An LED calibration system for the MINOS long baseline neutrino oscillation experiment

Philip Adamson

Submitted for the degree of Doctor of Philosophy at the

University of Sussex

March 2001

Abstract

The MINOS experiment aims to test the hypothesis that the observed deficit in the ratio of ν_μ to ν_e resulting from the impact of cosmic rays on the upper atmosphere is caused by ν_μ oscillating into another flavour in transit to the earth's surface.

MINOS uses a beam of ν_μ from the Fermilab Main Injector facility, and steel/plastic scintillator calorimiters, read out by wavelength shifting fibres and multi-anode photomultipliers, at Fermilab and 730 km away in the Soudan mine, Minnesota. A comparison of the neutrino flavour content in the two detectors provides a measurement of neutrino oscillations that is relatively independent of any flux uncertainties. However, to make an accurate spectrum comparison, accurate calibration is needed. MINOS has specified accuracies of 2% for relative calibration between the two detectors, and 5% for absolute energy calibration.

Calibration in MINOS has three major steps. A light injection system is used to map out the linearity of the phototubes and account for short-term variations in gain. Selected muons from cosmic rays provide a "standard candle" energy deposit that is the same in all locations, and monitor long-term deterioration of the scintillator. A small (1 m $\times$ 1 m $\times$ 3 m) calibration detector will be exposed to known hadron and electron beams to relate the response to different particle species to that of the "standard" muon.

This thesis describes an LED-based light injection calibration system that has been developed for MINOS. LED "pulser boxes" produce pulses of light with widths $10-40$ns; pulse widths and heights are controlled electronically, avoiding the moving parts required in filter-wheel based systems. The LEDs have no pulse-to-pulse jitter; the pulse heights are monitored by PIN diodes to account for the (small) variation in LED light output over time. Each LED illuminates up to 640 wavelength-shifting fibres; the production system produces a maximum of at least 200 photoelectrons in each PMT channel. The system has been tested in the lab and in a CERN testbeam, and is shown to be capable of the 2% accuracy required.

Preface

If no man is an island, then nor is a D. Phil. thesis. This work would not have been possible without the support and assistance of many people. I would like to thank my supervisor, Philip Harris, for his advice and support throughout my time as a graduate student, for prodding me when I needed prodding, and for not prodding when I didn't. I would also like to thank Jenny Thomas, for advice and assistance particularly during the final stages of this work.

I would like to thank my MINOS collaborators for being a thoroughly pleasant group of people to work with, and for many of the plots in chapter 3. Much of the work on the LED pulser described in chapter 6 was done in conjunction with Robert Morse, who I would like to thank particularly for his assistance with some of the plots in that chapter. I am grateful to Ryan Nichol for assistance with the PMT spectrum fits in chapter 7.

During my research and the preparation of this thesis I have had the pleasure of sharing an office with Rob Morse and Jim Richardson, and with Paul Crosby and Ryan Nichol. I am grateful to them for making it such a pleasant experience.

Finally, I would like to thank my parents, who have been a continual source of support throughout my sometimes seemingly endless journey as a student, and my wife Rachel, whose continuing patience with both me and this thesis in the face of repeated provocation is surely deserving of a medal.

Declaration

I hereby declare that this thesis has not been submitted, either in the same or different form, to this or any other University for a degree.

Philip Adamson

Contents

1. Introduction

Everything starts somewhere, although many physicists disagree.

— Terry Pratchett, "Hogfather"

The standard model of particle physics is the single most successful physical theory ever. It describes the electromagnetic, weak and strong forces, and their interactions with the fundamental particles. No definite measurement has been made which conflicts with the standard model. However, it cannot be the final theory of particle physics—it does not describe gravity, and contains a large number of arbitrarily-set parameters, which is deeply unsatisfying from an æsthetic point of view.

