Designing a 3.8-GeV/c muon-decay ring and experiment sensitive to electronvolt-scale sterile neutrinos

C.D. Tunnell

Lady Margaret Hall

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Hilary 2013

To my Dad

Acknowledgments

I could never have gotten this far without the help of numerous people. This work could not have been done without my supervisor John Cobb. His eternal patience and different world view have made me and my thesis better. Working with him has been enormously educational, though he never taught me to appreciate cricket. I still reminisce about our chats and the clarity he had when deconstructing problems. He is a true master of the *back of the envelope* calculation.

I'd also like to thank Alan Bross. His vast knowledge of the practical details of experimental physics taught me quite a bit about actually doing experiments. Also, my examiners Mark Thomson and Alfons Weber gave a wonderfuly critical reading of this thesis, which provided for a humbling viva.

I would also like to thank the mentors who helped me along the way: Bernard Gilroy, Josh Klein, Stan Seibert, Monica Dunford, Jim Pilcher, Victoria Blackmore, Linda Coney, and Alain Blondel. In addition to physics help, their tolerance of shenanigans requires special note (e.g., no pants day, air-hose safety, etc.). At each step of the way, these people cared for me and helped me develop as a physicist.

And most importantly my parents without whom, for more than biological reasons, I could not have done any of this.

Abstract

The liquid-scintillator neutrino-detector (LSND) and mini booster neutrino experiment (MiniBooNE) experiments claim to observe the oscillation $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$, which can only be explained by additional neutrinos and is a claim that must be further tested. This thesis proposes a new accelerator and experiment called neutrinos from stored muons (νSTORM) to refute or confirm the oscillation these claims by studying the CPT-equivalent channel $\nu_e \rightarrow \nu_\mu$. A 3.8-GeV/c muon decay ring is proposed with neutrino detectors placed 20 m and 2000 m from the decay ring. The detector technology would be a magnetized iron sampling calorimeter, where the magnetic field is induced by a superconducting transmission line. In a frequentist study, the sensitivity of this experiment after 5 years would be $>10\sigma$. The range of the thesis discussion starts with the proton front-end design and ends with neutrino parameter estimation. After describing the phenomenology of sterile neutrinos, the facility and detector performance work is presented. Finally, the systematics are explained before the sensitivity and parameter-estimation works are explained.

Contents

List of Figures

List of Tables

Acronyms

νSTORM neutrinos from stored muons. 4, 7, 10, 46, 58–60, 62, 64, 66–68, 71, 72, 76, 79–81, 89–92, 95, 119, 120, 122, 130, 138, 140, 142–145, 150

0νDBD neutrinoless double-β-decay. 24–26

AIDA advanced European infrastructure for detectors at accelerators. 142

APD avalanche photodiode. 86

ASIC application-specific integrated circuit. 86

ATLAS LHC experiment. 140

BCT beam-current transformers. 81, 144

BooNE Booster neutrino experiment. 17, 35, 37, 38, 40, 41, 45, 135, 157, 158

BPM beam-position monitor. 144

CalDet calibration detector. 140–142, 159

CC charged current. 16, 20, 22, 41, 50, 90, 92, 98, 101, 102, 106, 110, 111, 127, 129, 130, 135, 145, 150, 170

CKM Cabibbo–Kobayashi–Maskawa. 28, 169

CL confidence level. 37

CMB cosmic microwave background. 42

CSDA continually-slowing-down approximation. 52

DAB D0 assembly building. 82, 149, 151

DAG directed acyclic graph. 105

DAQ data acquisition. 155

DIS deep inelastic scattering. 38, 91, 96, 98, 101, 106, 127, 144, 145

ES elastic scattering. 41

FFAG fixed-field alternating gradient. 62, 76

GUT grand unified theory. 16, 29

ILL Institut Laue-Langevin. 42

LAMPF Los Alamos meson physics facility. 33

LAr liquid argon. 40, 130, 166, 170

LBNE long-baseline neutrino experiment. 82, 170

LBNO long-baseline neutrino observatory. 82, 170

LEP Large Electron–Positron Collider. 20

LOI letter of intent. 60, 77

LSND liquid scintillator neutrino detector. 17, 33, 90

MC Monte Carlo. 77, 90, 115, 120, 123, 163

MI main injector. 54, 66–68, 70, 81, 150

MIND magnetized iron neutrino detector. 91, 129, 142, 145, 151, 160, 166, 169, 170

MIP minimally-ionizing particle. 148

MS multiple scattering. 51

NC neutral current. 16, 41, 50, 92–94, 98, 106, 107, 127, 129, 130, 135, 170

NF neutrino factory. 62, 64, 67, 80, 92, 117, 119, 122, 124, 129, 138, 139, 144

PIP proton improvement plan. 68

PMNS Pontecorvo–Maki–Nakagawa–Sakata. 21, 28, 64, 169

PMT photomultiplier tube. 34, 35, 86

POT protons on target. 59, 66, 68, 71, 80, 124, 126, 130, 131, 135, 137, 157, 158

QE quasielastic. 96, 98, 106, 144, 145

QED quantum electrodynamics. 61, 62

RES resonant production. 96, 98, 144

S-matrix scattering matrix. 24

SiPM silicon photomultiplier. 86, 148

SNO Sudbury Neutrino Observatory. 16

STL superconducting transmission line. 48–50, 87, 89, 143

Super-K Super-Kamiokande. 17, 24, 26

TASD totally-active scintillator detector. 49

TPC time projection chamber. 40, 130, 166, 170

USD US dollar. 48

VLENF very-low-energy neutrino factory. 62

VLHC very-large hadron collider. 49, 50, 87

WLS wavelength shifting. 85, 86

Chapter 1

Introductory historical preamble: mass-splittings and their anomalies

No one ever lost money betting that a neutrino experiment was wrong!

D.H. Perkins

I lost 20 bucks on that #$%@ing Opera speed-of-light thing!

C.D. Tunnell

This section is a non-technical layman's introduction for historical context, and is written in a Scientific American style. It can be skipped by physicists.

Neutrinos—fundamental building blocks of nature—seemingly *disappear*. At least that was what experiments were seeing in the 1960s: a mere twenty years after the neutrino was discovered. Neutrinos are difficult to detect, and to neutrinos even something the size of the Earth is more transparent than the clearest glass. The excitement of being able to study neutrino interactions despite this difficulty was soon counteracted by utter confusion: why were less neutrinos interacting than expected (i.e., *deficits*)?

Whether neutrinos were produced in the sun or atmosphere, there were measured deficits, criticisms of the experiment, and skepticism in the theoretical models associated with neutrino production. Each of these deficits—solar and atmospheric—were later understood in terms of neutrino oscillations and serve as examples of the chaotic beauty of the scientific process. These mass splittings cause neutrinos to spontaneously—or rather, quantum mechanically—change their species, which in the jargon is called "oscillations." The history of neutrino oscillation physics is the history of these mass splittings.

The study of solar neutrinos eventually led to the realization that neutrinos oscillate. The fusion processes in the sun create electron neutrinos through (primarily) the proton–proton solar cycle: protons fuse into helium and other unstable elements. This solar neutrino flux was measured by Davis for the majority of his career, where he continually measured the "wrong" answer compared to theory. His first limit on solar neutrinos—around the time when Reines and Cowan discovered the neutrino—was actually at a reactor antineutrino experiment. He was trying to observe a chlorine transition allowed with neutrinos, but not antineutrinos, under the assumption that neutrinos were their own antiparticle. He observed no effect, but also was worried about solar neutrino backgrounds; therefore, he buried his detector a few meters underground to place a limit on the solar-neutrino flux. Upon submitting his publication, his reviewer said this was "like standing on a hill, waving your arms around, then saying the moon is more than a few meters away."

Davis went deeper underground and built bigger detectors. The experimental technique involved searching for the production of exotic radioactive isotopes—that only neutrinos could produce—within a large vat of freon. A chemistry technique was used to count the number of radioactive isotopes. His experiment, which should have been sensitive to the solar standard model neutrino flux, observed no significant solar neutrino flux. This raised doubts about his experimental technique—counting neutrinos is much harder than counting antineutrinos—and reenforced what physicists think about

chemists. He continuously improved his experiment until he eventually measured the flux at a third of its expected value. After 30 years, the Kamiokande experiment confirmed Davis's deficits. However, Kamiokande also demonstrate that these neutrinos were coming from the sun; the electrons emerging from the neutrino interaction point always pointed away from the sun.

It took 30 years before these deficits were confirmed by the Kamiokande experiment. By using the directionality information of the neutrino interaction events, Kamiokande showed that these events pointed back to the sun—i.e., came from the sun. However, Kamiokande measured a different fraction of missing neutrinos than Davis, further increasing the confusion.

50 years after Davis's original experiment, the Sudbury Neutrino Observatory (SNO) experiment established that Davis's deficit was caused by neutrino oscillations. The experiment used a different experimental technique, conceived by the late Herb Chen, that allowed for counting the total number of neutrinos at the same time as counting just electron neutrinos. The insight was that the neutral current (NC) interaction occurs for the three known neutrinos, whereas at these energies the charged current (CC) interaction only occurs for electron neutrinos. The discovery was that the neutrinos were still there, but were no longer electron neutrinos. In the sun, they had oscillated away, hence the deficits.

The other known mass splitting—the atmospheric—had a similar story: years of confusions. In the late 1970s and 1980s, various experiments searched for proton decay, which is a signal of various grand unified theories (GUTs), where the backgrounds are primarily atmospheric neutrinos. The atmospheric neutrino flux uncertainty was ≈20%; however, because the production mechanism ($\pi \to \mu + \bar{\nu}_\mu$ then $\mu \to e + \nu_\mu + \bar{\nu}_e$) was understood, the ratio of neutrino species was well known. The IMB and Kamiokande data showed that $\nu_\mu/\nu_e \approx 1$ rather than the expected value of 2.

Similar to the solar neutrino case, the models and experiments were independently

blamed. For instance, Frejus and NUSEX observed no effect. Later experiments like Soudan 2 and MACRO confirmed the initial experiments that indicated a deficit. It was not until Super-Kamiokande (Super-K) took data that the anomaly was determined to be neutrino oscillations, when Super-K (1998) demonstrated a difference between the upward and downward neutrino rates. All doubt about neutrino production being the cause of the anomaly was resolved when the MINOS experiment (2006)—using a detector similar to what will be discussed later—used neutrinos produced from an accelerator and confirmed the Super-K results.

Neutrino oscillations are firmly established by roughly a dozen experiments with different sources, detectors, and at different baselines and energies. Much is known about neutrinos, such as that they are—sadly—not massive enough to be dark matter. What is less firmly established is whether or not these two mass splittings are the entire story.

A current question in neutrino physics is whether there is another large mass splitting ($\approx$1 eV) that results in the existence of another yet-undiscovered "sterile" neutrino with no known interactions. Evidence from the liquid scintillator neutrino detector (LSND) experiment suggests this may be the case. Some experiments confirm it, yet other experiments refute the effect. A definitive experiment called Booster neutrino experiment (BooNE) was to be built to address this problem. Due to budget constraints it became MiniBooNE, and was unable to make definitive conclusions.

Regardless of the LSND anomaly, how does one probe the large mass splitting regime to determine if there are oscillations not accounted for in the standard paradigm? One should not ignore experimental anomalies. A definite experiment is required and proposed in the remainder of this thesis.

In Chap. 2, the theoretical and experimental status of neutrino oscillations is presented such that a nonspecialist particle physicist can understand. In Chap. 3, a new experiment is designed and proposed using simple arguments, which are strengthened

in the subsequent chapters. In Chap. 4, the details of the accelerator and detector are described. The detector performance is described in Chap. 5, where a detailed Monte Carlo study is developed. The sensitivity to sterile neutrinos is demonstrated in Chap. 6 followed by a discussion of systematics and backgrounds in Chap. 7. The precision is discussed in Chap. 8. Lastly, conclusions are included in Chap. 9.

Chapter 2

Phenomenology of neutrino oscillations

> There are two ways to slide easily through life; to believe everything or to doubt everything. Both ways save us from thinking.
>
> A. Korzybski

This chapter is intended for particle physicists who do not specialize in neutrinos, and can be skipped by neutrino physicists.

The standard neutrino framework has two main features. First, neutrinos have a relatively small mass and their mixing is different from that of quarks. Second, the mechanism for generating small neutrino masses is unknown. The main challenge to this model is the developing evidence for additional neutrinos without Standard Model interactions (i.e., *sterile neutrinos*).

In the first section, the phenomenology and experimental summary of three-neutrino oscillations are presented. After that, to accommodate sterile neutrinos, the three-neutrino formalism is extended and various anomalies are discussed.

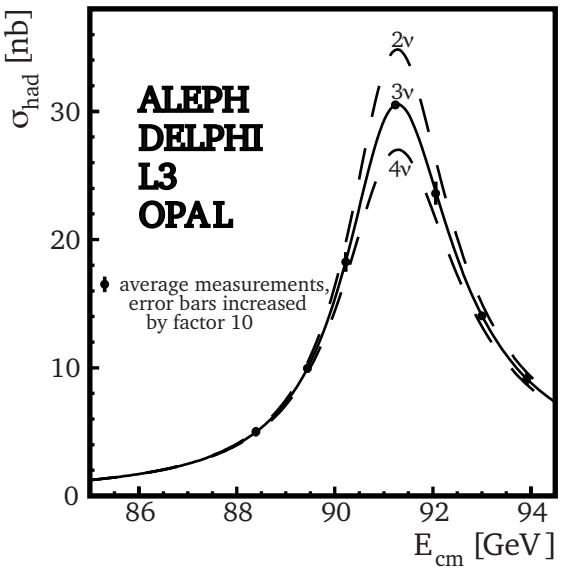

Figure 2.1: Measurements of the hadron-production cross section around the Z^0 resonance. Shown are the data and predicted cross sections for two, three, and four active neutrinos [1].

2.1 Three neutrinos

2.1.1 Neutrino flavors

The three known neutrinos interact via the electroweak interaction (i.e., are charged under the electroweak gauge); therefore, studying the Z^0 and $W^\pm$ bosons provides information about neutrinos. The invisible width of Z^0 decays was measured at the Large Electron–Positron Collider (LEP) and is consistent (Fig. 2.1) with three neutrinos—more specifically, three neutrinos with electroweak interactions and masses below 45 GeV $= m_Z/2$. An example of an active neutrino would be the one in the decay $W^+ \to \mu^+ \nu$. The neutrino in the final state is defined to be a muon neutrino ν_μ, and if the neutrino were forced to interact instantaneously (i.e., before "oscillation"), the neutrino would always produce a μ in CC interactions. Similar arguments can "label" neutrinos associated with the other charged leptons e and τ.

2.1.2 Mixing

Neutrino mixing occurs when the flavor-basis eigenstates are different from the mass-basis eigenstates (see the review in [2]). The decay $W^+ \to \mu^+ \nu$ results in a flavor-state neutrino, and may be written in the mass-basis as $W^+ \to l^+ \sum_i U_{li} \nu_i$ with amplitude U_{li}, where i identifies the mass eigenstate. These amplitudes form a matrix that may be written explicitly for three neutrinos as

$$U_{3\times3} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu1} & U_{\mu2} & U_{\mu3} \\ U_{\tau1} & U_{\tau2} & U_{\tau3} \end{pmatrix}, \tag{2.1}$$

where $U_{3\times3}$ is called the Pontecorvo–Maki–Nakagawa–Sakata (PMNS) matrix.

Neutrinos can mix if U is not diagonal. This matrix is frequently presented in terms of the product of rotation matrices: $U_{3\times3} = R_{23}(\theta_{23}, 0)R_{13}(\theta_{13}, \delta)R_{12}(\theta_{12}, 0)$, where R_{ij} represents a rotation between the ith and jth mass eigenstates and δ is a nonfactorizable Dirac CP-violating phase. The rotation order is just convention. More formally, the rotation is

$$[R_{ij}(\theta_{ij}, \delta_l)]_{pq} = \begin{cases} \cos\theta_{ij} & \text{if } p = q = i, j \\ 1 & \text{if } p = q \neq i, j \\ \sin\theta_{ij} e^{-i\delta_l} & \text{if } p = i; q = j \\ -\sin\theta_{ij} e^{i\delta_l} & \text{if } p = j; q = i \\ 0 & \text{otherwise.} \end{cases} \tag{2.2}$$

All the mixing angles in $U(\theta_{13}, \theta_{13}, \theta_{13}, \delta)$ are known to be nonzero at greater than 5σ (as will be discussed in Sec. 2.1.4); however, the phase δ is unknown and is ignored in this thesis unless otherwise stated. Experiments are being planned to measure this phase to complete our understanding of U.

2.1.3 Oscillations

As neutrinos propagate, mixing results in oscillations between eigenstates. In the previous example of $W^+ \to \mu^+ \nu_\mu$ decay, if the neutrino was not detected instantaneously, then it is possible for a ν_e to be detected at a later time due to the evolution of the ν_μ state in time and space. This evolution is well described by the Schrödinger view of quantum mechanics using plane waves (reviewed in [3]). Treatments using wave packets [4–8] and field theory [9] are indistinguishable experimentally at present and therefore not discussed. An illustrative treatment is presented, but the previously cited works are more complete.

Neutrinos produced by CC interactions are initially in pure flavor states, which later evolve. At $t = 0$, the initial state vector $|\nu_\alpha(t=0)\rangle$ is

$$|\nu_\alpha(t=0)\rangle = \sum_k U_{\alpha k} |\nu_k(t=0)\rangle , \tag{2.3}$$

where $\alpha \in \{e, \mu, \tau\}$ and k is summed over mass eigenstates. There must be at least three mass eigenstates since three active light neutrino flavors exist, and U must be Unitary.

The Schrödinger equation describes the evolution of $|\nu_\alpha(t=0)\rangle$ in time,

$$\frac{d}{dt} |\nu_k(t)\rangle = E_k |\nu_k(t)\rangle , \tag{2.4}$$

where E_k are the energy eigenvalues associated with the Hamiltonian in vacuum (i.e., $H|\nu_k\rangle = E_k|\nu_k\rangle$). The plane-wave solution results in

$$|\nu_k(t)\rangle = e^{-iE_k t} |\nu_k(0)\rangle . \tag{2.5}$$

Using Eq. (2.3), it can be seen that

$$|\nu_\alpha(t)\rangle = \sum_k U_{\alpha k} e^{-iE_k t} |\nu_k(0)\rangle , \tag{2.6}$$

where the flavor state is represented by the evolution of mass eigenstates. However, the mass eigenstates on the right can be represented as flavor states by inverting Eq. (2.3):

$$|\nu_k\rangle = \sum_\alpha U_{\alpha k} |\nu_\alpha\rangle . \tag{2.7}$$

Assuming that U is real, which is approximately true for our purposes, then substituting this equation into the right-hand side of Eq. (2.6) yields

$$|\nu_\alpha(t)\rangle = \sum_k \sum_\beta U_{\alpha k} e^{-iE_k t} U_{\beta k} |\nu_\beta\rangle . \tag{2.8}$$

Therefore, an initial flavor state $|\nu_\alpha(0)\rangle$, which is just a superposition of mass states $|\nu_k(0)\rangle$, becomes a superposition of different flavor states as it travels. This requires that U not be diagonal.

Quantum amplitudes are unobservable, and detectable neutrinos are relativistic. Neutrino experiments measure the probability of a neutrino oscillating between two flavor states $P_{\alpha\beta}$, where the probability is related to the amplitude by

$$\begin{aligned} P_{\alpha\beta} &= |A_{\alpha\beta}|^2 = |\left\langle \nu_\alpha \middle| \nu_\beta(t) \right\rangle|^2 && (2.9)\\ &= |\sum_k U_{\alpha k} e^{-iE_k t} U_{\beta k}|^2 && (2.10)\\ &= \sum_{k,j} U_{\alpha k} U_{\beta k} U_{\alpha j} U_{\beta j} e^{-i(E_k - E_j)t} && (2.11)\\ &= \sum_{k,j} U_{\alpha k} U_{\beta k} U_{\alpha j} U_{\beta j} e^{\frac{-i\Delta m^2_{ij}}{2E} L}, && (2.12) \end{aligned}$$

where $\Delta m^2_{ij} = m^2_j - m^2_i$. In the last line, a relativistic approximation is used ($t \simeq L$); therefore, the dispersion relation is expanded such that

$$E_k = \sqrt{p^2 + m^2} \simeq p_k + m_k^2/2E \tag{2.13}$$

because $m \ll p$ for any detectable neutrino.

The prediction in Eq. (2.12) is that neutrinos can turn into other neutrinos. Or rather, when the quantum wave function $|\nu_\alpha(t)\rangle$ in Eq. (2.8) collapses, Eq. (2.12) describes the probability of transforming into another neutrino flavor. If U is not diagonal, then this probability of Eq. (2.12) is nonzero when $\alpha \neq \beta$. The implication is the violation of lepton number; if oscillations is viewed—very roughly—as a black box, electrons can turn into muons. As mentioned before, the derivation can be performed in numerous ways and even more directly (e.g., scattering matrix (S-matrix) formalism); regardless, neutrinos oscillate among flavors.

When mixing between three neutrinos is considered, the sums in Eq. (2.12) can be tedious. Simplifications can be made, but do not elucidate much. Computer programs are typically used to compute 3×3 oscillation probabilities, mainly to help bookkeeping. In certain circumstances, it is possible to make assumptions that result in the oscillation probabilities simplifying. For example, if there are effectively two neutrinos due to the experiment not being sensitive to other mass splittings—as is the case later in Eq. (2.16)—then the appearance probability can be written, if $\alpha \neq \beta$, as the more familiar expression

$$P_{\alpha\beta} \quad = \quad \sin(2\theta_{\alpha\beta}) \sin^2\left(\frac{\Delta m^2 L}{4E}\right). \tag{2.14}$$

2.1.4 Experimental status of three-neutrino mixing

The Super-K experiment [10–12] provided concrete evidence that neutrinos oscillate and thus have mass, so for the last 15 years experiments have been studying neutrino mass. Three methods are used to study neutrino masses and mixing: end-point measurements of β-decay spectra, neutrinoless double-β-decay (0νDBD) experiments, and neutrino oscillation experiments. At present, the absolute scale of the heaviest mass eigenstate is constrained between 1 eV from end-point measurements and 0.05 eV $= \sqrt{\Delta m^2_{31}}$ from oscillation physics. Double-β-decay experiments have yet to produce conclusive results.

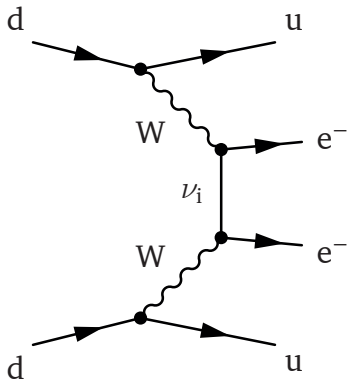

Figure 2.2: 0νDBD Feynman diagram. Two decays $n \rightarrow p^+ \bar{\nu}_e e^-$, but with "annihilation" of the neutrinos.

In β decay ($n \rightarrow p^+ \bar{\nu}_e e^-$), measurements of the electron energy lead to the best understanding of the absolute mass scale of neutrinos. These experiments compare the maximum electron energy E_e with the free energy Q of the decay process to determine the neutrino energy $E_\nu = Q - E_e$. This type of experiment also has the longest history: the shape of the β spectrum led Pauli to postulate that neutrinos had a mass smaller than the electron. As experimenters tried to measure the neutrino mass and failed, the upper limits on the allowed value became more stringent. Eventually, after parity violation was discovered in β decays, the neutrino was assumed to be massless in the two-component theory by Landau [13], Lee and Yang [14], and Salam [15]. Current limits on the squared electron neutrino mass include 2.05 eV^2 from β-decay results and 5.7 eV^2 from the Supernova 1987A [2]. By extending the work from the Mainz and Troitsk experiments, KATRIN experiment should be sensitive above 0.2 eV^2 [16]. However, at present the exact mass of the neutrino is unknown.

Neutrinos having mass may allow for interesting processes such as 0νDBD (Fig. 2.2), where the neutrinos in double-β decay seemingly annihilate. Depending on which

equation of motion (Dirac or Majorana) neutrinos obey, neutrinos may—as is said colloquially—be their own antiparticles. The experimental signature of 0νDBD is two electrons whose kinetic energy is equal to the nuclei binding energy difference. EXO, Kamland-Zen, SNO+, SuperNEMO, and others are searching for this process. If this process is observed, it could help explain the baryon asymmetry in the Universe; this process motivates a heavy neutrino with above the energies that can be explored terrestrially ($\gg$ TeV) and that could have decayed in the early Universe and violated CP. Majorana neutrino experiments provide an important input to understanding of neutrinos and the Universe.

Neutrino oscillation experiments provide information on the differences between neutrino mass eigenstates and can be sensitive to CP violation in the lepton sector. Precision measurements of oscillation parameters may reveal structure that advance the theoretical models of neutrino mass generation. Nevertheless, a dominant theory at present is called *anarchy* [17], which suggests that the mixing matrix is random and nothing can be learned by studying it further. The limited experimental precision of the current oscillation parameters prevents finding any structure within the mixing matrix.

Super-K (1998) and SNO (2002) established neutrino oscillations with atmospheric ($|\Delta m^2_{32}| = 2.7 \times 10^{-3}$ eV^2) and solar neutrinos ($\Delta m^2_{21} = 7 \times 10^{-5}$ eV^2), respectively. Successive experiments (Fig. 2.3) have confirmed their results using neutrino sources ranging from cosmic-ray interactions with the atmosphere, the Sun, supernovae, reactors, radioactive sources, and accelerator beams. Super-K and MINOS primarily measure $|\Delta m^2_{23}|$ and $\sin^2\theta_{23}$, whereas SNO and Kamland constrain $\sin^2\theta_{12}$ and Δm^2_{21}, respectively. Recently, the Daya Bay reactor neutrino experiment collaboration discovered that $\sin^2\theta_{13} \neq 0$. This result has been confirmed by the T2K and RENO experiments—T2K published a 2.5σ result before Daya Bay, which would have had better statistics if not for the 2011 Tōhoku earthquake (and tsunami) that damaged their accelerator facility J-PARC. The current understanding of 3×3 mixing is shown in Table 2.1.

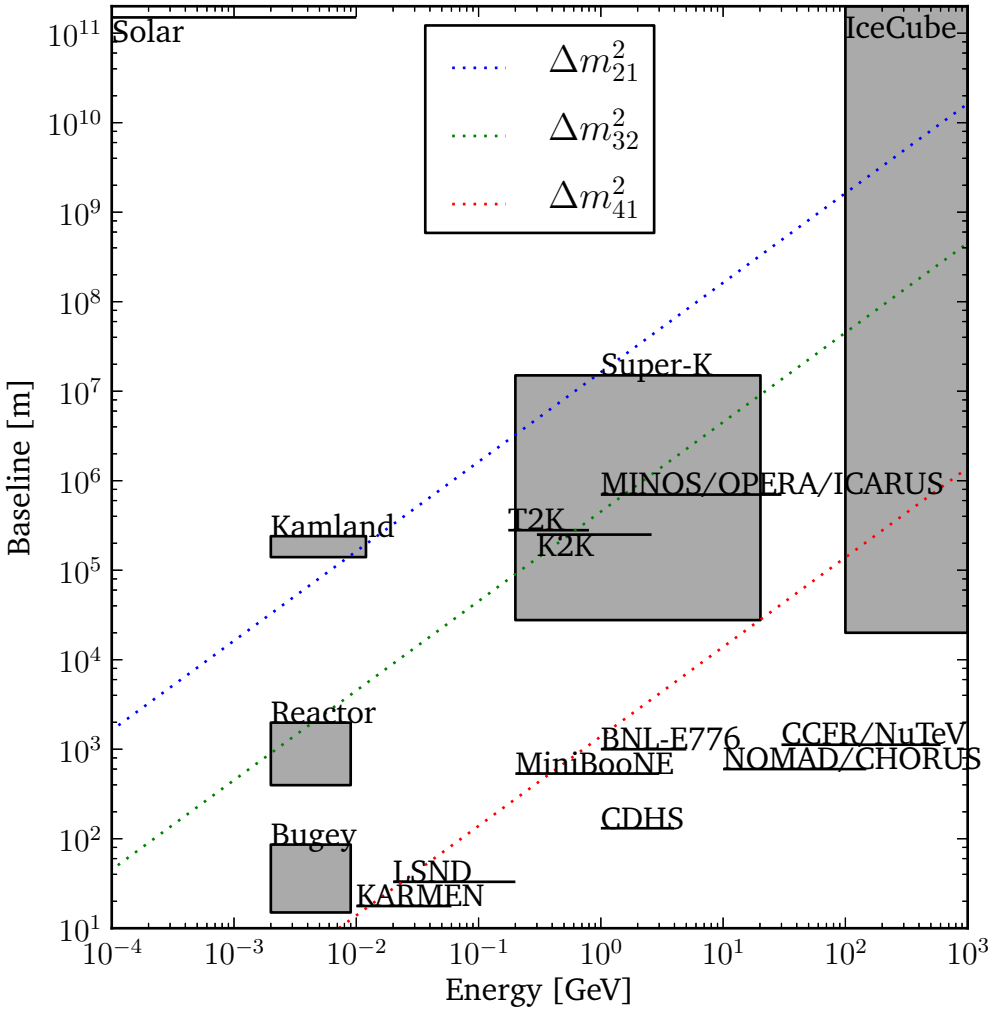

Figure 2.3: The baseline vs energy dependence of various neutrino experiments.

Table 2.1: Neutrino oscillation parameters. For Δm^2_{31}, $\sin^2\theta_{23}$, $\sin^2\theta_{13}$, and δ the upper (lower) row corresponds to normal (inverted) neutrino mass hierarchy [19].

Parameter	3σ range
$\Delta m^2_{21}\,[10^{-5}\,\mathrm{eV}^2]$	7.12–8.20
$\lvert\Delta m^2_{31}\rvert\,[10^{-3}\,\mathrm{eV}^2]$	2.31–2.74 2.21–2.64
$\sin^2\theta_{12}$	0.27–0.37
$\sin^2\theta_{23}$	0.36–0.68 0.37–0.67
$\sin^2\theta_{13}$	0.017–0.033
δ	0–2 π

At present, the mass hierarchy (i.e., $m_3 > m_1$ or $m_3 < m_1$) and value of the CP phase δ are unknown. Numerous proposals are being considered that can make hierarchy measurements using matter effects—the earth contains electrons but not positrons, and this discriminates the mass eigenstates by changing the effective neutrino mass. This CP phase is interesting due to its theoretical relations to leptogenesis; however, $\delta \neq 0$ does not imply leptogenesis, and vice versa. The ultimate goal for many neutrino physicists is to constrain the 3×3 PMNS matrix U to levels comparable to those of the Cabibbo–Kobayashi–Maskawa (CKM) matrix. The main motivation is to understand the mechanism that gives neutrinos mass. This mechanism—hopefully—arises from new physics at energies unreachable by colliders [18].

2.1.5 Conclusion

Within the standard theory of neutrino mixing and oscillations, only three neutrinos are assumed and most data is understood in this framework. Over a dozen experiments have been performed, and future experiments will improve upon this understanding.

2.2 Sterile neutrinos

> Before I came here I was confused about this subject. Having listened to your lecture I am still confused. But on a higher level.
>
> Enrico Fermi

In this section, the motivation for the study of sterile neutrinos is discussed, as well as the various astrophysical and terrestrial evidence for them. The status of global sterile fits is presented and future sterile neutrino experiments reviewed. A comprehensive literature review focusing on sterile neutrinos is available in [20].

Sterile neutrino results indicating the existence of an electronvolt-scale neutrino are contentious. An electronvolt-scale neutrino would affect solar, atmospheric, supernova, β-decay, and cosmological (nucleosynthesis, extra radiation of the universe, large scale structure formation) experiments. However, only a subset of these experiments claim an effect. No model fits all experimental results well. At present, it is not known whether or not sterile neutrinos exist.

2.2.1 Theory of sterile neutrinos

There is no strong theoretical motivation for sterile neutrinos with masses observable by neutrino oscillation experiments. There are nevertheless theoretical ideas. For example, a low-energy seesaw with a light Majorana neutrino on the order of 1 eV explains all sterile neutrino data [21]. The problem is that this is at the expense of the reason for the seesaw mechanism: explaining why neutrinos are light. Most GUT models either predict sterile neutrinos or easily accommodated them, but give no further guidance as to their mass. The difficulty is that there is no *a priori* mass scale but are then unobservable at neutrino oscillation experiments. Nevertheless, if there is a light sterile neutrino, then the oscillation phenomenology that follows applies.

2.2.2 Sterile neutrino oscillations

The formalism that was used for 3×3 neutrino oscillations can easily be extended to accommodate more neutrinos. In the previous formalism in Eq. (2.3) on p.p. 22, additional mass eigenstates are already accommodated by extending the sum over the k mass eigenstates. The 3×3 unitarity mixing matrix has a convenient representation in terms of rotations and the same is possible for $N \times N$ unitarity matrices. For 3×3 mixing matrices, it was shown that $U_{3\times3} = R_{23}R_{13}R_{12}$. Using a similar convention for ordering the rotations, higher dimensional matrices can be defined: $U_{4\times4} = R_{34}R_{24}R_{14}U_{3\times3}$, $U_{5\times5} = R_{45}R_{35}R_{25}R_{15}U_{4\times4}$, and by induction any $U_{N\times N}$. Rotations are defined such that

$$R_{ij} = \begin{pmatrix} 1 & \dots & 0 & \dots & 0 & \dots & 0 \\ \vdots & & \vdots & & \vdots & & \vdots \\ 0 & \dots & \cos\theta_{ij} & \dots & \sin\theta_{ij} & \dots & 0 \\ \vdots & & \vdots & & \vdots & & \vdots \\ 0 & \dots & -\sin\theta_{ij} & \dots & \cos\theta_{ij} & \dots & 0 \\ \vdots & & \vdots & & \vdots & & \vdots \\ 0 & \dots & 0 & \dots & 0 & \dots & 1 \end{pmatrix}. \tag{2.15}$$

Most data analyses and studies related to sterile neutrinos are sensitive to only one additional neutrino; therefore, only one additional neutrino is assumed here. These so-called (3+1) models have an additional mass eigenstate m_4 that is much heavier than the other mass eigenstates ($m_4 \gg m_3, m_2, m_1$), resulting in the ability to only probe one mass splitting (i.e., single-hierarchy dominance) and the inability to observe CP violation at short baselines (see, e.g., p. 273 of [3]). Two sterile neutrinos are required [i.e. (3+2) model] to observe CP violation at short baselines.

The evidence, that will be discussed shortly, for the existence of sterile neutrinos favors a sterile neutrino at 1 eV. Therefore, $\Delta m^2_{41} \simeq 1$ eV^2. The L/E where the effects of this additional neutrino become evidence is around $L/E \sim$ km/GeV; short baselines

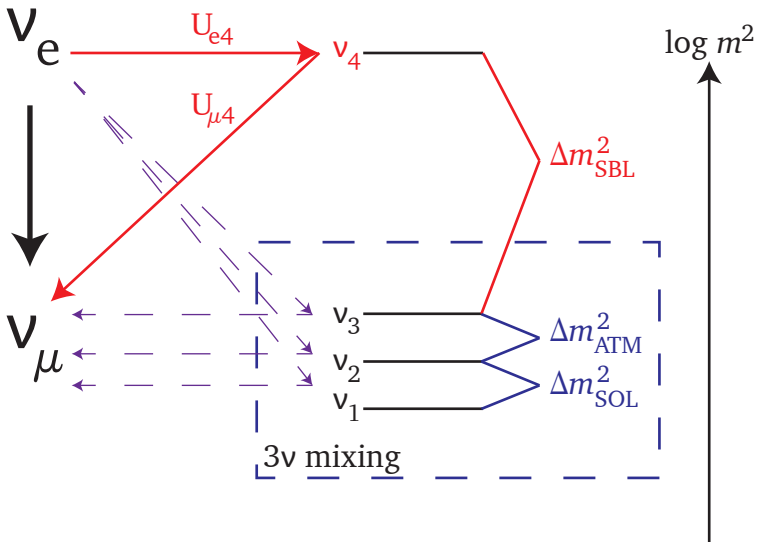

Figure 2.4: Pictorial motivation for $\nu_e \to \nu_\mu$ oscillations.

relative to most other neutrino experiments must be used to probe mixing elements relevant to an electron-volt scale neutrino. In the *short baseline approximation* (i.e., $\Delta m^2_{41} \gg \Delta m^2_{31}$), the submatrix $U_{3\times3}$ can be approximated as the unit matrix since it does not influence short-baseline oscillations. Therefore, $U_{4\times4}$ becomes

$$U_{4\times4} = \begin{pmatrix} \cos\theta_{14} & 0 & 0 & \sin\theta_{14} \\ -\sin\theta_{14}\sin\theta_{24} & \cos\theta_{24} & 0 & \sin\theta_{24}\cos\theta_{14} \\ -\sin\theta_{14}\sin\theta_{34}\cos\theta_{24} & -\sin\theta_{24}\sin\theta_{34} & \cos\theta_{34} & \sin\theta_{34}\cos\theta_{14}\cos\theta_{24} \\ -\sin\theta_{14}\cos\theta_{24}\cos\theta_{34} & -\sin\theta_{24}\cos\theta_{34} & -\sin\theta_{34} & \cos\theta_{14}\cos\theta_{24}\cos\theta_{34} \end{pmatrix},$$

where it is useful to note that $U_{e4} = \sin(\theta_{14})$ and $U_{\mu4} = \sin(\theta_{24})\cos(\theta_{14})$.

Oscillation probabilities can be computed using Eq. (2.12). All sterile neutrino data is understood in terms of three types of transitions (assuming CP is conserved): $\nu_e \to \nu_e$, $\nu_\mu \to \nu_\mu$, and $\nu_e \leftrightarrow \nu_\mu$. Interestingly, the addition of a massive sterile neutrino flavor state can affect the oscillations between $\nu_e \to \nu_\mu$: the ν_e oscillates into the sterile ν_s state, then into the ν_μ state (Fig. 2.4). Explicitly, in a (3+1) framework, where CP must

be conserved at short baselines, the probabilities are

$$
\begin{aligned}
P_{\nu_e \to \nu_e} &= 1 - 4\left(1 - |U_{e4}|^2\right)|U_{e4}|^2 \sin^2\left(\frac{\Delta m^2_{41} L}{4E}\right) && (2.16)\\
&= 1 - \sin^2(2\theta_{ee}) \sin^2\left(\frac{\Delta m^2_{41} L}{4E}\right), && (2.17)\\
P_{\nu_\mu \to \nu_\mu} &= 1 - 4\left(1 - |U_{\mu 4}|^2\right)|U_{\mu 4}|^2 \sin^2\left(\frac{\Delta m^2_{41} L}{4E}\right) && (2.18)\\
&= 1 - \sin^2(2\theta_{\mu\mu}) \sin^2\left(\frac{\Delta m^2_{41} L}{4E}\right), && (2.19)\\
P_{\nu_e \leftrightarrow \nu_\mu} &= 4|U_{e4}|^2 |U_{\mu 4}|^2 \sin^2\left(\frac{\Delta m^2_{41} L}{4E}\right) && (2.20)\\
&= \sin^2(2\theta_{e\mu}) \sin^2\left(\frac{\Delta m^2_{41} L}{4E}\right), && (2.21)
\end{aligned}
$$

in the short baseline limit ($\Delta m^2_{41} \frac{L}{E} \sim 1$ and $\Delta m^2_{31} \frac{L}{E} \ll 1$), where the definitions for $\sin^2(2\theta_{ee})$, $\sin^2(2\theta_{\mu\mu})$, and $\sin^2(2\theta_{e\mu})$ have been introduced.

A few observations can be made that will help understand the relations between the above oscillation probabilities. All three oscillation channels are coupled since each depends on the mass splitting Δm^2_{41}, although the channels differ in their dependence on U. Note that the disappearance channels depend only upon one matrix element $P_{\nu_e \to \nu_e} \sim |U_{e4}|^2$ and $P_{\nu_\mu \to \nu_\mu} \sim |U_{\mu 4}|^2$, whereas the appearance channel $P_{\nu_e \leftrightarrow \nu_\mu}$ depends on the product $|U_{e4}|^2 |U_{\mu 4}|^2$. Therefore, any nonzero appearance signal implies two nonzero disappearance signals:

$$
P_{\nu_e \leftrightarrow \nu_\mu} \neq 0 \iff \left(P_{\nu_\mu \to \nu_\mu} \neq 0 \text{ and } P_{\nu_e \to \nu_e} \neq 0\right). \qquad (2.22)
$$

By unitarity, every matrix element is required by be less than 1, and both $|U_{e4}|^2$ and $|U_{\mu 4}|^2$ must be small due to numerous experimental results. A comparison can be made to Cabibbo suppression. Disappearance deficits are *singly suppressed* by a small number $|U_{l4}|^2$ and appearance measurements are *doubly suppressed*; hence small appearance

probabilities should result in larger disappearance deficits.

This simple relation [Eq. (2.22)]—as explained later in greater detail—results in tension in global sterile neutrino fits. Deficits are seen in electron disappearance and appearance measurements, but not in muon disappearance. Either new physics has been discovered, or some experiments are wrong.

2.2.3 Evidence for sterile neutrinos

The strongest evidence for sterile neutrinos comes from the LSND experiment [22], which was an oscillation experiment that started taking data in 1993. It used neutrinos produced from a proton target and beam stop at the Los Alamos meson physics facility (LAMPF). The experiment was sensitive to neutrinos at ≈1 eV^2 (i.e., $L/E \sim$ m/MeV), with ≈60 MeV neutrinos at a baseline of 30 m. The intention was to measure $\bar{\nu}_\mu \to \bar{\nu}_e$ oscillations by the following method:

1. Proton–nucleon interactions produce π^+ and π^- in a target.
2. Both pion species are stopped in the beam stop.
3. π^- are captured in a nucleus because they have the same charge as the electron. π^+ are not captured.
4. The decay $\pi^+ \to \mu^+ \nu_\mu$ occurs.
5. The μ^+ stops as well and decays $\mu^+ \to \bar{\nu}_\mu \nu_e e^+$, where all neutrinos are at energies below the muon rest mass of $m_\mu = 106$ MeV.
6. Identify $\bar{\nu}_e$.

Three neutrino flavors are produced: ν_μ, $\bar{\nu}_\mu$, and ν_e. The proton pulse is 0.25 μs thus with adequate detector timing, the ν_μ from pion decay can be separated from the neutrinos from muon decay because it will be produced earlier. Therefore, if any $\bar{\nu}_e$ are

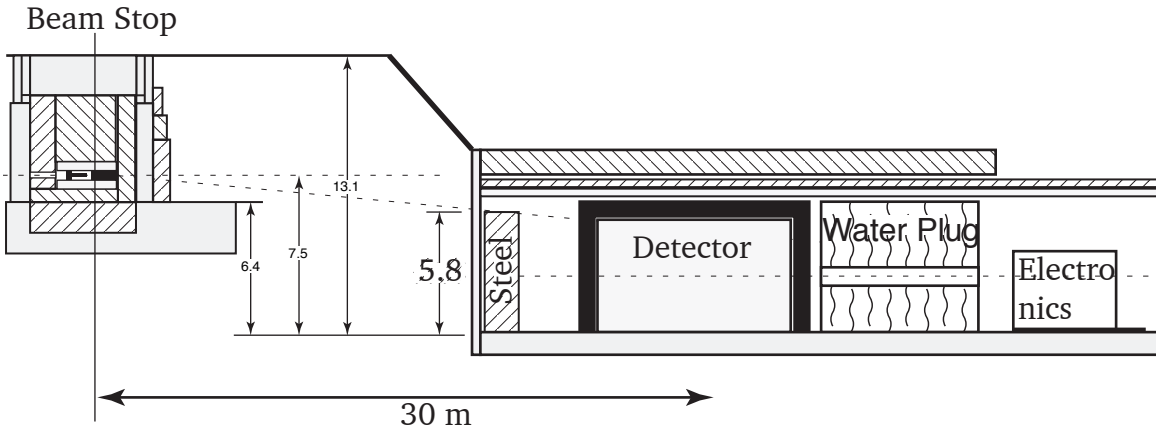

Figure 2.5: LSND experimental setup. Slightly clarified figure from [22].

detected, they must come from $\bar{\nu}_\mu \to \bar{\nu}_e$ transitions because early reactor experiments have shown that $\nu_e \to \bar{\nu}_e$ transitions are not possible [23]. Therefore, the detector was required to have good discrimination between ν_e from muon decay and $\bar{\nu}_e$.

The detection of $\bar{\nu}_e$ at these energies is straightforward experimentally—unlike ν_e and $\overset{(-)}{\nu}_\mu$—using the inverse-β-decay reaction $\bar{\nu}_e + p \to n + e^+$. Detecting this reaction relies upon the double coincidence of having a positron that annihilates quickly followed by the delayed capture of the neutron. Both the annihilation and neutron capture produce γs and release >1.0 MeV and 2.2 MeV of energy, respectively.

The detector was located 30 m from the beam stop and was composed of 167 t of mineral oil doped with scintillant. Any energy deposited resulted in scintillation light that was detected by 1220 8" photomultiplier tubes (PMTs) covering 25% of the surface area of the tank. To reduce external backgrounds, the detector was shielded by steel. A veto was used to reject cosmic ray muons, which traversed the detector at 4 Hz. These muons were also used to calibrate the e^+ and e^- response.

An analysis based on timing was performed to determine the $\bar{\nu}_e$ rates. LSND observed an excess of 87.9 ± 22.4 (stat.) $\pm$ 6.0 (sys.) $\bar{\nu}_e$ interactions above a background of 30 events, providing evidence for the transition $\bar{\nu}_\mu \to \bar{\nu}_e$ at 3.8σ.

Table 2.2: Parameters of the LSND and KARMEN experiments, a shortened version of [20].

Property	LSND	KARMEN
Proton energy (MeV)	798	800
Protons on target	1.8×10^{23}	5.8×10^{22}
Total mass (t)	167	56
Neutrino distance (m)	30	17.7
Events for 100% $\bar{\nu}_\mu \to \bar{\nu}_e$ transmutation	33,300	14,000

However, results from the KARMEN experiment do not agree with the LSND excess, which had a similar experimental setup and detector (Table 2.2). In a similar analysis to LSND, KARMEN observed no excess: 15 $\bar{\nu}_e$ events were observed with an expected background of 15.8. The experimental setups had slightly different sensitivities, due to their different baselines. A joint fit of KARMEN and LSND data shows that they may be compatible for mass splittings of 7 eV^2 or below 1 eV^2 [24].

The situation is further confused by results from the MiniBooNE—formerly BooNE—experiment, which both refuted [25, 26] and confirmed [27] the LSND signal. The experimental goal was to explore the LSND anomaly at higher energies while maintaining the same L/E. The experimental setup was similar (i.e., mineral oil, scintillator and PMTs); however, the beam was different. Instead of using a decay-at-rest beam, the experimenters decided to use a conventional pion beam $\pi^\pm \to \mu^\pm + \overset{(-)}{\nu}_\mu$.

However, external considerations resulted in an unanticipated loss of physics sensitivity for BooNE due to the absence of a near detector (how BooNE became MiniBooNE). Initially, the appearance [26] and disappearance [25] data excluded the $(3+1)$ model of the LSND anomaly; however, an unexpected electronlike low-energy excess was observed. Numerous fits were performed including or excluding the low-energy excess, with many fits suggesting a second sterile neutrino. (See [28] and references therein.) Brief tension in the sterile neutrino fits between neutrino and antineutrino data resolved itself with more statistics. At present, neutrino and antineutrino data from MiniBooNE favor a sterile neutrino at 3.8σ [27].

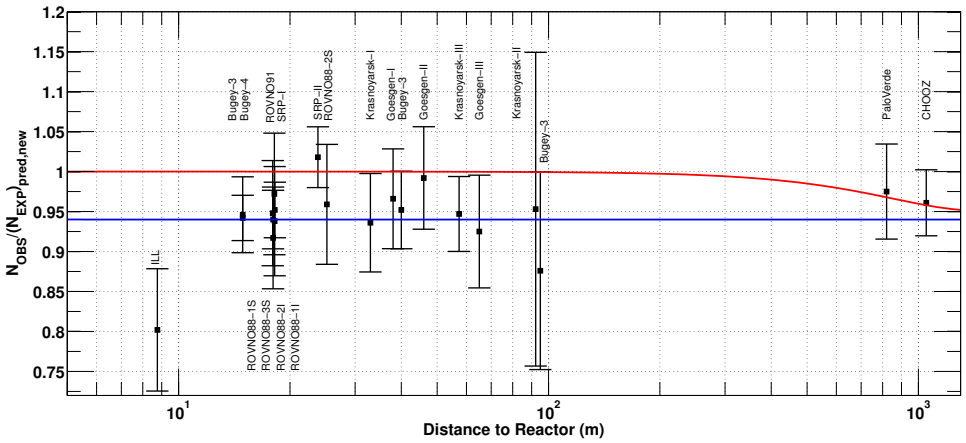

Figure 2.6: The short-baseline reactor antineutrino anomaly using the newly computed fluxes. The red line shows a possible three-active-neutrino mixing solution, with $\sin^2(2\theta_{13}) = 0.06$. The blue line displays a solution including a new neutrino mass state [$|\Delta m^2_{41}| \gg 1$ eV^2 and $\sin^2(2\theta_{ee}) = 0.12$] but does not include 3×3 oscillations . (Text and figure partially reproduced from [29].)