The standard model contains the observed fundamental fermions: six quarks, six leptons (three charged, three neutrinos) together with their respective antiparticles. It also contains the gauge bosons for the weak, strong and electromagnetic forces ($\mathrm{W}^{\pm}$, Z^0, eight gluons, and the photon). The scalar Higgs ϕ^0, whilst not a part of the original standard model, is generally held today to be "standard"—the standard model clearly requires a mechanism to provide particles with mass, and the Higgs mechanism is the generally accepted way of doing so. In the standard model, the neutrino is massless. The only other massless particle in the standard model, the photon, is required to be massless to produce a long-range electromagnetic force. There is no similar compelling reason for the neutrino's masslessness—the neutrino is massless *de facto* as no-one has measured it to have a mass. The fact that the neutrino is massless does explain the maximal parity violation in the weak interaction; however neutrino masses at the eV scale are also consistent with the observed parity violation.

Ever since the neutrino was first postulated by Wolfgang Pauli [1] in 1930, it has proved to be the most elusive of particles. It was not until 1956 that Reines and Cowan [2] finally detected the neutrino, with their pioneering experiment at the Savannah River nuclear reactor in South Carolina.

Since then, comparatively little has actually been discovered about the neutrino. The famous experiment of Wu et al. [3] demonstrated that the weak interaction maximally violated the parity symmetry, and so the massless neutrino was either left or right-handed. In 1958, Goldhaber et al. [4] demonstrated that the neutrino was left-handed, and, following the discovery of three generations of charged leptons and the muon neutrino, measurements of the width of the Z^0 resonance at CERN [5, 6, 7, 8, 9] and at SLC [10] demonstrated that there were precisely three neutrinos that couple to the Z^0 with mass less than $\frac{1}{2}\mathrm{m_Z}$.

However, a full understanding of Pauli's 'desperate remedy' was still proving elusive. Realising that the nuclear processes in the Sun must be a potent source of neutrinos, Ray Davis set out to measure them. His pioneering Homestake experiment [11] succeeded, but half the neutrinos seemed to be missing. Other experiments began to look at neutrinos formed in the particle showers caused by cosmic rays hitting the atmosphere, and this time saw too few ν_μ. Where could the neutrinos have gone?

A possible solution had already been proposed by Pontecorvo [12, 13, 14]. We know that in the quark system, the eigenstates of the weak force are different from the mass eigenstates, and are related by the Cabbibo-Kobayashi-Maskawa matrix. We also know that the neutral K mesons are formed as K^0 and $\overline{\mathrm{K}^0}$, but decay as the near CP eigenstates $\mathrm{K_L}$ and $\mathrm{K_S}$. This leads to the phenomenon of strangeness oscillations. Pontecorvo realised that if neutrinos were a superposition of mass eigenstates, they would be formed and detected in weak eigenstates, but propagate as mass eigenstates. Just as in the K-system, this would lead to flavour-changing oscillations. One could then explain the apparent disappearance of neutrinos by assuming that they had merely oscillated into another flavour state.

This would require the neutrinos to have mass, whereas everyone had assumed that they were massless. Attempts to measure neutrino mass directly have to date yielded only upper limits[1]. Precision measurements of the beta decay spectrum of atomic tritium near the endpoint have produced an upper limit on the electron neutrino mass of 2.5 eV (Troitsk)[15] and 2.2 eV (Mainz)[16]. A similar experiment for the muon neutrino [17] measures the energy spectrum of muons from pion decay at rest, and has upper limit 170 keV. The tau neutrino mass is limited [18] by analysis of the decays $\tau \longrightarrow 4\pi + \nu_\tau$ and $\tau \longrightarrow 5\pi + \nu_\tau$ to less than 18.2 MeV.

A completely independent limit on the electron neutrino mass was obtained from analysis of the time structure of neutrinos from the supernova SN1987A. The Kamiokande and IMB detectors observed 19 neutrinos over a period of 10 s; Bachall and Glashow [19] obtain a conservative upper limit on the electron neutrino mass of 11 eV.

Although these mass limits are of interest to cosmologists, as they preclude neutrinos from being the sole constituent of the 'dark matter' that is thought to make up some 90% of the universe's mass, they present no difficulty for the neutrino oscillation hypothesis. We shall see that we can explain the observed neutrino deficiencies with masses several orders of magnitude below these limits.

[1] Apart, that is, from a couple of temporary glitches. The Troitsk beta decay spectrum experiment did for a while report a mass for ν_e in the region of a few eV, but this was caused by an anomalous bump towards the end of the spectrum, which has since disappeared. There was also a brief flurry of interest in the late '80s and early '90s in a neutrino state with mass 17 keV, provoked by distortions in the measured beta-decay spectrum of ^{63}Ni and ^{35}S. These distortions have since been shown to be systematic effects arising from energy losses in the apparatus which had not been accounted for.

2. Neutrino theory

Two men say they're Jesus—one of them must be wrong.

— Dire Straits, "Industrial Disease"

The standard model neutrino is massless. In this chapter we see how mass enters in to the standard model, and consider extensions to the standard model that could allow neutrino mass. We then go on to consider the consequences of a massive neutrino.

2.1. Mass in the standard model

The standard model of particle physics is an $\mathrm{SU}(3)\otimes\mathrm{SU}(2)\otimes\mathrm{U}(1)$ gauge theory. We are concerned here only with the leptonic sector; leptons do not carry SU(3) colour quantum numbers, so it will suffice for our purposes to consider only the $\mathrm{SU}(2)_L \otimes \mathrm{U}(1)_Y$ of the electroweak theory. For clarity, we shall look at a model with one leptonic generation, and point out the differences when one moves to a three generation model.

Chirality is preserved by gauge interactions, so it is appropriate to decouple the fermion fields into independent left and right chiral projections. In order to produce W bosons that only couple to left chiral fermion fields, the left handed fields must form SU(2) doublets:

$$l_L = \begin{pmatrix} \nu_{eL} \\ e_L \end{pmatrix}, \quad q_L = \begin{pmatrix} u_L \\ d_L \end{pmatrix} \tag{2.1}$$

and the right handed fields form SU(2) singlets e_R, u_R, d_R.

The standard model Higgs also forms an SU(2) doublet:

$$\Phi = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix} \tag{2.2}$$

which will contribute the three extra degrees of freedom required for the massive $\mathrm{W}^{\pm}$ and Z bosons, leaving the physical scalar Higgs. We can then write down the Lagrangian for the Higgs-fermion couplings:

$$\mathcal{L} = f^e \overline{l}_L \Phi e_R + f^u \overline{q}_L \tilde{\Phi} u_R + f^d \overline{q}_L \Phi d_R + h.c. \tag{2.3}$$

with $\tilde{\Phi} = \mathrm{i}\tau_2 \Phi^*$, and the f are the Yukawa coupling constants for the fermion-Higgs interaction.

After the symmetry is broken, we have

$$\langle \Phi \rangle = \begin{pmatrix} 0 \\ v/\sqrt{2} \end{pmatrix} \tag{2.4}$$

and so produce mass terms:

$$\mathcal{L} = f^e \frac{v}{\sqrt{2}} \overline{e}_L e_R + f^u \frac{v}{\sqrt{2}} \overline{u}_L u_R + f^d \frac{v}{\sqrt{2}} \overline{d}_L d_R + h.c. \tag{2.5}$$

When the three leptonic generations are included, equation 2.3 can contain cross-couplings between members of different generations, so in general, equation 2.5 will contain 3×3 mass matrices viz:

$$\mathcal{L} = \begin{pmatrix} \overline{e}_L & \overline{\mu}_L & \overline{\tau}_L \end{pmatrix} \begin{pmatrix} \odot & \odot & \odot \\ \odot & \odot & \odot \\ \odot & \odot & \odot \end{pmatrix} \begin{pmatrix} e_R \\ \mu_R \\ \tau_R \end{pmatrix} + h.c. \tag{2.6}$$

2.2. Dirac and Majorana mass

The mass terms produced as above couple left and right handed fields together; they are known as Dirac mass terms. They are the only sort of mass term that is allowed for charged particles. However, neutrinos, being neutral, can also have a Majorana mass term, which couples a particle to its antiparticle.

Define the charge conjugation operator C so that

$$\psi^C = C\gamma^0\psi^* = \mathrm{i}\gamma^2\psi^*. \tag{2.7}$$

$$\text{Define chiral fields } \psi_L = \frac{1}{2}(1-\gamma_5)\psi, \quad \psi_R = \frac{1}{2}(1+\gamma_5)\psi \tag{2.8}$$

$$\text{and write } \psi_L^C = (\psi_L)^C = \frac{1}{2}(1+\gamma_5)\psi^C = (\psi^C)_R \tag{2.9}$$

so that the charge conjugation of a left-handed field is a right-handed (anti-)field.