The appearance anomaly observed by LSND has implications for other experiments [Eq. (2.22)], some of which may have been observed. For example, an electronvolt-scale sterile neutrino requires electron neutrino disappearance at short baselines, and reactor neutrino experiments observe this effect.

Until 2011, experiments at distances of <100 m from the reactor core—ILL-Grenoble, Goesgen, Rovno, Krasnoyarsk, Savannah River, and Bugey—agreed with the predicted reactor antineutrino flux with a measured/expected ratio of 0.976 ± 0.024. However, as part of the preparations for the reactor neutrino experiment Double Chooz, these fluxes were recomputed [29]. Previously, the model of the neutrino flux was phenomenological and based on 30 effective branches, whereas the new analysis took into account the detailed knowledge of the decays of thousands of fission products. Both the new and old models are constrained by the flux measurements at ILL-Grenoble. Over the past 15 years, the measured neutron lifetime has evolved from 926 ± 11 s [30] to the current PDG average value of 880.1 ± 1.1 s, resulting in an increase of the predicted neutrino cross section. The result was a larger deficit of 5.7%, where the measured/expected

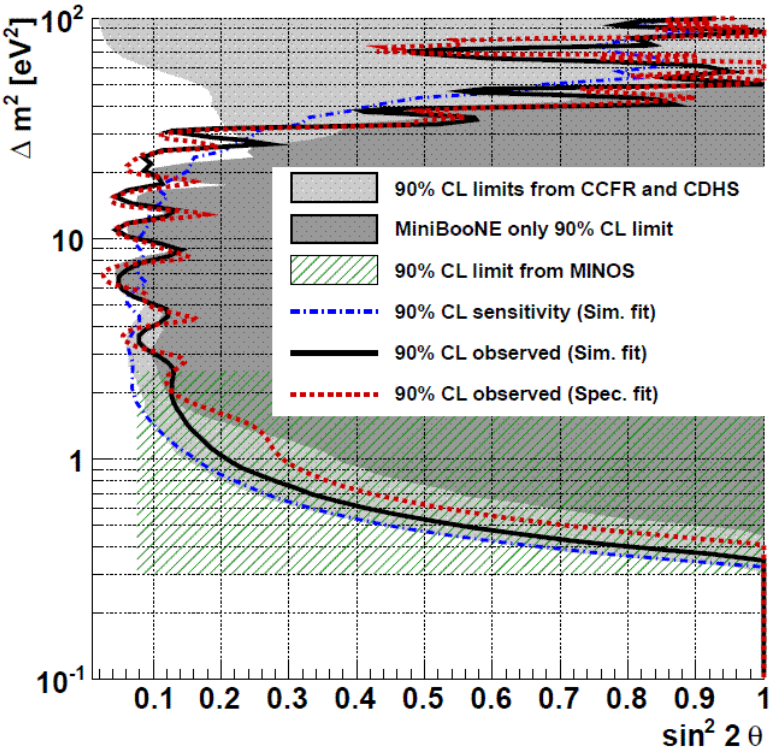

Figure 2.7: 90% confidence level (CL) sensitivity (dot-dash curve) and 90% CL limit (solid black curve) from simultaneous MiniBooNE/SciBooNE fit, and 90% CL limit from the spectrum fit method (red dashed curve). Previous limits from CCFR, CDHSW, MINOS, and MiniBooNE are also shown. (Text and figure reproduced from [20].)

ratio was 0.943 ± 0.023 (Fig. 2.6). This calculation was confirmed by Huber [31]. One explanation for this *reactor neutrino anomaly* is the short-baseline disappearance of neutrinos; however, the evidence is unconvincing on its own.

Similar anomalies are observed in radioactive decay experiments. For example, in the GALLEX and SAGE solar neutrino experiments, ^{51}Cr and ^{37}Ar radioactive sources were used as calibrations. For all the combinations of source-type and experiment, deficits were observed, although the significance arises only when combining data from the experiments. Assuming Gaussian errors, a 2.8σ effect (motivated in [32]) led to the *gallium anomaly*, which may be explained by the existence of sterile neutrinos.

However, not all disappearance results indicate sterile neutrinos: the CDHSW [33], CCFR [34], and MINOS [35] experiments do not observe a corroborating muon-neutrino disappearance signal at large L/E [required by Eq. (2.22)]. At large L/E, the oscillation

probability should be in the *averaging regime* and result in additional deficits beyond θ_{23} on the order of a few percent. The CDHSW experiment had a peak neutrino energy of 3 GeV—it used CERN's PS accelerator rather than the SPS used for deep inelastic scattering (DIS) studies—and a baseline of about 100 m; however, no oscillations were observed. The CCFR collaboration performed experiments using a ν_μ beam from protons with energies in the range 40–230 GeV and a $\bar{\nu}_\mu$ beam with a peak proton energy of 165 GeV. They performed oscillation studies where the combined neutrino and antineutrino analysis had a χ^2 for the null hypothesis of 15.7 for 18 degrees of freedom. These experiment's results are not compatible with any number of sterile neutrinos (Fig. 2.7).

2.2.4 Global fit

The data from the experiments discussed above can be understood within a single parameter-estimation fit, taking advantage of the wide range of oscillation channels, baselines, and energies. Different scenarios assuming a different number of sterile neutrinos can be considered. However, fits assuming either a $(3+1)$ [20], $(3+2)$ [36], or $(3+3)$ [37] model result in fits of comparable quality.

Fits have also been performed using different subsets of experiments, which is useful for assessing data consistencies within the same channel. The data was divided into three channels: the two disappearance channels $\overset{(-)}{\nu}_e \to \overset{(-)}{\nu}_e$ and $\overset{(-)}{\nu}_\mu \to \overset{(-)}{\nu}_\mu$ and the appearance channel $\overset{(-)}{\nu}_\mu \to \overset{(-)}{\nu}_e$. Combining the electron appearance and disappearance data always results in tension with muon disappearance data. In [20], a global $(3+1)$ fit was performed and described—paraphrasing their paper for this thesis—including:

- The short-baseline $\overset{(-)}{\nu}_\mu \to \overset{(-)}{\nu}_e$ data of the LSND, KARMEN, NOMAD, and MiniBooNE (neutrino and antineutrino) experiments.

- The short-baseline $\bar{\nu}_e$ disappearance data of the Bugey-3, Bugey-4, ROVNO91, Gosgen, ILL, and Krasnoyarsk reactor antineutrino experiments and the KamLAND

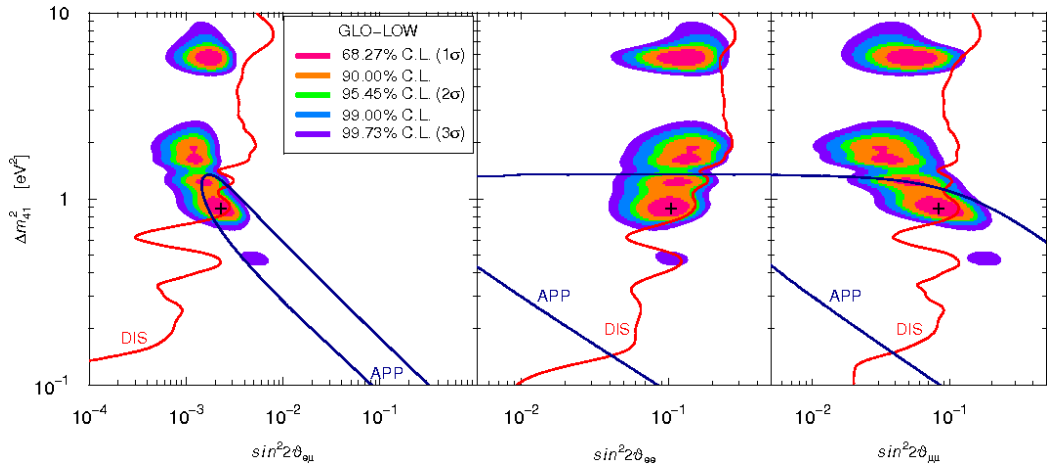

Figure 2.8: Allowed regions in the $\sin^2 2\theta_{e\mu}$–Δm^2_{41}, $\sin^2 2\theta_{ee}$–Δm^2_{41}, and $\sin^2 2\theta_{\mu\mu}$–Δm^2_{41} planes obtained by global analyses of short-baseline neutrino oscillation data. The best-fit points are indicated by crosses. The thick solid-blue lines with the label APP show the 3σ allowed regions obtained from the analysis of $\overset{(-)}{\nu}_\mu \to \overset{(-)}{\nu}_e$ appearance data. The thick solid-red lines with the label DIS show the 3σ allowed regions obtained from the analysis of disappearance data. (Text and figure reproduced from [20].)

bound on $|U_{e4}|^2$.

- The short-baseline ν_μ disappearance data of the CDHSW experiment, the constraints on $|U_{\mu 4}|^2$ from the analysis of the data of atmospheric neutrino oscillation experiments, and the information on $|U_{\mu 4}|^2$ obtained from MINOS neutral-current data.

- The data of gallium radioactive source experiments (GALLEX and SAGE).

- The $\nu_e + {}^{12}\mathrm{C} \to {}^{12}\mathrm{N}_{\text{ground state}} + e^-$ scattering data of the KARMEN and LSND experiments, which constrain the short-baseline ν_e disappearance.

The χ^2 of a spectral fit to the above experiments is minimized at 1 eV^2 when combining both appearance and disappearance data, agreeing with the initial claims of LSND (Fig. 2.8). However, when fit independently, the appearance and disappearance data favor different regions of parameter space. This inconsistency is not solved by adding more than one additional sterile neutrino.

If the data can be explained by neutrino oscillations alone, rather than some other new physics, then one or more experiment is wrong. An experiment that can probe all of these channels would be advantageous, given the complications of combining data from different experiments.

2.2.5 Future experiments

Future experiments must be sensitive to the LSND anomaly, while also resolving the differences between appearance and disappearance results. The history of having had numerous nondefinitive experiments means that future experiments must be sensitive to this effect at greater than 5σ and should rely on proven technologies.

Funded experiments will take data shortly (see [38] and references therein) that may marginally increase our understanding of sterile neutrinos. The MINOS+ experiment will take data in the high-energy NuMI beam, concurrently with NoVA, and should provide more data for the muon disappearance channel. The μBooNE experiment is a liquid argon (LAr) time projection chamber (TPC) that will be placed in a Fermilab booster beamline and should determine the origin of the MiniBooNE low-energy excess, but it will not resolve the tensions in oscillation fits.

Countless experiments have been—and continue to be—proposed. A partial list is given in Table 2.3. In addition to the topic of this thesis, various experimental techniques may be sensitive to sterile neutrinos at 5σ.

Two LAr TPCs are being proposed at either CERN (ICARUS/NESSiE) or Fermilab (LAr1). Both experiments intend to use a conventional beam aided by having a near detector. These experiments have the advantage that they also provide useful detector R&D.

Also under consideration is the deployment of an MCi radioactive source in the detector halls of either Borexino, SNO+, or Daya Bay, which are low-energy ($\approx$MeV) neutrino detectors. Radioactive sources provide megaelectronvolt neutrinos and, if sterile neu-

Table 2.3: Nonexhaustive list of past and future planned experiments able to probe oscillations to sterile neutrinos. Details include the type of experiment (decay-at-rest, short baseline, reactor, etc.), how they are sensitive to oscillations (rate, energy spectrum, and/or length dependence), and the reaction channel probed (either CC, NC, or elastic scattering (ES)). Past experiments are denoted by †. Names in italics are proposed experiments. (Table reproduced and modified from [20].)

Experiment(s)	Source	Type	Sensitivity	Channel
LSND†	Decay-at-rest	$\bar{\nu}_\mu \to \bar{\nu}_e$	Rate, energy	CC
MiniBooNE†	Short baseline	$\bar{\nu}_\mu \to \bar{\nu}_e$	Rate, energy	CC
Reactor measurements†	Reactor	$\bar{\nu}_e$ dis.	Rate	CC
Gallium anomaly†	EC source	ν_e dis.	Rate	CC
Future decay-at-rest				
(*OscSNS*)	Decay-at-rest	$\bar{\nu}_\mu \to \bar{\nu}_e$	Rate, energy	CC
		ν dis.	Rate	NC
Future short baseline				
(*μBooNE, BooNE,*	Short baseline	$\overset{(-)}{\nu}_\mu \to \overset{(-)}{\nu}_e$	Rate, energy	CC
NESSiE, LArLAr)		$\overset{(-)}{\nu}_\mu$ dis.	Rate, energy	CC
(*νSTORM; this thesis*)	Short baseline	$\overset{(-)}{\nu}_e \to \overset{(-)}{\nu}_\mu$	Rate, energy	CC
		$\overset{(-)}{\nu}_e, \overset{(-)}{\nu}_\mu$ dis.	Rate, energy	CC, NC
Future reactor measurements				
(*Nucifer, Stereo, Solid*)	Reactor	$\bar{\nu}_e$ dis.	Rate, length	CC
Future source experiments				
(*Borexino, Ce-LAND, Daya Bay*)	β^- Source	$\bar{\nu}_e$ dis.	Rate, length	CC
(*Borexino, SNO+Cr*)	EC Source	ν_e dis.	Rate, length	ES
(*LENS, Baksan*)	EC Source	ν_e dis.	Rate, length	CC
(*RICOCHET*)	EC Source	ν_e dis.	Rate, length	NC

trinos exist, should result in an oscillation length for $\Delta m^2 \sim 1$ eV^2 comparable to the detector size. Various different isotopes are being considered. The main experimental complication is that hot radioactive sources have short half-lives and that radioactive source presents radiological concerns.

The radioactive source concept is similar to those using nuclear reactors to search for electron disappearance (e.g., Nucifer, Solid). These ideas require small research reactors that emit $\bar{\nu}_e$, such as at the Institut Laue-Langevin (ILL) in Grenoble, and baselines less than 10 m. Backgrounds are reduced by using coincidences between e^+ and n from inverse β decay—much like LSND. These experiments are on the surface to be near the reactor core, which increases the cosmic ray background rate; therefore, ideally, a submerged nuclear submarine on a six month deployment would be used instead of the ILL, but this presents other issues. Most importantly, these reactor experiments provide a cheap means for confirming or refuting the reactor neutrino anomaly—cheap experiments help optimize baselines for expensive ones.

Planck is a next-generation cosmic microwave background (CMB) satellite that aims to improve upon WMAP and KOBE's measurements of the CMB. It also helps determine the number of effective neutrinos N_{eff} that contribute to the formation of the Universe. A light sterile neutrino may have thermalized and affected the formation of the Universe; however, the observation of N_{eff} greater than $\sim$3 would not exclusively mean the existence of a sterile neutrino since other new physics (e.g., axions) could have been the cause. The prediction of N_{eff} without sterile neutrinos from big bang nucleosynthesis is 3.046 due to various corrections that need be applied (See, e.g., [39] and references therein).

At the time of writing this thesis, the Planck satellite space mission [39] released its first major result that confirmed, from a particle physicist's perspective, that the Universe is exactly like we expected. Although the Planck collaboration should be congratulated on having performed exceptional science, the results are nevertheless disappointing

since they do not provide any anomalies that could lead to new physics. Specifically, their sterile neutrino result measured $N_{\text{eff}} = 3.30 \pm 0.27$ [40]. Without introducing large primordial neutrino asymmetries, these results are hard—but not impossible—to reconcile with terrestrial sterile neutrino experiments [41]. Nevertheless, a 10 eV^2 sterile neutrino is still consistent—albeit with tension—with both short-baseline experiments and cosmology [42]. Cosmological experiments fit the CMB using the physics models developed terrestrially; therefore, if sterile neutrinos were shown to agree, then our cosmological models would be updated accordingly—much like how the discovery of neutrino mass changed the models.

2.3 Conclusions

Neutrino oscillations are firmly established. However, numerous anomalies challenge the standard 3ν oscillation framework. Experiments will continue to be performed to elucidate these effects, but no funded experiment will conclusively confirm or reject the sterile neutrino hypothesis. Conventional beams have been unable to confirm the effect at $>5\sigma$. A new experiment is needed that can measure appearance and disappearance signals.

Chapter 3

Designing a new sterile neutrino experiment

> If you see a snake, just kill it—don't appoint a committee on snakes.
>
> Ross Perot

This chapter shows the arguments that were used to design the VLENF (now νSTORM) in the summer and fall of 2011, and was the initial and original sterile neutrino work for νSTORM. At this point, nobody else had worked or was working on the sterile neutrino reach of this type of facility since VLENF/νSTORM was originally conceived for cross-section measurements. After this work, it became clear that there was a strong sterile neutrino case. This chapter pedagogically discusses the design of the proposed experiment. The arguments are illustrative and are estimates that need further confirmation. The chapters following this one expand arguments, confirm assumptions, and give well-defined descriptions.

A new experiment is needed to determine whether or not sterile neutrinos exist. What would be the best experiment to build? And what are the requirements? Firstly, the experiment must be decisive and model independent: the LSND $\bar{\nu}_\mu \to \bar{\nu}_e$ channel should be directly studied at $L/E \sim$ km/GeV with a sensitivity of at least 5σ. Further-

more, studying several oscillation channels in the same detector is advantageous. These requirements limit the experimental possibilities.

Any new experiment needs—obviously—neutrinos. Physical processes that produce neutrinos include nuclear, π, and μ decay. However, each neutrino source has unique experimental difficulties.

Nuclear β decays (i.e., the process $u \rightarrow d + e^+ + \bar{\nu}_e$) have been extensively used throughout the history of neutrino physics, whether occurring in solar fusion or nuclear reactors. For example, reactors allowed the discovery of neutrinos, precision measurements of Δm^2_{21}, and the measurement of the mixing angle θ_{13}. However, β-decay processes are only useful for disappearance measurements. The appearance of the second lightest lepton—the muon—requires energies greater than m_μ =106 MeV; however, the Q value of β decay is around 1 MeV. It has been proposed to accelerate radioactive nuclei so that the electron neutrinos have energies above the muon production threshold—the so-called β beam concept—but this poses significant technical difficulties. Studying electron disappearance is crucial for studying sterile neutrinos, but β decay does not allow for directly testing the LSND signal.

The LSND signal $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ can be studied using the process $\pi^+ \rightarrow \mu^+ + \nu_\mu$, where the LSND collaboration used a stopped-π beam and the MiniBooNE collaboration used conventional beams with a π-decay pipe. A tremendous amount of experience using these beams has been accumulated. Unfortunately, appearance measurements are also difficult—not impossible—with π beams, but for different reasons than nuclear-decay beams. Pions are produced when protons collide with a fixed target, and this production depends upon the quark and gluon content within the nuclear target. Hadronic models have large uncertainties that translate into $\pi^\pm$ and $K^\pm$ production uncertainties, where background $\bar{\nu}_\mu$ and $\overset{(-)}{\nu}_e$ are produced at the level of ≈1%. Typically, a near detector is built to provide an external constraint on the product of the flux and cross section $\phi \times \sigma$, and also measure, e.g., the single-π-production background. This σ is actually the cross

section for muon-neutrino interactions, σ_{ν_μ}. However, knowledge of σ_{ν_e} is required for these appearance measurements, and the cross-section ratio ν_e/ν_μ has theoretical uncertainties and no experimental measurements. Also, the short 7-m decay length of the pion makes it difficult to measure and manipulate the pion beam.

In contrast, a $\mu^+ \to e^+ + \nu_e + \bar{\nu}_\mu$ decay beam has smaller uncertainties and can be used to measure $\sigma_{\nu_e}/\sigma_{\nu_\mu}$. The decay length of a muon is 600 m; therefore, the beam is long lived, and can be stored in a ring and measured. For conventional pion beams, the complication is the uncertainty in the target production, and—due to the short lifetime—the beam cannot be precisely measured before it decays. With muons stored in a ring, the same production uncertainties exist because the muons result from π decay; however, the long lifetime allows for precise measurements of the muon beam using conventional beam instrumentation.

Muons will be used for this proposed experiment, which is called the neutrinos from stored muons (νSTORM)[1] experiment. Nuclear and π decays produce only one neutrino flavor; yet μ decay produces both electron and muon neutrinos. All $\overset{(-)}{\nu}_{e/\mu} \leftrightarrow \overset{(-)}{\nu}_{e/\mu}$ channels are experimentally accessible at νSTORM (Table 3.1). Therefore, many channels can be studied in one experiment, which helps reduce the systematic uncertainties. In a μ^+ decay, $\bar{\nu}_\mu$ and ν_e are produced and, if there are no oscillations, only a μ^+ can be produced by neutrino interactions—the so-called "right-sign" disappearance channel $\bar{\nu}_\mu \to \bar{\nu}_\mu$. However, if there are oscillations, then a μ^- (i.e., "wrong-sign" muon) can be produced by $\nu_e \to \nu_\mu$ oscillations. Therefore, if the experiment is viewed as a black box, then a nonzero appearance probability converts a μ^+ into a μ^-, creating this "wrong-sign" muon (i.e., $\mu^+ \to \nu_e \rightsquigarrow \nu_\mu \to \mu^-$). Just to be clear: charge is still conserved because the positron and other neutrino were ignored.

There is, nevertheless, a downside: the appearance signals require a high degree of lepton charge discrimination. For example, if μ^+ are stored, the oscillation $\nu_e \to \nu_\mu$

[1] Joachim Kopp invented this clever name.

Table 3.1: Table of possible oscillation channels from muon decay rings. All oscillation channels below the τ threshold are present. Channels are either called "right sign," if the detected neutrino is produced in muon decay, or "wrong sign."

Parent decay	Oscillation channel	Final-state lepton	Colloquialism
$\mu^+ \to e^+ + \nu_e + \bar{\nu}_\mu$	$\bar{\nu}_\mu \to \bar{\nu}_\mu$	μ^+	right-sign
	$\nu_e \to \nu_\mu$	μ^-	wrong-sign
	$\nu_e \to \nu_e$	e^-	right-sign
	$\bar{\nu}_\mu \to \bar{\nu}_e$	e^+	wrong-sign
$\mu^- \to e^- + \bar{\nu}_e + \nu_\mu$	$\nu_\mu \to \nu_\mu$	μ^-	right-sign
	$\bar{\nu}_e \to \bar{\nu}_\mu$	μ^+	wrong-sign
	$\bar{\nu}_e \to \bar{\nu}_e$	e^+	right-sign
	$\nu_\mu \to \nu_e$	e^-	wrong-sign

occurs concurrently with $\bar{\nu}_\mu \to \bar{\nu}_\mu$, which produce μ^- and μ^+, respectively. These muons must be differentiated.

The required level of discrimination can be estimated. The number of events is

$$N = \epsilon \times [\sigma_{\text{nucleon}} \times N_A \times m_{\text{det.}}] \times P_{e\mu} \times \Phi, \tag{3.1}$$

where ϵ is the efficiency, σ_{nucleon} is the appropriately-defined neutrino–nucleon cross section for an isoscalar target, N_A is Avagadro's constant 6.02×10^{23}, $m_{\text{det.}}$ is the target mass in grams, $P_{e\mu}$ is an oscillation probability from an electron flavor state to a muon flavor state, and Φ is the flux. Until more of the facility has been designed later in this thesis, it is not possible to make proper statistical arguments about the required detector performance since the interaction rates are unknown. However, the following pedagogical relation

$$\frac{N_{\text{sig.}}}{N_{\text{bkg.}}} = \frac{\epsilon_{\text{sig.}} \times \sigma_{\text{sig.}} \times P_{e\mu} \times \Phi_{\text{sig.}}}{\epsilon_{\text{bkg.}} \times \sigma_{\text{bkg.}} \times P_{\mu\mu} \times \Phi_{\text{bkg.}}} \sim 1 \tag{3.2}$$

can be used for the moment. The neutrino cross-section ratio $\sigma_{\text{sig.}}/\sigma_{\text{bkg.}}$, flux ratio $\phi_{\text{sig.}}/\phi_{\text{bkg.}}$, and disappearance probability are all approximately 1. Rearranging Eq. (3.2)

results in

$$P_{e\mu} = \frac{\epsilon_{\text{bkg.}}}{\epsilon_{\text{sig.}}}. \tag{3.3}$$

The appearance probability is $\mathcal{O}\left(10^{-3}\right)$. If the detector is only—conservatively—10% efficient at detecting muons, then the background rejection must be $\mathcal{O}\left(10^{-4}\right)$.

The only way to determine a particle's charge at this level is to magnetize the detector. Only muons are considered because it is difficult to determine an electron's charge due to its short interaction length. This would be worrying given that LSND's signal was electron appearance if it were not for the conservation of CPT. It is true that $\bar{\nu}_\mu \to \bar{\nu}_e$ cannot be easily explored using muon-decay beams; nevertheless, the CPT-equivalent channel $\nu_e \to \nu_\mu$ *can* be explored. (If CPT does not hold, then QFT is wrong and it is a bigger discovery than sterile neutrinos.) There are advantages to this channel: neutrino cross sections are larger than the antineutrino cross sections and muons are easier to detect than electrons.

The two design requirements are that the detector must identify muons and be magnetized. Identifying muons from NC events is best done using range—muons go farther than, for example, pions—because reconstructing interaction kinematics or looking for decay "kinks" is challenging. The complication with using range is that magnetizing the large volumes, which is required to stop the muon, is expensive (large volumes are also required to increase the detector target mass).

Ideally, the detector would be totally active and inside of a large magnet. Superconducting magnets provide more field than conventional magnets, therefore would be advantageous and are used for νSTORM. However, producing magnetic fields over large volumes is difficult if the detector is not ferromagnetic. For example, a totally active scintillator detectors could work, but a conventional superconducting magnet would cost[2] ≈40 million dollars. Less conventional magnets such as superconducting trans-

[2]A model of superconducting magnet costs was developed in 1993 [43], $C = 0.4(BV)^{0.635}$, where C is in millions of US dollars (USDs), B is in Tesla, and V is in m^3.

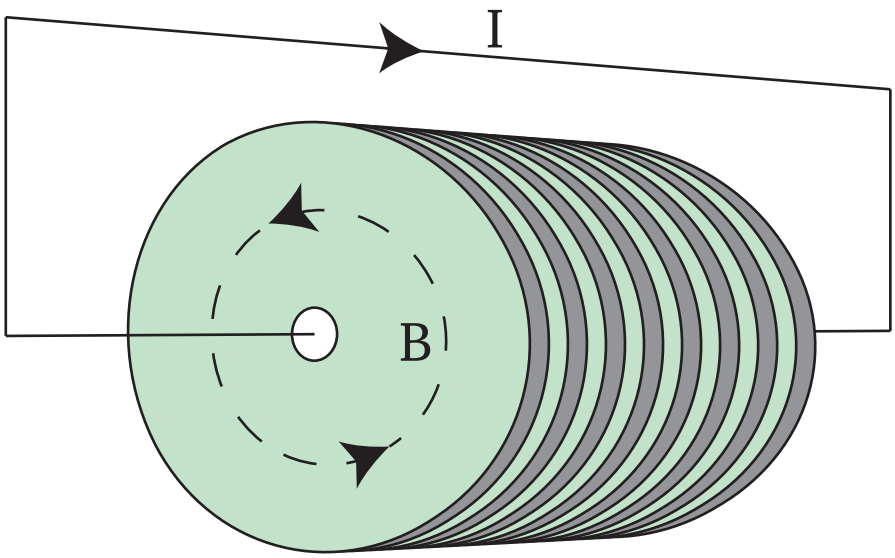

Figure 3.1: Diagram of a magnetized iron sampling calorimeter. Scintillator planes are green and iron planes are gray. The solid line is a conducting cable, where the arrow indicates the direction of the flowing current. The dashed line indicates the toroidal magnetic field lines.

mission lines (STLs)—which are just superconducting cables—introduce technical risk since they were not designed to be wrapped. The radius of curvature for the STL from very-large hadron collider (VLHC) R&D is 15 m, and opinions differ as to how difficult it would be to adapt it to, for example, a 5-m totally-active scintillator detector (TASD).

Another alternative is to use a magnetized-iron sampling calorimeter, much like MINOS (Fig. 3.1). This technology is well proven for neutrino detection and has a long history [33,44]. Iron and plastic scintillator planes are stacked (iron, plastic, iron, plastic, and so on), where the planes are assumed to have equal thickness. A current carrying cable passes through a hole in the center. The current magnetizes the ferromagnetic iron and, in the case of a STL, saturates the iron at 2 T with a toroidal field—compared to 1.16 T in the middle of the MINOS near detector target volume [45] using a normal-conducting magnet.

Current in a conductor passing through the detector center produces a toroidal field in the iron. To obtain the maximum field in the iron—which is advantageous for charge

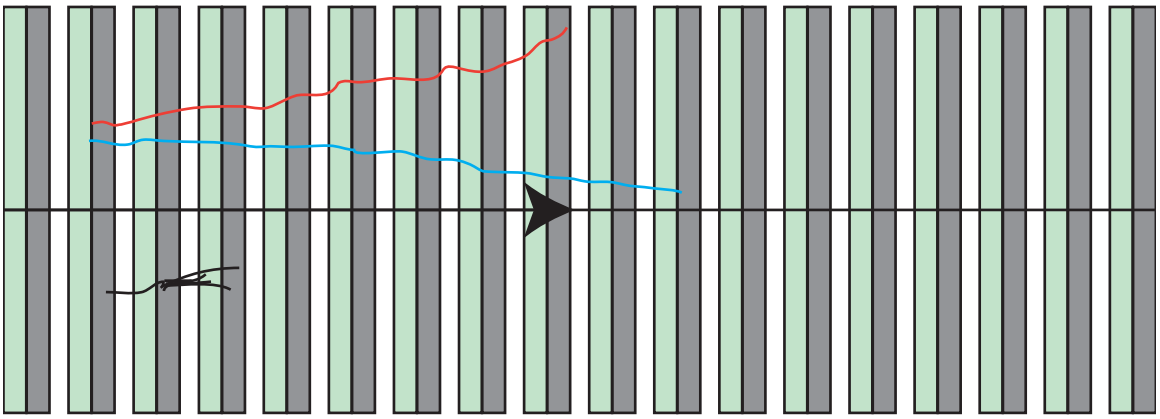

Figure 3.2: Sketch of how the detector rejects backgrounds. The solid line is the conducting cable, where the arrow indicates the direction of current flow. The blue line in the middle is a μ^- that is focused towards the conducting cable. The red is a "wrong-sign" muon (i.e., μ^+) that is defocused. The short black lines below the superconducting cable is a NC event that does not penetrate much steel.

identification—a superconducting cable must be used. A suitable cable was developed for the VLHC: the STL. It can carry a maximum current of 100 kA; however, higher currents can be achieved by having multiple conductors. A current of 250 kA will saturate the iron (B =2T) up to a radius of 2.5 m.

As mentioned before, the two design constraints for the detector are that it must differentiate muons and non-muon events, and that it must be magnetized. Simple arguments can be used to demonstrate the performance of a magnetized-iron sampling calorimeter

Separating NC and ν_e CC events from ν_μ CC events is straightforward: muons will penetrate more steel than electrons or hadrons. Requiring that particles go through 2 m of steel will reject most NC and ν_e CC backgrounds. The hadronic interaction length in steel is $X_0^{\text{Int.}} = 20$ cm, so 2 m corresponds to ten interaction lengths. The probability that a pion is able to travel farther than 2 m without interacting is then at most $e^{-10} = 4\times10^{-5}$. The radiation length in steel is much shorter ($X_0^{\text{Rad.}} = 1.75$ cm); therefore, ν_e CC can be

ignored. The background levels from NC and ν_e CC satisfy the requirement of Eq. (3.3).

Demonstrating that μ^+ and μ^- can be separated requires a small derivation. The muons must be bent in a magnetic field to determine the charge. However, multiple scattering (MS) can appear to bend the muon the "wrong way". Achieving the requisite level of background rejection requires the bending due to the magnetic field to exceed the apparent bending due to MS.

If the steel is saturated, how far must a muon travel for the charge to be known at the required level? Both bending and MS can be represented as momentum kicks transverse to the direction of the muon, where the momentum kick from bending should exceed that of multiple scattering. The momentum kick for bending in a magnetic field is

$$p_{\perp}^{B}\ [\text{MeV/c}] = 300Bx\ [\text{Tesla meters}], \tag{3.4}$$

where B is the magnetic field and x is the distance traveled [46]. The competing MS momentum is $p_{\perp}^{\text{MS}}$ and is defined in terms of the RMS angular spread σ_θ from MS. For high-Z material such as steel,

$$\sigma_\theta = \frac{13.6\ \text{MeV}}{\beta cp}\sqrt{x/X_0}, \tag{3.5}$$

where p, βc, and x/X_0 are, respectively, the momentum, velocity, and the distance in units of radiation length [2]. The X_0 for the composite material can be found by taking the weighted-average of the reciprocals,

$$\frac{100\%}{X_0^{\text{composite}}} = \frac{50\%}{X_0^{\text{iron}}} + \frac{50\%}{X_0^{\text{scintillator}}}, \tag{3.6}$$

where $X_0^{\text{composite}}$ is found to be 3.38 cm (Table 3.2), and the 50% comes from half the thickness of a layer being either iron or scintillator. Using the small-angle approximation

Table 3.2: Detector parameters including radiation lengths.

Region	Parameter	Value
Iron	X_0	1.76 cm
	Density	7.874 g cm^{-3}
	Magnetic field	2 T
Scintillator (Polystyrene $[C_6H_5CHCH_2]_n$)	X_0	43 cm
	Density	1.06 g cm^{-3}
	Magnetic field	0
Effective	X_0	3.38 cm
	Density	4.467 g cm^{-3}
	Magnetic field	1 T

it is possible to find $p_\perp^{\rm MS}$:

$$p_\perp^{\rm MS} = \frac{13.6 \text{ MeV}}{\beta}\sqrt{x/X_0}. \tag{3.7}$$

The ratio r of $p_\perp$ kicks determines the importance of scattering vs magnetic bending:

$$r = \frac{p_\perp^B}{p_\perp^{\rm MS}} = \frac{300Bx \times \beta}{13.6\sqrt{x/X_0}} \tag{3.8}$$

$$\simeq 22B\sqrt{x}\sqrt{X_0}\beta. \tag{3.9}$$

$$\implies x = \frac{r^2}{484B^2X_0\beta^2} \tag{3.10}$$

The ratio r must exceed 4.3, which corresponds to a probability of 10^{-5} if the distribution is Gaussian. From Table 3.2, for the composite material $X_0 = 3.38$ cm, $B = 1$ T, and $\beta \sim 1$; therefore, the muon must penetrate at least 1 m. The range requirement for rejecting NC events provides more than enough bending to separate muon polarities.

Imposing a cut that requires a muon to penetrate 4 m of detector (i.e., 2 m of steel) will limit the neutrino energies where this detector technology is useful, which determines the required accelerator energy. The relationship between muon momentum and range can be determined by using the continually-slowing-down approximation

(CSDA). For a muon with kinetic energy T, the integral for the range R is

$$R = \int_{T_0}^{T} \langle dE/dx \rangle^{-1} \, dE, \tag{3.11}$$

where the mean stopping power $\langle dE/dx \rangle$ arises from ionization loss and is given by the Bethe-Bloch expression:

$$\langle dE/dx \rangle = \frac{4\pi}{m_e c^2} \frac{nz^2}{\beta^2} \left(\frac{e^2}{4\pi\varepsilon_0} \right)^2 \left[\ln\left(\frac{2m_e c^2 \beta^2}{I(1-\beta^2)} \right) - \beta^2 - \frac{\delta}{2} \right], \tag{3.12}$$

where $\beta = v/c$, z is the particle's charge, e is the electron charge, m_e is the electron mass, n is the electron density, I is the mean excitation energy, ε_0 is the vacuum permittivity, and δ is a density correction [2]. The computation in [2] of $\langle dE/dx \rangle$ for iron and scintillator is used when computing the range in the composite material. The effect of radiative corrections and multiple scattering is ignored. Integrating R for the composite material, T_0 must be greater than 3 GeV for the muon to penetrate 2 m of steel.

The energy threshold of the detector requires that the neutrinos from the accelerator be above 3 GeV. The neutrino spectrum from muon decay is continuous—it's a wide-band beam. Ideally, the accelerator muon energy would be higher than the detector energy threshold; however, this is expensive and it is useful to know how well the 4 GeV muons perform. A 4 GeV muon will result—roughly—in a quarter of the neutrinos having energies above 3 GeV; therefore, the assumed signal efficiency is 10% (p.p. 48).

For many reasons, higher stored-muon energies are better than lower ones: cross sections, detector efficiency, and beam collimation. However, the beam energy cannot be increased indefinitely. The muons are a tertiary beam; therefore, the pion and proton energies must be high enough to create high-energy muons. The relation between proton

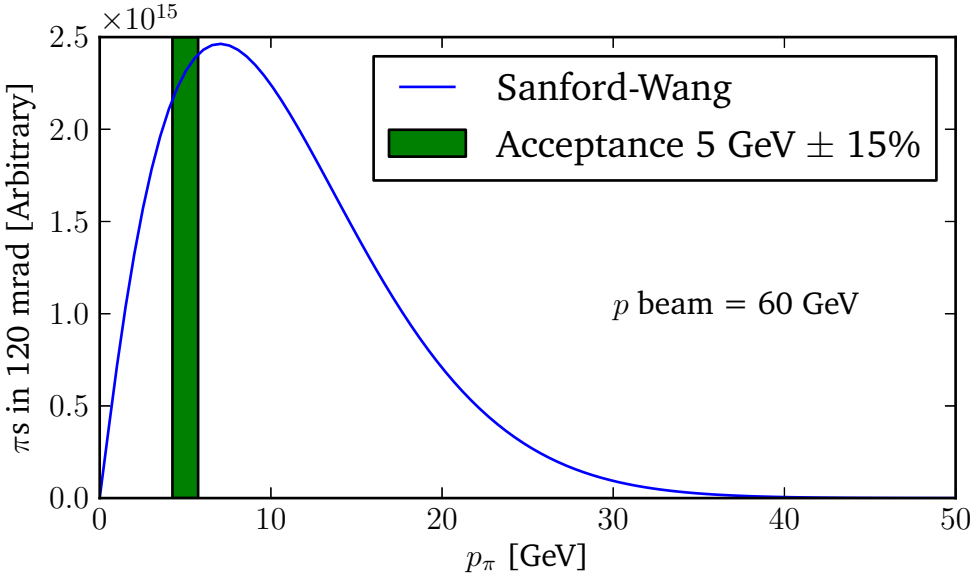

Figure 3.3: Sanford-Wang parameterization for π production using 60 GeV protons for a thin target.

and pion energies is given by the Sanford-Wang parameterization for a thin target,

$$\frac{d^2\sigma(p+A\rightarrow\pi^+ +X)}{dpd\Omega}(p,\theta) = c_1 p^{c_2}\left(1-\frac{p}{p_{beam}}\right)\exp\left[-c_3\frac{p^{c_4}}{p_{beam}^{c_5}}-c_6\theta(p-c_7 p_{beam}(\cos\theta)^{c_8})\right], \quad (3.13)$$

where p is the pion momentum, Ω is solid angle, p_{beam} is the beam momentum, and the rest (i.e., c_1, c_2, ..., c_8) correspond to fit parameters given in Table 3.3 [47]. The highest energy proton machine that could deliver enough protons for a new neutrino experiment is the Fermilab main injector (MI), which can have a proton energy of 60 GeV; however, it typically runs at 120 GeV. Pion production is maximal at approximately 10% of the proton energy, which means that peak pion production occurs at 5–6 GeV. These pions will decay lower energy muons; therefore, the muon energy is assumed to be 3.8 GeV.

The baseline can be estimated, assuming this muon energy. The oscillation probability

Table 3.3: Sanford-Wang parameterization for Be targets.

c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8
196	1.08	2.15	2.31	1.98	5.73	0.137	24.1

is

$$\text{Prob}[\,\nu_e \to \nu_\mu] \sim \sin^2\left(\frac{L}{4\langle E_\nu\rangle}\Delta m^2\right), \tag{3.14}$$

where L is the baseline, E is the neutrino energy, and Δm^2 is the neutrino parameter. The guesstimate L is then

$$\begin{aligned}
\frac{L}{4\langle E_\nu\rangle}\Delta m^2 &= \frac{\pi}{2} && (3.15)\\
L &= \frac{4\langle E_\nu\rangle\pi}{2\Delta m^2} && (3.16)\\
&= \frac{4\langle E_\nu\rangle\pi}{2\Delta m^2} && (3.17)\\
&= 7.6\times 10^{18}\ \text{GeV}^{-1} && (3.18)\\
&= 2.3\ \text{km} && (3.19)
\end{aligned}$$

where $\langle E_\nu\rangle = 2$ GeV and the last step used the relation, in natural units, $\hbar c = 197$ MeV fm. A more sophisticated baseline optimization will be performed in later sections. In the meantime, even though the energy threshold of the detector is 3 GeV, choosing $\langle E_\nu\rangle =$ 2 GeV helps optimize the disappearance searches. The disappearance signal is larger, and energy cuts can be relaxed for this analysis. Solid angle effects have also been ignored, but are included later.

The final parameter to determine is how many neutrino interactions are required, which depends upon the number of stored muons and also the target mass. Recalling Eq. (3.1) on p.p. 47, the number of interactions can be written as

$$N = \epsilon \times [\sigma_{\text{nucleon}} \times N_A \times m_{\text{det.}}] \times P_{e\mu} \times \Phi, \tag{3.20}$$

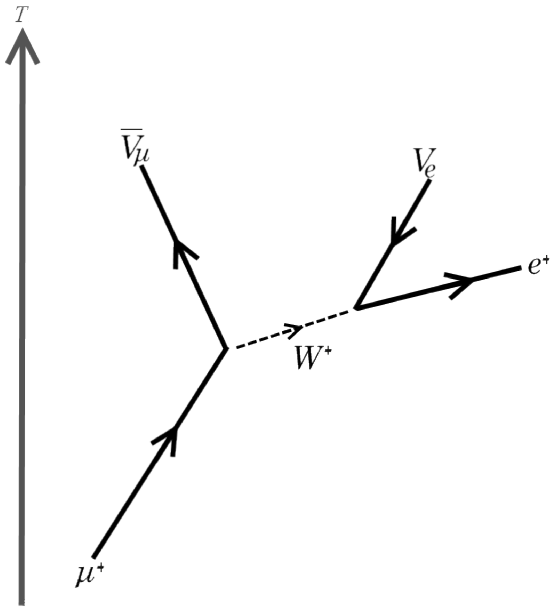

Figure 3.4: Feynman cartoon of muon decay. Specifically, the decay $\mu^+ \to \bar{\nu}_\mu \nu_e e^+$.

where σ_{nucleon} is the appropriately-defined neutrino–nucleon cross section for an isoscalar target, N_A is Avagadro's constant 6.02×10^{23}, and $m_{\text{det.}}$ is the target mass in grams. Assuming a 1.5 kt detector could be built, what flux ϕ is needed?

The neutrino flux arises from the electroweak decay $\mu^+ \to \bar{\nu}_\mu \nu_e e^+$ (see Feynman cartoon in Fig. 3.4). The electron-neutrino spectrum is

$$\frac{dn_{\nu_e}}{dx\, d\Omega} = \frac{1}{4\pi}\left[12x^2(1-x)\right](1 + \mathscr{P}\cos\theta), \tag{3.21}$$

where $x = 2E_\nu^{\text{c.o.m.}}/m_\mu \in [0,1]$ and Ω are the scaled neutrino energy and solid angle in the muon rest frame, $\mathscr{P}$ is the polarization, and θ is the neutrino angle. (See, for example, Chap. 6 of [48] for a derivation.) Electron and neutrino masses are negligible for this process and can be ignored, hence the inclusive range for values of x. Polarization $\mathscr{P}$ is ignored due to the incoherent Thomas precessions of the beam (studied later in Sec. 6.1.1).

The boosted electron-neutrino double-differential flux at distance L in the laboratory

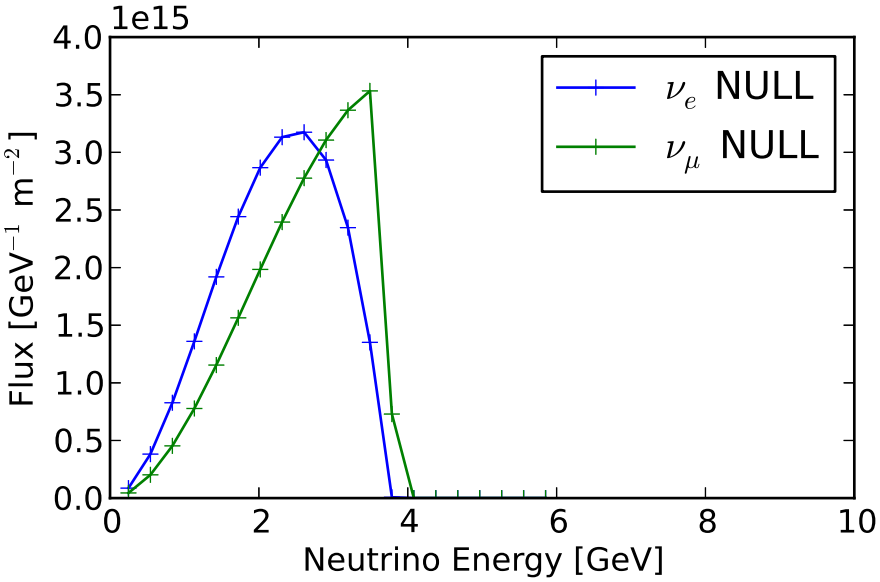

Figure 3.5: The flux of ν_e and $\bar{\nu}_\mu$ for a 3.8-GeV/c muon decay without oscillations at 2000 m. No smearing due to accelerator effects has been performed.

frame is

$$\frac{d^2N_{\nu_e}}{dy\,dA} = \frac{24n_\mu}{\pi L^2 m_\mu^6} E_\mu^4 y^2 (1-\beta\cos\phi)\Big[m_\mu^2 - 2E_\mu^2 y(1-\beta\cos\phi)\Big], \quad (3.22)$$

where $y = E_\nu/E_\mu$ is the scaled neutrino energy in the laboratory frame, $\beta = \sqrt{1-m_\mu^2/E_\mu^2}$, A is an area, n_μ is the number of muons, and ϕ is the angle between the muon and neutrino. This distribution (Fig. 3.5) is for a point source, which is a suitable approximation if the baseline is longer than the accelerator.

Integrating the flux over energy, assuming that ϕ is small, and that the cross section $\sigma = 0.67\times10^{-42}$ m$^2/E$ GeV yields the event rate $N_e = 1.4\times10^{-13}\times n_\mu$. As mentioned before, the goal is to have a 5σ measurement and the oscillation probability is 10^{-4}. Roughly, 25 events are required to have a 5σ sensitivity (i.e., $25/\sqrt{25} = 5$). Given the 10% efficiency, 250 interactions are required. The number of muons can be determined:

$$n_\mu = \frac{10\%\times250}{1.4\times10^{-13}\times10^{-4}} = 1.8\times10^{18}. \quad (3.23)$$

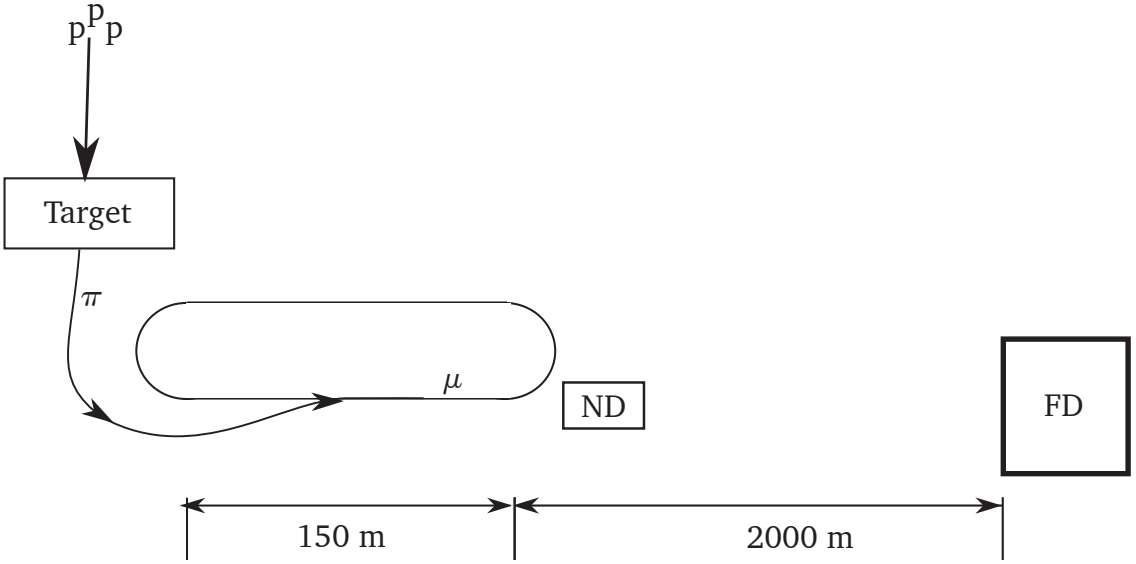

Figure 3.6: A diagram of the proposed νSTORM experiment and detector. Protons are collided onto a fixed target. The resulting pions are then collected and injected into the decay ring where they decay. The muons circulate before decaying to neutrinos that travel toward the near and far detectors.

Based on the estimates in [49] of the muon per proton on target ratio, obtaining this n_μ would require 10^{21} protons on target.

A summary of the experimental parameters is in Table 3.4, and these parameters are the baseline for the proposed νSTORM experiment (Fig. 3.6). The remainder of this thesis further discusses, motivates, and defends the parameters discussed above. How νSTORM could be built is explained in the next chapter. The detector performance is further motivated in Chap. 5, before the physics reach is discussed.

Table 3.4: νSTORM parameters.

Stored muon energy	3.8 GeV
Number of stored muons n_μ	1.8×10^{18}
protons on target (POT)	10^{21}
Baseline	2 km
Detector mass	1.5 kt
Efficiency	10%
Oscillation probability	10^{-3}
Background rejection level	10^{-4}
Magnetic field	Toroidal at 2 T
Muon range cut	2 m
Energy threshold	3 GeV

Chapter 4

νSTORM accelerator and detector technical designs

> If an idea's worth having once, it's worth having twice.
>
> Tom Stoppard

A large fraction of the work in this chapter were done by engineers and other physics. Where specific, about half of the plots in the chapter were created by myself.

In this chapter, a brief history of neutrino beams is discussed followed by a description of how νSTORM could be built. The historical background is intended to give context and show that muon-decay beams are—and have been—feasible for decades. By describing the facility in more depth than the previous chapter, the claim of little new R&D being needed is strengthened further. (More information is available in the letter of intent (LOI) that was submitted to Fermilab [50].)

4.1 Muon-decay neutrino beam review

Neutrino beams have led to a wide range of discoveries, despite not changing much since their initial conception. The pion-decay $\overset{(-)}{\nu}_\mu$ beam was invented in 1961 by Steinberger *et al.* [51], and resulted in the discovery of the muon neutrino. It is also the only reasonable method for creating a "pure" ν_μ beam because pions mainly decay to muon neutrinos, with backgrounds from other decay channels and produced particles at the percent level. These backgrounds arise from muons decaying despite the decay pipe being short, and from the suppressed decay of pions and kaons to electrons. Every accelerator neutrino beam has been made in a similar way, with only slight variations (e.g., focusing horns, or having the muon decay at rest).

However, 13 years after the invention of pion-decay beams (i.e., 1974), neutrino beams from particles other than pions were first being considered [52]. Muons are an ideal candidate because they are unstable, charged, and relatively long lived. Neutrino oscillations were not yet established; however, various accelerators and detectors were being designed to investigate the possibility of oscillations. For example, Neuffer proposed using a 4-GeV/c muon-decay ring to probe neutrino oscillations at the electronvolt scale [49]. However, his idea was not well received within the neutrino community and was promptly forgotten; particle physicists prefer to solve problems by developing new detector technologies, while accelerator physicists prefer building accelerators. Luckily, 30 years later when there was an independent resurgence of interest in probing electronvolt scale neutrinos using muon-decay beams, he was able to reveal that he had done much of the initial accelerator work. His original concept was the starting point for the accelerator design described in this thesis—specifically, his injection technique.

Muon storage rings have never been used for neutrino oscillation physics; however, muons have been stored in accelerators for precision quantum electrodynamics (QED) studies. The first muons were stored in an accelerator in 1961 [53] to test the developing theory of QED. One of the predictions of QED is that the gyromagnetic ratio g is not

exactly 2 due to loop corrections, therefore there is an anomalous magnetic moment $a = (g-2)/2$. Shortly after parity violation was discovered (1956), it was realized that muons from pion decay were polarized and that a nonzero g—2 would result in a measurable spin precession as the muons circulated in a ring. The E821 experiment at Brookhaven National Lab (BNL) was able to measure this effect and claims that $a = (11659208.9 \pm 5.4\ (\text{stat}) \pm 3.3\ (\text{sys})) \times 10^{-10}$, which agrees with theoretical calculations [2]. However these beams are not good for neutrino physics due to poor neutrino intensity; the design requirements of small energy spreads and uniform magnetic fields results in an accelerator that is not ideal for neutrino physics.

Therefore, in the 1980s, little had been demonstrated about muon acceleration for neutrino physics. The beam concept stagnated until the modern idea in the 1990s of the neutrino factory (NF) by Geer [54], where beam cooling and fixed-field alternating gradient (FFAG) accelerators were considered to increase neutrino brilliance in various NF designs. These NFs were designed to perform the difficult measurements of the unknown neutrino parameters: θ_{13} and δ_{CP} (recall Sec. 2.1.4). Even though muon-decay beams can be applied to many different neutrino problems, most NF designs are ambitious and expensive (e.g., 5.3–7.3 billion EUR [55]).

However, cheaper muon-decay beams are possible. Cooling and FFAGs, which constitute a large fraction of the NF accelerator cost, are not required for, e.g., cross-section physics. More recently, Rubbia *et al*. have independently considered neutrino experiments using muon-decay rings to perform cross-section measurements in a similar energy range [56] to what is proposed in this thesis. The νSTORM idea, which was previously called the very-low-energy neutrino factory (VLENF), was also originally conceived for cross section measurements, before this work. The recent focus of muon acceleration programs toward R&D for next-generation beams—or the even more complex Higgs factory—has skipped the intermediate step of a facility like νSTORM, which is more complex than muon rings for QED studies but less complex than a NF.

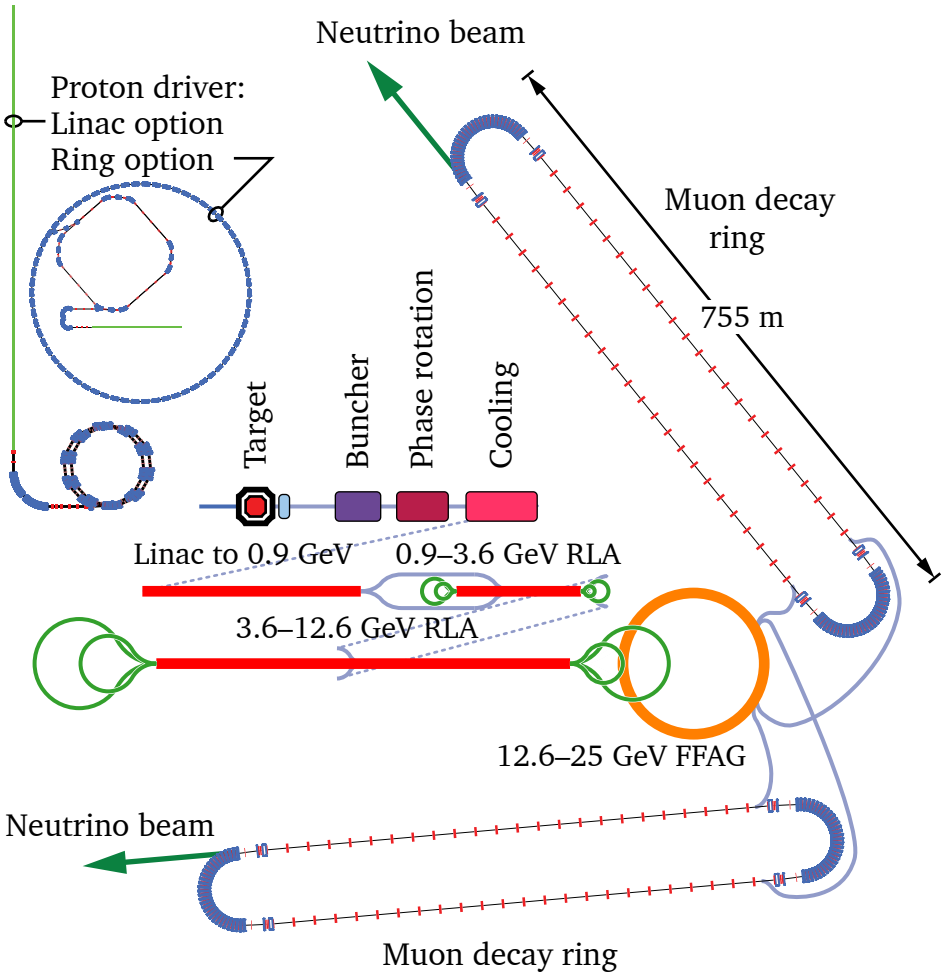

Figure 4.1: High-energy high-luminosity neutrino factory with 10^{21} stored muons with a 25–50 GeV momentum. Two options are shown for the proton driver: a linear accelerator and a synchrotron. Whichever is used would collide protons with a target, where the emerging muons are manipulated in the buncher, phase rotation, and cooling stages. After cooling, the muons are accelerated to 12.6 GeV in a series of linear accelerators (that include a recirculating linear accelerator). Afterward, an FFAG is used to accelerate the muons to 25 GeV. At this point, the beam can be extracted into decay rings such that neutrinos will be directed at detectors. [57]

R&D from the design of high-energy high-luminosity NFs (25–50 GeV, 10^{21} stored muons, Fig. 4.1) can nevertheless be applied to νSTORM (3.8 GeV, 10^{18} muons) [57]. For example, the detectors and beam instrumentation are the same. Efforts toward the construction of NF prototypes relating to cooling magnets and RF cavities (MICE, MUCOOL), acceleration (EMMA), and detector technology (AIDA) are underway. No construction R&D effort, however, has been undertaken to demonstrate the effectiveness of muon-decay beams for neutrino physics. νSTORM can be viewed as a neutrino-factory test facility where beam instrumentation can be tested and detector analyses prototyped.

This νSTORM test-beam facility can be the initial stage of a long-term, staged physics program using muon-decay beams. For example, νSTORM could aid next-generation long-baseline experiments and perform sterile neutrino measurements while fitting on either FNAL's or CERN's site. The next stage would be a facility with sensitivity to >80% of the possible values for the neutrino CP-violating phase, where this CP phase is a benchmark for sensitivity to new physics. There is also merit in overconstraining the PMNS matrix to help model builders. Later stages could include muon colliders at either the Higg's resonance or any new energy scale of interest indicated by LHC experiments. All of these stages would be aided by the successful demonstration of using muon-decay beams for neutrino physics. Muon-accelerator R&D provides the only logically staged program where the benefits of each stage advance accelerator, neutrino, and collider physics.

4.2 νSTORM Accelerator

νSTORM is the simplest possible conception of a "neutrino factory." It consists only of a target and a decay ring: if the target is removed, then there is no beam, and if the decay ring is removed, then it is a conventional pion-decay beam. A technically feasible accelerator design is elaborated in this section. Optimizations of the accelerator complex

are ongoing, and various designs exist for different components that result in similar accelerator performance.

The facility has three high-level components: the accelerator, near detector, and far detector. This accelerator consists of six major components

1. existing proton facility up to 60 GeV/c,
2. extraction from a proton storage ring,
3. a proton target station,
4. pion collection and transport,
5. injection into a muon decay ring, and
6. a muon decay ring.

The accelerator accounts for the majority cost and effort for this experiment—the accelerator physicists work hard so the particle physicists do not have to. The near and far detectors are placed at their requisite distances to receive "useful" neutrinos, where "useful" is defined later. For simplicity, only the Fermilab site will be discussed despite a CERN design existing [58]. Lastly, the term "accelerator" could be more accurately replaced with "beam line" because no RF cavities are used (i.e., no acceleration); however, by convention this machine is still called an accelerator.

A figure of merit that takes into account the performance of each accelerator component is developed to allow for an overall optimization.The accelerator performance metric is the number of muons N_μ that decay toward the detectors, which can be represented by the performance of each component of the accelerator as follows

$$\begin{aligned} N_\mu = (\text{POT}) \times (\pi \text{ per POT}) \times \epsilon_{\text{collection}} \times \epsilon_{\text{transport}} \\ \times \epsilon_{\text{injection}} \times (\mu \text{ per } \pi) \times A_{\text{dynamic}} \times \Omega, \end{aligned} \tag{4.1}$$

where (POT) is the number of protons on target (POT), (π per POT) is the pion multiplicity per proton, $\epsilon_{\text{collection}}$ is the collection efficiency, $\epsilon_{\text{transport}}$ is the transport efficiency, $\epsilon_{\text{injection}}$ is the injection efficiency, (μ per π) is the chance that an injected pion results in a muon within the acceptance, A_{dynamic} is the probability that a muon within the physical aperture is within the dynamic aperture, and Ω is the fraction of the ring that aims muons at the far detector. The values of these parameters are derived in this chapter.

4.2.1 Host ring and extraction

The proton energy is an important parameter for νSTORM. For instance, the proton energy creates an upper limit on E_ν; neutrino interaction rates scale as E_ν^3 (E from cross sections, and E^2 from decay kinematics $1/\gamma$), and detector efficiencies scale roughly as E_ν. However, the proton energy—and accordingly, pion and muon energies—cannot be raised indefinitely. Specifically, a limitation arises from using synchrotron rings to store protons and the limited gradients of RF cavities in the proton front end. The beam takes time to accelerate, yet dipole magnets only transport one beam momentum. Therefore, new beam from the initial proton source cannot be injected until the previous beam has finished accelerating and extracting.

νSTORM is designed to be parasitic off a proton ring, and the Fermilab MI (Fig. 4.2) has been chosen as the host ring. It is the highest energy accelerator (8–150 GeV/c) operating at Fermilab, but also provides a reasonable proton current as it serves other experiments. The MI operates at 120 GeV/c for the MINOS, NOνA, and MINERνA neutrino experiments, but it is possible to extract at 60 GeV/c. For a fixed pion energy, increasing the extracted proton energy increases the pion multiplicity dN_π/dE_π (See, e.g., the Sanford-Wang formalism in [59]). By not cooling or accelerating the beam of muons, the method employed to increase the neutrino flux is to increase the proton energy. Proton energy cannot be increased indefinitely, since it is linearly proportional to beam power, and proton targetry above 400 kW is expensive due to potential radi-

Figure 4.2: A possible νSTORM siting on the Fermilab site using the MI.

ological concerns. The proton energy is already higher than is typically considered in NF designs, which use 6–8 GeV protons—but those "cool" and accelerate muons. This entails taking a low-momentum range of the beam, reducing its emittance so it can be accelerated, then accelerating the beam up to some higher energy.

As an aside, the proton requirements for this machine are compatible with allocations to future experiments. The proton improvement plan will have been completed at FNAL (700 kW at 120 GeV/c), and NOνA will have finished taking data by the time νSTORM is ready to take data. The main beam users will be the muon experiments Mu2e and g-2, which have proton limitations to avoid pileup. LBNE should occur later than nuSTORM.

Extraction may be performed in either of two ways: either an extraction of two MI bunches or extracting the entire MI once every seven fills. Quick extraction would use a fast kicker and start in empty buckets since only 12 of the 15 MI buckets are full. Extracting the entire MI complicates injection because the circumference of the MI is larger than the decay rings being considered. It is important to note that the beam current is higher if the entire MI is extracted, which is preferable for the decay-ring beam instrumentation.

The MI can deliver up to 700 kW, but a modest 100-kW target station is assumed to allow for the operation of other experiments—though the target hall will have enough concrete for 400 kW to allow upgradability. Good sterile neutrino sensitivity requires 10^{21} POT [motivated from Eq. (3.23)]. Power is defined as $P = NE$, where P is the power, N is the number of particles per second, and $E = 60$ GeV is the energy of the beam. Accordingly, the target station could receive $N = 10^{13}$ protons per second on average and, assuming a 2×10^7 s operational year, will result in 2×10^{20} POT per year. After five calendar years νSTORM should deliver its design sensitivity to sterile neutrinos.

The number of protons per second is reasonable in the FNAL proton improvement plan (PIP) era. The booster's linac delivers 4×10^{12} protons per pulse and the repetition rate is 15 Hz [60]. There are 1.2×10^{21} protons per calendar year at 8 GeV/c, which is an upper limit for the number of usable protons, since other accelerators are downstream. For the 60-GeV/c MI, only 80% of the protons that are in the booster can be accelerated to 60 GeV/c, yielding 5×10^{13} protons per second at 60 GeV/c.

4.2.2 Targeting

Targeting entails impeding a proton beam onto a stationary target in order to generate pions. A 100-kW target station for 60-GeV/c protons using a 21-cm tantalum target is considered at present (Fig. 4.3). Histograms of the number of pions in the window 5 GeV ± 10% for various materials (e.g., diamond, Be, Au) of different lengths was provided by Striganov [61] using MARS, and various observations made. The pion yields of these targets are comparable: the target designs agree within the 30% uncertainty within MARS at these energies [62]. Accordingly, the target can be chosen by considering engineering constraints such as heat loads, mechanical stresses, and other engineering constraints, which will most likely result in a preference for a low-Z target.

There is an energy mismatch between the incident proton energy and the collected

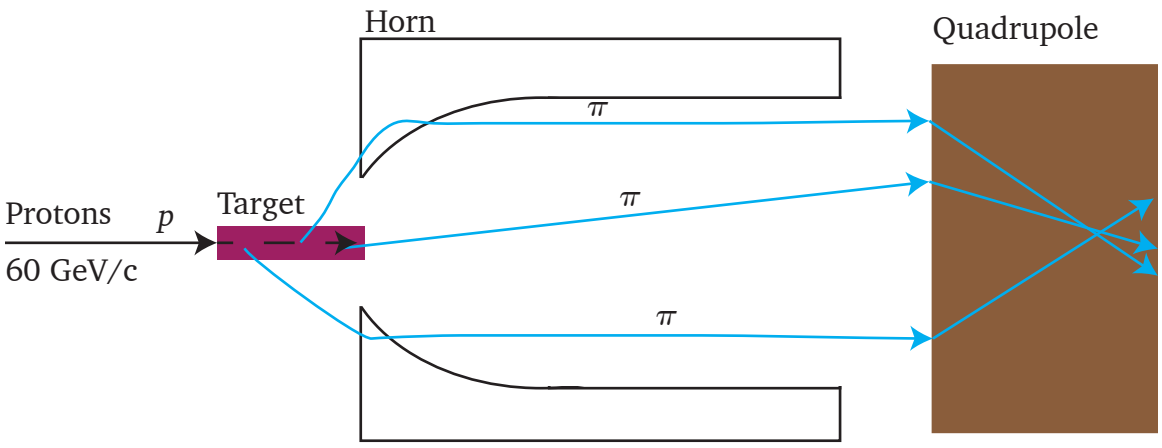

Figure 4.3: Target, horn, and collection.

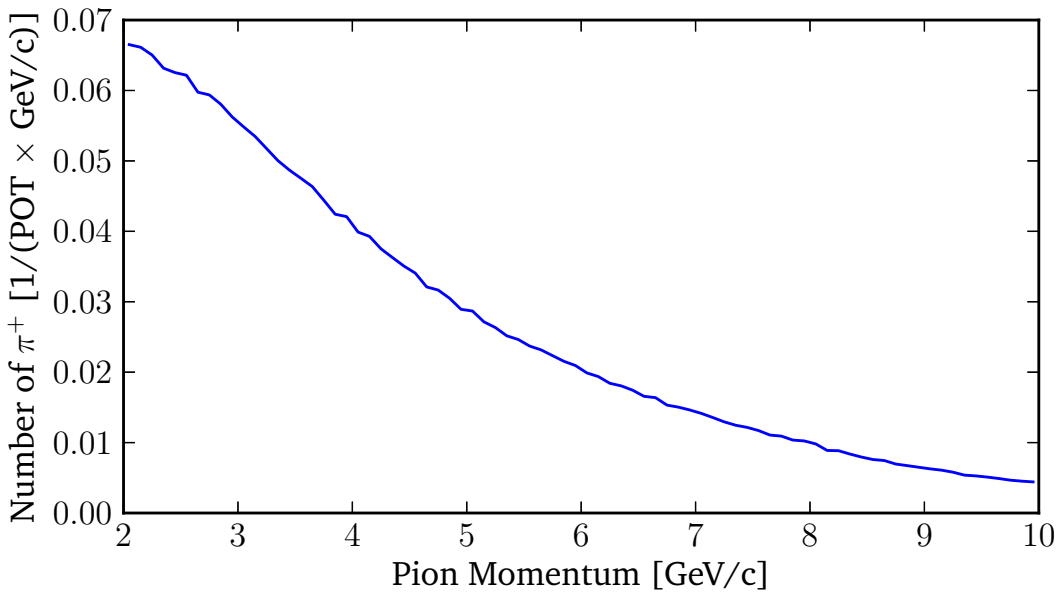

Figure 4.4: Proton 60 GeV, 70-cm Be target, 1-cm radius, $\theta < 120$ mrad, MARS [61].

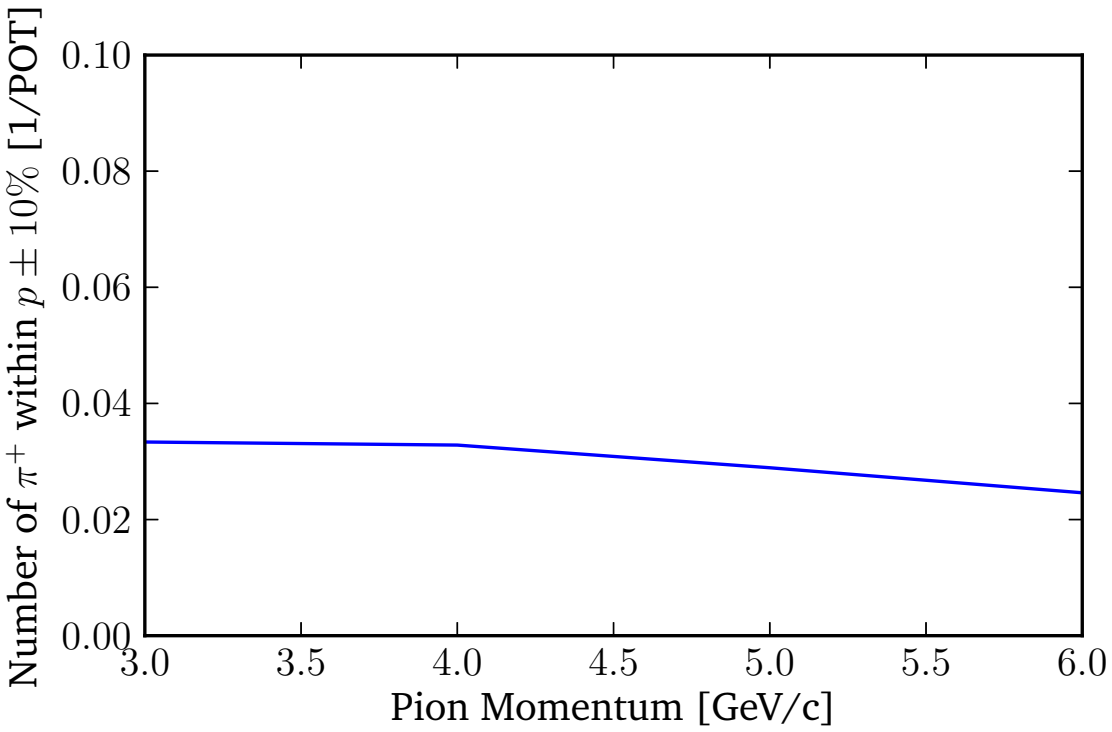

Figure 4.5: Pion production integrated over dynamic aperture.

pion energy. Pion production will increase monotonically with proton energy. Accordingly, the 60-GeV MI is favored to the 8-GeV booster to allow for more pions in 5 GeV ± 10%. The pion multiplicity within 120 mrad and the energy acceptance is shown in Fig. 4.4.

The momentum acceptance of the collection and transport channel is ±10% by design, and this relative acceptance results in $\int_{p_\pi \pm 10\%} (dN_\pi/dp_\pi)dp$ being constant for different momenta as shown by Fig. 4.5. This result is due to the approximate cancellation of two effects as the pion collection momentum increases: the pion production decreases (Fig. 4.4) while the absolute acceptance increases. The function dN_π/dp_π, where p is the pion momentum, is the pion multiplicity within an energy range. Pions with a momentum spread of 10% will be collected from the target. Using linear

interpolation and integrating dN_π/dp_π (i.e., Fig. 4.4), one finds

$$(\pi \text{ per POT}) = \int_{p\pm 10\%} \frac{dN_\pi}{dp_\pi} dp \simeq 0.09, \tag{4.2}$$

where (π per POT). Note that the result is independent of the central pion momentum from 3–6 GeV due to the ±10% momentum acceptance.

4.2.3 Collection and transport

The 4.5–5.0 GeV/c pions emerging from the target must be collected and then transported to the decay ring (the horn and quadrupole in Fig. 4.3). Various techniques were considered (e.g., lithium lenses, solenoids, and horns); however, a horn much like the one used by NuMI gives the desired performance.

Horns are a hollow coaxial structures of conductor through which large currents flow to produce an azimuthal field. They were invented by Simon van der Meer to increase the intensity of the S$p\bar{p}$S accelerator. They have since been used by many accelerator-based neutrino experiments and are a standard technology in collecting secondary beams. The horn gives adequate optics into the collection channel and is the current favored option for νSTORM.

For pions within 5 ± 0.5 GeV and the desired phasespace, a conservative 10% loss of pions during the collection phase is a reasonable estimate [50]; therefore, an efficiency of 90% is used in subsequent calculations. These are pions that are within the design acceptance since most pions (e.g., a 10 GeV pion) are lost. If the transport channel is short enough (<100 m) such that muons do not decay, then the transport efficiency is ≈1.

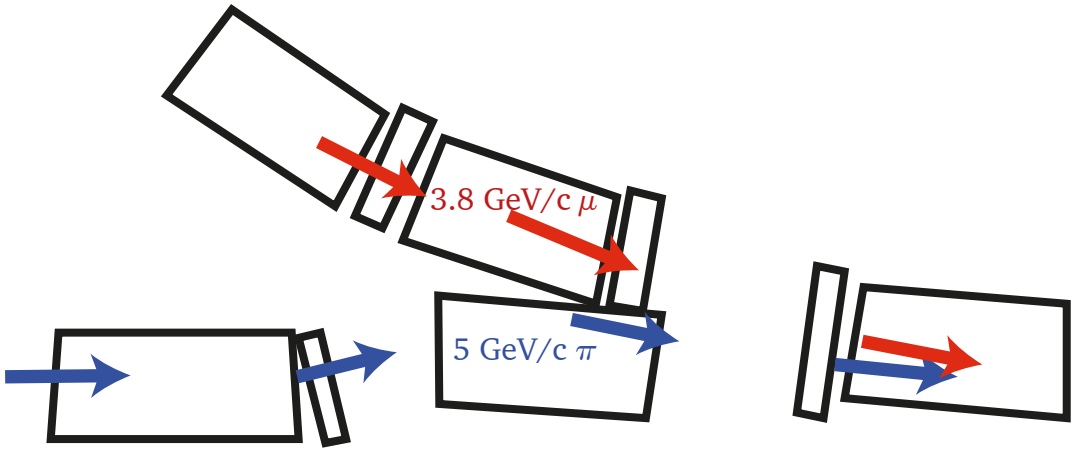

Figure 4.6: A cartoon showing the general principle of stochastic injection.Ó

4.2.4 Injection

Injection into the decay ring will be accomplished by stochastic injection. Stochastic injection was the original concept envisioned by Neuffer and was conceived in the 1980s when designing accelerators *stochastically* was in fashion [49, 63]. For example, stochastic cooling—also invented by Simon van der Meer—was employed at the S$p\bar{p}$S to reduce secondary $\bar{p}$-beam emittance, providing R&D that led to future successful $p\bar{p}$ colliders. Stochastic processes allow for manipulating the beam emittance, despite Liouville's theorem that states the phase space volume of a beam is a conserved quantity.

The circumference of the MI (3319 m) is longer than the decay rings being considered (<500 m), which presents a problem for injection. Beams cannot be injected into the phase space of a circulating beam, and a usual technique such as H^- injection with stripping foils does not work with muons. Kickers would disrupt the circulating beam.

Stochastic injection (Fig. 4.6 and 4.7) relies upon the injected particle species being different from the circulating particle species. For νSTORM, pions could be injected into the ring with the circulating muons by directing the pion beam toward the edge of a quadrupole magnet, which creates a dipole field (Fig. 4.8). The pion and muon occupy different phase space volumes, since they are spatially separated once injected, and also

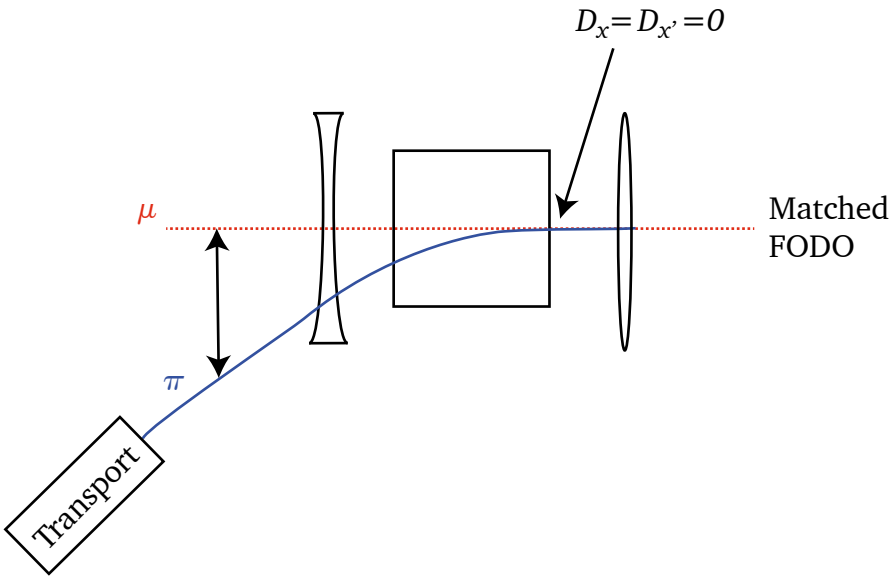

Figure 4.7: Stochastic injection, as understood by accelerator physicists. The beam is dispersion matched (i.e., $D_x = D_{x'} = 0$ for both beams), where dispersion is roughly dp/dx.

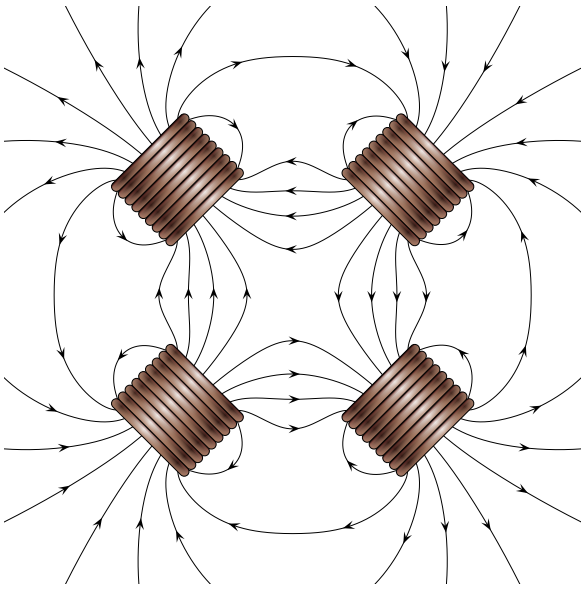

Figure 4.8: Quadrupole field. Note that a particle off-center will see a dipole field.

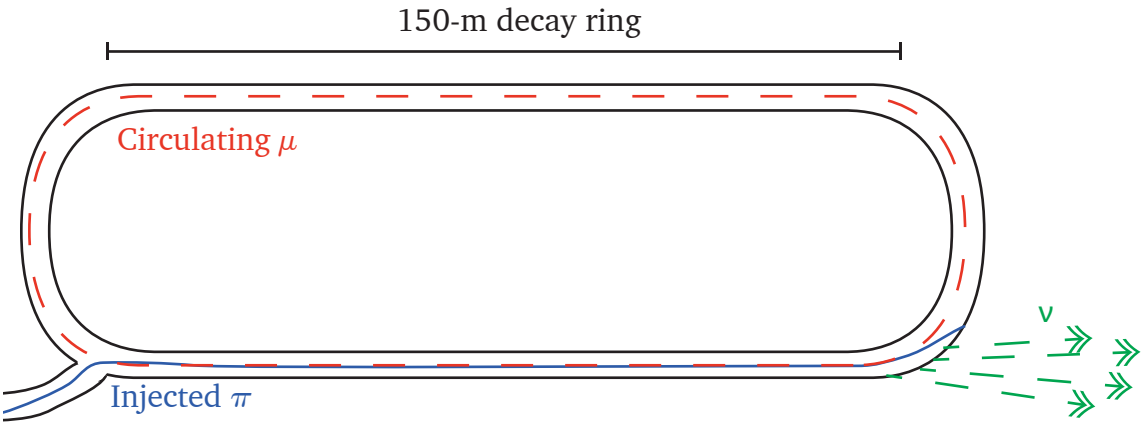

Figure 4.9: A schematic of the decay ring. Pions are injected in the bottom left of the figure and decay in the bottom straight. Muons from these decays circulate in the ring to produce neutrinos.

have different momenta. Some of the muons from pion decay will be within the dynamic aperture of the machine.

Injection can be designed to have pions decay toward or away from the far detector. The favored option is to inject into the accelerator portion that results in pions decaying such that neutrinos from this decay can be observed at the far detector—called the *pion-decay flash*. These pion-decay neutrinos may provide interesting physics, do not result in dead time (<1%) whereby pion-decay neutrinos interact at the same time as muon-decay neutrinos, and only marginally complicate the far detector analysis.

Simulation work has been performed to demonstrate the efficiency of various injection schemes [64, 65]. Given these results, it is conservative to assume an injection efficiency of 90%. Further optimizations are still ongoing.

4.2.5 Decay ring

A racetrack ring has two long, straight sections (Fig. 4.9). Circulating muons will decay, producing neutrinos along both straights. One straight points toward the detector and is called the *decay straight*. The key difference between a muon-decay straight and a pion-

decay pipe—how accelerator neutrinos are normally made—is that muons which arrive at the end of the straight without decaying are "recycled" instead of lost. Dipoles steer the beam back to the beginning of the straight to allow the muons another possibility to decay toward the detector. This is possible because muons ($c\tau = 658$ m) are longer lived than pions ($c\tau = 8$ m). Recycling the beam is an efficient way to get bright neutrino beams.

The accelerator optimization depends upon three different aspects of the design. First, the number of circulating muons in the ring must be maximized. Due to pion decay kinematics, the muon beam is larger than any reasonable ring dynamic aperture. Therefore, more muons are accepted when the magnet apertures are increased. Once the beam is within the dynamic aperture, it is transported until it decays because the beam is too short lived to hit a resonance. Next, the ratio of the decay straight length to the circumference determines how many muons decay toward the detectors; therefore, this must be maximized. The final constraint is cost. To increase the statistics in a detector, either the accelerator can be improved to deliver more muon decays or the detector can be made more massive. However, if the facility will serve many experiments, then it is better to spend money on the accelerator. Costing exercises are ongoing and are not discussed further. (An optimization is presented in Appendix A.)

Designs exist for both a classic FODO (Fig. 4.10) and a scaling FFAG lattice, where a FODO lattice uses mainly quadrupoles for vertical stability and dipoles in bends, while a FFAG uses a more complicated field to ensure the same bending for many muon momenta. The most interesting feature of the FODO design is the asymmetric β function, which was an idea conceived for this thesis and implemented by Bogacz [66]. The benefit is that magnet costs are minimized in the straight that does not aim beam at the detectors (beam size $\sigma_x \sim \sqrt{\beta} \sim$ magnet size); however, the beam divergence, $\sigma_{x'} \sim 1/\sqrt{\beta}$, is also minimized. By having small divergence, the beam is aimed at the detector. Ideally, the beam divergence should not affect the flux predictions. To accomplish this, $\sigma_{x'}$ should

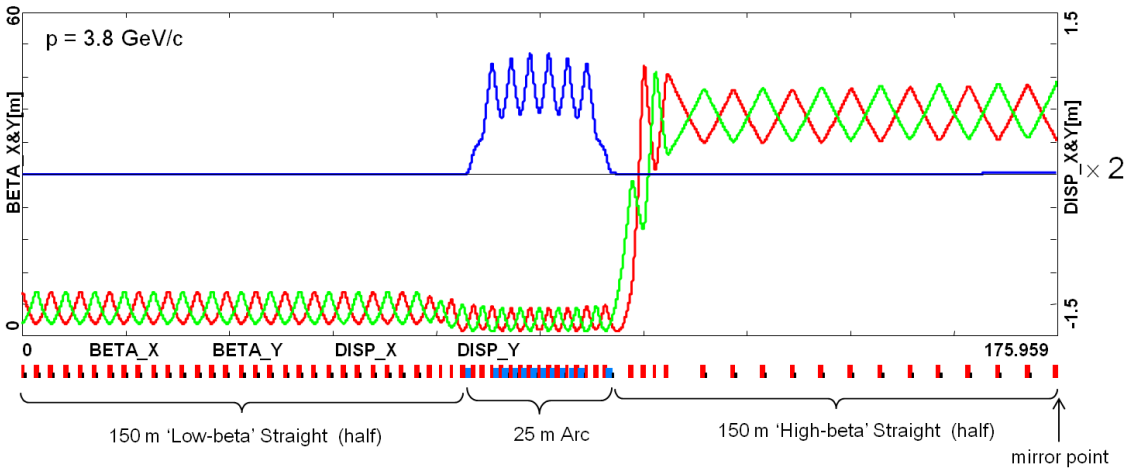

Figure 4.10: The decay-ring lattice. The β functions are shown for a FODO, which means uses focusing (F) and defocusing (D) quadrupoles, decay-ring design with small beam size β ($\beta = \sqrt{\sigma_x \epsilon}$) in the straight that is not the decay straight [50, 66].

be less than the "natural" divergence of the beam $1/\gamma$ arising from the decay kinematics. The geometric emittance is 2 mm; therefore, $\beta = 40$ m results in $\sigma_{x'} = 7$ mrad, where the divergence (i.e., opening angle) is $1/\gamma = 30$ mrad for a 3-GeV muon.

The main distinction between the FODO and FFAG accelerators is their longitudinal momentum acceptances: 10% and 16% for the FODO and FFAG, respectively. Maximizing the longitudinal acceptance is the main design consideration for increasing the number of useful muons due to the broad energy spectrum of muons that could potentially be transported.

Different lattice designs achieve large momentum acceptance in different ways. FFAGs are intended for high-current machines that have bunches of different energies circulating simultaneously (e.g., neutron spallation sources and muon accelerators). Specific to scaling FFAGs is that the optics are the same regardless of energy, allowing for tuned lattices at a wide range of reference particle momenta. It is possible to get $p_z \pm 16\%$ for νSTORM [50]. A FODO lattice requires larger magnets to increase the aperture, which results in higher magnet costs since costs scale with magnetic volume [67]. This can be partially mitigated by making β small in the nondecay straight. The details

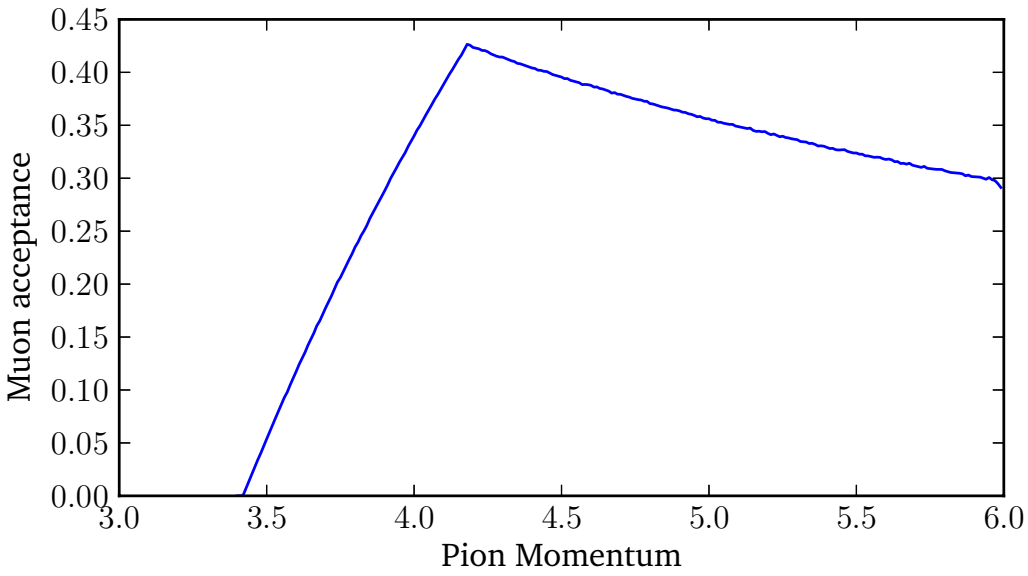

Figure 4.11: Number of muons from the pion decays within the decay-ring aperture of p_z between 3.42 GeV and 4.18 GeV.

of the design implementation are beyond the scope of this thesis and can be seen in the LOI [50].

The probability of the pion decaying depends on the straight length and pion energy. The ring has a 150-m straight section. The $c\tau\gamma$ of a pion with momentum 3 GeV/c is 280 m; therefore, the probability of a pion decaying in a 150-m straight is $1 - \exp[-150/280] = 40\%$, which forms an upper limit on the number of resulting muons.

The pion decay kinematics must be taken into account because only 3.8 ± 0.38 GeV/c muons are transported in the ring. The muon acceptance is the fraction of muons from pion decay that are within the dynamic aperture of the ring. A Monte Carlo (MC) was written for this thesis to compute the acceptance by sampling the phase space of the

resulting muons in the pion rest frame, then boosting to the laboratory frame. Only 40% of the pions decay. Only 35% of the pion decays result in a muon within the longitudinal acceptance (Fig. 4.11). Lastly, only half of the muons within the longitudinal acceptance are within the dynamic aperture. Thus the combined probability that a decaying pion results in a useful muon (μ per π) is about 7%.

Studies were performed [50] for both styles of ring (FFAG and FODO) to determine the efficiency of a muon within the dynamic aperture, A_{dynamic}, being transported around the ring. Particles within the dynamic aperture can be lost due to nonlinearities in the lattice. Every study has shown a greater than 90% efficiency after 100 turns for muons ignoring decays; 90% will be taken as a lower limit.

Only muons that decay toward the far detector result in neutrinos that may be used for oscillation physics. The decay length of the muon is much longer than the size of the accelerator; therefore the decay probability distribution is uniform throughout the ring. The ratio Ω of the decay straight length to the ring circumference must be maximized to increase the neutrino flux. The arc lengths are fixed and determined by the dipole magnet fields. As an estimate, assume that the field is 2 T and, with $p[\text{MeV}] = 300Br[\text{T-m}]$, where B is the magnetic field and r the bending radius, then $2\pi r \approx 40$ m. The two straights are 150 m each; therefore, the circumference is 350 m. The fraction of useful muons is $\Omega = 43\%$. The diminishing returns on the investment of buying more magnets results in it being inefficient to increase the straight length indefinitely.

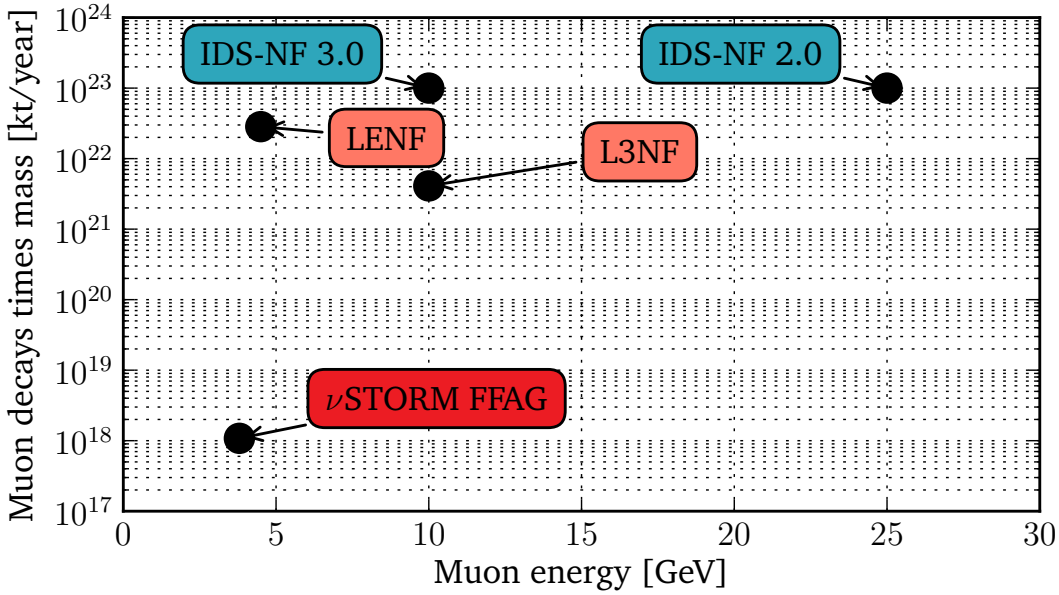

Figure 4.12: Comparison of νSTORM to other muon-decay facilities. In terms of muon-decay multiplied by the mass of the detector, νSTORM would have 10^5 times less exposure, but is at a much shorter baseline. Data points reflect the information in [50,57,68].

4.3 Number of useful muon decays

Under the assumptions discussed, the number of useful muon decays can be determined:

$$N_\mu = (\text{POT})(\pi \text{ per POT})\epsilon_{\text{collection}}\epsilon_{\text{transport}}\epsilon_{\text{injection}}(\mu \text{ per } \pi)A_{\text{dynamic}}\Omega \quad (4.3)$$

$$= 10^{21} \times 0.09 \times 0.9 \times 1 \times 0.9 \times 0.07 \times 0.9 \times 0.375 \quad (4.4)$$

$$= 1.8 \times 10^{18} \text{ good muon decays.} \quad (4.5)$$

νSTORM has a smaller exposure (i.e., muon decays multiplied by the mass of the detector) than other NF facilities (Fig. 4.12). As a comparison, the standard neutrino-factory design (IDS-NF 2.0) has 10^{21} muons per year. The standard neutrino factory and νSTORM differ by 5 orders of magnitude in exposure: approximately 2 from detector size, 1.5 from collection and cooling, and 1.5 from targeting. These are significantly different facilities in terms of their requirements.

4.3.1 Beam instrumentation

One motivation to use muon-decay beams is that they should have smaller flux uncertainties than pion-decay beams. The accelerator neutrino beam that discovered the muon neutrino had 30% flux uncertainties [51] that did not affect the experimental results. Over the past 50 years, the same technique has been used to produce all accelerator neutrino beams. The main differences are the number of POT and that the inclusion of horns has helped beam brilliances. In this time, flux uncertainties have dropped to only 15%. These uncertainties are mainly dominated by hadron production uncertainties, despite the efforts of HARP [69], MIPP [70], and NA61/SHINE (and predecessors) [71].

The claimed flux uncertainty for the 25-GeV/c neutrino factory is 0.1% from beam instrumentation [57] but has never been demonstrated experimentally. The beam instrumentation accomplishes two goals beyond accelerator operations: First, the beam

instrumentation can be used to normalize the oscillation physics analyses by precisely determining the flux. But second, νSTORM will be the first practical demonstration that neutrino factory beam instrumentation can actually measure the beam to the precision claimed by simulation.

Various beam instrumentation systems are being considered. The two most important are beam-current transformerss (BCTs) and polarimeters, which measure the number of circulating muons and their energy spectrum, respectively. The beam current being considered is about 10 μA, which is small compared to typical electron or proton storage rings. However, techniques that are used for measuring beam currents during slow proton-beam extractions can be applied to the νSTORM decay ring. Another important feature of the beam is that it has the bunch structure of the MI, which—even though there is no RF—does not decohere since particles of different energy are still going the same speed. However, measurements, similar to what has been achieved in the MI and at J-PARC, should be possible, although more simulation is needed [72].

A polarimeter to measure the energy spectrum is also envisioned. A design exists that has been thoroughly studied and simulated within the IDS NF effort. (See Sec. 6.1.1 for a detailed discussion of polarization.) The energy must be well known for the oscillation physics, because neutrino interaction rates are proportional to E_μ^3—E from the cross section times E^2 from the opening angle. A thorough systematic study has not yet been performed, but the energy uncertainties are expected to be below 1% based on the arguments in [57] which rely upon experience with existing (different) accelerators. However, if the energy distribution is not well known from accelerator instrumentation, then the near detector can measure it.

The concept behind the polarimeter is that muons are polarized by the V-A interaction in pion decay—the muons need the "wrong" helicity state—and this polarization can be measured with electrons from the subsequent muon decay. If the time evolution of the polarization is well known, then the energy is also known from the physics of Thomas

Figure 4.13: D0 assembly building. This building serves as an ideal location for the far detector.

precession. The polarimeter is a scintillator detector downstream of a dipole magnet: electrons will have less momentum since they are decay products and will bend into the polarimeter, while muons will continue on down the beam pipe.

4.4 Detectors

At least two detectors are envisioned: a near detector at 50 m and a far detector at 2000 m from closet point in the decay-ring straight. The near detector will be in a large hall to serve future experiments and detectors with a well-understood neutrino test beam. For example, experiments from the long-baseline neutrino experiment (LBNE) and long-baseline neutrino observatory (LBNO) collaborations could make ancillary measurements at this test beam to aid their future programs. These future long-baseline experiments are unable to measure their electron neutrino response in any conventional test beam, which also results in electron-neutrino cross-section uncertainties. The far detector at 2000 m could be situated on the surface in the D0 assembly building (DAB) (Fig. 4.13) in the Fermilab siting plan (Fig. 4.2).

As mentioned before, muon-decay neutrino-beam experiments require magnetization to perform appearance measurements. The favored detector technology is a sampling iron-scintillator calorimeter, similar to the MINOS/CDHS designs. Engineering such a detector is straightforward and has been performed [50] (Fig. 4.14 and Fig. 4.15).