Then, writing down all possible couplings between pairs of our uncharged fields ψ, we can have both Dirac and Majorana mass terms:

$$\text{Dirac Mass: } \mathcal{L}_D \quad = m_D(\overline{\psi}_L\psi_R + \overline{\psi}_R\psi_L) \quad = m_D\overline{\psi}\psi \tag{2.10}$$

$$\text{Majorana Mass: } \mathcal{L}_{M_a} \quad = m_a(\overline{\psi}_L^C\psi_L + \overline{\psi}_L\psi_L^C) \quad = m_a\overline{\chi}\chi \tag{2.11}$$

$$\mathcal{L}_{M_b} \quad = m_b(\overline{\psi}_R^C\psi_R + \overline{\psi}_R\psi_R^C) \quad = m_b\overline{\omega}\omega \tag{2.12}$$

with the self-conjugate fields χ and ω defined

$$\chi = \psi_L + \psi_L^C = \chi^C, \quad \omega = \psi_R + \psi_R^C = \omega^C. \tag{2.13}$$

It is obvious that Majorana mass terms are only possible where the field is uncharged.

So the most general mass term will occur when all three of the terms above are present:

$$\mathcal{L}_{DM} = m_D\overline{\psi}_L\psi_R + m_a\overline{\psi}_L^C\psi_L + m_b\overline{\psi}_R^C\psi_R + h.c. \tag{2.14}$$

$$= \frac{1}{2}m_D(\overline{\chi}\omega + \overline{\omega}\chi) + m_a\overline{\chi}\chi + m_b\overline{\omega}\omega \tag{2.15}$$

$$= \begin{pmatrix} \overline{\chi} & \overline{\omega} \end{pmatrix} \begin{pmatrix} m_a & \frac{1}{2}m_D \\ \frac{1}{2}m_D & m_b \end{pmatrix} \begin{pmatrix} \chi \\ \omega \end{pmatrix}. \tag{2.16}$$

Diagonalising this produces two mass eigenstates

$$\eta_1 = \cos\theta\,\chi - \sin\theta\,\omega \tag{2.17}$$

$$\eta_2 = \sin\theta\,\chi + \cos\theta\,\omega \tag{2.18}$$

It can be seen that these are both Majorana states, with masses

$$m_{1,2} = \frac{1}{2}\left(m_a + m_b \pm \sqrt{(m_a - m_b)^2 + m_D^2}\right). \qquad (2.19)$$

So in the general case, the neutrino would be a Majorana particle. To produce a Dirac neutrino, one would like to impose a symmetry principle that forbids the Majorana coupling.

If the neutrino does have a Majorana mass, one should be able to observe neutrinoless double beta decay. A small number of isotopes can decay by emitting two betas simultaneously (i.e. two neutrons decay into two protons, as in figure 2.1).

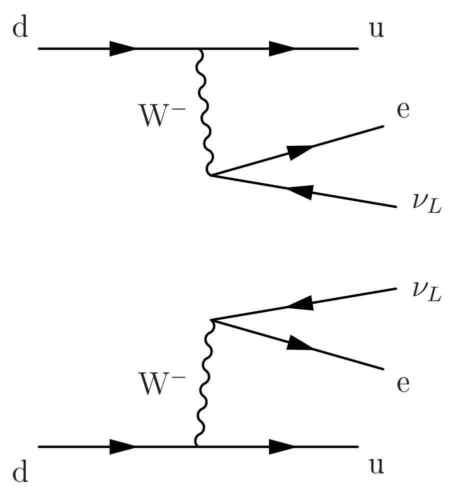

Figure 2.1.: Standard double beta decay, with two neutrinos produced.

If the neutrino has a Majorana mass, the right-handed anti-neutrino emitted by one decay can be transformed by the Majorana mass term to become a left-handed neutrino, which is then absorbed in the second decay, as in figure 2.2. The net result is a double beta decay that produces two electrons and no neutrinos, so all the energy lost must be carried by the electrons. Summing the electron energies would then give a delta function, rather than the continuous distribution obtained when 'invisible' neutrinos carry away some of the energy. This would be a clear signal that neutrinos had a Majorana mass. Experiments [20] have studied double beta decays to look for such a process, but so far haven't seen a signal.