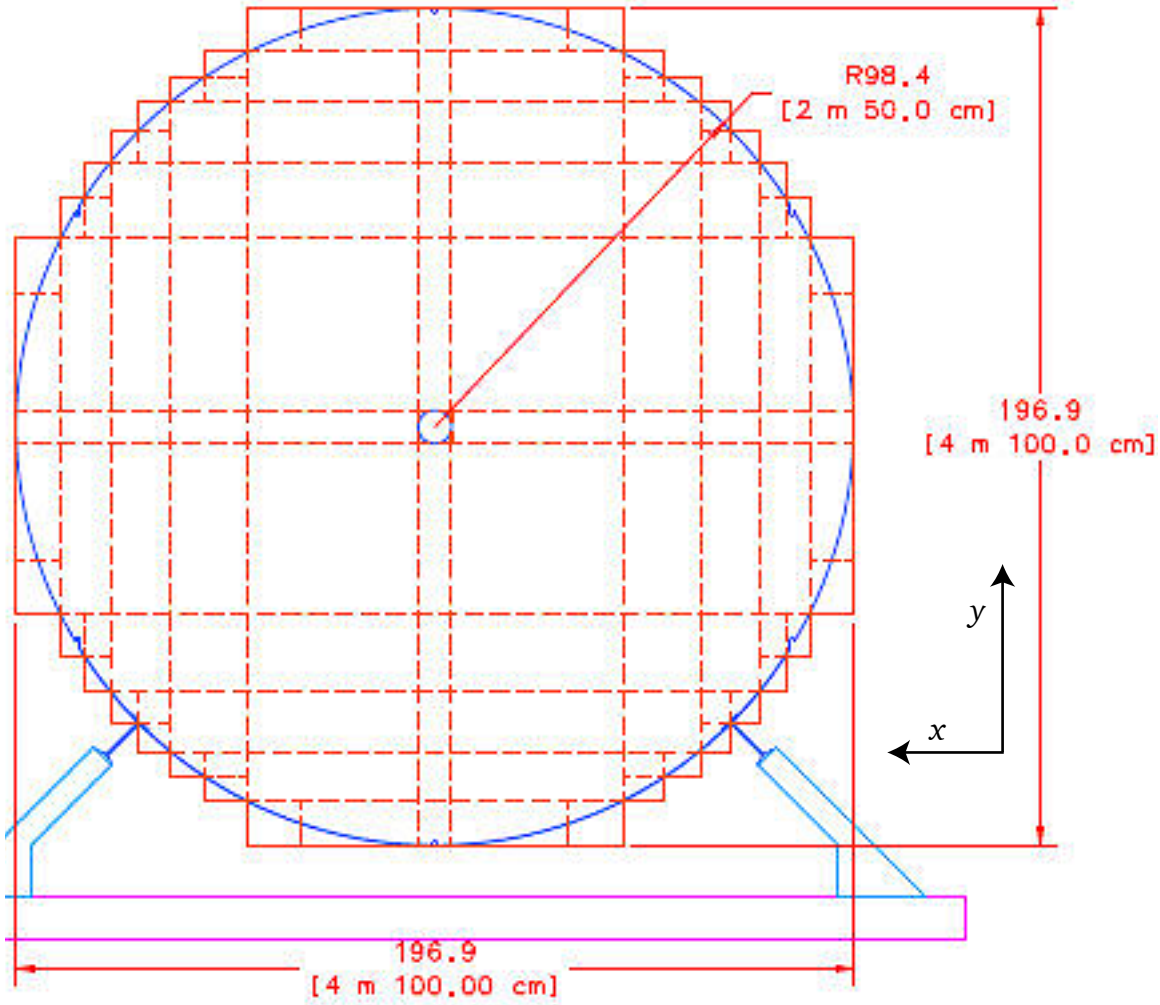

Figure 4.14: Engineering diagram of far detector as viewed by the beam [50].

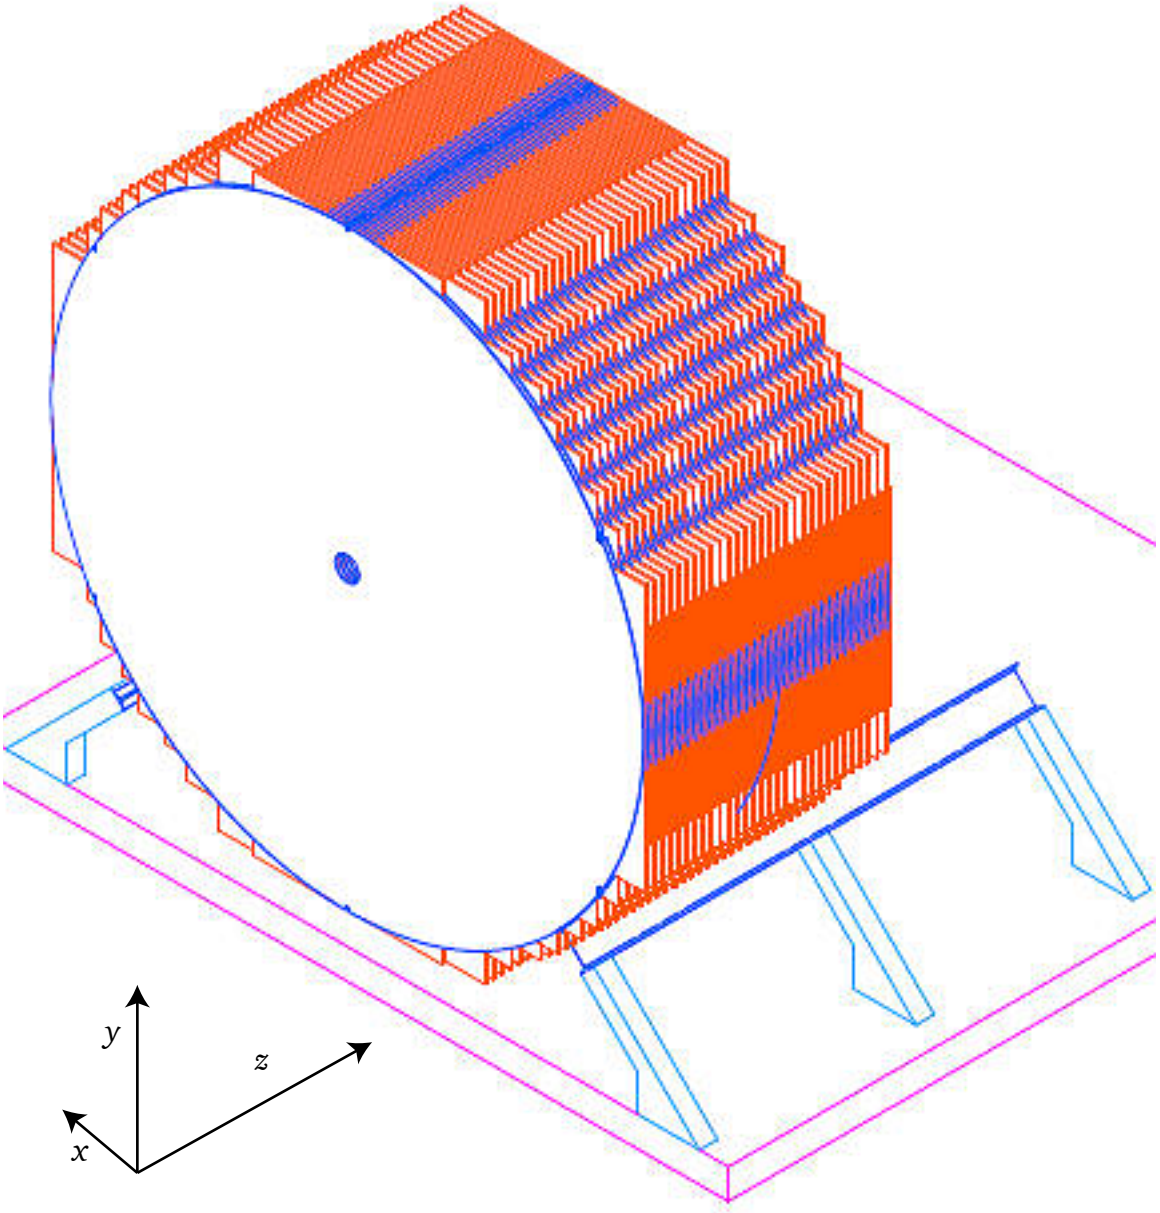

Figure 4.15: Engineering of far detector as viewed off-axis [50].

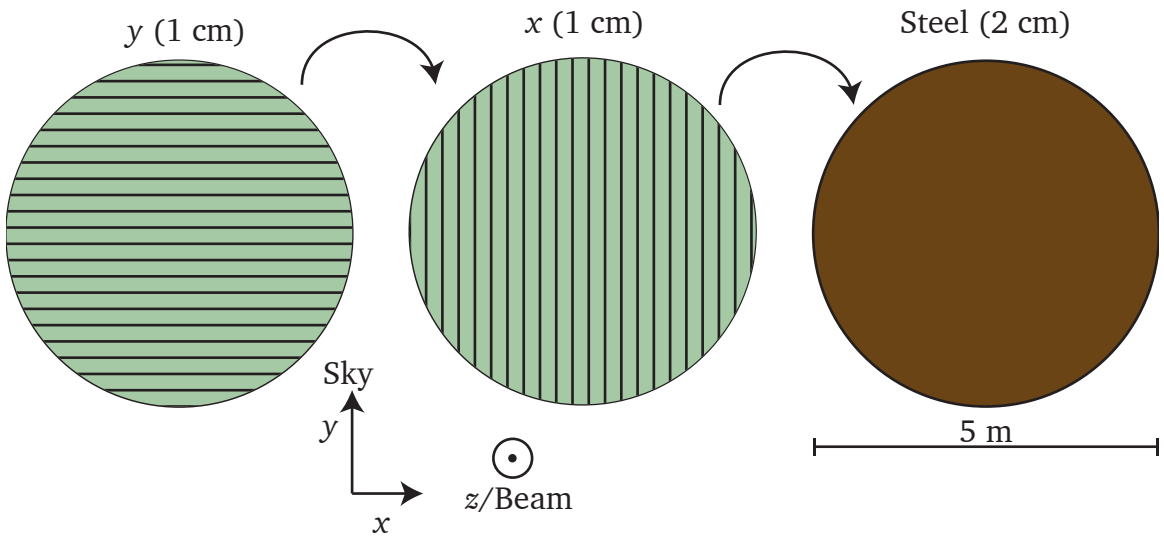

Figure 4.16: Layer assembly. A y view is placed on top of an x view, which is stacked on top of a steel sheet.

These calorimeters are built from alternating circular sheets of steel and scintillator (shown before in Fig. 3.1 and Fig. 3.2). The steel plates will be 2 cm thick, have a 5 m diameter, and be skip welded from two half-moon steel pieces. By having a simpler plate construction and welding than MINOS, the field uncertainties should be smaller.

A *view* is formed by stacked 1 cm^2 scintillator bars. Two views measure the position transverse to the z beam-direction axis, where the y axis is local vertical and the x axis is defined so that x, y, and z form a right-handed coordinate system. The y view is mounted to the x view, which is supported by the steel plane (Fig. 4.16). The two views and steel plane form a *layer*. By having the x and y views adjacent, it is possible to form 3 dimensional space points if a muon traverses the layer.

The active regions consist of scintillator bars with embedded wavelength shifting (WLS) fibers (Fig. 4.17), which is a standard detector technology in particle physics. Scintillation light is capture by a WLS fiber where the photon is reemitted with a different wavelength which has a longer attenuation length. The scintillator bars have a cross section of 1 × 1 cm and consist of extruded polystyrene doped with blue-emitting

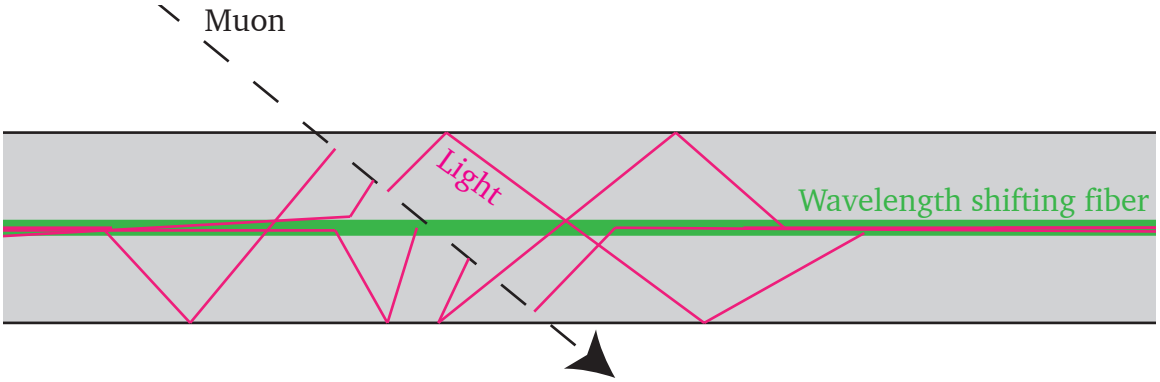

Figure 4.17: Wavelength shifting fiber within a scintillating bar. When a muon traverses the bar (dotted line), it deposits energy that is converted into light (pink). This light reflects within the bar until it hits the fiber (green), after which the light total internally reflects within the fiber. Most reflection is spectral (i.e., angle in equals the angle out).

fluorescent compounds, a coextruded TiO_2 outer layer for reflectivity, and a hole for the WLS fiber. The WLS fiber uses PPO and PPOP as the wavelength shifter. This technique has been used by numerous experiments. (See [50] for more scintillator details.)

The light can be detected by silicon photomultipliers (SiPMs). (See [73] for review of SiPMs.) These are now a proven technology that is being used at experiments such as T2K [74]. SiPMs consist of a 2D array of avalanche photodiodes (APDs) working in Geiger mode on a single silicon chip. Each APD is sensitive to single photons, but by connecting the APDs in parallel it is possible to measure the number of incident photons, assuming that the APD density is high enough (i.e., $N_{\text{photon}} < N_{\text{pixel}}$). The size of these devices is 20–100 μm, so a WLS fiber can be glued directly to it. APDs provide adequate gains and quantum efficiencies ($\sim$80%) for particle physics purposes. Compared to PMTs, they are cheaper, do not age, and work in magnetic fields. However, other technologies could be used instead. The only requirement for the appearance analysis is that it is known when bars are traversed by a minimum ionizing particle.

The readout is still to be decided, but there are many possibilities. The main consideration is price per channel rather than physics. The most expensive option would be to design an application-specific integrated circuit (ASIC), but it may be less effort—and

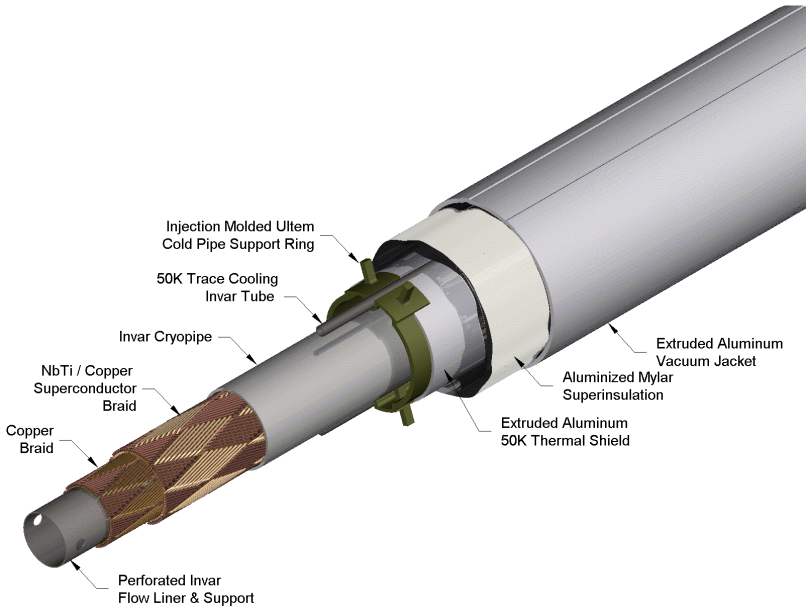

Figure 4.18: The superconducting transmission line used for detector magnetization [75].

cheaper—to use a preexisting design from another experiment or industry. Since these are just silicon chips, they are cheap to produce in bulk but expensive to design.

A STL is used to magnetize the detector in order to saturate the steel ($B \approx 2$ T). As seen in Fig. 4.19, the steel is saturated out to the edges of the plates, unlike the MINOS ND which has a field of about 1 T in their analysis region. The STL was developed for the design study of the VLHC, and a 17-m functioning prototype was built at Fermilab [75] that carried 100 kA. NbTi superconductor and copper braiding is wrapped around an Invar cylinder (Fig. 4.18). The 2.5-cm bore of the Invar carries superfluid helium at 4.5–6 K. The helium cools the superconductor through the Invar. A small coaxial cryostat surrounds the superconductor. The advantage to having a small conducting cable to magnetize the steel detector planes is that it can be fed through a hole in the detector planes (Fig. 4.15 and Fig. 3.1). In MINOS—which has a larger regular conducting cable to magnetize the steel—muons sometimes traverse the cable, which leads to reconstruction uncertainties. Having a small coil reduces this source of uncertainty.

A total current of 250 kA is assumed to flow through the STL at the center of the

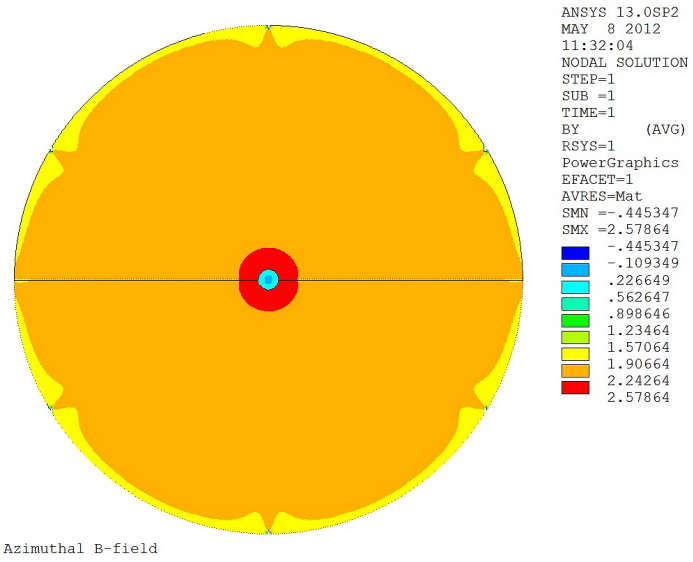

(a) *ν*STORM.

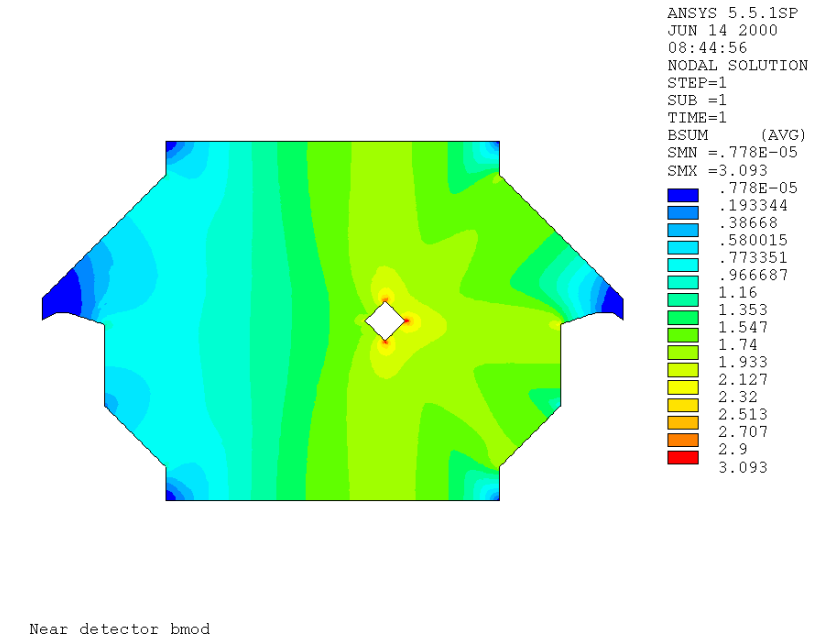

(b) MINOS ND [45].

Figure 4.19: Magnetic field in *ν*STORM compared to the MINOS ND, where the color scale is in units of tesla. For *ν*STORM, the steel is saturated throughout most of the steal at 2 T. However, for the MINOS ND, the field varies from 1 to 2 T throughout their plates.

detector, which based on finite element analyses saturates the steel [50]. The STL could reasonably carry 50 kA per turn, requiring five turns for saturation.

As an aside, additional detectors—beyond the sterile neutrino program—could be added to further utilize the beam. Given the cost of construction the νSTORM accelerator, is should be expected that there will be numerous other experiments operating in the same beam: νSTORM could be viewed as a user facility, much like the light source Diamond. For example, cross section measurements could be performed with this well-understood beam, as well as the first ν_e cross section measurements. Also, detector prototyping could be in a ν_e beam—the typical signal for long-baseline experiments. Lastly, there are all the experiments that have not been thought of that may want the next generation of neutrino beam. However, all of these applications are beyond the scope of the thesis question: "are there sterile neutrinos?"

4.5 Conclusion

In this chapter, the high-level design and engineering of νSTORM was described. The accelerator forms the majority of the facility but requires little technical R&D. The stochastic injection, FFAG, and beam instrumentation provide interesting accelerator physics. For the detectors, a classic design is used for the oscillation physics program, while leaving room for other detectors to take advantage of the beam line.

Chapter 5

Detector performance

> Science is what we understand well enough to explain to a computer. Art is everything else we do.
>
> Donald Knuth

In this chapter, I present the MC that I wrote from scratch and also the cuts I developed for the appearance analysis.

Up to now, only back-of-the-envelope calculations have been used to design νSTORM (Chap. 3); more thorough calculations are now presented to confirm that νSTORM is able to resolve the LSND anomaly at a suitable level of confidence. After a reminder of the far detector technology choice, a way of extracting the appearance signal CC interactions $\nu_e \to \nu_\mu$—the CPT conjugate of the LSND channel—from the data is discussed. The physically motivated cuts that were developed earlier are extended to depend only upon measurable quantities such as reconstructed track length. A GEANT4 MC was written to demonstrate that this small muon-appearance signal is observable at 3 GeV. The work in this chapter summarizes the appearance analysis, MC design, and detector performance.

The detector must be magnetized to differentiate "right-sign" and "wrong-sign" muons resulting from the channels $\bar{\nu}_\mu \to \bar{\nu}_\mu$ and $\nu_e \to \nu_\mu$, respectively (Table 3.1). As was dis-

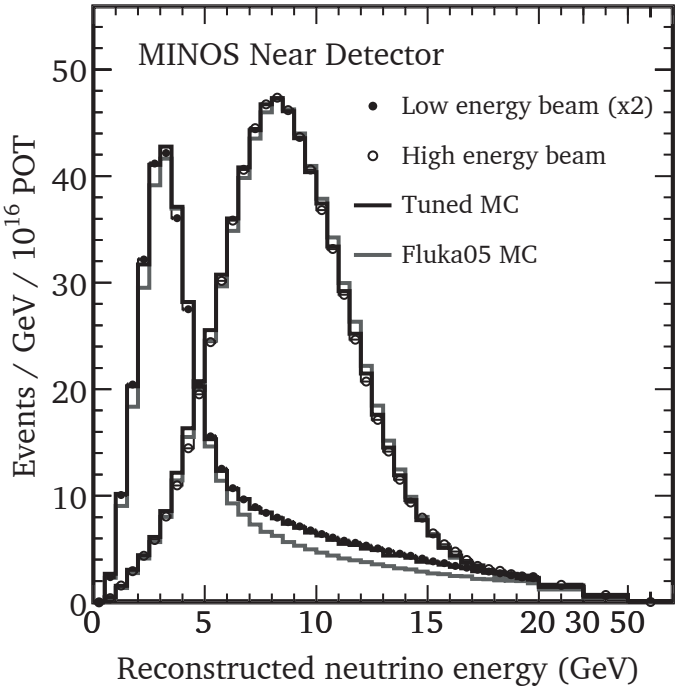

Figure 5.1: The MINOS neutrino energy spectrum for neutrino-mode runs. [76]

cussed in Chap. 3, magnetized iron sampling calorimeters—also called magnetized iron neutrino detectors (MINDs)—are the chosen detector technology. MINDs were invented in the 1970s and are a well-proven technology (e.g., CDHS, MINOS experiments) for identifying muon final states, and are the only economic way of building large magnetized detectors. The steel acts as both a muon ranger and spectrometer. (Brief technical information is discussed in Sec. 4.4, or [50] and references therein.)

MINDs were originally designed for DIS studies in the energy range $E_\nu = 30$–70 GeV [77], but MINOS used a MIND at energies less than 10 GeV. Specifically, the MINOS near-detector energy spectrum in the low-energy configuration is peaked at 3 GeV (Fig. 5.1) and similar to what is expected for νSTORM; therefore, MINOS has already demonstrated the validity of particle interaction models in the νSTORM energy range—more on this is discussed later when the MINOS calibration detector is discussed in the context of systematics.

The biggest difference between νSTORM and MINOS is the analysis. The analy-

Table 5.1: Channels, interactions, and cuts for the appearance analysis. The oscillation channels $\overset{(-)}{\nu}_e \to \overset{(-)}{\nu}_e$, $\overset{(-)}{\nu}_\mu \to \overset{(-)}{\nu}_\mu$, and $\overset{(-)}{\nu}_e \leftrightarrow \overset{(-)}{\nu}_\mu$ (i.e. all channels below τ threshold) can be explored. Both NC and CC interaction types are presented.

Decaying particle	Channel	Interaction	Cut
μ^+	$\nu_e \to \nu_\mu$	CC	*(Signal; do not cut)*
	$\bar{\nu}_\mu \to \bar{\nu}_\mu$	CC	Curvature
	$\bar{\nu}_\mu \to \bar{\nu}_\mu$	NC	Range
	$\nu_e \to \nu_e$	CC/NC	
	$\bar{\nu}_\mu \to \bar{\nu}_e$	CC/NC	
	$\nu_e \to \nu_\mu$	NC	
π^+	$\nu_\mu \to (\nu_\mu$ or $\nu_e)$	CC/NC	Timing
μ^-	$\bar{\nu}_e \to (\bar{\nu}_e$ or $\bar{\nu}_\mu)$	CC/NC	Known from magnet polarity
	$\nu_\mu \to (\nu_\mu$ or $\nu_e)$	CC/NC	
π^-	$\bar{\nu}_\mu \to (\bar{\nu}_\mu$ or $\bar{\nu}_e)$	CC/NC	

sis that MINOS was designed for involved a ν_μ disappearance spectral fit to measure atmospheric-neutrino parameters. MINOS does do appearance searches, but only electron appearance [78] to attempt to measure θ_{13}. However, unlike MINOS, νSTORM is designed for muon appearance.

5.1 ν_μ CC appearance requirements

The wrong-sign muon appearance signal needs to be separated from seven other event topologies: having two initial neutrinos from muon decay, with two possible final states, and two interaction types results in $2 \times 2 \times 2 = 8$ event topologies (Table 5.1). All of these other channels need to be separated from the appearance channel CC $\nu_e \to \nu_\mu$. Existing work on isolating CC ν_μ appearance events—the so-called *golden channel*—has been studied extensively for NF R&D studies [57], and can be adapted to νSTORM. As an aside, it is interesting to note that the pion-decay neutrinos in νSTORM are not typical of NFs, and could possibly be used as, e.g., a charge-ID calibration.

Accelerator operators know the injected beam polarity since the horn current can be changed to only focus either positive or negative particles; therefore, neutrinos from

π^- and μ^- decay—remember, the signal comes from μ^+ decay—can be ignored in this analysis. The polarity of the horn, chicane, and injection magnets will be known because it is the polarity of a power supply. The beam polarity is known at any given time; therefore, only half of the event topologies in Table 5.1 need to be separated from the CC $\nu_e \to \nu_\mu$ signal.

In positive polarity running, π^+s are injected into the ring and decay to μ^+. These decays only occur within one straight of the ring, and within 500 ns of injection. The pion momentum is 5 ± 0.5 GeV, where the tails are cut off by the beam line magnets. Consequently, no pions have momenta within 3.8 ± 0.38 GeV; collimators can be added to ensure this is true. The decay ring is designed to only transport 3.8 ± 0.38 GeV. No pions will be transported around the first dipole in the decay ring, and even if they did, the $\gamma c\tau$ for a 5-GeV pion is 280 m; therefore pions travel less than one turn. Muons on average decay in ≈100 turns, and cutting the first turn will reject all pion backgrounds. A timing cut of 500 ns is technically feasible and will reject all π^+ backgrounds thereby leaving only neutrinos from muon decay.

The channels that remain are from μ^+ decay, and cuts need to be designed to isolate the appearance signal. For example, the final-state muon will penetrate more steel than NC events or electrons. All nonmuon final states can be removed with a range cut. Differences between the ν_μ appearance signal and the $\bar{\nu}_\mu$ background include muon charge, energy distributions, and whether the Michel electron annihilates. In the following analysis, bending in a magnetic field is used to determine the charge.

Oscillation parameters need to be assumed when designing the cuts; however, the cut must be effective for a wide range of oscillation parameters. Cuts are designed assuming that the oscillation probabilities are in the "averaging regime" of large Δm^2 [see Eq. (2.16) to (2.21) on p.p. 32], which results in the oscillation probability being independent of energy. The normalization is determined at a later step, which removes any dependence on $\sin^2(2\theta)$. Therefore, for disappearance signals, the cuts are tuned

on an unoscillated interaction spectrum. For the appearance signal, tuning is performed assuming that all ν_e oscillate into ν_μ, so the signal ν_μ flux is identical to the ν_e spectrum from muon decay. Depending on the value of Δm^2, the shape of these distributions can change, which is why the cuts are studied later for a wide range of possible neutrino parameters.

Before going into the details of the appearance analysis, it is worth making a side point; it is possible to be more frugal with the neutrinos than wasting those from π^+ decay. By analyzing data up to 500 ns after injection stops, it is possible to use the π-decay neutrinos for different physics analyses. Primarily, the muon disappearance signal could be used as a calibration since there are as many neutrinos from pion decay as muon decay. The current favored injection option is to inject into the straight aimed at the far detector, which will result in a "pure" muon-neutrino beam before the muon-decay beam. If these pion decays provide difficulty for the muon-decay analysis, they can be avoided by injecting the beam into the other decay straight.

In addition to calibrations, appearance physics may be possible using these pion decays. Probing the electron appearance signal—à la LSND—from pion decay is difficult; the oscillation probability is about 10^{-3}, but there is a 1.2% background of electrons from muons decay since the accelerator straight is long. For 5 years exposure with a kilotonne detector, the unoscillated interactions have rates of about 100,000 events, resulting in $\sqrt{1.2\% \times 100000} = 31$ background event. There would be 300 expected electron-neutrino appearance interactions. Though possible, a special purpose liquid argon detector would be required to ensure discrimination between $\overset{(-)}{\nu}_e$ and NC events. Therefore, it is expected that these pion decays could be used to calibrate the muon charge-misidentification and energy scale.

Table 5.2: Simulation parameters within GNOMON.

Event generator	Code	GENIE
	Particles	$\overset{(-)}{\nu}_e$ and $\overset{(-)}{\nu}_\mu$
	Position	Uniform in (x, y, z)
	Energy	Uniform in $E \in [0, 5]$ GeV
Particle propagation	Code	GEANT4
	Physics models	QGSP_BERT
Magnetic Field	Model	2D finite element analysis
	Approximate field	2 T throughout steel
Geometry	Bar dimensions	5 m × 1 cm × 1 cm
	Steel dimension	5 m × 5 m × 2 cm
	Layer	Steel, x, y, then repeat
	Layer dimension	5 m × 5 m × 5 cm
	Layers	444
	Bars per view	500
	Steel	$I_e = 286$ eV ^{26}Fe g/mol = 58.85 Density = 7.87 g/cm^3
	Scint.	$I_e \simeq 64$ eV 91% ^{6}C , 9% ^{1}H g/mol = 12.01, 1.01, respectively Density = 1.06 g/cm^3
Electronics	Code	GNOMON
	Noise threshold	2 pe
	Energy scale	20 pe/(MeV deposited)

5.2 The Monte Carlo simulation

The initial design of νSTORM came from analytic arguments; however, code has been developed to confirm the detector performance estimates and study subtle effects. Code was developed based on the GEANT4 [79, 80] library to model particle transport and interactions. Also, a wrapper around GEANT4 was written called GNOMON that maintains the geometry, simulates the electronics (by photon statistics and attenuation), and performs reconstruction. The code was made open source, and further information about implementation details as well as a manual are available online [81].

The simulation parameters used for this analysis are shown in Table 5.2. Neutrino

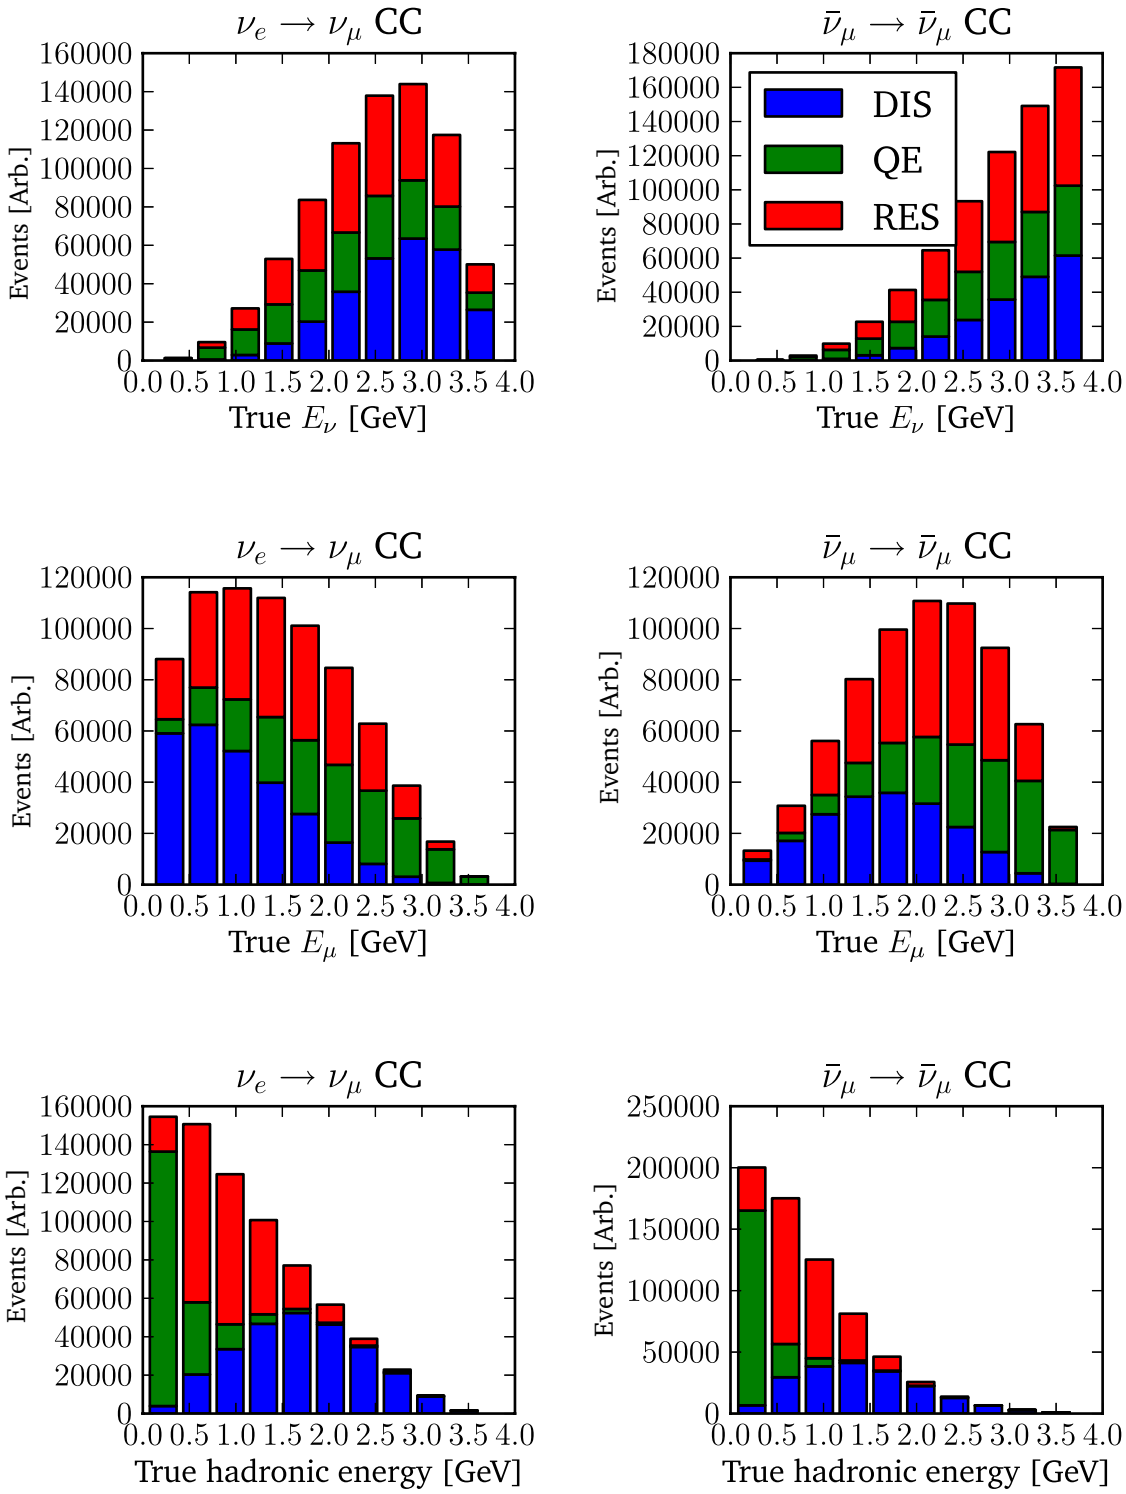

Figure 5.2: Kinematic distributions of E_ν, E_μ, and hadronic energy for averaging-regime ν_e and $\bar{\nu}_\mu$ fluxes. The stacked histograms show the DIS (blue), quasielastic (QE) (green), and resonant production (RES) (red) cross section components. The normalization corresponds to one million simulated events.

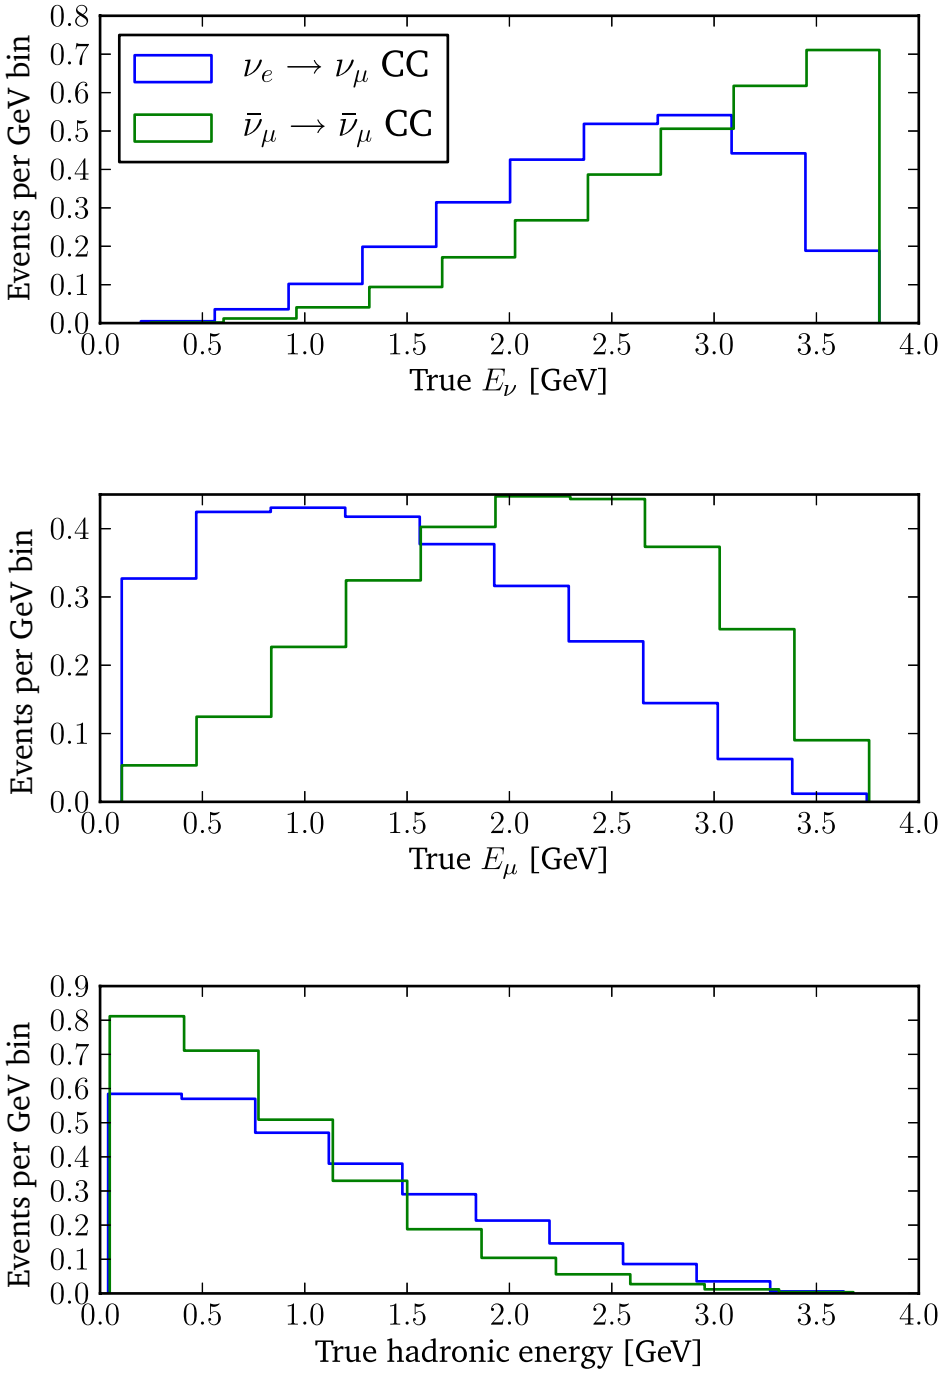

Figure 5.3: Overlaid kinematic distributions of E_ν, E_μ, and hadronic energy for averaging-regime ν_e and $\bar{\nu}_\mu$ fluxes. The distributions are normalized to unity.

interactions were generated by GENIE [82], distributed uniformly throughout the detector volume, and propagated with GEANT4 through the geometry, which was shown in Fig. 4.14 and 4.16. An example distribution of E_ν, E_μ, and the hadronic energy ($= E_\nu - E_\mu$) is shown in Fig. 5.2 and 5.3. The kinematic distributions are influenced by the QE, DIS, and RES cross sections—despite not being shown since insignificant, coherent single-pion production and neutrino–electron scattering are simulated. The electronics response within GNOMON simulates the expected SiPM performance by assuming 20 photoelectrons per MeV deposited in the bar. The details of the electronics has not yet been determined, yet is not required for this basic simulation.

In the next few sections, there are plots created with GNOMON that demonstrate the cut performance. Later in the chapter, this study is generalized such that the cut efficiency is determined as a function of neutrino energy, thereby allowing arbitrary fluxes to be used in the sensitivity analysis in the next chapter.

5.3 Range cut

As was shown in Chap. 3, a range cut is capable of separating NC and ν_e CC events from $\overset{(-)}{\nu}_\mu$ CC events. This range cut relied on hadron showers not traversing as much steel as muons; however, it used *true* path lengths. This cut must be redefined in terms of measurable path lengths from digitized MC hits.

Following previous experiments like MINOS, NuTeV, and CDHS, the range is defined in terms of the number of hit layers. These hit planes need not be contiguous, though imposing continuity requirements does not negatively impact the cut performance.

A range cut was implemented in GNOMON by counting the number of hit layers. Fig. 5.4 shows the efficiency for signal and background versus number of layers. More than 100 layers must be penetrated to achieve a NC/ν_e-CC background of $<10^{-5}$. After the 100-layer cut, the signal efficiency is 20% while the background efficiencies are

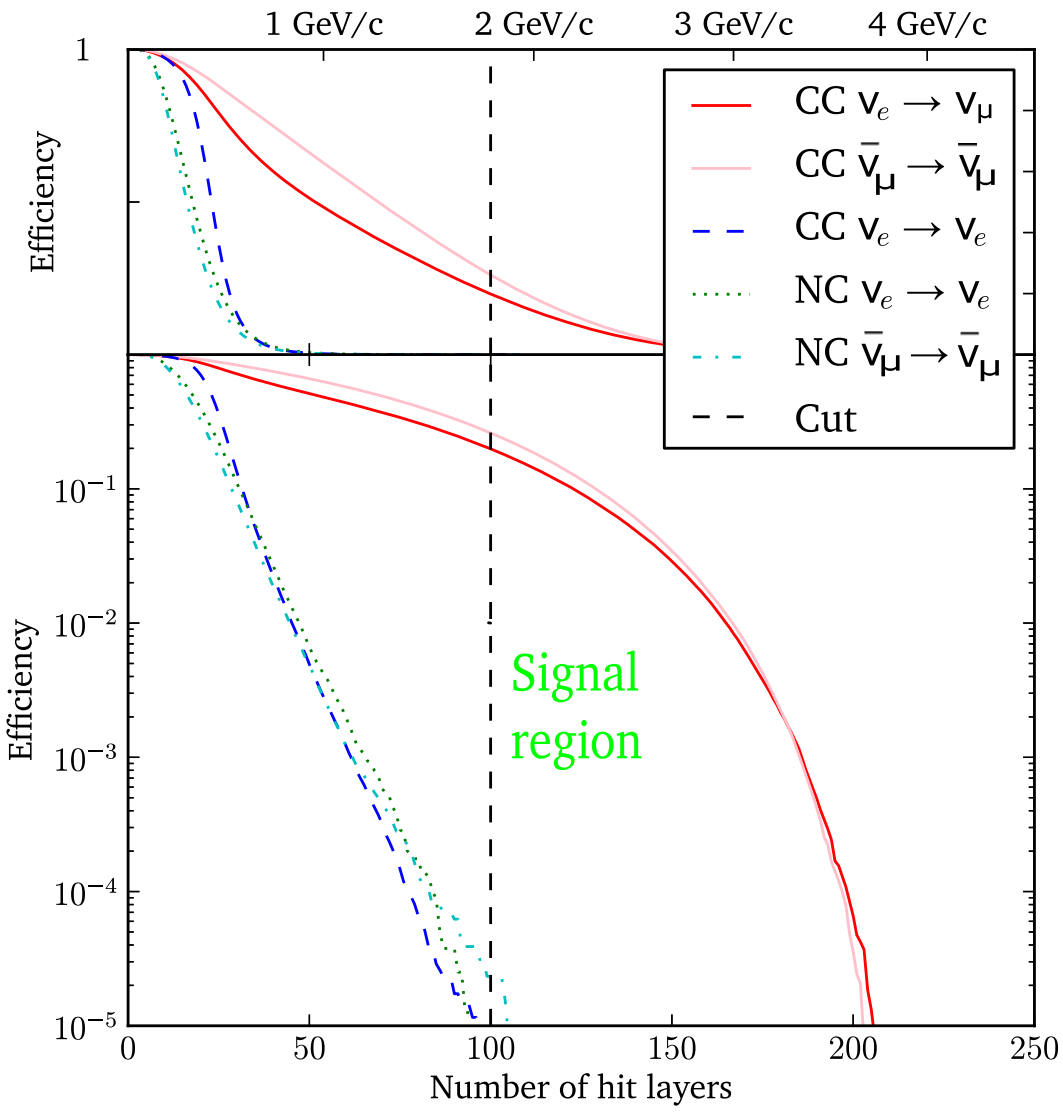

Figure 5.4: Performance of the range cut for LSND best fit of 8.9 eV^2. Efficiency is plotted on both a linear (top) and logarithmic (bottom) scale. All relevant channels from stored-μ^+ decay are shown. Requiring more than 100 hit layers, which a roughly 2 GeV muon-energy cut, removes 80% of the signal but results in backgrounds at the $\mathcal{O}(10^{-5})$ level, where 100 hit layers corresponds to 2 m (4 m) of steel (gapless detector). The continually-slowing-down approximation range for a muon is shown at the top.

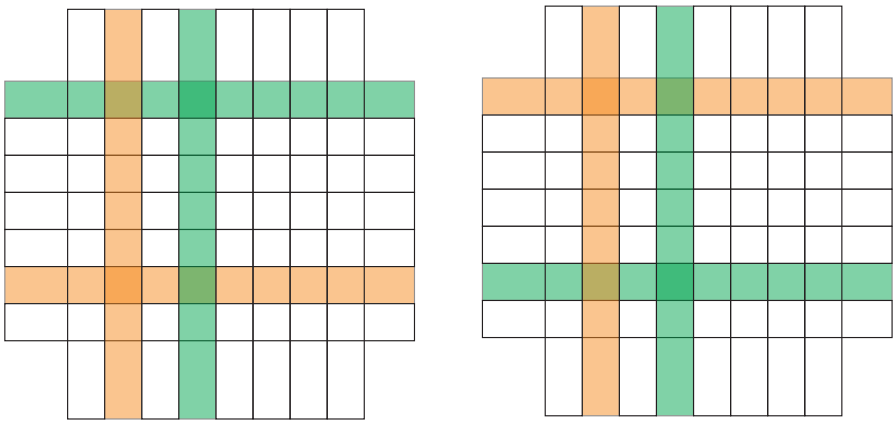

Figure 5.5: A diagram demonstrating hit *ghosting*. Shown are two views, x and y, where an orange and green particle have lit up two bars each. The situation on the left and right are indistinguishable using only position information.

at the 10^{-5} level. A substantial amount of signal needs to be removed to reject the backgrounds. The difference between muon appearance and disappearance is caused by the neutrinos having different energy distributions (Fig. 3.5) and different $y = (E_\nu - E_\mu)/E_n u$ distributions.

5.4 Pattern recognition algorithm

A curvature cut is required to determine if muons—the only particles that pass the length cut—are being focused or defocused, and is equivalent to the question of whether or not the muon track radius is increasing or decreasing in z. However, this requires being able to extract the muon track from a collection of hits, which is complicated by *ghosting* in the detector (Fig. 5.5), whereby it is impossible to match x and y views when two particles cross if only hit position information is used. Accordingly, the hits associated with the muon track must be extracted before it is possible to determine the curvature of the track.

An algorithm for extracting the muon track has been developed. Muons have long isolated tracks; therefore, it is easy for humans to identify muons in event displays, as can

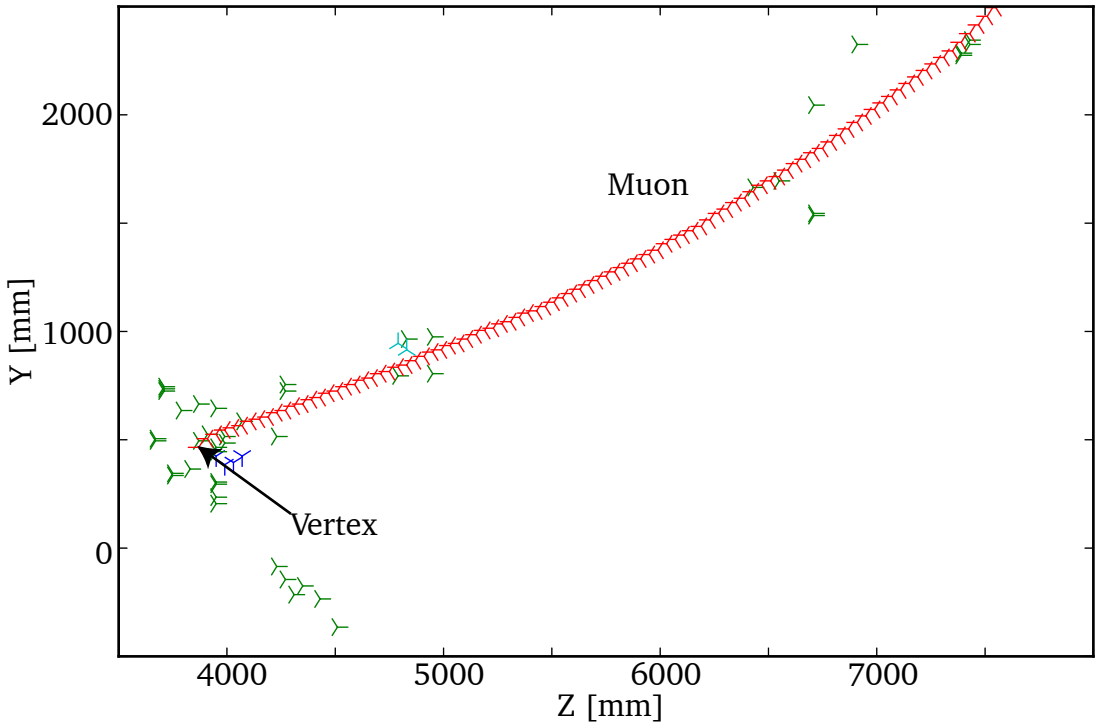

Figure 5.6: Event display in the y view of digitized hits corresponding to a ν_μ CC DIS event where the muon has been extracted. The algorithm described in Sec. 5.4 was used to extract the muon track (seen in red).

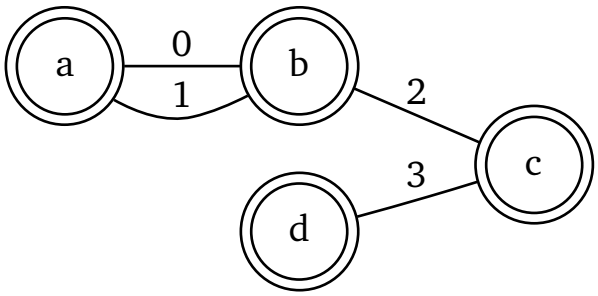

Figure 5.7: An example graph where vertices are circles, and the edges are lines. Each vertex is labeled with a letter. The graph is weighted by the number at each edge. The graph is not directed, but contains a cycle $a \to b \to a$. Starting from point a, the longest path is $a \to b \to c \to d$ with distance 6.

be seen in Fig. 5.6. However, developing a computer algorithm is less straightforward since computers—sadly—do not have innate pattern recognition abilities.

Numerous methods exist for finding muon tracks in $\overset{(-)}{\nu}_\mu$ CC events (reviewed in [83] and references therein). Most algorithms make the reasonable assumption that the muon is the longest track in the interaction, and that the beam interaction vertex is at the most upstream hits. For this thesis, a graph theoretic method is developed, which was independently developed in 1971 by Zahn for spark chambers [84].

Graph theory is a well studied field[1], where problems typically entail finding the extreme paths through an interconnected set of points. Example graphs are shown in Fig. 5.7 and 5.8, where the simplicity of these abstract objects can be seen. The vertices are labeled circles, and a path—and respective distance—exists between various vertices.

As previously mentioned, graphs can be used to find extreme paths; in the case of νSTORM, the extreme path is the muon. Graphs are used in a wide range of models: e.g., travel planning[2], power distributions, traffic flow, and even coloring books. In the case of computer networks, the extreme path is the quickest path between two devices that need to send packets to one another. However, in the case presented here, graphs are used to identify the muon.

[1] Part of why they are well studied is because they are generalizations of finite-state machines, trees, and automata, which are used extensively in computer science. Graph theory is reviewed in [85].

[2] See *traveling salesman problem*.

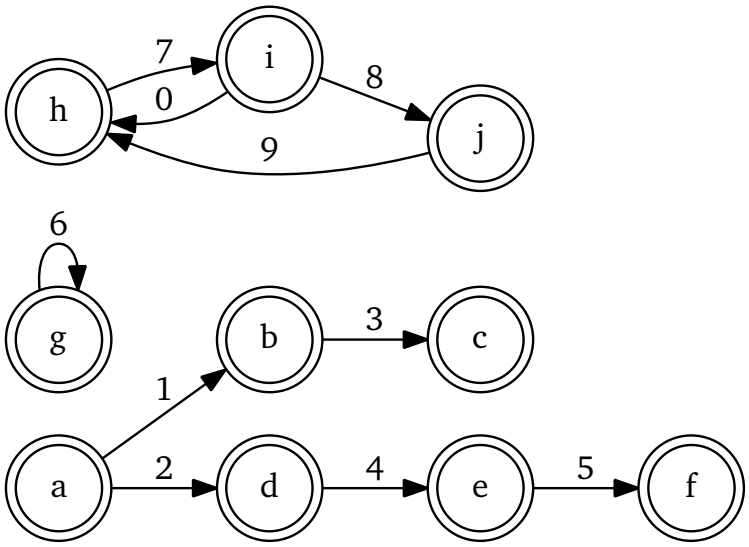

Figure 5.8: A directed graph, where the direction is given by arrows. The graph is not connected, and contains three connected subgraphs containing vertices a–f, g, and h–j. The top subgraph contains a cycle. The bottom connected graph is a directed acyclic graph. The longest path is the bottom subgraph starting from vertex a is $a \to d \to e \to f$ with length 11.

Table 5.3: Requirements for an edge between graph vertices (x_i, z_i) and (x_j, z_j).

Requirement 1	$z_i < z_j$
Requirement 2	$\lvert x_i - x_j \rvert < 5 \times$ (bar width)
Requirement 3	$z_j - z_i =$ (layer width)

As can be seen in the figures, a graph is a set of vertices and edges. Without loss of generality, only the x view is discussed at first. A *hit* in the x view has a location, which is represented as an ordered pair (x_i, z_i). The transverse coordinate x_i is determined from the bar number, and z_i, in the beam direction, is determined from the layer number.

Now that the vertices have been identified, the next step is to define how edges are created. Every edge between two vertices must satisfy three requirements. The direction of the edges "point" downstream, which is the path of muons from beam interactions (Requirement 1). Tracks are distinguishable if separated by five bars, so there are no edges between hits belonging to different tracks (Requirement 2). Lastly, edges should be between neighbors and there should not be 'shortcuts' to the end of the track, which could underestimate the track length (Requirement 3). More technically, this graph is transitively reduced since if there is an edge $a \to b$ and another edge $b \to c$, there cannot exist an edge $a \to c$. The transitive reduction ensures the path length is accurately estimated, including not having cycles that result in infinite path lengths. These requirements are summarized in Table 5.3.

Finally, the graph that will be used for pattern recognition is weighted. Every edge between two vertices has a numerical weight associated with the Cartesian distance between the respective hit bars. In other words, this is a Cartesian-weighted digraph with weighting $w(v_i, v_j) = \sqrt{(x_i - x_j)^2 + (z_i - z_j)^2}$. These weights are used when finding the longest path.

To identify the longest path, the neutrino interaction point is chosen to be the start of the longest path. This point is identifiable because it is the vertex in the graph that

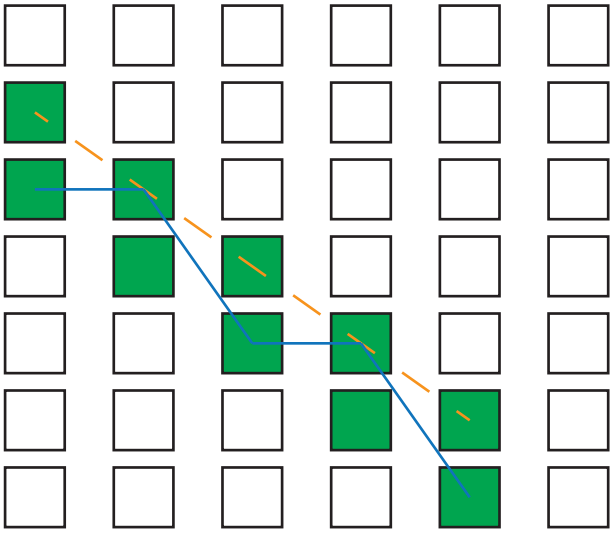

Figure 5.9: An illustration of why, after the extreme vertices have been identified, the shortest path between them must be computed. Scintillator bars are represented by boxes, and are hit when filled in green. The distance between the leftmost and rightmost hits can be calculated in different ways. For example, the solid blue path zigzags due to the hit doublets and overestimates the muon path length. However, the minimal distance shown in dashed orange corresponds to a path closer to the muon path.

has the most paths to other vertices in the graph. This is called the most accessible node[3] (e.g., vertex a in Fig. 5.8 is the most accessible node). By defining the interaction vertex to be the most accessible node, rather than the most upstream hit, ambiguities from the nuclear breakup releasing neutrons that create upstream "fake" hits is removed.

With an initial point identified, the longest *direct* path in the graph can be found. For every point in the graph, the distance between it and the interaction point is computed. To ensure a direct path back to the interaction point, the Bellman-Ford algorithm [86,87] is used[4]. The longest of these paths is called the muon. For example, in Fig. 5.8 it would be the path a to f with a cumulative weight of 11. More mathematically, the most direct path between any two vertices is also the shortest path. Therefore, if $P =$

[3]This vertex is found by taking the transitive closure of the graph and then counting the number of edges from each point.

[4]The Bellman-Ford algorithm requires a directed acyclic graph (DAG) because, otherwise, the complexity of the problem is NP-complete. The DAG requirement is satisfied by the problem being physical; therefore, edges only point downstream. For a DAG using this algorithm, the worst case runtime is $\mathcal{O}(|V||E|)$.

{set of shortest paths between every two hit bars}, then the muon is the path $\max(P)$. The computation of shortest paths prior to the selection of the muon path avoids zig-zagging artificially lengthening the track (Fig. 5.9). Hits immediately adjacent to a hit in the path (i.e., doublets) are added to the muon path.

Other "legs" of the interaction—e.g., if two muons emerge from the interaction point—can be found by removing the muon track and then repeating the extraction routine. This is particularly useful for for events with multiple pions and muons, such as from resonant production, but is not relevant for $\overset{(-)}{\nu}_\mu$ CC QE or DIS.

This algorithm identifies final-state muons. In the event display shown in Fig. 5.6, a muon is found using this algorithm. The muon track is identifiable as the track that extends from $z = 4$ m to $z = 8$ m, and is shown in red. The interaction vertex was well identified and the hadronic noise around the vertex did not affect the search for the interaction vertex. A second track extraction was attempted, but no other track could be found.

In contrast, in Fig. 5.10, which is a NC event, no long track (i.e., muon) is found. The longest track is 60 cm, and attempts at finding other tracks failed. Note that the interaction vertex was found. The algorithm still performs well, but did not find a muon candidate that satisfied the criterion of a 4 meter track.

5.5 Track fitting

Each view was dealt with independently but the views still need to be combined to form 3D space points, resulting in a set of vertices V_X and V_Y for the x and y views, respectively. At this point, since the longest-track muon has been found by the algorithm in the previous section, there is only one measurement per x- and y- detection plane. It is then possible to define a 3D space point without any concern for ghosting. At the end of the algorithm there is a set of (x_i, y_i, z_i) associated with the track.

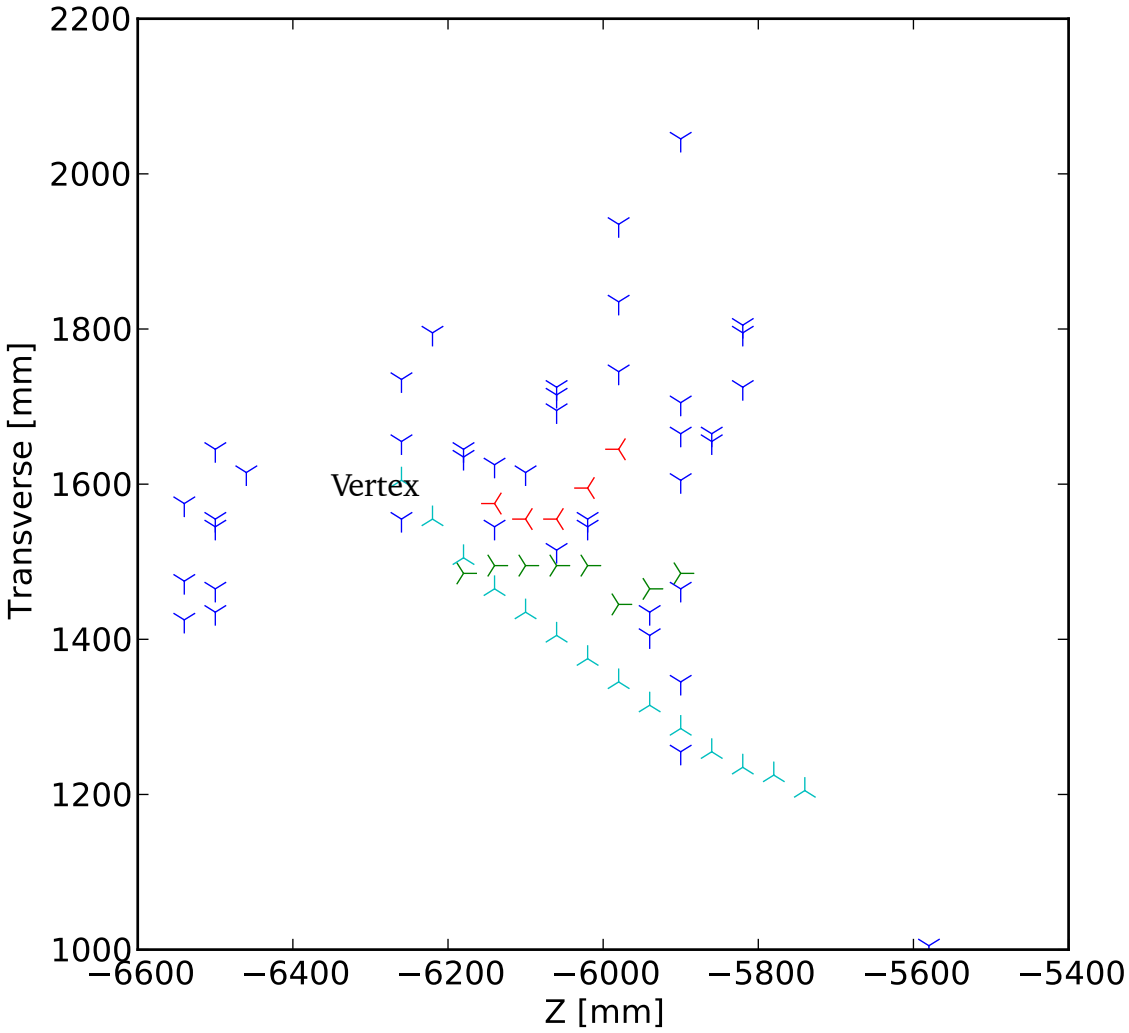

Figure 5.10: Event display in the x view of digitized hits corresponding to a ν_μ NC event.

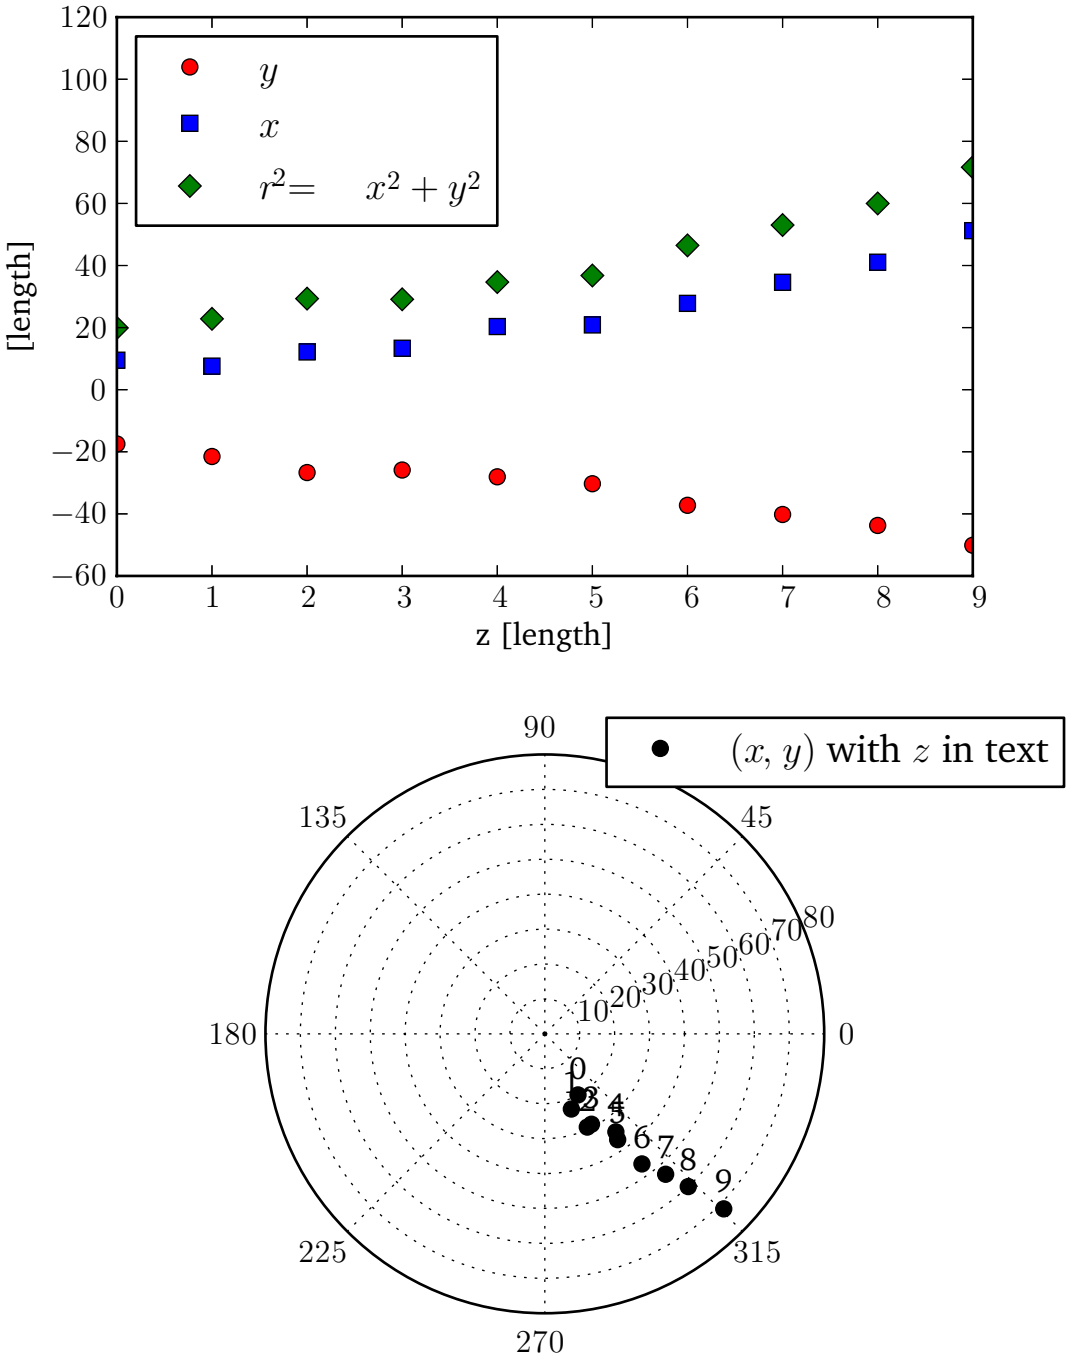

Figure 5.11: Sketch of fit procedure using fake data of a muon track. The data points do not correspond to a simulation or data but were rather chosen to demonstrate the fit procedure. The top plot shows the x vs z and y vs z projections of a potential muon track, where it can be seen that the radius of the track is increasing so the muon is being defocused. The bottom plot shows the polar projection of the same data, where it can be seen that the muon is being defocused and will exit the detector.

After the track has been extracted, there exists a set of (x_i, y_i, z_i) coordinates associated with the event track. A fit is performed to parameterize these points into three numbers. The track is fit using the simple second-order polynomial function

$$r(z) = a + bz + cz^2, \tag{5.1}$$

where a least-squares fit is performed in order to minimize,

$$R^2 = \sum_i \left(r(z_i) - \sqrt{x_i^2 + y_i^2} \right)^2 . \tag{5.2}$$

The result of the fit is values a, b, and c that correspond to an offset, slope, and curvature, respectively. The sign of the curvature c corresponds to the muon being focused or defocused. As illustration of the fit procedure is shown in Fig. 5.11, where r_i is determined for each z_i then fit to a second order polynomial. In this case, the particle is defocused since the radius is increasing.

This track fit is considerably simpler than those typically used in particle physics experiments. For example, the trajectory for the muon is assumed to follow the simple quadratic equation in Eq. (5.1), where this is an approximation of the true trajectory. Effects such as the muon slowing down or multiple scattering are ignored. As a result, much of the muon-momentum information is lost and muon range must be used instead. Numerous more advanced methods exist for track fitting (e.g., Kalman filtering); however, as will now be presented, this simple fit performs remarkably well.

5.6 Curvature cut

Track curvature can be used to differentiate μ^+ and μ^-, and the Lorentz force law applies to these muons,

$$F = \frac{d\vec{p}}{dt} = q\left(\vec{v} \times \vec{B}\right), \tag{5.3}$$

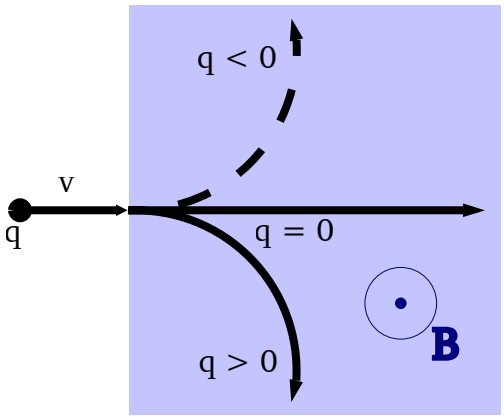

Figure 5.12: Lorentz force [88].

where F is the force, $\vec{p}$ is the momentum 3-vector, t is the time, q is the charge, $\vec{v}$ is the muon velocity, and $\vec{B}$ is the magnetic field (Fig. 5.12). The field is toroidal as was shown in Fig. 3.1, which results in $\vec{v}$ and $\vec{B}$ being at approximate right angles since, topologically, a toroidal field is locally dipole (i.e., the muon only 'sees' a dipole field if it does not travel too far). The formalism is nonrelativistic—the results are easily generalizable to relativistic muons—but all muons travel at roughly the same speed relativistic $\beta = 1$. Therefore, $\vec{v} \times \vec{B} = |\vec{B}|$ in natural units, and

$$F = \frac{dp_r}{dt} \sim q|\vec{B}|. \tag{5.4}$$

The field B is the same for μ^+ and μ^-; therefore, the only variable that affects the direction of the force is the charge q. If q is negative (e.g., μ^-), then the particle starts traveling toward the center of the detector (i.e., $dp_r/dt < 0$) and is said to be *focused* by the field. Conversely, if q>0, then the particle is said to be defocused.

The muon charge identification performance is determined. After requiring the range cut of 100 Fe planes outlined previously to be satisfied, it is possible to look at the fit parameter c for CC $\nu_e \to \nu_\mu$ and CC $\bar{\nu}_\mu \to \bar{\nu}_\mu$ events (Fig. 5.13), which will determine whether muons are being focused or defocused.

A simple cut of requiring $c < 0$ for μ^- and $c > 0$ for μ^+ is effective at determining

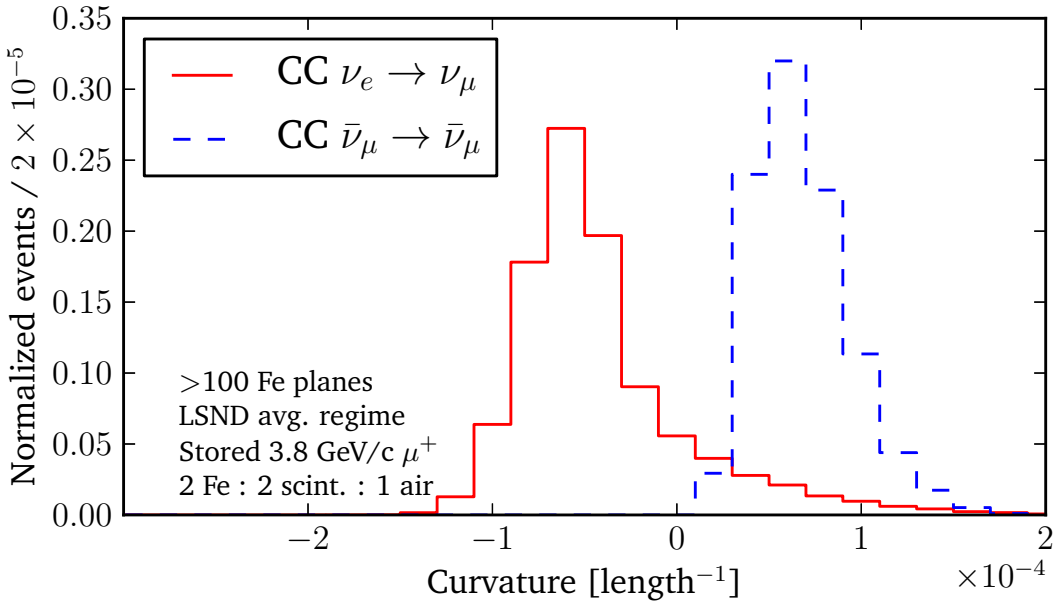

(a) Linear

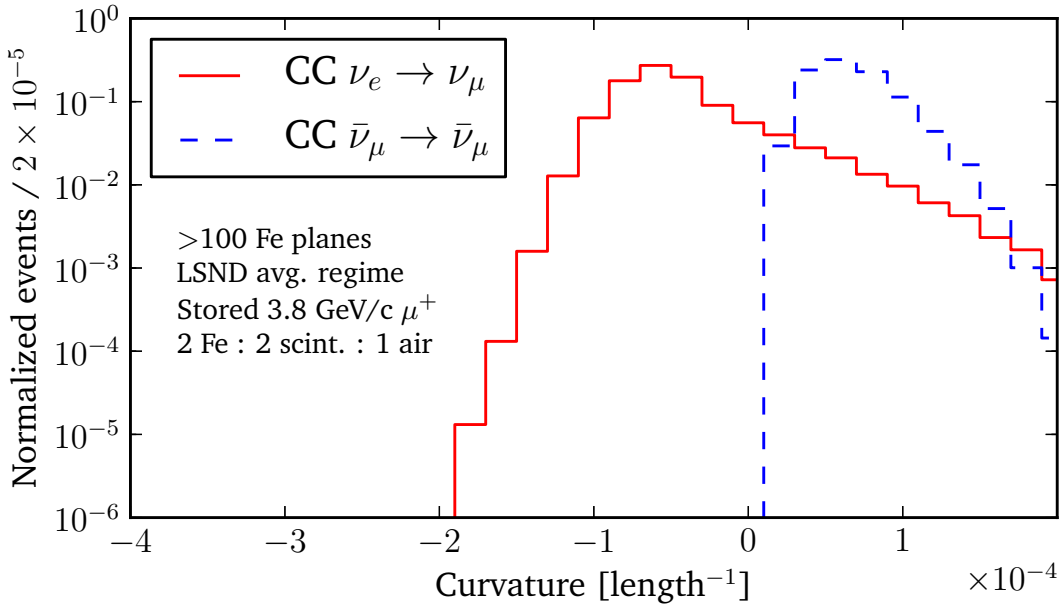

(b) Log

Figure 5.13: Linear and log distributions of curvature for CC $\nu_e \rightarrow \nu_\mu$ signal and CC $\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$ background after a 100-Fe-plane track-length cut for LSND best fit of $\Delta m^2 = 0.89$ eV^2. Negative values correspond to focusing toward the beam axis, while positive values are defocused. A cut at $c = 0$ separates the distributions. The long tail for CC $\nu_e \rightarrow \nu_\mu$ toward positive values is a geometric effect. The standard 2 cm of iron, 2 cm of scintillator, and 1 cm of steel was assumed.

the charge for most of the particles (Fig. 5.13a). The tail of the signal distribution in the positive direction is not matched by a tail of the background distribution in the negative direction due to the choice of coordinate system. The variable c is positive if the radius from the central coil increases, and negative if the radius decreases. Muons at small radii cannot have decreasing radii. Therefore, focused particles do not have tails of small values of c. This geometric effect reduces the tail of the background distribution into the signal window. When the magnetic field is flipped, the tail exclusively extends in the other direction. This is simply a statement about the definition of the coordinate system.

Fig. 5.13b shows the distribution of c, and the level to which the background distribution leaks into the signal window of $c < 0$ is determined. The 2D Opera field model mentioned in Table 5.2 was used, which results in a 2-T field throughout most of the detector. No background events have $c < 0$ and 531077 events were simulated; therefore, the Poisson upper limit is $2.3/531077 = 4 \times 10^{-6}$. The remaining signal efficiency for just the curvature cut is 87%. The cut is effective at reducing right-sign muons by having both a high purity and efficiency.

In order to show the dependence on the field strength, the field is reduced by 20% and is shown in Fig. 5.14 (compare to Fig. 5.13). Despite the saturation of the steel, finite-element analyses show that it is sufficient to approximate the lower field by applying a linear scaling [89]. If the field is lowered 20%, the charge misidentification rate will rise to $\mathcal{O}(10^{-4})$, which is within the specification from the previous chapters.

5.7 Performance after all cuts

The two cuts that have been designed give the requisite background rejection for $\nu_e \rightarrow \nu_\mu$ appearance searches. Fig. 5.15 shows the cumulative performance of the cuts as a function of the neutrino energy E_ν, which can later be used to determine the sterile neutrino

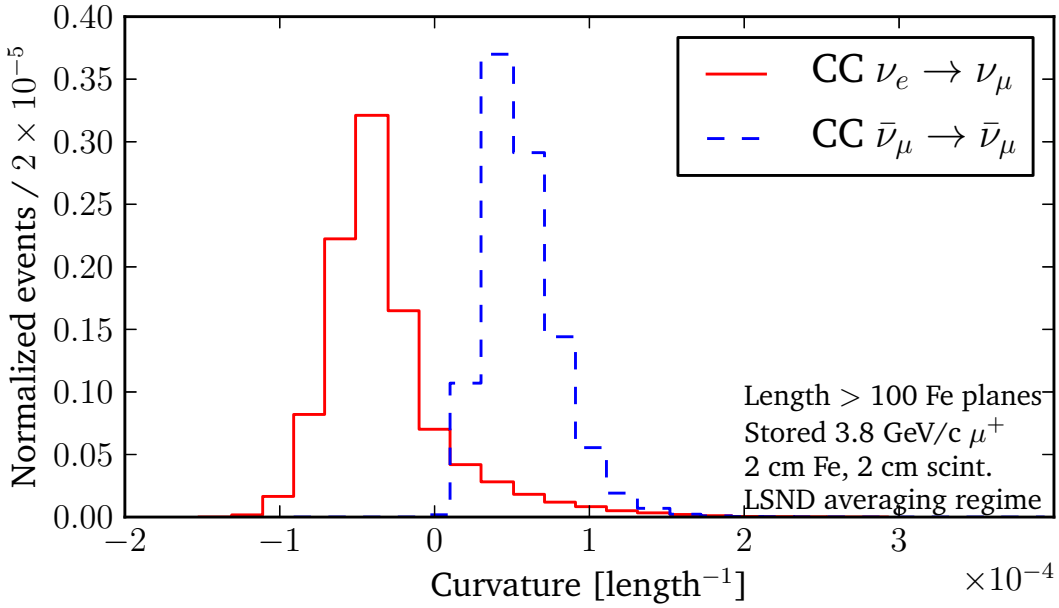

(a) Linear

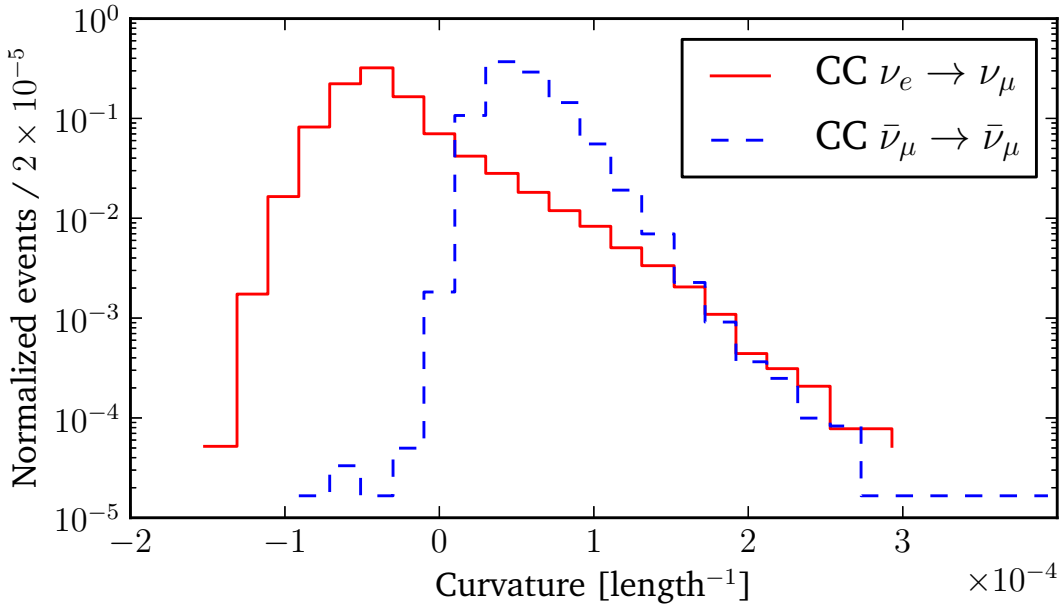

(b) Log

Figure 5.14: Distribution of curvatures for background and signal for 80% field strength with $\Delta m^2 = 0.89$ eV^2 and a range cut of 100 iron planes.

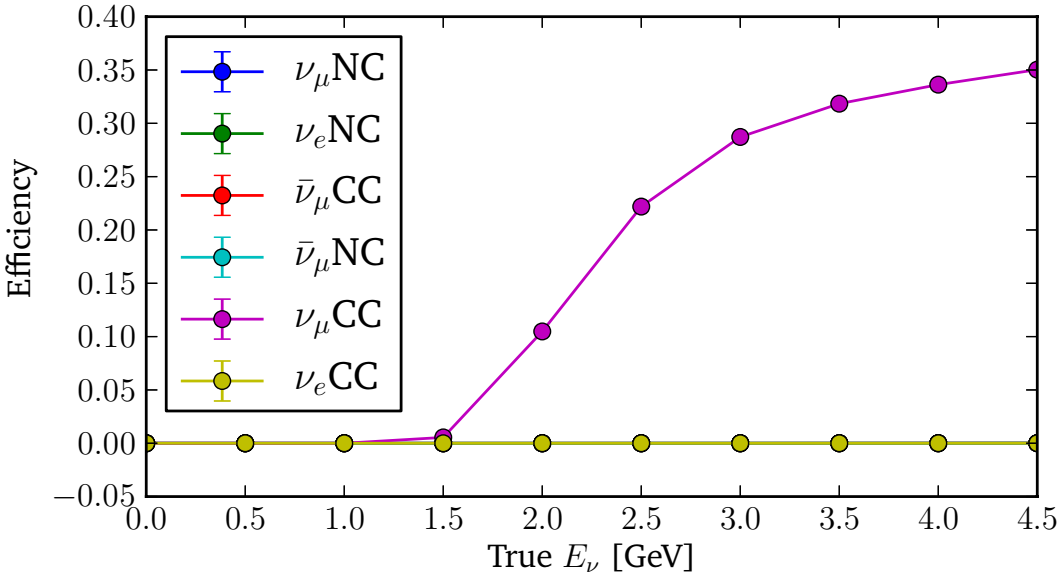

(a) Linear scale

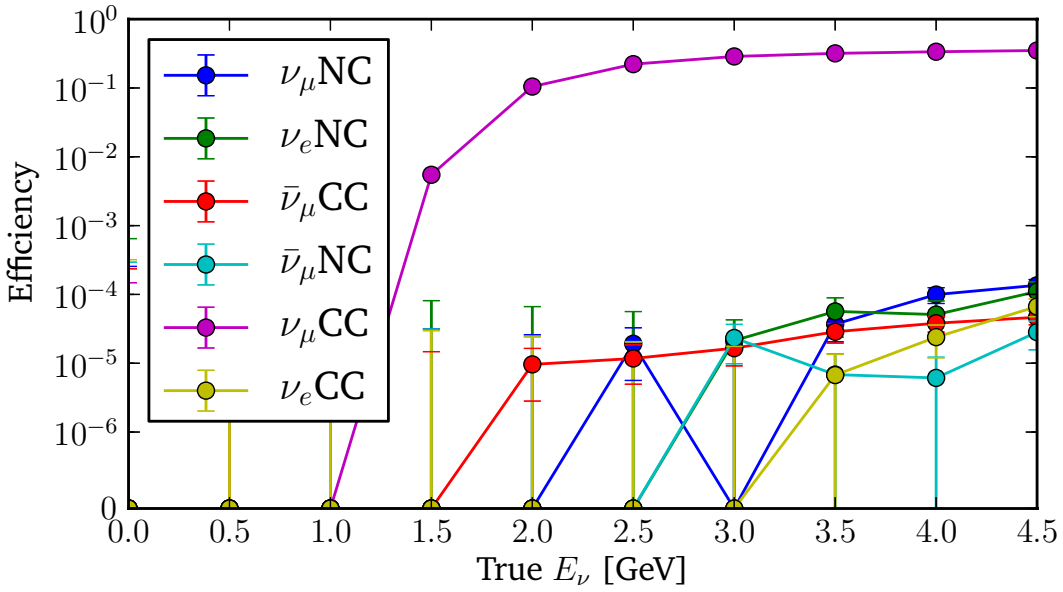

(b) Log scale

Figure 5.15: Efficiency curve for a simple cuts-based analysis. Error bars shown for all points and arise from the statistics from the simulation.

sensitivity with different neutrino fluxes. For neutrino energies of 3 GeV, this appearance search is 30% efficient. The signal efficiency plateaus at 35%. The background levels are at $\mathcal{O}(10^{-5})$, as required. This efficiency curve agrees with the later independent work presented in [50] that used Kalman filtering and likelihood cuts.

5.8 Conclusion

A MC has been implemented to strengthen the arguments that νSTORM is able to do muon appearance physics. An electronics simulation and pattern recognition algorithm were designed to derive these performance estimates. Lastly, two simple cuts have been designed that demonstrate the required background rejection levels.

Chapter 6

Appearance physics sensitivity

Chris. I don't want to see any more 5σ contours. I want 10σ.

Stephen Parke (theorist)

Chris. We don't believe you when you say 10σ.

Luca Stanco (experimentalist)

Having determined the detector performance in the previous chapter, it is now possible to determine the sterile neutrino sensitivity. The work presented here supersedes the work in [50, 90, 91]. (See [92] for work on disappearance measurements.) The method developed here has been used for all νSTORM sterile neutrino plots, at the time of writing.

Sensitivity contours of the neutrino parameter space have been calculated for over 20 years, so the methods used here are fairly standard. Even though the work presented is new, the principles are the same. However, there is one major difference; the flux computation developed here is the most sophisticated treatment of muon-decay-beams to date. This is primarily because this is a study that integrates over the phase space

Table 6.1: Electroweak muon-decay parameters.

Particle	$f_0(x)$	$f_1(x)$
ν_μ	$2x^2(3-2x)$	$2x^2(1-2x)$
ν_e	$12x^2(1-x)$	$12x^2(1-x)$

of the beam. Said differently, most NF studies (e.g., [57]) assume a point source of monoenergetic muons.

For any neutrino experiment, the calculation of sensitivity requires the following sequential well-established computations:

1. flux Φ,
2. oscillation probability (Prob),
3. cross section σ,
4. number of interactions $N_{\text{int}} = \Phi(\text{Prob})\sigma$, where $\sigma = \sigma_{\text{nucleon}} \times N_{\text{nucleons}}$,
5. number of events $N_{\text{evt.}} = N_{\text{int}}\epsilon$ for efficiency ϵ, and
6. test statistic χ^2,

where all of these are energy dependent. For example, as was shown in Fig. 3.5, monoenergetic muon decays produce a wide-band neutrino beam, so the flux depends on energy. Similarly, the other items in the list above depend on energy, as will be discussed in this chapter.

6.1 Neutrino flux: Φ

Interest in μ-decay beams is because they can have a two orders of magnitude smaller flux uncertainty (0.1%) than π-decay beams ($\sim$20–40%). This has never been shown experimentally—which is part of the motivation for νSTORM—but is motivated by two

key arguments: one can better measure a muon beam than pion beams since muons are longer lived, and avoid the production uncertainties from a fixed target beam. Measuring the beam is easier since the $c\tau$ of a muon is 600 m/7 m $\simeq$ 100 times longer than that of a pion; therefore, muons are stored in an accelerator, and can be measured using accelerator beam diagnostics. Also, the production uncertainties from π decay from the proton–nucleon production mechanism is avoided because, once again, the muon beam can be measured *in situ*. Also, uncertainties in the number of, e.g., kaons produced at the target station are irrelevant since kaons will decay ($c\tau = 3.7$ m) before the muons and pions do.

Once the muons have been separated from πs and Ks and measured, the muons will decay. This electroweak muon decay $\mu^+ \to \bar{\nu}_\mu \nu_e e^+$ is well understood, and produces equal numbers of $\bar{\nu}_\mu$ and ν_e. The neutrino flux from muon decay was already discussed, but this discussion will now be expanded here. Previously, only the $\overset{(-)}{\nu}_e$ flux was shown, but $\bar{\nu}_\mu$s from μ^+ decay result in backgrounds and must understood. More generally than Eq. (3.21) on p.p. 56, the neutrino spectrum for a $\mu^\pm \to e^\pm + \nu_e(\bar{\nu}_e) + \bar{\nu}_\mu(\nu_\mu)$ decay in the rest frame of the muon is

$$\frac{dn}{dx\, d\Omega} = \frac{1}{4\pi}\left[f_0(x) \mp \mathscr{P} f_1(x)\cos\theta\right], \tag{6.1}$$

where n is the number of muons, $x = 2E_\nu^{\text{c.o.m.}}/m_\mu \in [0,1]$ is the scaled neutrino energy in the rest frame, Ω is the solid angle in the rest frame, $f_0(x)$ and $f_1(x)$ are muon-decay parameters (see Table 6.1), and $\mathscr{P}$ is the polarization. The muon-decay parameters are different for ν_e and $\bar{\nu}_\mu$, which results in different neutrino spectra (see, e.g., Fig. 3.5).

6.1.1 Polarization

The beam polarization $\mathscr{P}$ affects neutrino fluxes, as can be seen from Eq. (6.1). For instance, if the polarization $\mathscr{P}$ is unity, then the ν_e flux from μ^+ decay is zero since

the ν_es are at rest in the laboratory frame [Eq. (6.1)]. (Polarization of muon beams is explored extensively in [93].)

Muons at νSTORM are tertiary beams and arise from pion decay. Pions, however, are long lived because they have zero spin ($J^P = 0^-$) and therefore have helicity suppressed decays. For the decays $\pi^+ \rightarrow l^+ \nu_l$ to conserve spin, the charged lepton l^+ must be righthanded since the neutrino is chiraly lefthanded. This helicity suppression is why the pion-decay mode $\pi^+ \rightarrow e^+ \nu_e$ is suppressed at $\mathcal{O}(10^{-4})$ compared to $\pi^+ \rightarrow \mu^+ \nu_\mu$. The helicity of the μ^+ is then –1 in the pion frame.

However, this spin is a projection onto the momentum vector of the muon, which changes as the muons circulate around the ring. Each particle within a bunch has a spin $\vec{S}$, which undergoes Thomas precession. Polarization is defined as the ensemble average of the spin vectors:

$$\mathscr{P} = \begin{pmatrix} \mathscr{P}_x \\ \mathscr{P}_y \\ \mathscr{P}_z \end{pmatrix} = \langle \vec{S} \rangle = \frac{1}{N} \sum_i^N \vec{S}_i. \tag{6.2}$$

The projection of the polarization vector on the beam axis $\mathscr{P} \cdot \hat{z} = \mathscr{P}_z$ is called the *longitudinal polarization*.

The initial polarization can be determined using similar methods employed for NF studies. The muon has a helicity in the laboratory frame,

$$h = \frac{EE^* - \gamma_\pi m^2}{pp^*}, \tag{6.3}$$

where m, p, and E are the muon's mass, momentum, and energy, respectively [94]. The starred and unstarred frames correspond to the pion-rest and laboratory frame, respectively. (For a more comprehensive treatment, see [95].) The polarization can now be determined by averaging over the decay kinematics.

Approximations can be made to determine h without using an MC accelerator simulation, such as G4BEAMLINE. The muon and pion energy distributions are different since the decay ring lattice is designed to only transport muons that decayed backward in the pion rest frame. The muon distribution in the pion rest frame is isotropic. Using a toy MC and only considering backward decaying muons, the average helicity can be computed using Eq. (6.3). The longitudinal polarization is −48% in the laboratory frame.

The evolution of the polarization in the ring can now be determined. The spin precession frequency for a polarized muon is

$$\nu = \frac{g_\mu - 2}{2}\frac{E_{\text{beam}}}{m_\mu} \tag{6.4}$$

$$= \frac{E_{\text{beam}}}{90.6233\ \text{GeV}}, \tag{6.5}$$

where $g-2$ is well measured to one part in 10^8. νSTORM is far from a spin resonance; the closest is $n = 24$. Without an energy spread, muons coherently Thomas precess. Since $g-2$ is not zero, the spin precesses at a different frequency than the circulating muons do. Also, an energy spread results in the spins precessing at different frequencies with respect to one another; therefore the polarization decoheres, and can average to zero—as will be shown.

Introducing a 10% Gaussian energy spread for a 3.8-GeV stored muon energy, the evolution of the polarization as a function of time can be computed (Fig. 6.1). A spread of 10%, compared to the 15–20% that may be possible in νSTORM, was used to overestimate the polarization. However, the polarization shown in the previous figure does not account for muon decay. This is important since the beam is more polarized earlier in time where the muon is more likely to decay. Therefore, for an accurate estimate of the polarization including the effect of muon decay, the average polarization can be determined. The average polarization is 2×10^{-3}—smaller than other uncertainties at

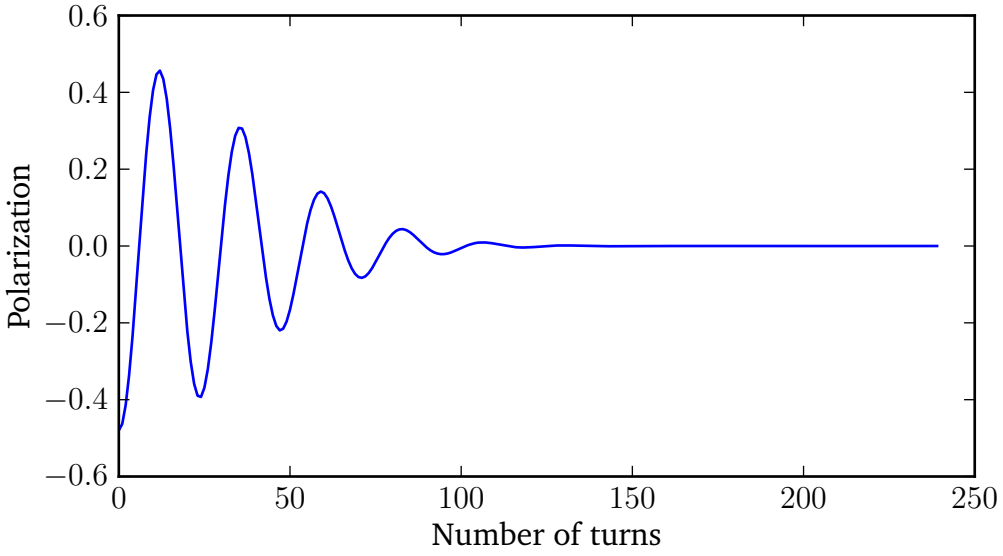

(a) Linear scale

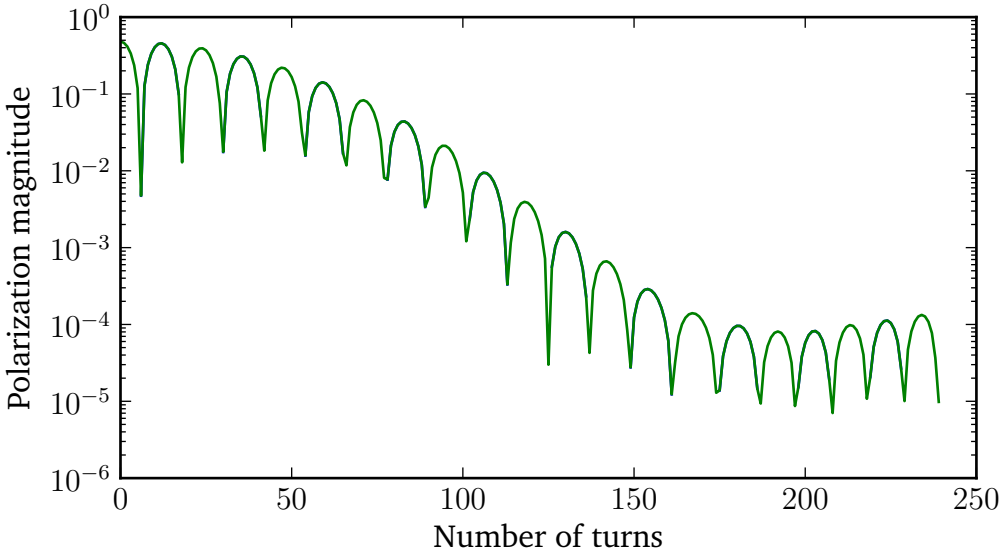

(b) Log scale

Figure 6.1: The muon polarization evolution if the initial muons have –48% initial polarization and an energy distribution of 3.8 GeV ± 10%. The decay-weighted average is 2×10^{-3} under this assumption, where 10^{6} muons were simulated.

νSTORM. Similar results are seen in other NF studies [57].

Beam polarization is typically ignored in NF studies [48, 57, 96] because, as was just seen, polarization does not affect the flux when averaged over the beam lifetime. However, importantly, these claims have now been verified for a lower-energy NF where these effects are more severe—but still negligible.

6.1.2 Phase-space integration

Once the neutrino spectra in Eq. (6.1) are boosted into the laboratory frame, the resulting distributions are

$$\frac{d^2N_{\nu_\mu}}{dy\,dA} = \frac{4n_\mu}{\pi L^2 m_\mu^6} E_\mu^4 y^2 (1-\beta\cos\phi)\Big[3m_\mu^2 - 4E_\mu^2 y(1-\beta\cos\phi)\Big], \tag{6.6}$$

$$\frac{d^2N_{\nu_e}}{dy\,dA} = \frac{24n_\mu}{\pi L^2 m_\mu^6} E_\mu^4 y^2 (1-\beta\cos\phi)\Big[m_\mu^2 - 2E_\mu^2 y(1-\beta\cos\phi)\Big], \tag{6.7}$$

where $y = E_\nu/E_\mu$ is the scaled neutrino energy in the laboratory frame, $\beta = \sqrt{1-m_\mu^2/E_\mu^2}$, A is an area, n_μ is the number of muons, and N is the number of neutrino interactions for a neutrino flavor [97].

When computing the flux for νSTORM, the far detector approximation of a point-source accelerator and detector is not necessarily applicable since the size of the detector and accelerator straight (150 m) are comparable to the baseline of 2000 m. The neutrino fluxes are computed by integrating over the decay straight, transverse beam phase space, and detector volume. The beam occupies a 6D phase space (x, y, z, p_x, p_y, p_z), and the detector has a 5×5 m^2 cross section with the depth set by the desired fiducial mass of 1.5 kt ($\sim$13 m, if no air gaps). Both transverse 2D phase spaces are represented by the Twiss parameters, which are convenient descriptions of accelerator beams. For example, by definition, the spread in, for example, x is $\sigma_x = \sqrt{\beta_\perp \epsilon}$ and the angular divergence in x is $\sigma_{x'} = \sqrt{\epsilon/\beta_\perp}$. The Twiss parameters for this ring are $\alpha = 0$ and

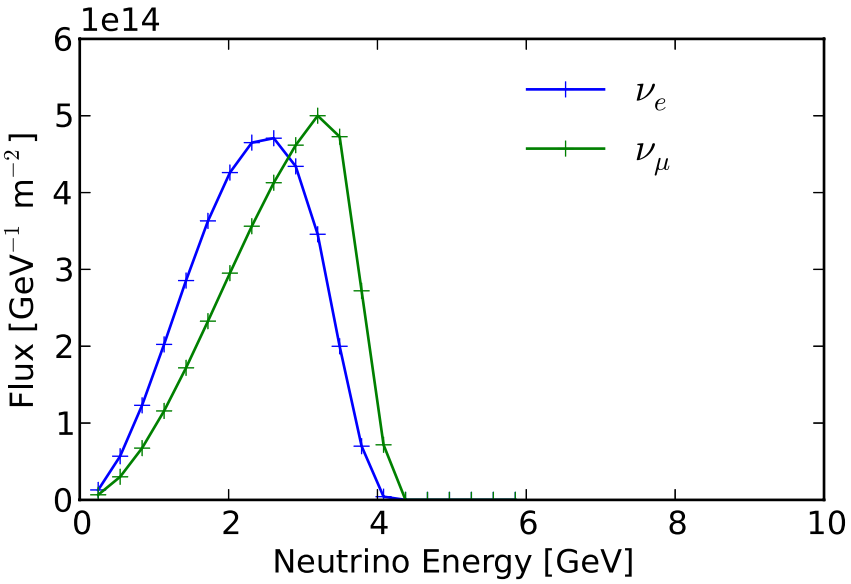

Figure 6.2: The unoscillated flux of ν_e and $\bar{\nu}_\mu$ for a (3.8 ± 0.38) GeV/c muon decay at 2000 m. Accelerator effects are included; see the text for details.

$\beta_\perp = 40$ m, where the 1σ Gaussian geometric emittance is assumed to be 2.1 mm (motivated from FODO study in [50]). These optics correspond to a 3.8-GeV/c muon, but dynamic aperture studies have shown efficient transport for off-momentum particles (i.e., no resonances) [64]. The longitudinal phase space (z and p_z) is described by assuming uniform distributions in both $z \in [0, 150\text{ m}]$—accurate to 0.5% based on the muon lifetime—and $p_z \in [3.8 \pm 0.38\text{ GeV/c}]$.

The flux is computed by MC integration: random points are chosen within the beam phase space and within the detector volume to determine the expected flux. This integration introduces a new computational requirement: the baseline is a variable that affects both the oscillation probability (L/E) and the flux (L^{-2} geometric factor). The GLoBES software (version 3.1.10) [98,99] that is used in neutrino factory phenomenology treats these as separable problems and had to be modified to compute this flux, and later the event rates and sensitivities. Specifically, GLoBES is modified such that both the flux and oscillation probability are computed in the *oscillation probability engine*. The code for the analysis is available [100] under the GPL license [101].

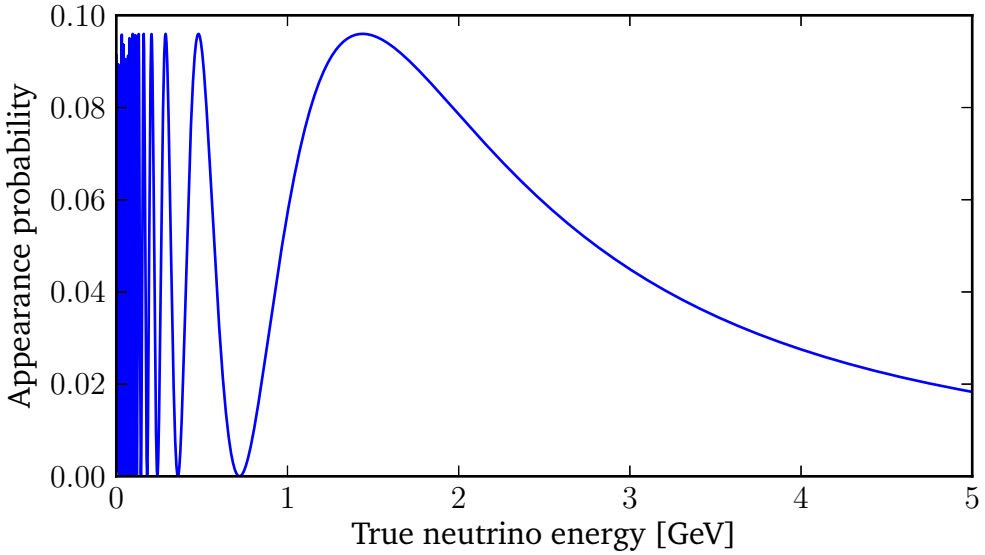

Figure 6.3: The appearance oscillation probability for the "golden channel" $\nu_e \rightarrow \nu_\mu$ from Eq. (2.21) using the (3+1) oscillation parameters in Table 6.2. A baseline of 2 km is assumed.

The resulting flux with accelerator effects is shown in Fig. 6.2, where $N_\mu = 1.8 \times 10^{18}$ [see Eq. (4.5)] and is based on 10^{21} POT at 60 GeV/c. The corrections introduced by integrating over the beam phase space are small (compare to 3.5). Given that these differences are small, it has now been demonstrated that at short 2-km baselines, relative to the long 7000-km baselines of other NFs, it is still reasonable to assume a point source despite the 150-m decay straight. These accelerator effects are nevertheless included in the subsequent calculations for completeness. (Note that integrating over the beam phase space *is* important for baselines of 20 m where, for example, cross-section physics is performed. This is beyond the scope of this thesis, but this work was published in [50].)

6.2 Oscillation probability

For a point-source baseline of 2000 m, it is possible to calculate the oscillation probability (Fig. 6.3) using Eq. (2.21) for any combination of L and E. Computationally, the SNU

Table 6.2: Best-fit oscillation parameters for the (3+1) sterile neutrino scenario using combined MB $\bar{\nu}$ and LSND $\bar{\nu}$ data [104].

Parameter	Value
Δm^2_{41} [eV^2]	0.89
$\lvert U_{e4}\rvert^2$	0.025
$\lvert U_{\mu 4}\rvert^2$	0.023

Table 6.3: Values for 3×3 oscillations used. Normal hierarchy is assumed. The values are in agreement with current best-fit values (Table 2.1).

$\sin^2\theta_{12}$	=	0.319
$\sin^2\theta_{23}$	=	0.462
$\sin^2\theta_{13}$	=	0.010
Δm^2_{21}	=	7.59×10^{-5} eV^2
Δm^2_{31}	=	2.46×10^{-3} eV^2

(version 1.1) add-on [102, 103] has been used to extend computations in GLoBES to 4×4 mixing matrices.

The best-fit parameters for the "short baseline anomaly" [104] (see Table 6.2) and 3×3 mixing [i.e., $\sin^2(2\theta_{13})$, Δm^2_{12}, etc.] (see Table 6.3) are used throughout the analysis, and agree with those published by the LSND Collaboration [105]. For completeness, oscillations between known mass eigenstates are included, despite not influencing the sensitivity; the correction was computed to be on the order of 10^{-5} in this analysis. The best-fit data from [106] is used to specify standard 3×3 oscillations. Without loss of generality, a normal hierarchy is assumed. As a cross check, errors on the 3×3 mixing matrix were included in the analysis and these are negligible; they affect the oscillation probability at much less than the 10^{-5} level. Though a more complete analysis is performed here, a two-flavor analysis would have been sufficient.

To understand the convolution of the flux with this oscillation probability, the unoscillated and oscillated fluxes are shown in Fig. 6.4. Interestingly, the shapes the ν_e unoscillated, ν_e disappearance, and $\nu_e \rightarrow \nu_\mu$ appearance at the LSND best fit spectra are similar. This is true throughout all LSND-motivated Δm^2_{41} values, and was due to (remember Fig. 6.3) the majority of the neutrino energy spectra being at the tail of

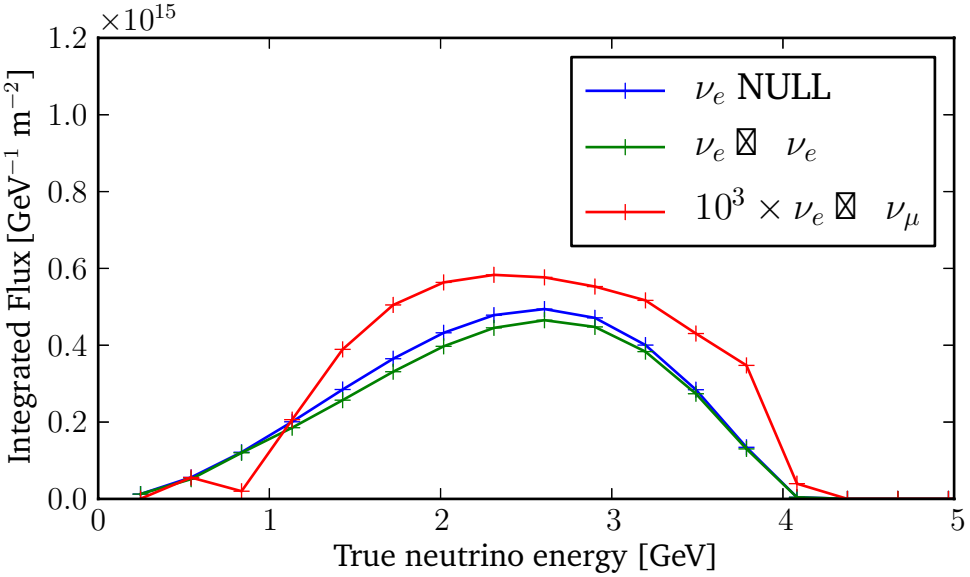

Figure 6.4: The flux at the far detector for a (3.8 ± 0.38) GeV/c stored μ^+ beam for initial ν_e states including averaging over the beam envelope and detector volume. Final states include ν_e without oscillations and both ν_e and $\bar{\nu}_\mu$ with best-fit, short-baseline oscillations. The normalization is 10^{21} POT. The dip in the oscillated appearance signal at 0.7 GeV corresponds to the first minimum of the oscillation probability (see Fig. 6.3).

the oscillation probability. At first glance, this would appear disadvantageous; however, there are more effects that depend on L and E than just the oscillation probability. For example, including all effects such as the L^2 solid angle, E^2 opening angle, and E from neutrino cross sections makes 2 km the optimal baseline and 3.8 GeV an optimal energy (as will be seen later).

6.3 Cross section: σ

Cross sections are required for each neutrino flavor (ν_μ, $\bar{\nu}_\mu$, ν_e, $\bar{\nu}_e$) and each interaction type (CC or NC). The nucleon cross sections (Fig. 6.5) for an isoscalar target are calculated in [12] and [107] for low and high energies, respectively. The νSTORM target is mostly iron, which is modeled as ^{56}Fe ($Z = 26, A = 56$) since this isotope constitutes 91% of the natural abundance, and can be approximated as isoscalar, which introduces a 5% error that is much less than the cross-section uncertainty discussed later. NC cross sections are flavor independent, and the neutrino NC cross section is about twice the antineutrino NC cross section. The CC cross sections are approximately flavor independent; *Fermi's Second Golden Rule* results in the same matrix elements and, at these energies, the phase spaces for the final-state electrons and muons are roughly equal. Differences do arise from structure functions being suppressed by powers of the lepton mass and, at lower energies, threshold and radiative effects [108], but these effects are small in this context. The CC cross section for neutrinos, in the DIS region, is twice that of antineutrinos, and is approximately three times the respective NC cross section.

At present, only the nucleon cross section σ_{nucleon} is known; therefore, the number of nuclei N_{nuclei} needs to be known to determine the total cross section $\sigma = N_{\text{nuclei}} \times \sigma_{\text{nucleon}}$ before the interaction rate $N_{\text{int}} = \sigma\phi$ can be determined. The fiducial mass of 1.5 kt $= 1.5\times10^9$ grams determines the number of nucleons via Avogadro's number $N_A = 6\times10^{23}$: $N_{\text{nuclei}} = 1.5 \times 10^9 \times 6 \times 10^{23} = 9 \times 10^{32}$. Fiducial mass uncertainties should be less than

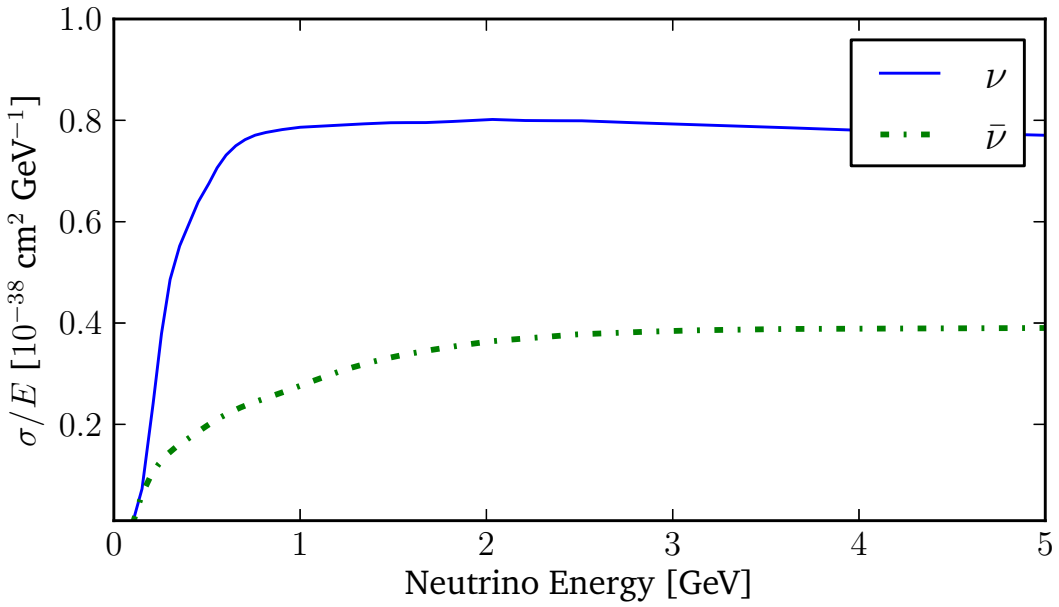

(a) CC interactions

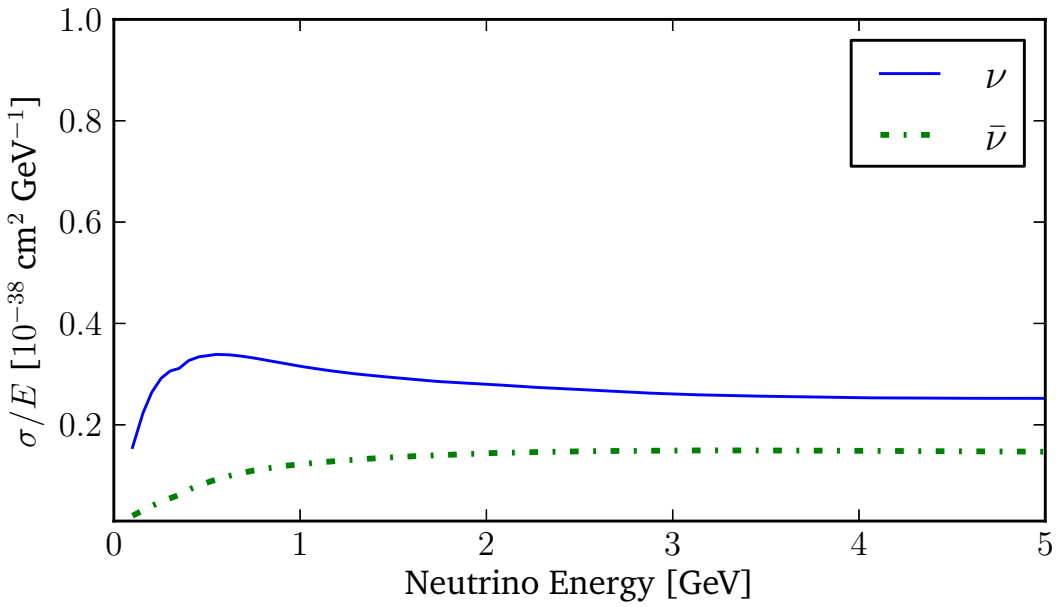

(b) NC interactions

Figure 6.5: Neutrino–nucleon cross sections for an isoscalar target as a function of neutrino energy for CC and NC interactions.

1%, and thus are ignored.

6.4 Interaction rates: $N_{\mathbf{int}}$

The number of neutrino interactions is computed, which does not require assumptions about the detector technology. The interaction rates can be computed by

$$N_{\text{int}} = \Phi \times P_{e\mu} \times \sigma \tag{6.8}$$

$$= \Phi \times P_{e\mu} \times \sigma_{\text{nucleon}} \times N_{\text{nuclei}} \tag{6.9}$$

for flux Φ, oscillation probability (Prob), and cross section σ, where all of these quantities have been computed in the previous sections.

Using the LSND anomaly best fit parameters (Table 6.2) as an example sterile neutrino signal, the event rates for μ^+ and μ^- decays are shown in Table 6.4. Assuming 1.8×10^{18} decays, with either stored μ^+s or stored μ^-s, the *statistical* significance of all channels is greater than 10σ. Combining the NC channels together results in a statistical significance of 20σ and 17σ for stored μ^+ and μ^-, respectively. However, these $\tilde{6}$% disappearance measurements will be *systematically* limited due to the 1% normalization uncertainty. Neither $\nu_e \to \nu_\mu$ CC nor $\bar{\nu}_e \to \bar{\nu}_\mu$ CC interactions have backgrounds from other oscillations with similar final states, except for negligible atmospheric-term oscillations; therefore the backgrounds will be misidentifying interactions arising from nonsignal oscillation channels.

With the disappearance signals, it is not possible in the experimental setup described above to make any useful measurements. As has been discussed, a near detector has not been assumed for the appearance analysis and the uncertainties this introduces makes it impossible to measure these 6% deficits. The flux uncertainties are conservatively expected to be a few percent, unlike the typical NF assumption of 0.1%. Therefore, it is not clear whether or not these disappearance channels could be probed in a MIND

Table 6.4: Interaction rates for 10^{21} POT and 1.5 kt target mass for the no-oscillations and short-baseline oscillations described by Table 6.2. The statistical significances are computed. The combined statistical significance of NC events are 20σ and 17σ for stored μ^+ and μ^-, respectively. No physics backgrounds exist for $\nu_e \to \nu_\mu$ CC interactions. Question marks indicate that there will be a nonzero number of events, where the exact value requires further work.

	Channel	N_{null}	$N_{\text{osc.}}$	Difference	$(N_{\text{osc.}} - N_{\text{null}})/\sqrt{N_{\text{null}}}$
μ^+	$\nu_e \to \nu_\mu$ CC	0	332	∞	∞
	$\bar{\nu}_\mu \to \bar{\nu}_\mu$ NC	50073	47679	–4.8%	–10.7
	$\nu_e \to \nu_e$ NC	78805	73941	–6.2%	–17.3
	$\bar{\nu}_\mu \to \bar{\nu}_\mu$ CC	128433	122322	–4.8%	–17.1
	$\nu_e \to \nu_e$ CC	230766	216657	–6.1%	–29.4
π^+	$\nu_\mu \to \nu_\mu$ CC	?	?	?	?
	$\nu_\mu \to \nu_e$ CC	?	?	?	?
μ^-	$\bar{\nu}_e \to \bar{\nu}_\mu$ CC	0	117	∞	∞
	$\bar{\nu}_e \to \bar{\nu}_e$ NC	32481	30511	–6.1%	–10.9
	$\nu_\mu \to \nu_\mu$ NC	69420	66037	–4.9%	–12.8
	$\bar{\nu}_e \to \bar{\nu}_e$ CC	82589	77600	–6.0%	–17.4
	$\nu_\mu \to \nu_\mu$ CC	207274	197284	–4.8%	–21.9
π^-	$\bar{\nu}_\mu \to \bar{\nu}_\mu$ CC	?	?	?	?
	$\bar{\nu}_\mu \to \bar{\nu}_e$ CC	?	?	?	?

detector. For channels other than muon disappearance, there is the added difficulty of separating ν_e CC and NC event; one could perform an analysis that only looked at short tracks, but even ignoring the muon CC backgrounds, the flux uncertainties once again make this measurement insensitive. Probing disappearance channels requires a near detector, and possibly also a LAr TPC, to extract all the potential physics from this beam. However, this is beyond the scope of this thesis since only the appearance channel is of interest here.

The number of events can also be determined as a function of energy since the evolution of ρ, σ, and (Prob) as a function of energy is known. These event distributions are shown in Fig. 6.6.

Numerous channels probe the sterile neutrino parameter space. Most other experiments have one or two channels to explore (see Table 2.3 for list of experiments), whereas in the best case νSTORM allows for 14 signals (Table 6.4).

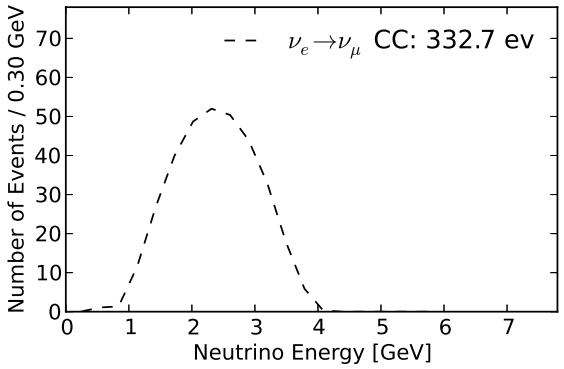

(a) μ^- appearance with stored μ^+.

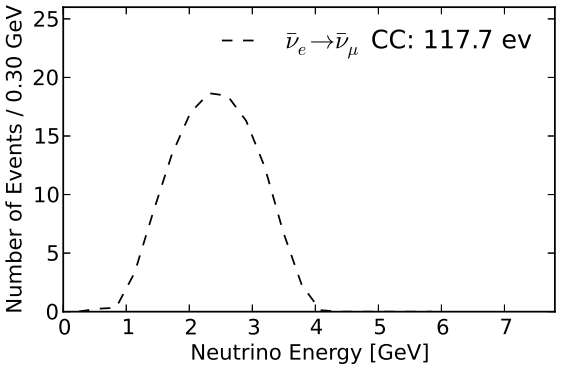

(b) μ^+ appearance with stored μ^-.

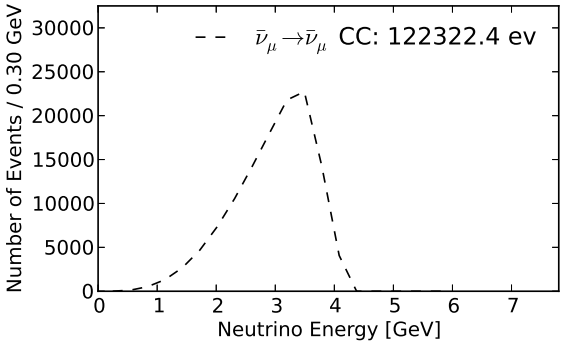

(c) μ^+ disappearance with stored μ^+.

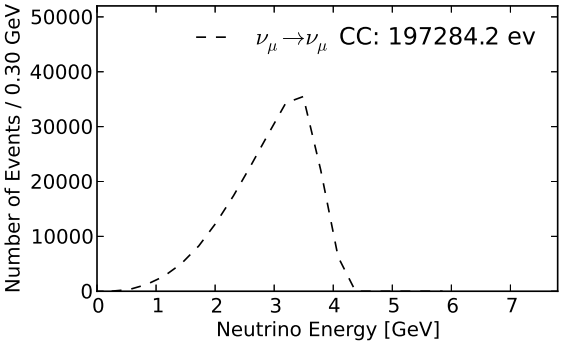

(d) μ^- disappearance with stored μ^-.

Figure 6.6: Interaction rate energy distributions assuming the LSND anomaly best-fit values (Table 6.2), 10^{21} POT, 1.5 kt of detector, a baseline of 2 km, and 3.8-GeV stored muons. The transitions $\nu_e \to \nu_\mu$, $\bar{\nu}_e \to \bar{\nu}_\mu$, $\bar{\nu}_\mu \to \bar{\nu}_\mu$, and $\nu_\mu \to \nu_\mu$ are shown.

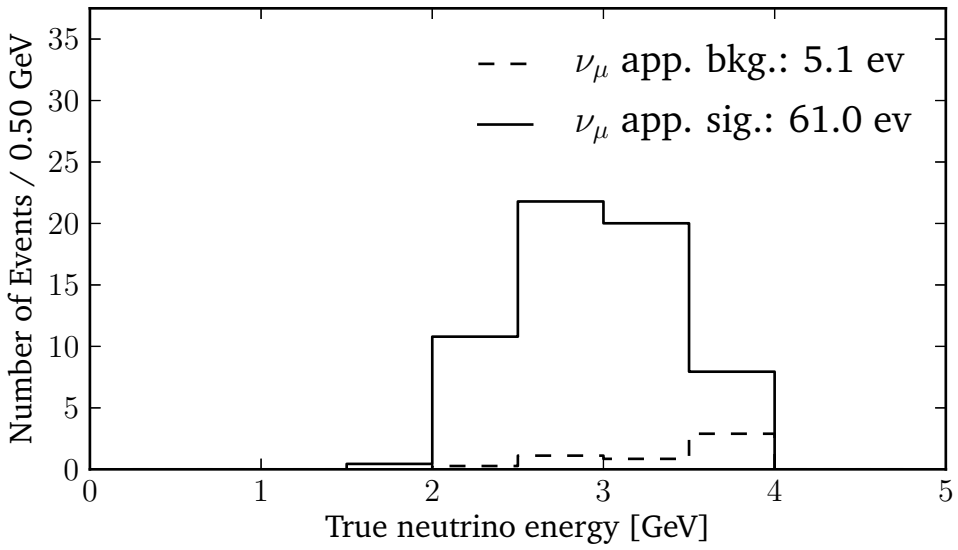

Figure 6.7: Event rates after cuts vs true neutrino energy.

6.5 Event rates after cuts

The energy-dependent interaction rates can be convolved with the detector performance to yield the measured neutrino total rates. The detector performance information is contained within the efficiency curve (Fig. 5.15). As seen in Fig. 6.7, applying these efficiency curves to the example interaction rates shown in Table 6.4 results in 61 signal and 5 backgrounds. The subsequent sensitivity analysis is a total rate analysis; an analysis that included shape information will not greatly improve the sensitivity since the oscillation maximum is below the energy threshold of the detector.

6.6 Statistics

It is necessary to determine if the number of events observed after cuts is statistically significant. In this section, statistics will be reviewed to explain later work. More specifically, a χ^2 test statistic is derived that can be used in the sensitivity analysis since the experiment must reject the null hypothesis, accounting for statistical fluctuations.

The hypothesis of no oscillations H_0 is the null hypothesis and H_1 the alternate hypothesis. These hypotheses have oscillation parameters associated with them: let $\Theta_0 = \{\Delta m^2_{41}, \sin^2 2\theta_{e\mu}\}$ be the oscillation parameters associated with H_0, and similarly Θ_1 for H_1.

The *test statistic* X is a function of the experimental observations and let w be the values of X that suggest that the null hypothesis H_0 is not true. It is desirable to have a small probability of X—by statistical fluctuations alone—taking a value in w when H_0 is true. A level of significance α can be defined

$$P(X \in w|H_0) = \alpha, \tag{6.10}$$

where α corresponds to, colloquially, "5σ" when $\alpha \simeq 2.8 \times 10^{-7}$ and "10σ" when $\alpha \simeq 7.6 \times 10^{-27}$. The number of "$\sigma$" corresponds to the p value of having a greater than $n\sigma$ upward fluctuation of a Gaussian centered at zero.

The test statistic used for hypothesis testing is the likelihood ratio test. For this total rates analysis, there is only one observation that is designated X and has a probability distribution function $f(X|\Theta)$. The likelihood function is

$$\begin{aligned} L(X|\Theta) &= f(X|\Theta) && (6.11)\\ &= e^{-\lambda}\lambda^X/X!, && (6.12) \end{aligned}$$

where λ is the expected number of events with X actual events, and depends on Θ. The distribution is Poisson because the background levels are small. The short-baseline parameters Θ_1 for H_1 are free to take any value, but the parameters Θ_0 are fixed to zero by the null hypothesis requiring no oscillations. The likelihood ratio test defines a test statistic λ such that

$$\Lambda = \frac{L(X|\Theta_0)}{\max_{\Theta_1} L(X|\Theta_1)}, \tag{6.13}$$

where the denominator is maximized with respect to Θ_1 and the numerator remains fixed. Using Eq. (6.12) leads to

$$\Lambda = e^{-\lambda+X}\,(\lambda/X)^X\,. \tag{6.14}$$

From Wilks theorem, with enough statistics, the χ^2 can be defined as $\chi^2 = -2\ln\Lambda$ (see [109]), which is preferable to using λ because of floating-point precision considerations—multiplying numbers in a computer compounds the floating point errors. Using this definition, one finds

$$\chi^2 = -2\ln\Lambda = 2\left[\lambda - X + X\ln\left(\frac{\lambda}{X}\right)\right], \tag{6.15}$$

which has as many degrees of freedom as Θ_1 has parameters since the numerator of Eq. (6.13) has no degrees of freedom, and the denominator allows Θ_1 to take any value.

Within this analysis, there are only two degrees of freedom given that $\Theta_0 = \{\Delta m^2_{41}, \sin^2 2\theta_{e\mu}\}$. Sometimes when a χ^2 is defined as above it is called a $\Delta\chi^2$ since it refers to how adding parameters affects the goodness of fit of the model. Adding parameters will always improve a fit because the "noise" within the data will be fit; however, adding sterile neutrino parameters should improve the fit beyond what would be expected from statistical fluctuations alone.

6.7 Appearance sensitivity

The parameters to be explored in the appearance analysis are Δm^2_{41} and $\sin^2(\theta_{e\mu}) = 4|U_{e4}|^2|U_{\mu 4}|^2$. Contours in the neutrino parameter space Δm^2_{41} vs $\sin^2(\theta_{e\mu})$ can be used to compare the sensitivities of various proposed short-baseline experiments.

Care must be taken when defining $\chi^2(\Delta m^2_{41}, \sin^2(2\theta_{e\mu}))$ to ensure that it is well defined. In the (3+1) scenario, the signal $\nu_e \to \nu_\mu$ depends on the amplitude $\sin^2(2\theta_{e\mu}) =$

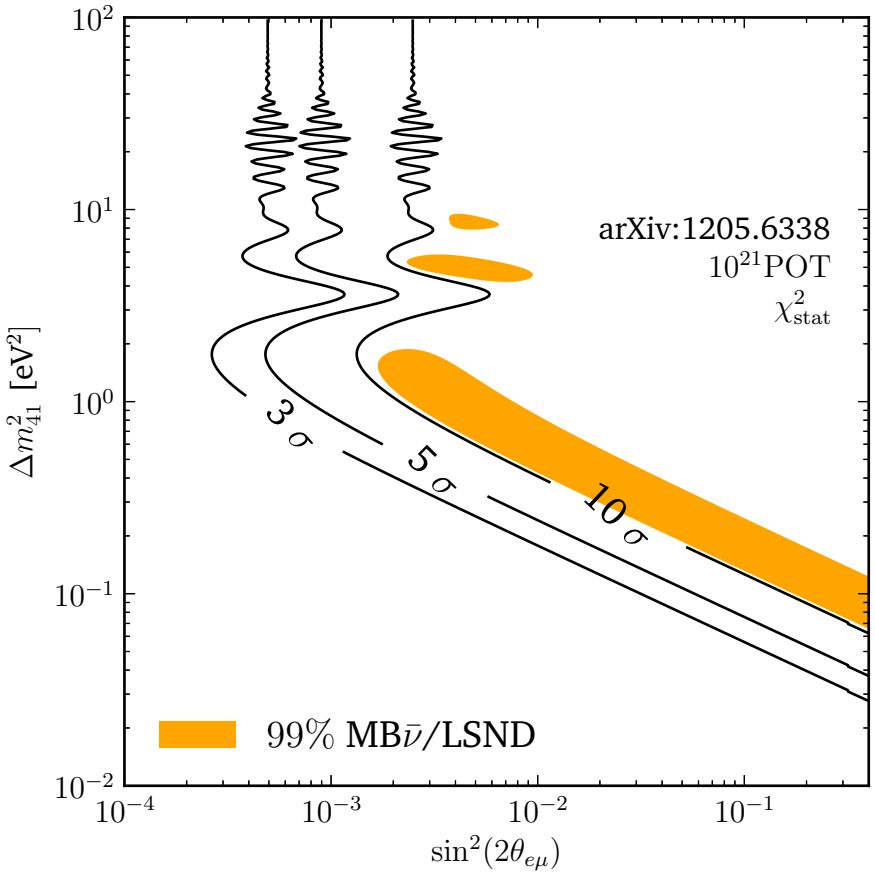

Figure 6.8: Sterile neutrino sensitivity using a statistical χ^2 of the appearance channel $\nu_e \to \nu_\mu$ assuming 10^{21} POT, a stored μ^+ beam, and 1.5 kt detector at 2 km. This channel is the CPT of the LSND anomaly $\bar{\nu}_\mu \to \bar{\nu}_e$. In this frequentist study, sensitivity to the LSND and MiniBooNE 99% confidence interval [104] is 10σ.

$4|U_{e4}|^2|U_{\mu 4}|^2$ and oscillation frequency Δm^2_{41} [Eq. (2.21)]. If there is an appearance signal, then $|U_{e4}|^2|U_{\mu 4}|^2 \neq 0$, which implies that both U_{e4} and $U_{\mu 4}$ are nonzero, resulting in disappearance. The disappearance of CC ν_e and NC backgrounds [Eq. (2.18)] affects the background estimation. However, by not oscillating the backgrounds (e.g., $\bar{\nu}_\mu \to \bar{\nu}_\mu$), the backgrounds are overestimated and the degeneracy avoided.

The statistical χ^2 is computed using the definition from the previous section. As can be seen in Fig. 6.8, there is over 10σ sensitivity to the entire LSND and MiniBooNE 99% confidence interval.

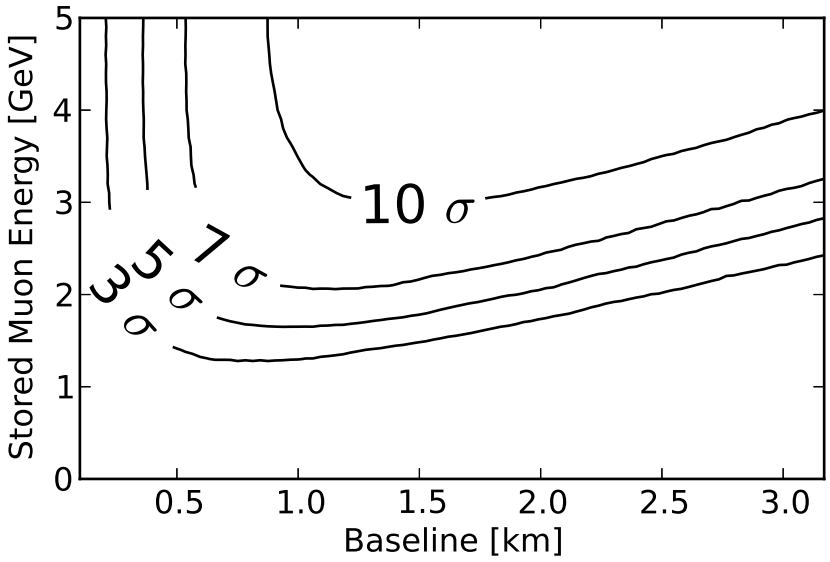

Figure 6.9: Baseline optimization using a total-rates statistics-only χ^2. On this plot, νSTORM's current baseline corresponds to 3.8 GeV and 2 km for the Fermilab siting.

6.8 Optimizations

It is important to understand how varying experimental parameters affects the sensitivity. For example, if there are site constraints, can the far detector be moved? Or if the charge misidentification is worse than expected, can this be compensated by building a better accelerator? These questions will be addressed now.

The optimization of baseline and energy seen in Fig. 6.9 allows one to change the baseline depending on site constraints or lower the energy of the ring if the accelerator gets too expensive. Interestingly, this shows that for a fixed baseline, increasing the muon energy is always advantageous. This effect arises because the maximum of the ν_e flux is not at the oscillation maximum but rather at a higher energy. As an illustrative exercise, at high energies the oscillation probability is

$$\Pr[\nu_e \to \nu_\mu] = \sin^2(2\theta_{e\mu}) \sin^2\left(\frac{\Delta m^2_{41} L}{4E}\right) \tag{6.16}$$

$$\simeq \sin^2(2\theta_{e\mu}) \left(\frac{\Delta m^2_{41} L}{4}\right)^2 E^{-2}. \tag{6.17}$$

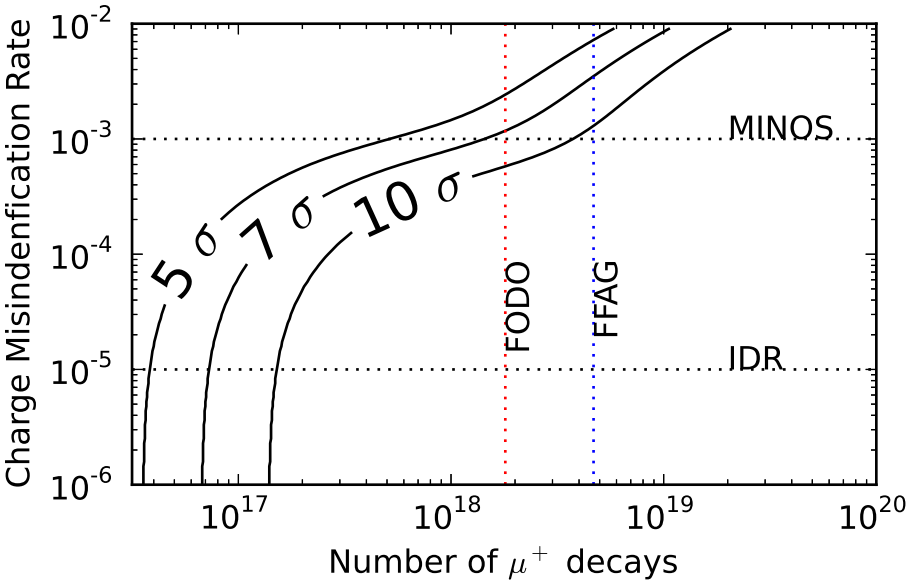

Figure 6.10: An optimization between the detector performance and accelerator performance using the charge misidentification rates and number of muon decays as the performance metric. IDR refers to the interim design report [57] expected detector performance. MINOS refers to their published performance. FODO refers to the FODO lattice design that gives 1.8×10^{18} useful muon decays, while FFAG refers to the FFAG design that gives 4.68×10^{18} useful muon decays. Both accelerators assume a front end of the main injector at 60 GeV/c for 10^{21} POT.

The oscillation probability decreases as E^{-2} for a fixed baseline. However, the signal rates increase as E^3: there is a factor of E^2 from the solid angle arising from the $1/\gamma$ opening angle, and another factor of E from the cross section. Or rather:

$$N = \sigma \times P_{e\mu} \times \Phi \tag{6.18}$$

$$\propto E \times E^{-2} \times E^2 \tag{6.19}$$

$$\propto E \tag{6.20}$$

The conclusion is that raising the stored muon energy will increase the event rates linearly with energy for a fixed baseline. As the event rates increase so do the background rates; therefore, the sensitivity only increases as $\sqrt{E}$. This effect has been noticed by similar analyses for other muon-decay-based facilities (see sensitivity work in [57]).

Even though the charge misidentification was derived on p.p. 51, understanding how the performance of νSTORM changes as this efficiency changes is important. Unlike the range cut, which has been used in numerous experiments, no neutrino experiment has been as reliant on the successful identification of muon charge—but collider experiments have demonstrated even better muon-sign discrimination, and even 10^{-4} electron-sign discrimination at these energies [110]. Accordingly, the charge misidentification rate is varied against the number of stored muons in Fig. 6.10. As can be seen, a reasonable number of muons using either the FODO or FFAG accelerator designs allows for enough statistics that there are reasonable requirements on the charge misidentification. For example, even if the charge misidentification rate was the same as MINOS [111], which has a much lower field and thicker plates, νSTORM could still make a 7σ detection of sterile neutrinos. However, based on the earlier arguments here and also estimates from the NF IDR, 10σ should be possible. These numbers will be further defended in Sec. 7.1.4, where uncertainties on the multiple scattering model are further discussed.

6.9 Summary

Using the detector performance work from the previous chapter and defining a χ^2, the sensitivity of νSTORM is 10σ across the LSND 99% confidence interval using only appearance information. Effects such as those related to the length of the accelerator straight have been included. Every step of the calculation—flux, oscillation probability, interaction rates—was shown to demonstrate the conservativeness of these estimates. Also, optimizations have been shown for the appearance channel sensitivity allows for variations on the experimental design. What remains to be done is ensuring that the systematic uncertainties and external background rates are as small as assumed, which will be shown in the subsequent chapter.

Chapter 7

Systematics and External Backgrounds

In the previous chapter, a statistical sensitivity analysis was performed, and demonstrated that νSTORM was sensitive to sterile neutrinos over the entire 99% C.L. indicated by LSND. However, the only claim made in the text so far is that the uncertainties should be small per other NF studies. Motivation for why this should be true was given, but these assertions need to be defended further. This is the first comprehensive estimate on what the systematic uncertainties should be for a muon-decay beam. In addition to the systematic discussion, external background rates from atmospheric neutrinos and cosmic rays will be computed. Afterward, a systematic χ^2_{sys} will be defined that includes both the systematics and external background information, before a sensitivity plot using χ^2_{sys} is shown.

7.1 Systematics

Systematic uncertainties must be identified and estimated, and the relevant ones here are:

1. hadronic and electromagnetic models,

2. magnetic field and steel, and

3. flux and cross sections.

After estimating the systematic uncertainty, the sensitivity was recomputed to determine the size of the effect. The physical nature of the cuts that have been used to extract the appearance signal eases assessing their impact. All of these systematics have been considered by MINOS and, in the end, they were shown to be small [112]. What must be shown is that, roughly, the five expected backgrounds do not have an uncertainty large enough to reduce the significance of the 61 signal; therefore, even background uncertainties of 100% will not affect the appearance analysis.

7.1.1 Hadron and electromagnetic models

Hadronic model uncertainties (e.g., $\pi + N$ reactions) are one of the dominant sources of systematic uncertainty for this type of detector and are not as well understood as electromagnetic models due to uncertainties arising from non-perturbative hadronization. Uncertainties in how, for example, pions are created from hadronizing quarks lead to uncertainties everywhere from targets for neutrino beams to calorimeters at a toroidal LHC apparatus (ATLAS).

Accordingly, experimenters typically put their calorimeters into test beams in order to measure their energy scales and particle response. For example, the Tile Calorimeter of ATLAS [113], CDF Endcap [114], CMS HCAL [115], and CDHS prototype [116] were all put in test beams. Both CERN and Fermilab have dedicated beam lines for this purpose. However, those experiments all operated at energies greater than 10 GeV—far beyond the energy range of interest for MINOS and thus νSTORM—which is why MINOS, whose energy range is 1–10 GeV, built the MINOS calibration detector (CalDet) at CERN. (Details can be found in [117–119].)

The CalDet was a scaled-down—yet unmagnetized—model of the MINOS detector that was built in order to understand the detector response in the 0.2–10 GeV/c momentum range for hadrons, electrons, and muons. This detector was put into the T11

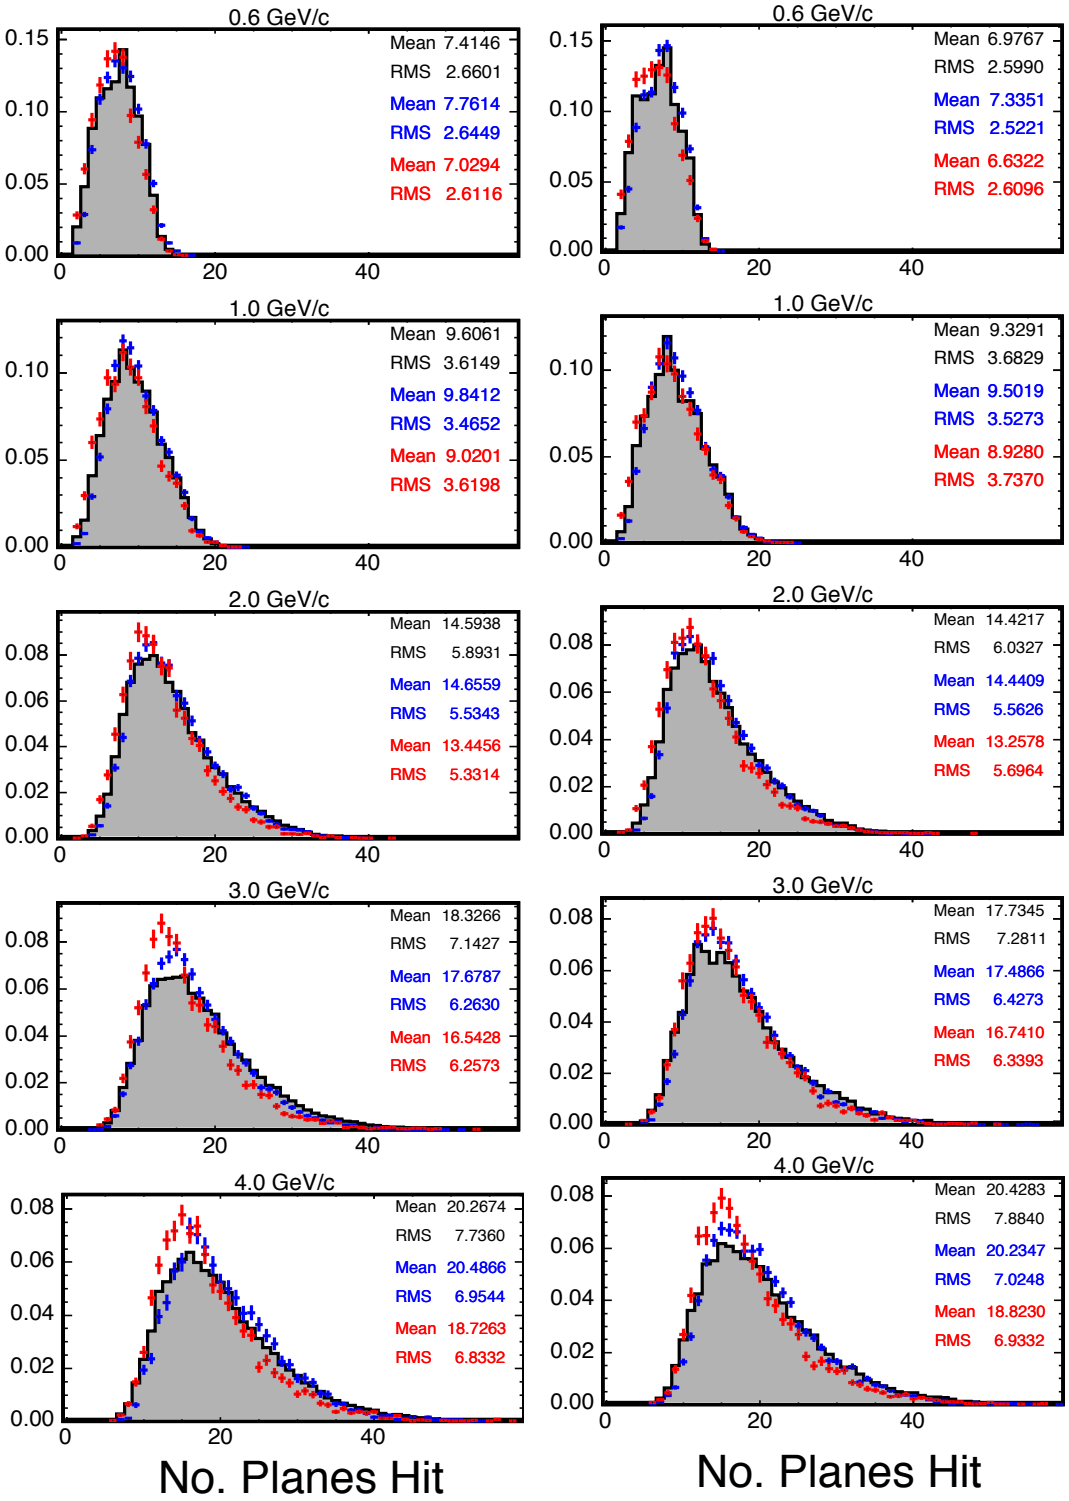

(a) π^- (b) π^+

Figure 7.1: Data and MC comparison for π^+ and π^- hit-plane multiplicity in CalDet. Data collected from T11 are shown shaded, pions simulated with GCALOR are shown in blue, and pions simulated with SLAC-GHEISHA are shown in red. Plots taken from thesis of Kordosky [118].

and T7 test beams in the CERN PS East Hall between 2001 and 2003. Time-of-flight and Čerenkov counters were used to identify the particle species, and the worst contamination was a 2.3% leakage of electrons into the pion sample at 10 GeV/c. The experimenters calibrated GMINOS—based on GEANT3—with this CalDet data. Two models were tested against the data: SLAC-GEISHA and GCALOR. The mean and RMS of the models agree with the data within ±15% over the energy range of interest for νSTORM (Fig. 7.1), which indicates that the simulation used for νSTORM should be correct to that level.

MINOS also noted a difference in the muon range between simulation and data. Part of the discrepancy was explained by a deficiency in how GEANT treated the density effect, which caused the $\langle dE/dx \rangle$ to be overestimated by 2%. Even accounting for this effect, in the data the muons appeared to be 3–4% longer than in the MC simulations.

Applying this information to νSTORM, the systematic errors indicated by the MINOS CalDet should result in uncertainties of ±4% for muon range and ±15% for pion ranges. A test-beam prototype, like CalDet, will be constructed for νSTORM as well. In addition to testing various physics models in the software, it is useful to determine calibration procedures and test the analysis chain. Also, which is less relevant to systematics, such an exercise will help understand the engineering integration. Before νSTORM is able to be built, these uncertainties should become smaller since the advanced European infrastructure for detectors at accelerators (AIDA) project at CERN will improve upon the physics models relevant to νSTORM by building a MIND-like detector.

7.1.2 Magnetic field and steel

The knowledge of magnetic field is important because $p_{\perp}^{B}$ is linear in B [Eq. (3.9)]. The systematic uncertainty arising from magnetic field errors is considered. Knowledge of the effective magnetic field is dependent on the knowledge of the steel, and experience from the MINOS experiment [45] can be extrapolated to νSTORM. The composition of the

Table 7.1: MINOS steel specifications per [120]. The flatness specification is half the ASTM A6 specification.

Quantity	Thin plates	Thick plates
Thickness [inch]	0.5	1
Thickness [mm]	12.7	25.4
Finished plate width [mm]	2000	3810
Finished width tolerance [mm]	±0.76	±0.76
Flatness over any 12 ft [mm]	8	14.5
Thickness tolerance [mm]	$^{+0.8}_{-0.254}$	$^{+1.8}_{-0.254}$
Waviness over full length [per 8 m]	8 waves	8 waves

steel will be well known. Any impurities can be below the 0.5% level by weight. MINOS made its steel plates from low-carbon, hot-rolled steel (AISI 1006 designation); the tolerances are satisfactory for νSTORM, with a flatness at half the ASTM A6 specification (Table 7.1). At each foundry run it will be possible to take block samples and measure an average steel density with an uncertainty of 0.3%. Each plane after construction can be weighted and the thickness measured. The variation of thickness should be negligible at ±0.3%, based on experiences with MINOS ultrasound thickness measurements. Also, the hysteresis (i.e., B-H) curve variation between steel plates was shown to be negligible.

Plane-to-plane magnetic field variations can occur from mechanical and chemical nonuniformity, but MINOS was able to get the field uncertainty at any given point to less than 3% [44]. Accordingly, the uncertainty in the integrated field which the muon encounters should be small. It is also possible to measure the B-H curve of the steel and perform a finite-element analysis with ANSYS. The current through the STL will be known to better than 10^{-4} by using current shunts at the power supply. Finally, a global absolute calibration may be possible by comparing the range vs curvature of cosmic-ray muons. Whatever magentic fluctuations exist should partially cancel, given that the muons will sample many plates.

7.1.3 Flux and cross sections

Flux and cross-section uncertainties provide a major source of uncertainty for conventional beams, but these uncertainties should be less at νSTORM. The optimistic assumption is that the flux can be measured to 0.1%, which is motivated from the 25-GeV NF work in [57]. Beam instrumentation—wire scanners, polarimeters, BCTs, and beam-position monitors (BPMs)— is used for flux measurements (previously discussed in Sec. 4.3.1). Ideally, the same instrumentation that is being considered for an NF should be used as a technology demonstration; however, the low beam current may provide difficulties. Therefore a conservative estimate of the flux normalization uncertainty is 1–3% based on conversations with beam instrumentation experts at CERN [121].

Cross-section uncertainties are infamous for complicating neutrino oscillation experiments. With νSTORM, the small flux uncertainty actually results in—in addition to the work presented here—a complementary experimental cross-section program that should be able to help νSTORM. However, such a program is not assumed for this analysis. The uncertainties without a near detector should be 10% for both QE and RES, and 5% for DIS [57, 122]. However, experiments such as MINOS measure the axial mass M_A, which is a parameter of the form factor, $G_A(q^2) = \frac{g_A}{(1-\frac{q^2}{M_A^2})^2}$, that encodes the nucleon structures, at the near detector and attempt to extrapolate the flux to their far detector. Axial mass uncertainties are $\mathcal{O}(10\%)$ [123], however this uncertainty only translates into a 2% normalization uncertainty on the observed neutrino-energy spectrum [124]. MINOS is able to predict the product of flux and cross section with an uncertainty of 4% [125] at their near detector.

The key advantage to νSTORM over other proposals is that it is able to reduce the cross-section uncertainties for appearance channels to the level of uncertainty typically associated with disappearance measurements—other neutrino experiments cannot measure their appearance signal at a near detector. At νSTORM, both $\overset{(-)}{\nu}_e$ and $\overset{(-)}{\nu}_\mu$ are in the beam, and if the decay ring can run in stored-μ^+ mode and stored-μ^- mode, then all

of these cross sections can be measured. Muon decay results in muon neutrinos, which also happens to be the signal from $\nu_e \rightarrow \nu_\mu$.

The flux uncertainty here is less than MINOS since the advantage of a muon-decay beam is the ability to better predict the flux, and cross section measurements better than MINOS should be possible for similar reasons. Therefore, it is conservative to assume a 4% uncertainty on the product of flux times cross section for the signal. For the appearance physics, the signal uncertainty is small because the physics behind CC QE and CC DIS is well understood.

However, background uncertainties are larger because backgrounds, such as pions, are produced by mechanisms such as resonant production, which are not as well understood. Therefore, background uncertainties will remain large and range from 20% to 40% [124], assuming they cannot be measured effectively in a MIND detector. These uncertainties may be further constrained with a complementary cross-section measurement program at νSTORM; however, these other experiments are not required.

A 4% and 40% normalization uncertainty is assumed for signal and background, respectively. The difference between signal and background is attributable to the difference in understanding of the neutrino–nucleon interaction physics. These uncertainties are derived from the experience and performance of existing experiments; therefore they may be decreased as future experiments improve the understanding of cross-section physics. However, both uncertainties are small relative to what is required of νSTORM.

7.1.4 Multiple scattering model

The level to which the muon charge can be identified is critical for sensitivity to sterile neutrinos. As was mentioned in Sec. 6.8, this level of background rejection has only been demonstrated experimentally at experiments such as ATLAS, where the muon energies are much higher than the 2 GeV considered here. However, the performance of the charge identification at 2 GeV using a MIND detector should be determined experimentally at

the νSTORM near detector. With 10^6 events at the near detector that should not have oscillated, there are enough statistics to set a 10^{-4} limit on the charge identification. By turning off the magnetic field, it should be possible to align the detection planes more than is required for this study.

There are two types of scattering that need to be addressed: hard scattering and multiple scattering. Hard scattering is also known as *Rutherford scattering*, where the muon scatters off of the nucleus (without seeing quark constituents), which has been modeled since 1911 [126]. This formalism assumes only one scatter. There can also be multiple Coloumb scattering, which can be derived from the Rutherford scattering model if there are numerous scatters [127]. Effects such as the electrons screening the charge of nucleus must be included. However, both of these models of scattering require only the classical picture of Rutherford scattering, which is well established.

These models were tested extensively in the 1950s with data on cosmic ray muon interactions [128]. The energy range that was explored included the range of interest at νSTORM: 1 to 4 GeV. Some anomalies were appearing in the larger q^2 data since it was not known that nuclei had quark constituents. Measurements of 2 GeV muons on carbon and lead were performed using an accelerator beam with a claimed purity of 10^{-6} [129], where no deviations from the Rutherford and Coloumb formalisms were found. Numerous other experiments have been performed (see, e.g., [130] for review) but no anomalies have been found in the low-q^2 region that νSTORM is considering.

It is nevertheless true that νSTORM requires better charge identification that has been previously performed with 2-GeV muons. Based on previous experiments, it is reasonable to assume that no new anomalies will be discovered (new physics?) that reduce the sterile sensitivity. However, if an anomaly is discovered then there are two possible situations. Either the scattering is soft but more than expected, and the near detector is able to measure this. Or the scattering is hard Rutherford scattering and will appear as an observable (and rejectable) kink in the path of the muon. There is every

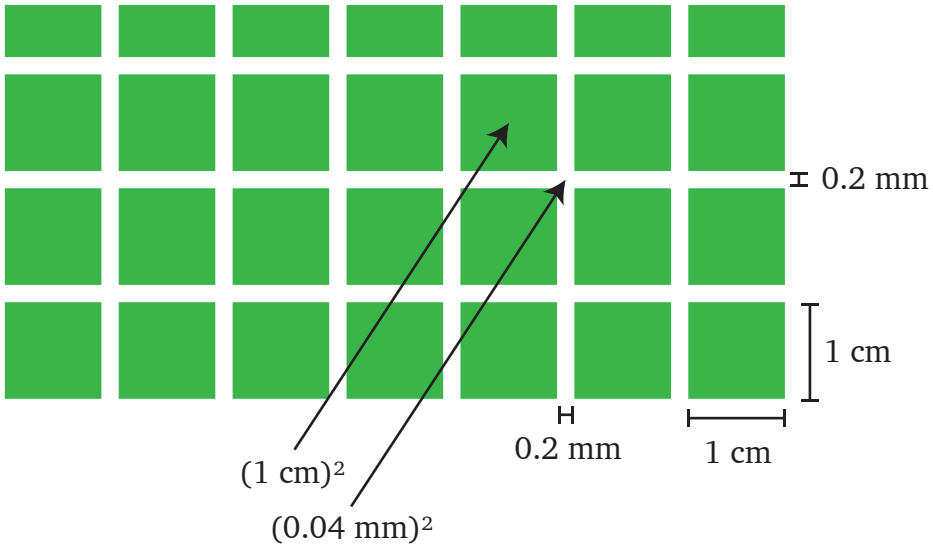

Figure 7.2: Detector cracks. Green corresponds to the pixels formed by 1-cm scintillator bars using both the x and y views. The white between the bars is a 0.2 mm gap.

reason to believe that these semi-classical models apply to νSTORM, and by measuring these models at νSTORM these assumptions can be checked.

7.2 External backgrounds

Backgrounds can arise from physics occurring outside the detector, which can potentially fake a wrong-sign muon. Three types of backgrounds are discussed: beam interactions with upstream material, atmospheric neutrinos, and cosmic-ray muons. Modest fiducial cuts are able to remove all of these backgrounds.

As seen in Fig. 7.2, there are "cracks" between the scintillator bars; therefore, before one can understand the efficacy of fiducial cuts, the detector cracks must be understood. There is a 100-μm layer around each scintillator bar that is reflective. Between any two 1-cm-wide bars is a 200-μm gap. Assuming muons had momentums perpendicular to the measurement plane, the area corresponding to the gap is $(200\ \mu\text{m})^2$. This as a percentage of the total area is $(200\ \mu\text{m})^2/(200\ \mu\text{m} + 1\ \text{cm})^2 \simeq 4 \times 10^{-4}$. Offsetting the bars between layers is required such that a muon that gets through the cracks will hit

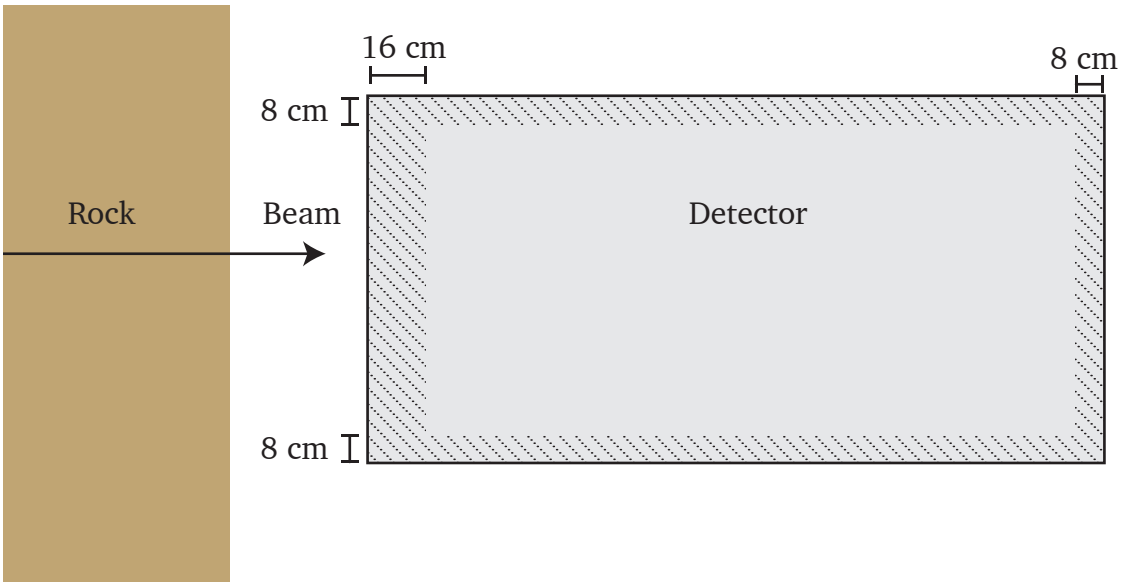

Figure 7.3: Sketch of fiducial cuts. Not to scale.

the center of the pixel in the next layer. The probability of a muon traversing two offset layers and going through the cracks is less than $(4 \times 10^{-4})^2 = 1.6 \times 10^{-7}$.

The SiPM hit efficiency is better than 90% for minimally-ionizing particles (MIPs). Assume a conservative 10% uncorrelated chance per bar that the particle is not detected. Requiring four modules to be traversed—eight scintillator bars—results in a probability of $(10\%)^8 = 10^{-8}$ and is sufficient.

The Poisson probability that a muon traverses a bar without the electronics registering a hit is also small. Assuming a conservative mean number of photoelectrons expected in a bar when a muon crosses is over 20 pe. The cut on noise is 2 pe; therefore the probability, assuming a Poisson distribution, of detecting less than or equal to 2 pe is 4×10^{-7}. Requiring eight scintillator bars to be hit makes this probability negligible.

Muons entering the detector can be rejected at the 10^{-8} level by vetoing on the eight outmost scintillator bars. As seen in Fig. 7.3, the upstream fiducial cut is 16 cm and the fiducial cuts on the sides are 8 cm.

Table 7.2: Material properties surrounding the DAB [2, 132]. Estimates for the range in till are shown.

	Liquid water	Till	Limestone	Std. rock
$\langle Z/A \rangle$	0.55		0.4955	0.5
Mean excitation energy [eV]	75		136.4	136.4
Density [g/cm^3]	1.0	2.12	2.8	2.65
Range 4 GeV μ [g/cm^2]	1.8×10^3	$\approx 2 \times 10^3$	2×10^3	2.1×10^3
Range 4 GeV μ [m]	18	≈ 9.5	7	8

7.2.1 Beam interactions with upstream material "rock muons"

Having computed the probability of an external muon penetrating the fiducial volume unnoticed, it is now possible to explore the "rock muon" example. The neutrino beam will interact with the material between the near and far detectors before emerging into the front of the far detector. If any of these muons slip through calorimeter cracks or hits are not registered, then these muons could be backgrounds for physics analyses.

The material composition on the Fermilab site is well known due to previous construction projects on the site. Below the topsoil is 15 m of glacial till and then the bedrock [131]. The far detector hall (i.e., the DAB) is surrounded by glacial till with density between 2.12 and 2.32 g/cm^3 [132]. The density-independent range in till should have the lower limit of water and upper limit of standard rock (Table 7.2). These density-independent ranges are similar and therefore the range 2×10^3 g/cm^2 is also used for till. Accordingly, 4-GeV muons are anticipated to penetrate 9.5 m of till.

Given this range, consider a cuboid of size 9.5 m $\times \pi \times$ (2.5 m)^2, which corresponds to a target mass of about 0.4 kt. Muons that interact within this target mass can penetrate the far detector, and to overestimate the background *all* muons are assumed to enter the detector. The till mass is the same order of magnitude as the detector mass; therefore, there will be comparable numbers of neutrino interactions. These interactions must be rejected to the same level to which other backgrounds were rejected: $\mathcal{O}(10^{-6})$. Given the 10^{-8} background rejection using the fiducial cut mentioned above, rock muons will

Table 7.3: Duty factor and atmospheric neutrino rates. The duty factor is understood under different assumptions as to how the proton beam from the MI is extracted. The muon beam in the decay ring is dumped after 76 μs (i.e., $3\gamma\tau$).

Extraction	Rep. rate	Duty factor	Atmos. ν events
Fast kicker; two MI bunches	1.2 s	6×10^{-5}	0.01
Slow kicker; entire MI fill	12 s	6×10^{-6}	0.001

not present a problem based on the fiducial cuts described above.

7.2.2 Atmospheric neutrinos

Cosmic rays, which are energetic extraterrestrial protons (90%), helium (9%), and heavy nuclei (< 1%), interact with the Earth's atmosphere to produce pions. Pions and muons, upon decaying, will produce atmospheric neutrinos in the energy range of interest.

Interactions of the atmospheric neutrinos within the detector can be rejected using timing and directional information. The Barr-Gaisser-Stanev model of the atmospheric neutrino flux is used [133]. For our energies of interest, the effects of the solar cycle can be ignored. Using the tables in [133], a 10-kt-years exposure (2-kt detector for 5 years) results in 74 $\bar{\nu}_\mu$ CC and 156 ν_μ CC interactions above 2 GeV. These rates agree with those seen by MINOS, which also resides in the American Midwest.

Atmospheric neutrinos should be negligible for νSTORM [111] once reasonable cuts are applied. A timing cut alone reduces the background rate to the <1 levels by taking advantage of the duty factor of the machine (Table 7.3). The duty factor depends on how the proton beam is extracted from the MI but is about 10^{-5} and 10^{-6} for fast and slow extraction, respectively, if the muon beam is only used for three lifetimes. Counter-intuitively, a slow kicker is preferred. The expected number of atmospheric neutrinos is at most 0.01 events.

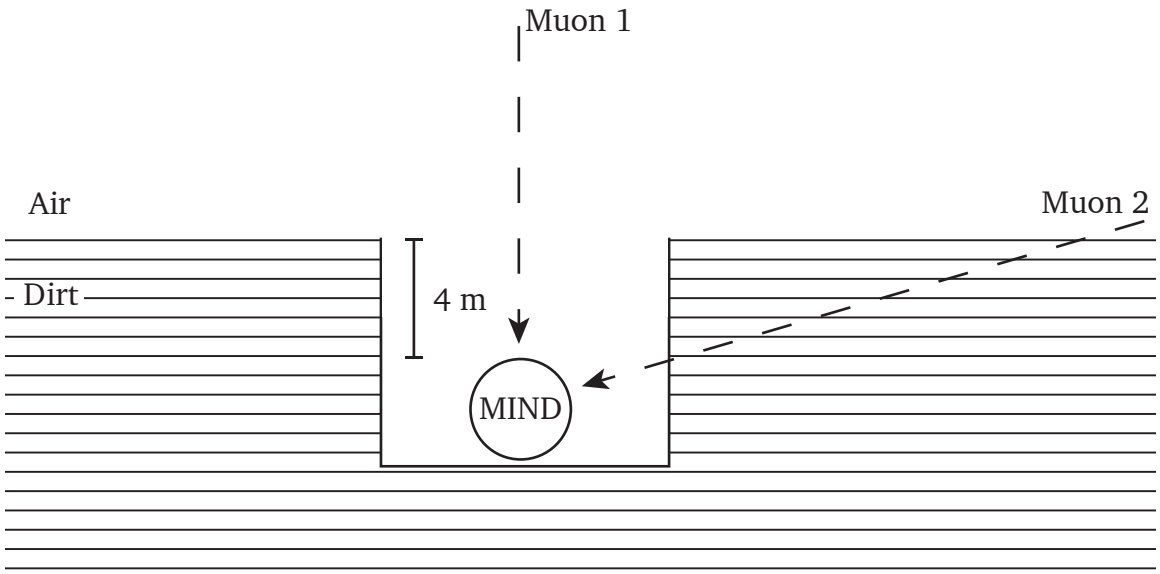

Figure 7.4: The overburden surrounding the DAB pit. The detector is assumed to be 4 m below the surface. Cosmic-ray muons may either enter the detector from the unexposed top (muon 1) or from the sides (muon 2). Muons entering from the side will travel through rock, which introduces an energy threshold for shallow angles.

7.2.3 Cosmic-ray muons

Only muons with energies below 2 GeV decay before arriving at the Earth's surface. Therefore, there are energetic muons (>2 GeV) that must be rejected from the appearance analysis. MIND detectors have, nevertheless, been used on the surface before by experiments such as CDHS. The calculation of cosmic-ray muon rates extends those for the MINOS near detector [134] and are an order of magnitude estimate.

The detector is in the DAB pit, where the top of the detector is approximately 4 m below the topsoil (Fig. 7.4). Given that the energy threshold of the detector (2 GeV) corresponds to 4 m of till (see Table 7.2 for further material information), the anticipated geometry of the ground around the detector can be approximated as if the detector were covered by this 4 m of till.

The Gaisser expression [136] is used to model the muon flux. However, this expression overpredicts the <10-GeV muon flux and can be corrected by using the modified Gaisser formalism (Fig. 7.5). For clarity and conservativeness, this study will use the

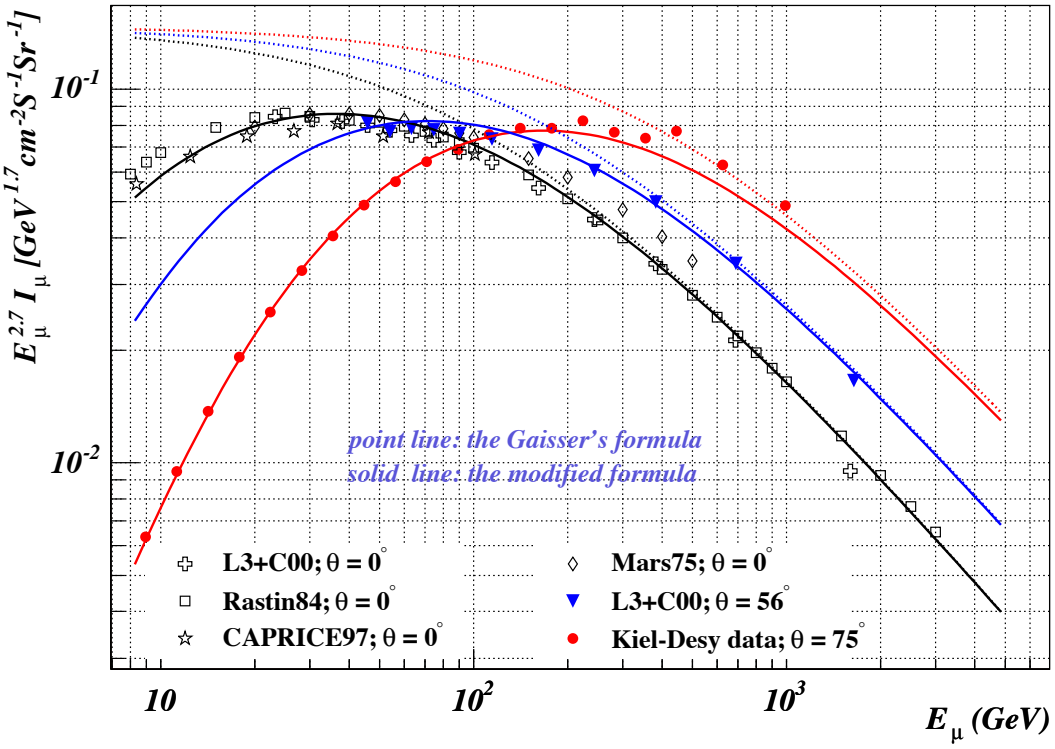

Figure 7.5: A comparison of the Gaisser formalism and the modified Gaisser formalism. Figure reproduced from [135].

Gaisser formalism that predicts twice the number of atmospheric muons at lower energies. The differential muon energy spectrum at the surface is

$$\frac{dN}{dE} = 0.14E^{-2.7}\left[\frac{1}{1+\frac{1.1E\cos\theta}{115\text{ GeV}}} + \frac{0.054}{1+\frac{1.1E\cos\theta}{850\text{ GeV}}}\right] \tag{7.1}$$

in units of muons/cm^2/sec/sr/GeV, where θ is the zenith angle of the muon and E is measured in GeV. The two terms in the brackets are approximately unity below 10 GeV:

$$\frac{dN}{dE} = 0.14E^{-2.7}. \tag{7.2}$$

The required muon energy to reach the detector without decaying will increase as the zenith angle increases, since the muon must penetrate more rock and atmosphere. Approximating the Earth as flat, the energy threshold is $E_{\text{min}}/\cos\theta$; therefore the angular distribution underground is

$$\begin{aligned}
\frac{dN}{d\Omega} &= \int_{E_{\text{min}}/\cos\theta}^{\infty} \frac{dN}{dE}\,dE & (7.3)\\
&= \frac{0.14}{1.7}\left(\frac{\cos\theta}{E_{\text{min}}}\right)^{1.7} & (7.4)\\
&= 0.082\left(\frac{\cos\theta}{E_{\text{min}}}\right)^{1.7}\ \text{muons/cm}^2\text{/sec/sr.} & (7.5)
\end{aligned}$$

To determine the number of muons that penetrate the detector, the integral must be calculated for each of the faces of the detector. Faces may either be horizontal or vertical

Table 7.4: Cuboid parameters for computing cosmic-ray rates.

Dimensions (m^3)	A_h (m^2)	A_v (m^2)	R_h (Hz)	R_v (Hz)	Total (Hz)
$5 \times 5 \times 13$	65	180	273	203	476

and have a flux dN/dA_h or dN/dA_v, respectively, where

$$\begin{aligned}
\frac{dN}{dA_h} &= \int_0^{2\pi}\int_0^1 \cos\theta \frac{dN}{d\Omega}\, d(\cos\theta)\, d\phi && (7.6)\\
&= 2\pi \int_0^{\pi/2} \cos\theta \frac{dN}{d\Omega} \sin\theta\, d(\theta) && (7.7)\\
&= 2\pi \int_0^{\pi/2} \cos\theta \left[0.082\left(\frac{\cos\theta}{E_{\text{min}}}\right)^{1.7}\right] \sin\theta\, d(\theta) && (7.8)\\
&= \frac{0.082 \times 2\pi}{E_{\text{min}}^{1.7}} \left[\frac{-1}{3.7}\cos^{3.7}(x)\right]_0^{\pi/2} && (7.9)\\
&= \frac{0.082 \times 2\pi}{E_{\text{min}}^{1.7}} \frac{1}{3.7} && (7.10)\\
&= 4.2 \text{ muons/m}^2\text{/sec} && (7.11)
\end{aligned}$$

and

$$\begin{aligned}
\frac{dN}{dA_v} &= \int_{-\pi/2}^{\pi/2} \cos\phi \int_0^1 \sin\theta \frac{dN}{d\Omega}\, d(\cos\theta)\, d\phi && (7.12)\\
&= 2\int_0^{\pi/2} \sin\theta \frac{dN}{d\Omega} \sin\theta\, d(\theta) && (7.13)\\
&= 2\int_0^{\pi/2} \sin\theta \left[0.082\left(\frac{\cos\theta}{E_{\text{min}}}\right)^{1.7}\right] \sin\theta\, d(\theta) && (7.14)\\
&= 1.13 \text{ muons/m}^2\text{/sec.} && (7.15)
\end{aligned}$$

where $E_{\text{min}} = 2$ GeV and dN/dA_v is computed numerically since the solution involves hypergeometric functions.

With the fluxes computed for vertical and horizontal surfaces it is possible to determine the number of muons incident upon the detector. Defining a cuboid containing a 2-kt detector with 2.5-m radius plates allows for the rates to be computed for each

surface the muon first penetrates (Table 7.4).

The rates are manageable at 476 Hz. Including the accelerator duty factor, the rates will be 0.05 and 0.5 Hz for fast and slow extraction, respectively. Over five operational years of 5×10^7 s, there will be at most 2.5×10^7 cosmic rays and background rejection of $1/(2.5 \times 10^6) = 4 \times 10^{-8}$ is required. Assuming cracks can be managed, then the fiducial cuts described for rejecting rock muons are sufficient.

The fiducial cuts already described will be more than adequate for rejecting cosmic-ray muons using an eight-outermost-bar fiducial cut (10^{-8} rejection). Cosmic-ray muons should never fake the muon signal. Independent work repeated this study using CRY and GEANT4, coming to the same conclusions within 50% [137], which is suitable for these purposes.

Cosmic-ray muons also provide useful information. In addition to being useful for locating dead channels, they may also be used to measure the detector's magnetic field by comparing momentum measurements from range vs bending. Current data acquisition (DAQ) technology should allow for recording all these events.

7.3 χ^2_{sys}

At present, various systematics have been discussed and quantified; however, what remains is incorporating these results in the χ^2 so sensitivity plots can be made using this information. The χ^2 defined in the previous chapter includes only statistical effects, but it is possible to include systematic uncertainties by using the so-called "pull method" [138], which has also been used to analyze data in experiments such as SNO's low-energy-threshold analysis [139]. So-called "penalty terms" are added to the χ^2 and marginalized over (i.e., minimized). A normalization systematic is introduced—shape systematics can be introduced if this were not a total rates analysis—and this is incorporated into the χ^2 by introducing a scale factor $(1 + \zeta)$, where ζ is minimized. The

Table 7.5: List of systematics considered and how much they affect appearance physics. Hadronic model uncertainties are how well NC interactions are modeled within the MC. Electromagnetic models are how well understood muon interactions with steel in scintillator are understood.

Source	$\nu_e \rightarrow \nu_\mu$ CC	$\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$ CC	$\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$ NC
Hadronic model	0%	15%	15%
Electromagnetic model	2%	0	0
Magnetic field	<1%	<1%	<1%
Steel	0.2%	0.2%	0.2%
Flux	1%	1%	1%
Cross section	4%	40%	40%

uncertainty on ζ is σ_ζ, which is corresponds to one of the systematic uncertainties derived earlier. Generalizing to k systematics, χ^2_{sys} is then defined as

$$\chi^2_{\text{sys}} = \min_{\zeta_i} \left[\chi^2(\Theta_0, \Theta_1, \zeta_i) + \sum_j \left(\frac{\zeta_i}{\sigma_{\zeta_i}} \right)^2 \right], \tag{7.16}$$

where Θ_0 and Θ_1 are the previously defined oscillation parameters to be compared. The model to be tested is made to agree as much as possible within errors before the goodness of fit is determined.

The χ^2 that has been defined allows for determining if an experiment is sensitive to various oscillation parameters. The systematics discussed previously are summarized in Table 7.5, and are all normalization systematics. The dominant systematic are the cross section uncertainties in this setup because a near detector has not been assumed, though may be present in the final design. As will be seen, not having a near detector is not an issue due to the signal to background ratio of 61:5. The signal errors are less than the background errors since the cross-section model uncertainties are smaller for muon final states than nonmuon final states. The external background rates seen in Table 7.6 are negligible compared to the half dozen backgrounds from disappearance channels.

Using the systematic errors as defined, a contour comparing the $\chi^2_{\text{sys.}}$ and $\chi^2_{\text{stat.}}$ is shown in Fig. 7.6. The contour only displaying systematics information is shown in

Table 7.6: Background register.

Source	Events
Cosmic-ray muons	$\ll 1$
Atmospheric neutrinos	≤ 0.14

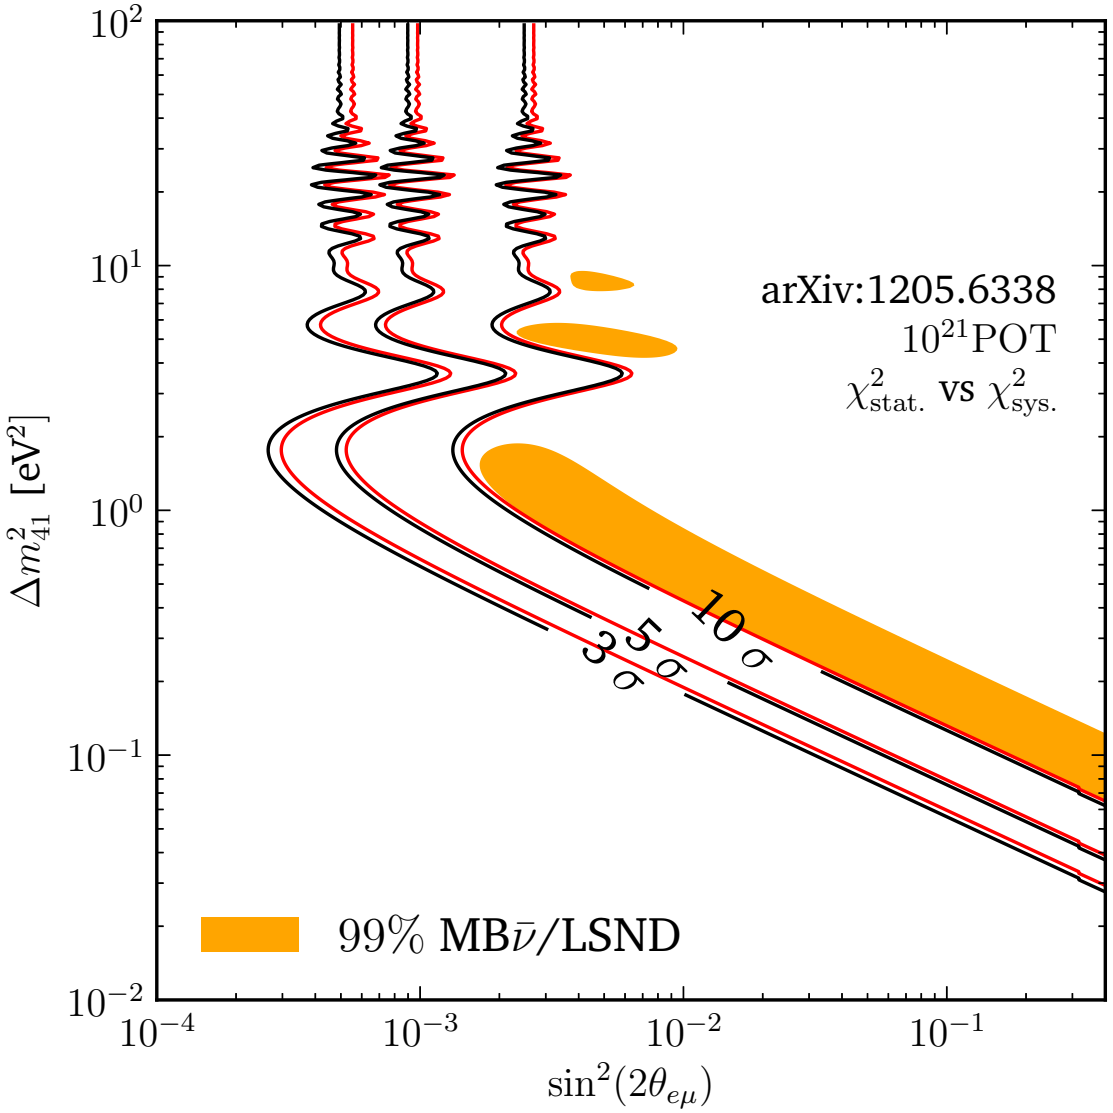

Figure 7.6: Sterile neutrino sensitivity comparing with (red) and without systematics (black) of the appearance channel $\nu_e \rightarrow \nu_\mu$ assuming 10^{21} POT, a stored μ^+ beam, and 1.5 kt detector at 2 km. For the systematic χ^2, a 4% signal and 40% background normalization uncertainty is assumed. This channel is the CPT of the LSND anomaly $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$. In this frequentist study, sensitivity to the LSND and MiniBooNE 99% confidence interval [104] is 10σ.

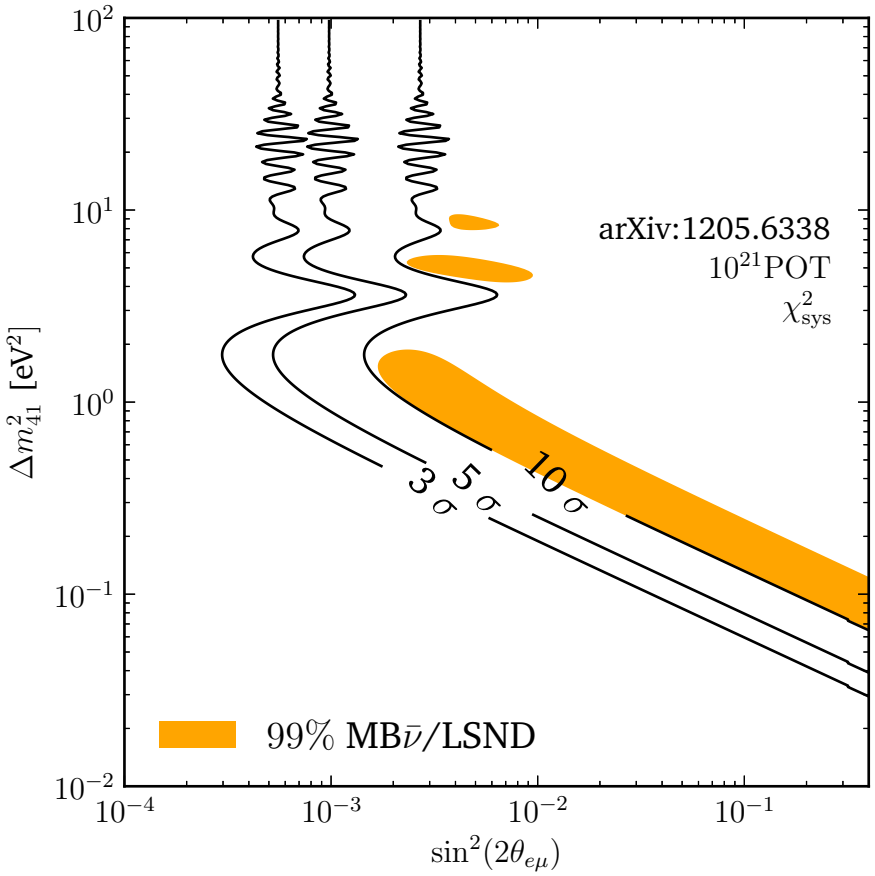

Figure 7.7: Sterile neutrino sensitivity using a systematics χ^2 of the appearance channel $\nu_e \rightarrow \nu_\mu$ assuming 10^{21} POT, a stored μ^+ beam, and 1.5 kt detector at 2 km. For the systematic χ^2, a 4% signal and 40% background normalization uncertainty is assumed. This channel is the CPT of the LSND anomaly $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$. In this frequentist study, sensitivity to the LSND and MiniBooNE 99% confidence interval [104] is 10σ.

Fig. 7.7. Due to the signal to background ratio of 61:5, introducing systematic errors does not affect the sensitivity of νSTORM (compare to Fig. 6.8). This can be understood as follows: if there are 61 signal—as is true throughout the LSND-favored region—and 5 background events, then even doubling the backgrounds does not affect the performance. If one expects 10 backgrounds, but sees 61 events, then it is clear that there is some new physics producing these additional appearance events.

7.4 Conclusions

Using the experience of MINOS, the systematics have been explored and shown to be negligible. Systematics related to the hadronic and electromagnetic interaction models were explored and shown to be small due to the work of MINOS and their CalDet. The field will be known within νSTORM to a better precision than is needed for the appearance physics since small field variations are averaged out over the path of the muon, and lowering the field by 20% results in similar detector performance. The composition and thickness of the steel can be measured accurately. Flux uncertainties will be small due to being able to measure the stored muon beam. Cross-section systematics did not affect the sensitivity. Cosmic ray and atmospheric neutrino backgrounds are negligible. Incorporating all of this information into a systematic χ^2 demonstrates that νSTORM is still sensitive at the 10σ level to the LSND effect.

Chapter 8

Oscillation-parameter estimation

Despite the central question of the thesis being how to determine if what LSND observes is sterile neutrinos, initial parameter-estimation calculations are presented. Measurement of the parameters Δm^2_{41} and $\theta_{e\mu}$ requires resolving their degeneracies in the oscillation probability. If only total rates are measured, then the oscillation probability $\Pr.[\nu_e \to \nu_\mu]$ is measured to be some value, but since

$$\Pr.[\nu_e \to \nu_\mu] \quad = \quad \sin^2(2\theta_{e\mu}) \sin^2\left(\frac{\Delta m^2_{41} L}{4E}\right) \tag{8.1}$$

it is impossible to determine Δm^2_{41} and $\theta_{e\mu}$ individually. The equation is underconstrained—there is one measurement, but two variables. To measure the parameter Δm^2_{41}, information in the shape of the neutrino energy distribution must be used; however, MINDs are sampling calorimeters with limited energy resolution. As will be seen, removing this parameter degeneracy requires using information from either the other oscillation channels or a better detector technology.

The determination of the precision requires understanding what the measured neutrino energy distribution will be for the signal. The neutrino energy resolution should be similar to previous experiments, which has been checked in [50]. The neutrino energy E_ν has two components: the muon energy E_μ and the hadronic energy E_{had}. From the experi-

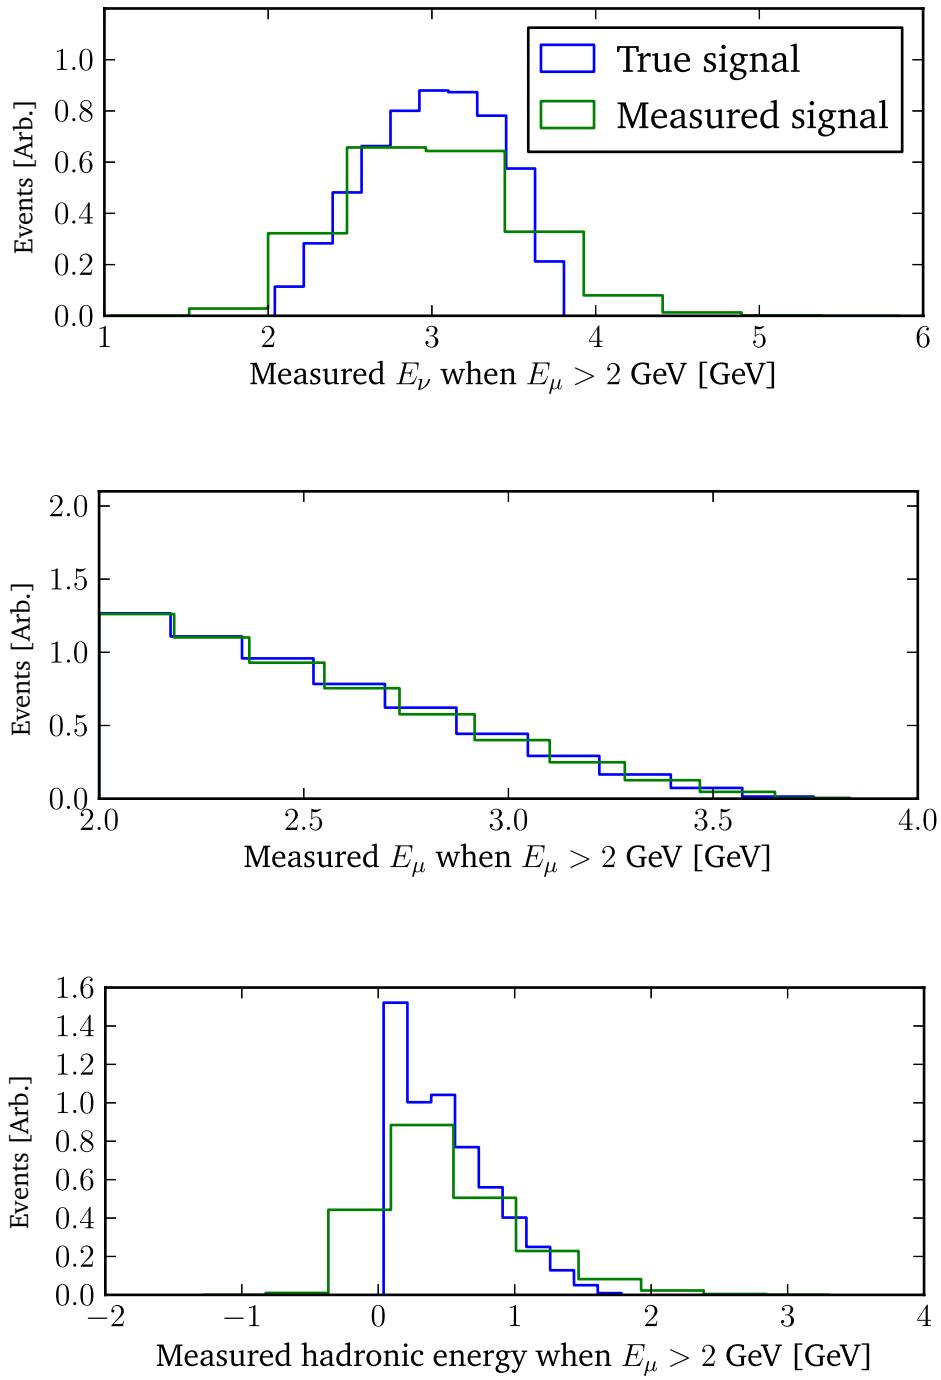

Figure 8.1: Overlaid kinematic distributions of E_ν, E_μ, and hadronic energy for averaging-regime ν_e and $\bar{\nu}_\mu$ fluxes. Negative values are retained to avoid any potential bias.

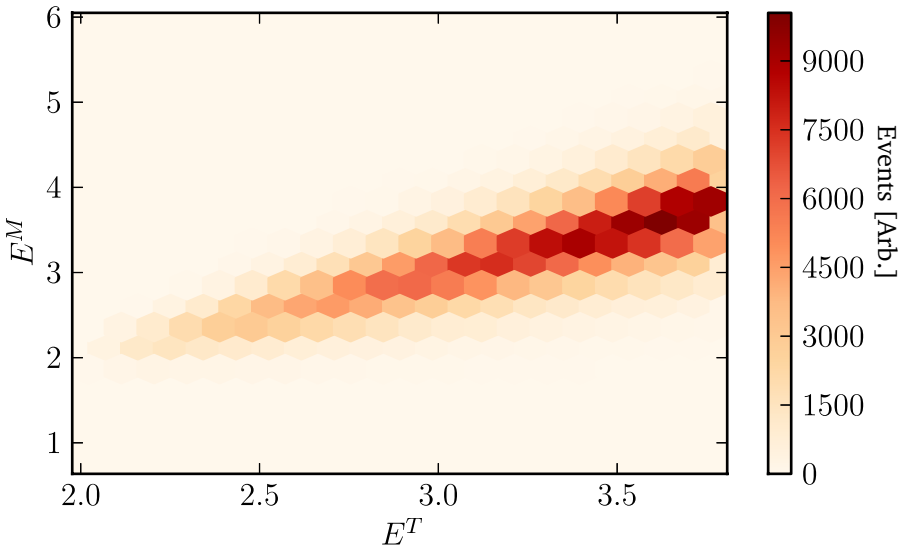

Figure 8.2: The measured energy E^M versus the true neutrino energy E^T.

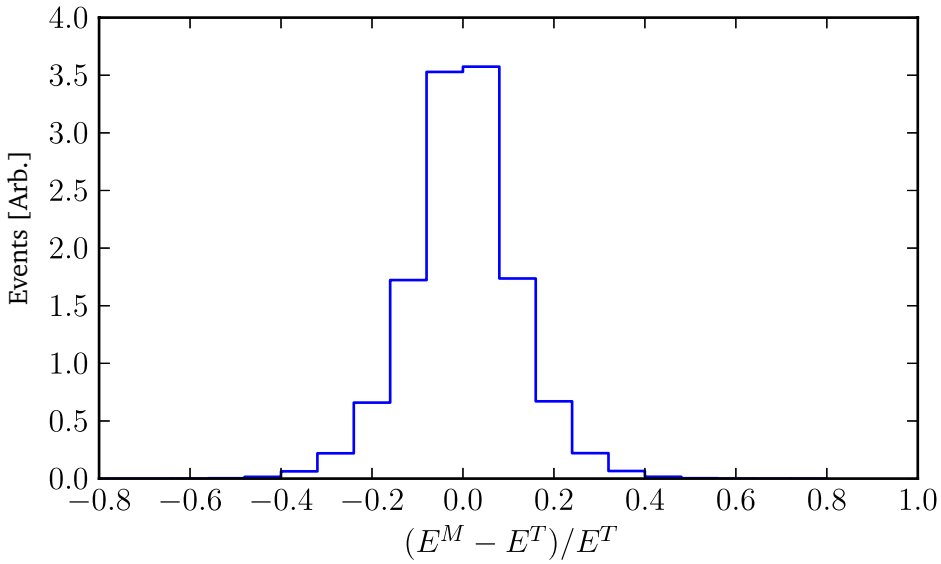

Figure 8.3: The relative error between the measured neutrino energy E^M and the true neutrino energy E^T. The RMS of the fractional difference is 12%.

ence of MINOS, the muon energy above threshold—remember the 2-GeV range cut leading to Fig. 5.15—can be measured to 5% based only on range. Also, summing all the hits in the hadronic shower should allow E_{had} to be measured to $50\%/\sqrt{E\ (\text{in GeV})}$ [140]. An estimate of the E_ν resolution can be made based on the E_μ and E_{had} resolutions and using a toy MC. In Fig. 8.1, the effect of convolving these resolutions with the unsmeared E_ν, E_μ, and E_{had} kinematic distributions presented in Fig. 5.3 is shown. The E_ν resolution arises from summing the smeared E_μ and E_{had} energy distributions. Comparing the unsmeared and smeared true neutrino energy in Fig. 8.2 and 8.3 shows that a 12% measurement of the neutrino energy should be possible.

Using this resolution, neutrino-energy distributions are shown in Fig. 8.4 for different Δm^2_{41}, and the significance of the difference is small due to low statistics. Nevertheless, a spectral likelihood fit is performed to determine the extend to which shape information aids parameter estimation.

The χ^2 discussion from p.p. 132 can be extended to include shape information. Specifically, instead of one measurement X, a set of $N = 16$ measurements $\mathbf{X} = \{X_0, X_1, ..., X_{15}\}$ is used, where each measurements corresponds to a 250 MeV bin in measured neutrino energy. For example, X_{15} corresponds to the number of events with measured energies between 3.75 and 4.0 GeV. The χ^2 definition follows from the generalization of the likelihood function and test statistic

$$L(\mathbf{X}|\Theta) = \prod_{i=1}^{N} f(X_i|\Theta) = \prod_i e^{-\lambda_i}\lambda_i^{X_i}/X_i! \tag{8.2}$$

$$\Lambda = \frac{L(\mathbf{X}|\Theta_0)}{\max_{\Theta_1} L(\mathbf{X}|\Theta_1)} = \prod_i e^{-\lambda_i + X_i}\left(\lambda_i/X_i\right)^{X_i}, \tag{8.3}$$

where λ_i is the expected number of events ($\lambda_i = P_{e\mu}(E) \times \sigma(E) \times \Phi(E) \times \epsilon(E)$) in the ith bin with X_i actual events. Therefore, as before, the χ^2 is:

$$\chi^2 = 2\sum_i \left[\lambda_i - X_i + X_i \ln\left(\frac{\lambda_i}{X_i}\right)\right]. \tag{8.4}$$

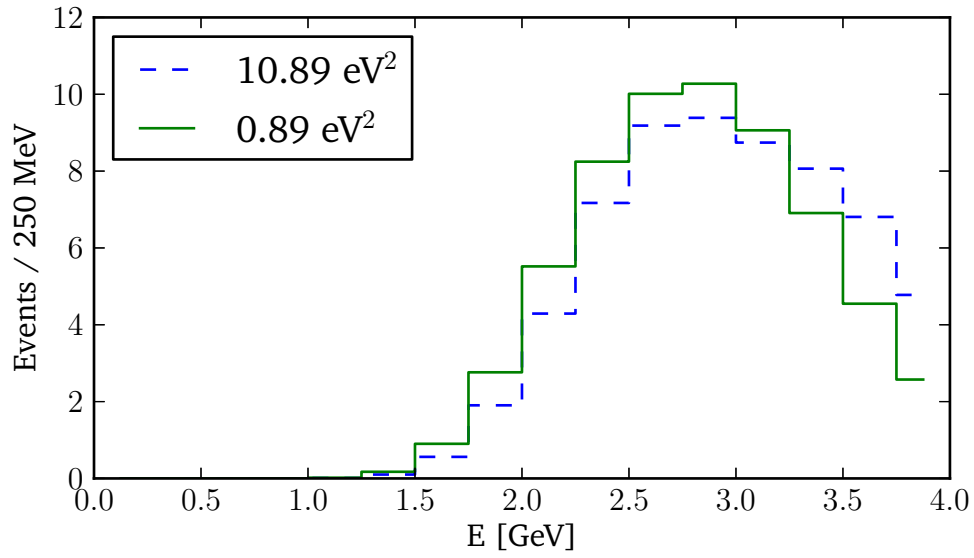

(a)

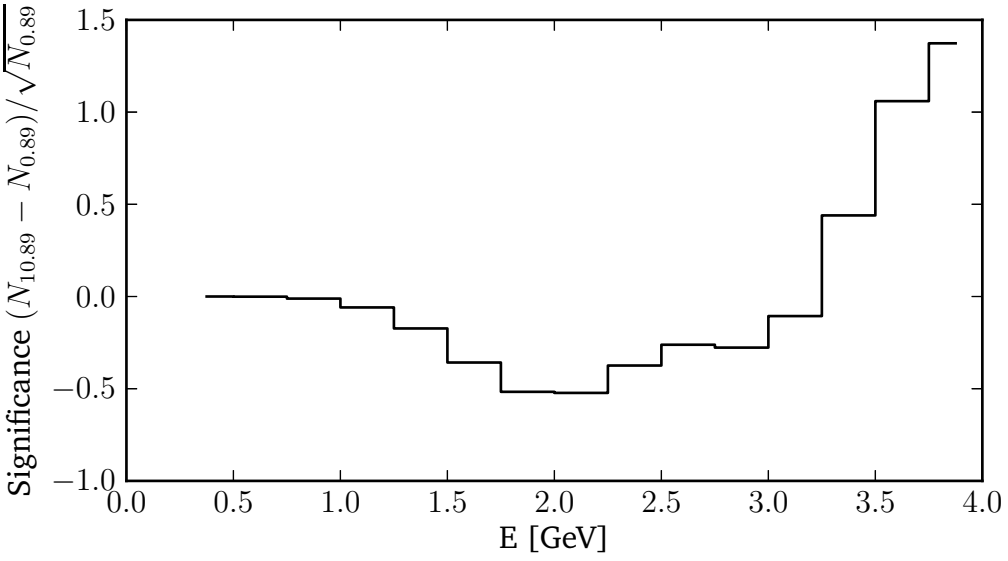

(b)

Figure 8.4: Comparison of measured neutrino-energy distributions for different values of Δm^2_{41}. The distributions for mass splittings of 0.89 eV^2 and 10.89 eV^2 are shown in (a). The 10.89-eV^2 distribution is normalized to the same area as the other distribution (i.e., 61 events). The significance of the shape differences between the distributions is shown in (b).

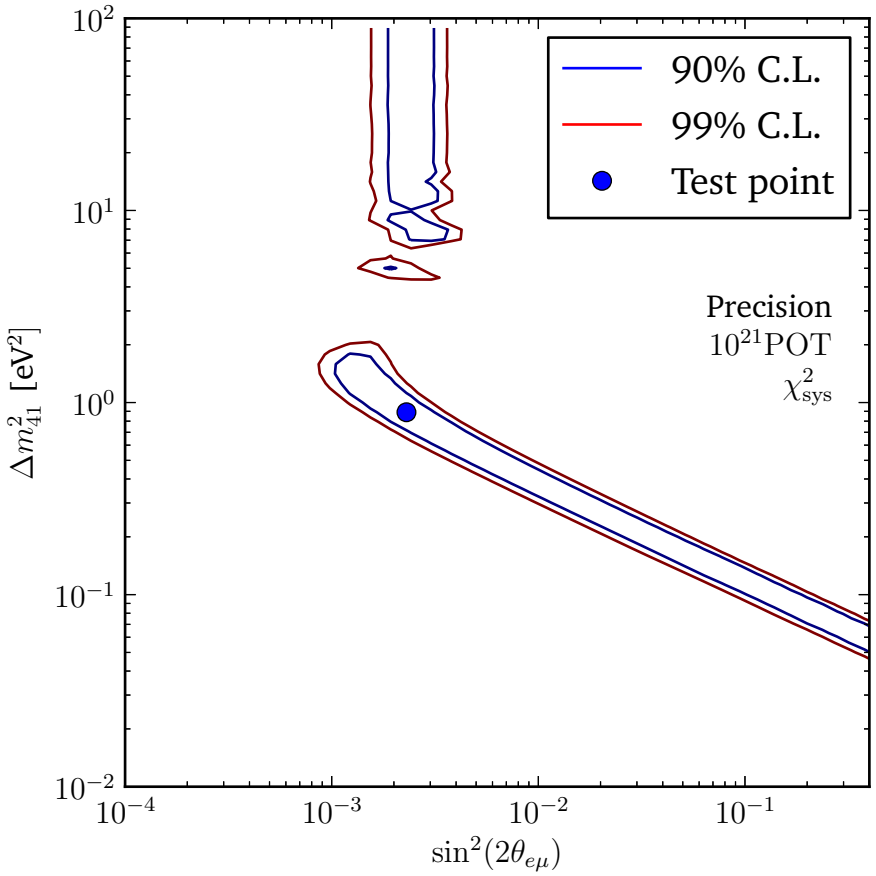

Figure 8.5: Precision study. The (3+1) best-fit point $\Delta m^2_{41} = 0.89$ eV^2, $|U_{e4}|^2 = 0.025$, and $|U_{\mu 4}|^2 = 0.023$ is used as the true test point. The 90% and 99% confidence levels are shown.

One important difference between the previous sensitivity study and this precision study is that the hypothesis Θ_0 no longer corresponds to 3×3 mixing, but now to the LSND best-fit point of $\Delta m^2_{41} = 0.89$ eV^2, $|U_{e4}|^2 = 0.025$, and $|U_{\mu 4}|^2 = 0.023$. Only the measured signal E_ν spectrum will be used for the precision study presented here since small background levels (e.g., Fig. 6.7) do not affect it.

A backgroundless parameter-estimation analysis is shown in Fig. 8.5. As was claimed above, νSTORM is unable to resolve the parameter degeneracy (i.e., no closed contour). The 4% systematic uncertainty (from Table 7.5) was used, but using a 1% uncertainty does not resolve the degeneracy. Similarly, the νSTORM target hall is designed for 400 kW; therefore, another study was performed using flux with a factor of 5 increase ($8.5 \times$

10^{18} muons; 4 years of 100 kW and 4 years of 400 kW) and the degeneracy remained. Lastly, changing the stored muon energy anywhere between 2 GeV and 10 GeV also is not beneficial. With a fixed baseline and this detector technology, the parameter-estimation abilities of νSTORM are limited.

The MIND detector is unable to resolve the parameter degeneracy in the conventional νSTORM design using exclusively appearance information. This is not an issue for refuting the LSND signal; however, if there are sterile neutrinos, then a magnetized LAr TPC should be commissioned to understand these neutrinos since the lower energy threshold would be closer to the oscillation maximum. The energy threshold for LAr TPCs is also better. Discovering sterile neutrinos is possible at νSTORM; however, precision measurements are more difficult.

There are, nevertheless, ways of designing a second phase to νSTORM if sterile neutrinos are discovered. To probe Δm^2_{41}, one needs $L = 4E\pi/(2\Delta m^2_{41}) = 5$ km if $E = 3$ GeV. (As an aside, the Fermilab site boundary is at 4 km in the current design; therefore, the decay ring may have to be repositioned.) Previously, the baseline was set to be shorter because longer baselines result in less flux (i.e., L^{-2}). However, by assuming a new 400 kW target-horn module for 8.5×10^{18} stored muons, the loses from the detector having a smaller solid-angle area is counteracted by the increased beam power. In Fig. 8.6, it can be seen that the parameter-estimation precision is increased. If a 3-kt fiducial volume is assumed, rather than the 1.3 kt, then Fig. 8.7 shows that closed contours can be made.

The parameter-estimation abilities of νSTORM have now been presented. With the experimental setup described to discovery sterile neutrinos (i.e., 2 km), the degeneracy between Δm^2_{41} and $\theta_{e\mu}$ cannot be resolved. However, a second phase of νSTORM at 5 km with a new 400 kW target-module would allow for precision measurements of the sterile neutrino parameter space.

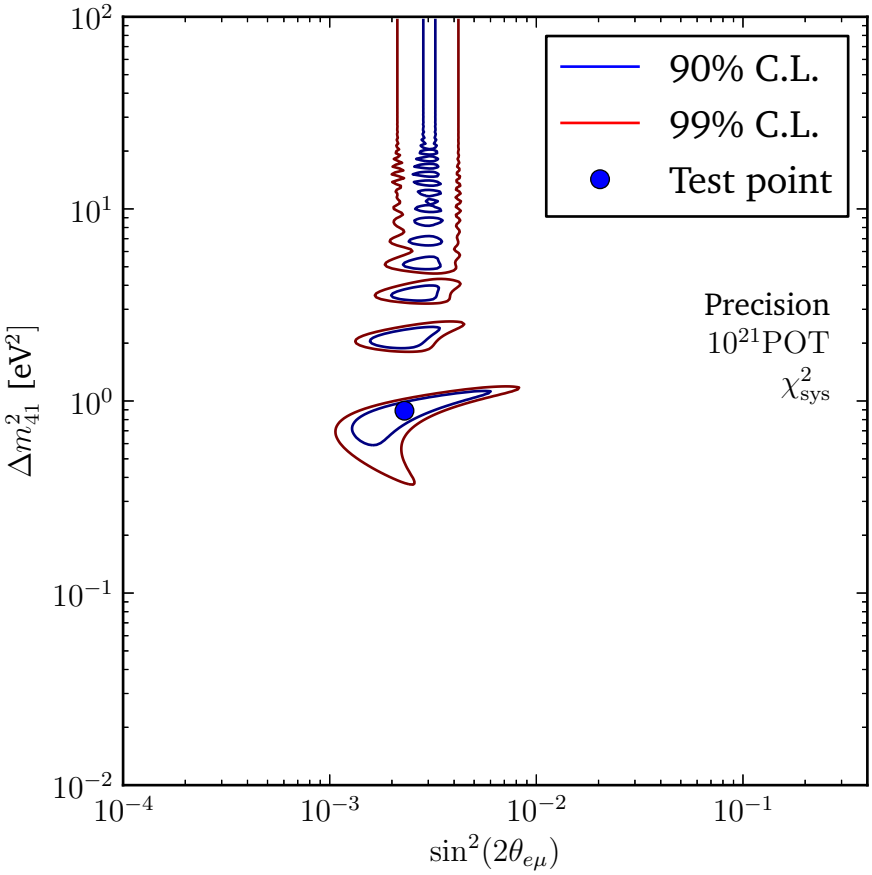

Figure 8.6: Precision study at 5 km with 8.5×10^{18} muons. A new 400 kW target-horn module was assumed in a second νSTORM phase. The (3+1) best-fit point $\Delta m^2_{41} = 0.89$ eV^2, $|U_{e4}|^2 = 0.025$, and $|U_{\mu 4}|^2 = 0.023$ is used as the true test point. The 90% and 99% confidence levels are shown.

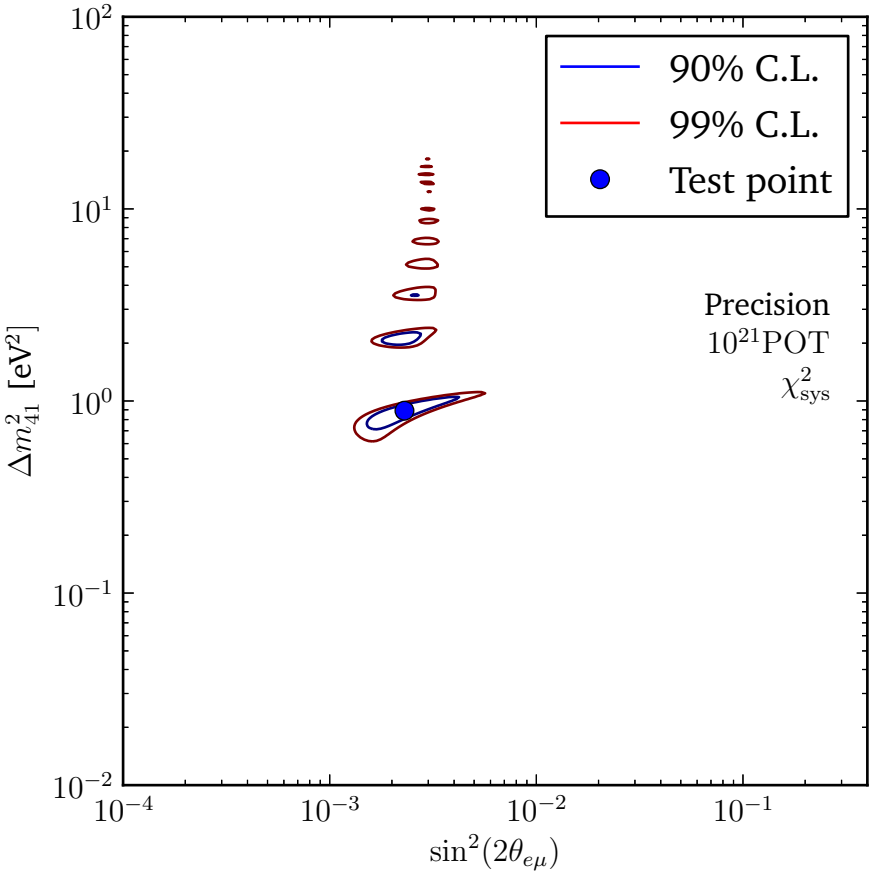

Figure 8.7: Precision study at 5 km with 8.5×10^{18} muons and a larger 3-kt detector. A new 400 kW target-horn module was assumed in a second νSTORM phase. The (3+1) best-fit point $\Delta m^2_{41} = 0.89$ eV^2, $|U_{e4}|^2 = 0.025$, and $|U_{\mu 4}|^2 = 0.023$ is used as the true test point. The 90% and 99% confidence levels are shown.

Chapter 9

Conclusion and future work

Throughout the course of this thesis, it has been demonstrated that the construction of a 3.8-GeV/c muon-decay ring could address many current physics questions, with a MIND at 2 km allowing the study of sterile neutrinos. The everpresent challenge in particle physics is determining how best to find new physics. The Standard Model is mankind's most precise theory, but it neither includes the unknown 95% of the Universe's energy density nor has been reconciled with gravity. The theory is incomplete. More generally than just νSTORM, utilizing the precise nature of muon-decay beams will allow explorations of the neutrino sector for new physics. Hopefully, studying the neutrino sector will reveal even more surprises than were found in the quark sector. If we start building new beams now, we can use them in the future to help determine if there are more surprises in the neutrino sector.

The accelerator and experiment described here provide a unique opportunity for exploring a wide range of phenomena. For example, in the long term, demonstrating the muon-decay beam concept allows for measuring the PMNS matrix to the precision of the CKM matrix, and hopefully reveal hints of new physics, in addition to being a new technology that may have both foreseen (i.e., μ collider) and unforeseen uses. In the medium term, cross-section measurements will increase the sensitivity of experiments

such as LBNO and LBNE to leptonic CP violation—primarily by constraining the $\sigma_{\nu_e}/\sigma_{\nu_\mu}$ ratio—thereby possibly helping explain why this Universe is compromised of matter rather than antimatter. It also would provide the only ν_e test beam.

The primary near-term physics case is the sensitivity to light sterile neutrinos, which come naturally from the discovery of neutrino mass. No proposed facility has comparable sterile neutrino sensitivity to νSTORM. The facility does not require any new technology. Detectors are being considered that are similar to previously built detectors. The sterile sensitivity analysis is simple by requiring only a range and curvature cut. By using the CPT of the LSND anomalous channel, the muon final state allows for negligible systematic uncertainties. Therefore, νSTORM is able to exclude sterile neutrinos at 10σ in a frequentist study.

Even though the physics case has been demonstrated, future work must still be performed in order to fully optimize the facility for physics other than the appearance channel. For disappearance analyses, the level to which NC and ν_e CC can be separated must be determined, and whether a dedicated LAr TPC is needed with MIND as the muon ranger. The cross-section measurements should be possible by introducing special purpose near detectors, but the detector-technology decision has not been made yet. However, this thesis describes a well-understood neutrino source that contains both $\overset{(-)}{\nu}_e$ and $\overset{(-)}{\nu}_\mu$; therefore it is expected that there will be numerous users of the beam, much like commercial light and neutron sources (e.g., Diamond, ISIS, ESS).

The physics case is diverse. More importantly, if this experiment is built, physicists may finally be able to explain the LSND results.

Appendix A

Optimization of baseline and straight length for near detector

The number of events at a near detector can be optimized against the decay straight length and near detector baseline. The arc length is fixed to $l_a = 60$ m. Therefore there are two other free parameters: the baseline l_b and the straight length l_s. The ring has a circumference $C = 2 \times (l_a + l_s)$.

The number of muons that decay toward a detector is proportional to l_b/C. The flux at the detector is—due to solid angle effects—proportional to L^{-2}, where L is the averaged effective baseline since muons can decay anywhere in the straight.

Define $L = \sqrt{l_s(l_s + l_b)}$. The two quantities l_b and l_s can be maximized. The maximum of $C \times L$ is shown in Table A.1. However, this is not applicable to the stochastic

Table A.1: Optimization with $l_a = 60$ m, assuming $\pi \to \mu$ decay channel.

Baseline l_b [m]	Optimal straight length l_s [m]
20	$20\sqrt{3} \simeq 24$
50	$10\sqrt{30} \simeq 54$
100	$20\sqrt{15} \simeq 77$
1500	300
2000	$200\sqrt{3} \simeq 346$

Table A.2: Optimize with $l_a = 60$ m, assuming $l_\pi = 250$ m and stochastic injection.

Baseline l_b [m]	Optimal straight length l_s [m]
20	191
50	228
100	272
1500	641
2000	701

injection scheme where pions decay in the straight; therefore another constraint is introduced. The charged pion mass is 139.57018 MeV, and accordingly, a 5-GeV/c pion has a boost of about $\gamma = 35$. The $c\tau$ of the pion is 7 m. The decay length of a pion in the laboratory frame is $l_\pi = 250$ m. Another term must be introduced: $1 - \exp(-l_s/l_\pi)$.

After including this new pion-decay term, a close near detector at 20 m favors a 200-m decay straight, whereas longer baselines are optimal for longer decay straight lengths (Table A.2). The reason longer straight lengths are optimal for longer baselines is that more pions are able to decay in the straight.

For a near detector at 20 m, there is a wide range of acceptable values. For example, using a 272-m straight length—the optimal value for a 100-m baseline—results in less than a 5% shift for the 20-m baseline. However, the shift is greater for longer baselines: there is a 30% shift by choosing a small straight length for baselines of 1.5 km. This is acceptable because these longer baselines desire straight lengths of over 600 m, which is not economically optimal; the money would be better spent building a bigger detector. Therefore the optimal near detector baseline is between 20 and 100 m, and the optimal straight length is between 250 and 300 m.

Bibliography

[1] ALEPH Collaboration, DELPHI Collaboration, L3 Collaboration, OPAL Collaboration, and SLD Collaboration. Precision electroweak measurements on the Z resonance. *Phys. Rept.*, 427:257–454, 2006.

[2] J. Beringer et al. Review of particle physics. *Phys. Rev. D*, page 10001, 2012.

[3] C. Giunti and C. Kim. *Fundamentals of neutrino physics and astrophysics*. Oxford University Press, USA, 2007.

[4] C. Giunti, C.W. Kim, and U.W. Lee. When do neutrinos really oscillate? Quantum mechanics of neutrino oscillations. *Phys. Rev. D*, 44:3635–3640, 1991.

[5] C. Giunti, C.W. Kim, and U.W. Lee. Coherence of neutrino oscillations in vacuum and matter in the wave packet treatment. *Phys. Lett. B*, 274(1):87–94, 1992.

[6] K. Kiers, S. Nussinov, and N. Weiss. Coherence effects in neutrino oscillations. *Phys. Rev. D*, 53:537–547, 1996.

[7] C. Giunti and C.W. Kim. Coherence of neutrino oscillations in the wave packet approach. *Phys. Rev. D*, 58:17301, 1998.

[8] C. Giunti. Coherence and wave packets in neutrino oscillations. *Found. Phys. Lett.*, 17:103–124, 2004.

[9] E.Kh. Akhmedov and J. Kopp. Neutrino oscillations: Quantum mechanics vs. quantum field theory. *JHEP*, 1004:8, 2010. arXiv:hep-ph/1001.4815.

[10] R. Wendell et al. Atmospheric neutrino oscillation analysis with sub-leading effects in Super-Kamiokande I, II, and III. *Phys. Rev. D*, 81:92004, 2010. arXiv:hep-ex/1002.3471.

[11] Y. Ashie et al. A measurement of atmospheric neutrino oscillation parameters by Super-Kamiokande I. *Phys. Rev. D*, 71:112005, 2005. arXiv:hep-ex/0501064.

[12] M. Messier. *Evidence for neutrino mass from observations of atmospheric neutrinos with Super-Kamiokande*. PhD thesis, Boston University, 1999. UMI-99-23965.

[13] L. Landau. On the conservation laws for weak interactions. *Nucl. Phys.*, 3(1):127–131, 1957.

[14] T.D. Lee and C.N. Yang. Parity nonconservation and a two-component theory of the neutrino. *Phys. Rev.*, 105:1671–1675, 1957.

[15] A. Salam. On parity conversation and neutrino mass. *Nuovo Cimento*, 5:299–301, 1957.

[16] Ch. Weinheimer. KATRIN: A next generation tritium β-decay experiment in search for the absolute neutrino mass scale. *Prog. Part. Nucl. Phys.*, 48(1):141–150, 2002.

[17] L. Hall, H. Murayama, and N. Weiner. Neutrino mass anarchy. *Phys. Rev. Lett.*, 84:2572–2575, 2000.

[18] F. Bonnet, D. Hernandez, T. Ota, and W. Winter. Neutrino masses from higher than d = 5 effective operators. *J. High Energy Phys.*, 2009(10):076, 2009.

[19] D.V. Forero, M. Tortola, and J.W.F. Valle. Global status of neutrino oscillation parameters after Neutrino 2012. arXiv:hep-ph/1205.4018, 2012.

[20] K.N. Abazajian et al. Light sterile neutrinos: A white paper. arXiv:hep-ph/1204.5379, 2012.

[21] A. de Gouvea. Seesaw energy scale and the LSND anomaly. *Phys. Rev. D*, 72:33005, 2005.

[22] C. Athanassopoulos et al. The liquid scintillator neutrino detector and LAMPF neutrino source. *Nucl. Instrum. Meth.*, A388:149–172, 1997.

[23] R. Davis. Attempt to detect the antineutrinos from a nuclear reactor by the $Cl^{37}(\overline{\nu}, e^{-})A^{37}$ reaction. *Phys. Rev.*, 97:766–769, Feb 1955.

[24] E.D. Church, K. Eitel, G.B. Mills, and M. Steidl. Statistical analysis of different muon-anti-neutrino to electron-anti-neutrino searches. *Phys. Rev. D*, 66:13001, 2002.

[25] A.A. Aguilar-Arevalo et al. A search for muon neutrino and antineutrino disappearance in MiniBooNE. *Phys. Rev. Lett.*, 103:61802, 2009.

[26] A.A. Aguilar-Arevalo et al. Unexplained excess of electron-like events from a 1-GeV neutrino beam. *Phys. Rev. Lett.*, 102:101802, 2009.

[27] A.A. Aguilar-Arevalo et al. A combined $\nu_\mu \to \nu_e$ and $\bar{\nu}_\mu \to \bar{\nu}_e$ oscillation analysis of the MiniBooNE Excesses. arXiv:hep-ex/1207.4809, 2012.

[28] C. Giunti and M. Laveder. 3+1 and 3+2 sterile neutrino fits. *Phys. Rev. D*, 84:73008, 2011.

[29] Th. Lasserre et al. Reactor antineutrino anomaly. *Phys. Rev. D*, 83:73006, 2011.

[30] H. Kwon et al. Search for neutrino oscillations at a fission reactor. *Phys. Rev. D*, 24:1097–1111, Sep 1981.

[31] P. Huber. Determination of antineutrino spectra from nuclear reactors. *Phys. Rev. C*, 84:24617, 2011.

[32] C. Giunti and M. Laveder. Statistical significance of the gallium anomaly. *Phys. Rev. C*, 83:65504, 2011.

[33] CDHSW Collaboration. A search for ν_μ oscillations in the Δm^2 range 0.3–90 eV^2. *Phys. Lett. B*, 134(3–4):281–286, 1984.

[34] I.E. Stockdale et al. Limits on muon-neutrino oscillations in the mass range 30 $< \Delta m^2 <$ 1000 eV^2. *Phys. Rev. Lett.*, 52:1384–1388, 1984.

[35] P. Adamson et al. Active to sterile neutrino mixing limits from neutral-current interactions in MINOS. *Phys. Rev. Lett.*, 107:11802, 2011.

[36] J. Kopp, M. Maltoni, and T. Schwetz. Are there sterile neutrinos at the eV scale? *Phys. Rev. Lett.*, 107:91801, 2011. arXiv:hep-ph/1103.4570.

[37] J.M. Conrad, C.M. Ignarra, G. Karagiorgi, M.H. Shaevitz, and J. Spitz. Sterile neutrino fits to short baseline neutrino oscillation measurements. arXiv:hep-ex/1207.4765, 2012.

[38] J.L. Hewett et al. Fundamental physics at the intensity frontier. arXiv:hep-ex/1205.2671, 2012.

[39] P.A.R. Ade et al. Planck early results. I. The Planck mission. *Astron. Astrophys.*, 536:16464, 2011. arXiv:astro-ph.IM/1101.2022.

[40] P.A.R. Ade et al. Planck 2013 results. XVI. Cosmological parameters. 2013.

[41] Alessandro Mirizzi, Gianpiero Mangano, Ninetta Saviano, Enrico Borriello, Carlo Giunti, et al. The strongest bounds on active-sterile neutrino mixing after Planck data. 2013.

[42] J. R. Kristiansen, Ø. Elgarøy, C. Giunti, and M. Laveder. Cosmology with sterile neutrino masses from oscillation experiments. 2013. arXiv/astro-ph:1303.4654.

[43] M.A. Green and S.J. St. Lorant. Estimating the cost of large superconducting thin solenoid magnets. *Adv. Cryog. Eng.*, 39A:271–276, 1994.

[44] D.G. Michael et al. The magnetized steel and scintillator calorimeters of the MINOS experiment. *Nucl. Instrum. Meth. A*, 596:190–228, 2008. arXiv:0805.3170.

[45] J.K. Nelson. The MINOS magnets. *Int. J. Mod. Phys. A*, 16:1181–1184, 2001.

[46] J. Jackson. *Classical electrodynamics*. Wiley, 3rd edition, 1998. Equation 12.42.

[47] C. L. Wang. Pion, kaon, and antiproton production between 10 and 70 bev. *Phys. Rev. Lett.*, 25:1068–1072, Oct 1970.

[48] P. Renton. *Electroweak interactions: An introduction to the physics of quarks and leptons*. Cambridge University Press, Cambridge, 1990.

[49] D. Neuffer. Design considerations for a muon storage ring. Telemark Workshop on Neutrino Mass, 1980.

[50] νSTORM Collaboration. nuSTORM: Neutrinos from stored muons, 2012. arXiv:hep-ex/1206.0294.

[51] G. Danby, J-M. Gaillard, K. Goulianos, L.M. Lederman, N. Mistry, M. Schwartz, and J. Steinberger. Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos. *Phys. Rev. Lett.*, 9:36–44, 1962.

[52] D.G. Koshkarev. Proposal for a decay ring to produce intense secondary particle beams at the SPS. Technical Report ISR-DI/74-62, CERN, 1974.

[53] G. Charpak, F. J. M. Farley, R. L. Garwin, T. Muller, J. C. Sens, V. L. Telegdi, and A. Zichichi. Measurement of the anomalous magnetic moment of the muon. *Phys. Rev. Lett.*, 6:128–132, Feb 1961.

[54] S. Geer. Neutrino beams from muon storage rings: Characteristics and physics potential. *Phys. Rev. D*, 57:6989–6997, Jun 1998.

[55] E. Wilder et al. EUROnu costing report. 2012.

[56] M. Campanelli, S. Navas, and A. Rubbia. Neutrino cross-section measurement with neutrinos from muon decay. *Nucl. Instrum. Meth.*, A503:151–153, 2001.

[57] S. Choubey et al. International design study for the neutrino factory, interim design report. arXiv:hep-ex/1112.2853, 2011.

[58] C.D. Tunnell. Possible muon-decay beams for sterile physics in the 2010s: Pion decay beams are *so* 1990s, 2012. CERN Neutrino town meeting.

[59] C.L. Wang. Pion production in high-energy collisions. *Phys. Rev. D*, 10:3876–3878, 1974.

[60] M. Popovic. Personal communication, 2012.

[61] S. Striganov. Personal communication, 2012.

[62] N.I. Bozhko et al. Hadron production in inclusive processes in proton-nuclear collisions at 67 GeV/c. *Yad. Fiz.*, 31:1494–1500, 1980.

[63] S. van der Meer et al. Physics and technique of stochastic cooling. *Phys. Rep.*, 58(2):73–102, 1980.

[64] A. Liu. Personal communication, 2012.

[65] B. Pasternak. Personal communication, 2011.

[66] A. Bogacz. Personal communication, 2012.

[67] M.A. Green and B.P. Strauss. The cost of superconducting magnets as a function of stored energy and design magnetic induction times the field volume. *IEEE Trans. Supercond.*, 18(2):248–251, 2008.

[68] S. Geer, O. Mena, and S. Pascoli. A low-energy neutrino factory for large θ_{13}. *Phys. Rev. D*, 75:93001, 2007. arXiv:hep-ph/0701258.

[69] M.G. Catanesi et al. The HARP detector at the CERN PS. *Nucl. Instrum. Meth.*, A571:527–561, 2007.

[70] M. Sadler et al. Proposal to upgrade the MIPP experiment. Technical Report Proposal P-960, Fermilab, 2010.

[71] M. Posiadala. Status of the NA61 (SHINE) experiment at CERN. arXiv:hep-ex/0901.3332, 2009.

[72] J. Crisp. Personal communication, 2012.

[73] D. Renker. New developments on photosensors for particle physics. *Nucl. Instrum. Meth. A*, 598(1):202–212, 2009. doi:10.1016/j.nima.2008.08.023.

[74] A. Ichikawa and K. Atsuko. Status of the T2K long baseline neutrino oscillation experiment. *J. Phys. Conf. Ser.*, 203:12104, 2010.

[75] G.W. Foster, V.S. Kashikhin, E. Malamud, P. Mazur, A. Oleck, H. Piekarz, J. Fuerst, R. Rabehl, P. Schlabach, and J. Volk. The 100 kA VLHC transmission line magnet superconducting cable test facility. *IEEE Trans. Supercond.*, 10(1):318–321, 2000.

[76] P. Adamson et al. Measurement of neutrino oscillations with the MINOS detectors in the NuMI beam. *Phys. Rev. Lett.*, 101:131802, 2008. hep-ex/0806.2237.

[77] M. Holder et al. Is there a high-y anomaly in antineutrino interactions? *Phys. Rev. Lett.*, 39:433–436, 1977.

[78] P. Adamson et al. Improved search for muon-neutrino to electron-neutrino oscillations in MINOS. *Phys. Rev. Lett.*, 107:181802, 2011.

[79] S. Agostinelli et al. Geant4: A simulation toolkit. *Nucl. Instrum. Meth. A*, 506(3):250–303, 2003.

[80] J. Allison et al. Geant4 developments and applications. *IEEE Trans. Nucl. Sci*, 53(1):270–278, 2006.

[81] `http://github.com/nuSTORM/gnomon`.

[82] C. Andreopoulos, A. Bell, D. Bhattacharya, F. Cavanna, J. Dobson, et al. The GENIE neutrino Monte Carlo generator. *Nucl. Instrum. Meth.*, A614:87–104, 2010.

[83] R. Mankel. Pattern recognition and event reconstruction in particle physics experiments. *Rep. Prog. Phys.*, 67(4):553, 2004.

[84] C.T. Zahn. Graph-theoretical methods for detecting and describing gestalt clusters. *IEEE Trans. Comput.*, C-20(1):68–86, 1971.

[85] J. Bang-Jensen and G. Gutin. *Digraphs: theory, algorithms and applications*. Springer, 2008.

[86] R. Bellman. On a routing problem. *Quart. Appl. Math.*, 16:87–90, 1958.

[87] J. Ford, D.R. Fulkerson, and R.G. Bland. *Flows in networks*. Princeton Landmarks in Mathematics. Princeton University Press, 2010.

[88] Wikipedia. Lorentz force — Wikipedia, the free encyclopedia, 2013. [Online; accessed 14-Feb-2013].

[89] B. Wands. Personal communication, 2012.

[90] C.D. Tunnell, J.H. Cobb, and A.D. Bross. Sensitivity to eV-scale neutrinos of experiments at a very-low-energy neutrino factory. arXiv:hep-ph/1111.6550, 2011.

[91] C.D. Tunnell. Sterile neutrino sensitivity with wrong-sign muon appearance at nuSTORM. arXiv:hep-ph/1205.6338, 2012.

[92] W. Winter. Optimization of a very-low-energy neutrino factory for the disappearance into sterile neutrinos. arXiv:hep-ph/1204.2671, 2012.

[93] A. Blondel. Muon polarisation in the neutrino factory. *Nucl. Instrum. Meth. A*, 451(1):131–137, 2000.

[94] Y. Fukui, R. Fernow, and J. Gallardo. Muon polarization in a front-end channel of a neutrino factory. *Proc. Part. Accel. Conf., 19th*, 3:2126–2128, 2001.

[95] G. Penn. Beam envelope equations in a solenoidal field. MUCOOL note 71.

[96] D. Meloni, J. Tang, and W. Winter. Sterile neutrinos beyond LSND at the neutrino factory. *Phys. Rev. D*, 82:93008, 2010. arXiv:hep-ph/1007.2419.

[97] A. Cervera et al. Golden measurements at a neutrino factory. *Nucl. Phys.*, B579:17–55, 2000. hep-ph/0002108.

[98] P. Huber, M. Lindner, and W. Winter. Simulation of long-baseline neutrino oscillation experiments with GLoBES. *Comput. Phys. Commun.*, 167:195, 2005. hep-ph/0407333.

[99] P. Huber, J. Kopp, M. Lindner, M. Rolinec, and W. Winter. New features in the simulation of neutrino oscillation experiments with GLoBES 3.0: General long baseline experiment simulator. *Comput. Phys. Commun.*, 177:432–438, 2007. arXiv:hep-ph/0701187.

[100] C.D. Tunnell. https://code.launchpad.net/∼c-tunnell1/+junk/vlenf_scripts, 2012.

[101] GNU. GPL 3.0.

[102] J. Kopp. Efficient numerical diagonalization of hermitian 3×3 matrices. *Int. J. Mod. Phys.*, C19:523–548, 2008. arXiv:physics/0610206.

[103] J. Kopp, M. Lindner, T. Ota, and J. Sato. Non-standard neutrino interactions in reactor and superbeam experiments. *Phys. Rev. D*, 77:13007, 2008. arXiv:hep-ph/0708.0152.

[104] C. Giunti and M. Laveder. Towards 3+1 neutrino mixing, 2011. arXiv:hep-ph/1109.4033.

[105] C. Athanassopoulos et al. Evidence for $\nu_\mu \to \nu_e$ neutrino oscillations from LSND. *Phys. Rev. Lett.*, 81:1774–1777, 1998. arXiv:nucl-ex/9709006.

[106] M.C. Gonzalez-Garcia, M. Maltoni, and J. Salvado. Updated global fit to three neutrino mixing: Status of the hints of $\theta_{13} > 0$. *JHEP*, 1004:56, 2010. arXiv:hep-ph/1001.4524.

[107] E.A. Paschos and J.Y. Yu. Neutrino interactions in oscillation experiments. *Phys. Rev. D*, 65:33002, 2002. arXiv:hep-ph/0107261.

[108] M. Day and K. McFarland. Differences in quasi-elastic cross sections of muon and electron neutrinos. *Phys. Rev. D*, 86:53003, 2012.

[109] S. Baker and R.D. Cousins. Clarification of the use of χ^2 and likelihood functions in fits to histograms. *Nucl. Instrum. Meth.*, 221(2):437–442, 1984.

[110] N. Ryder. Personal communication, 2013.

[111] P. Adamson et al. Measurements of atmospheric neutrinos and antineutrinos in the MINOS Far Detector. *Phys. Rev. D*, 86:52007, 2012.

[112] MINOS collaboration. MINOS neutrino and antineutrino charged-current disappearance analysis. `http://www-numi.fnal.gov/pr_plots/CC_Brag_Slide_Jun2012.pdf`.

[113] ATLAS Collaboration. Readiness of the ATLAS tile calorimeter for LHC collisions. *Eur. Phys. J. C*, 70:1193–1236, 2010.

[114] P.A. Poole, R.C. Ruchti, A.D. Bross, and M.R. Wayne, editors. *Test beam performance of the CDF plug upgrade hadron calorimeter*, volume 450. AIP, 1998.

[115] J. Damgov. CMS HCAL testbeam results and comparison with GEANT 4 simulation. In S.R. Magill and R. Yoshida, editors, *Calorimetry in High Energy Physics: XII*, volume 867 of *AIP Conf. Proc.*, pages 471–478, 2006.

[116] E. Hughes. Study of hadronic and electromagnetic shower development between 10 and 140 GeV by an iron-scintillator calorimeter. Technical report, SLAC, 1990. SLAC-PUB-5404.

[117] R. Ospanov. *A measurement of muon neutrino disappearance with the MINOS detectors and NuMI beam*. PhD thesis, University of Texas at Austin, 2008.

[118] M. Kordosky. *Hadronic interactions in the MINOS detector*. PhD thesis, University of Texas at Austin, 2004.

[119] A. Cabrera et al. Comparisons of the MINOS near and far detector readout systems at a test beam. *Nucl. Instrum. Meth. A*, 609(2):106–113, 2009.

[120] J. Kilmer. Specification for carbon steel plate for MINOS detectors. Technical Report NuMI-L-422, Fermilab, 1998.

[121] L. Soby. Personal communication, 2013.

[122] P. Coloma, P. Huber, J. Kopp, and W. Winter. Systematic uncertainties in long-baseline neutrino oscillations for large θ_{13}. arXiv:hep-ph/1209.5973, 2012.

[123] N. Mayer. Results for CCQE scattering with the MINOS near detector. *J. Phys. Conf. Ser.*, 312(7):72010, 2011.

[124] P. Adamson et al. Neutrino and antineutrino inclusive charged-current cross section measurements with the MINOS near detector. *Phys. Rev. D*, 81:72002, 2010.

[125] P. Adamson et al. A study of muon neutrino disappearance using the Fermilab Main Injector neutrino beam. *Phys. Rev. D*, 77:72002, 2008.

[126] E. Rutherford. The scattering of α and β particles by matter and the structure of the atom. *Philo. Mag.*, 6(21):669–688, 1911.

[127] Leon N. Cooper and James Rainwater. Theory of multiple coulomb scattering from extended nuclei. *Phys. Rev.*, 97:492–504, Jan 1955.

[128] A.W. Wolfendale. *Cosmic Rays*. Technical survey series. Philosophical Library, 1963.

[129] G. E. Masek, L. D. Heggie, Y. B. Kim, and R. W. Williams. Scattering of 2-bev/c muons in carbon and lead. *Phys. Rev.*, 122:937–948, May 1961.

[130] Gerald Feinberg and Leon M. Lederman. The physics of muons and muon neutrinos. *Annual Review of Nuclear Science*, 13(1):431–504, 1963.

[131] S. Dixon. Personal communication, 2012.

[132] K. Terzaghi, R.B. Peck, and G. Mesri. *Soil mechanics in engineering practice*. A Wiley Interscience Publication. John Wiley & Sons, 1996.

[133] G. Barr, T.K. Gaisser, and T. Stanev. Flux of atmospheric neutrinos. *Phys. Rev. D*, 39:3532–3534, 1989.

[134] J. Cobb. Cosmic ray muon rates in the MINOS near detector. NuMI-L-232, 1996.

[135] G. Mengyun. Muon simulation at the Daya Bay site. Technical report, LBNL, 2011. 4262E.

[136] T.K. Gaisser. *Cosmic rays and particle physics*. Cambridge University Press, 1991.

[137] R. Bayes. Personal communication, 2012.

[138] G.L. Fogli et al. Getting the most from the statistical analysis of solar neutrino oscillations. *Phys. Rev. D*, 66:53010, 2002.

[139] B. Aharmim et al. Low energy threshold analysis of the phase I and phase II data sets of the Sudbury neutrino observatory. *Phys. Rev. C*, 2009. nucl-ex/0910.2984.

[140] M.A. Thomson. Status of the MINOS experiment. *Nucl.Phys.Proc.Suppl.*, 143:249–256, 2005.

Wovon man nicht sprechen kann, darüber muss man schweigen.