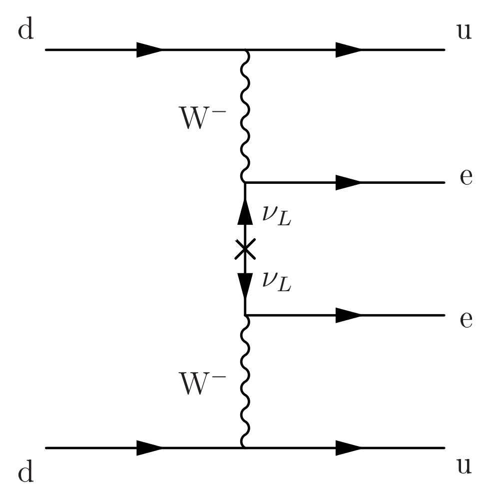

Figure 2.2.: Neutrinoless double beta decay. The cross represents the action of the Majorana mass term.

2.3. Adding a neutrino mass to the standard model

As we have seen, the standard model generates fermion masses from the coupling with the Higgs. Recall equation 2.3:

$$\mathcal{L} = f^e \bar{l}_L \Phi e_R + f^u \overline{q}_L \tilde{\Phi} u_R + f^d q_L \Phi d_R + h.c. \tag{2.20}$$

The absence of a right-handed neutrino field means that we cannot write down a term like those above involving the neutrino, so the standard model neutrino is massless. To generate neutrino mass, we must add extra leptons or extra Higgs bosons.

Extra leptons

The simplest and most obvious solution is to add a right-handed neutrino field. One may then add extra terms to the standard model Lagrangian:

$$\mathcal{L} = f^\nu \bar{l}_L \nu_R \tilde{\Phi} + B \overline{\nu}_R^C \nu_R + h.c. \tag{2.21}$$

$$\rightarrow m_D \overline{\nu}_L \nu_R + B \overline{\nu}_R^C \nu_R + h.c. \tag{2.22}$$

where the bare Majorana mass term is allowed, as ν_R does not carry any $\mathrm{SU}(2) \otimes \mathrm{U}(1)$ quantum numbers. Ultimately, in a unified theory, one would like a larger gauge group which also contains the ν_R and so B would arise from a Higgs-like symmetry breaking mechanism at the unification scale.

Using the notation of equation 2.16, we have

$$\mathcal{L} = \begin{pmatrix} \overline{\chi} & \overline{\omega} \end{pmatrix} \begin{pmatrix} 0 & \frac{1}{2}m_D \\ \frac{1}{2}m_D & B \end{pmatrix} \begin{pmatrix} \chi \\ \omega \end{pmatrix}. \tag{2.23}$$

There are three scenarios that can be considered:

$\boldsymbol{B = 0}$ Here we recover the Dirac neutrino. A unified theory must find an explanation for why the neutrino mass is so small compared to the other fermion masses.

$\boldsymbol{B \gg m_D}$ We produce a light neutrino with mass $\sim m_D^2/B$ which is predominantly ν_L, and a heavy neutrino with mass $\sim B$ which is predominantly ν_R. This is an example of the seesaw mechanism, which attempts to explain the smallness of the neutrino masses. We start with the coupling of the neutrino to the SM Higgs to be the same order as the Higgs-electron coupling (and hence m_D of order the electron mass). The bare coupling B is assumed to arise from a Higgs-like symmetry breaking at the unification scale, and so is large. This large mass then drives down the mass of the light neutrino to its physical value.

$\boldsymbol{B \sim m_D}$ Here the two mass eigenstates have similar masses (see equation 2.19). One then produces neutrino-antineutrino oscillations. These lie beyond the scope of this thesis—one could however note in passing that the atmospheric neutrino measurements ignore the sign of the lepton, so $\nu \longrightarrow \overline{\nu}$ oscillations do not help us here. As MINOS is magnetised, and so is able to measure the charge of a muon, it is able to measure $\nu_\mu \longrightarrow \overline{\nu_\mu}$ oscillations.

Extra Higgs

Consider the group structure of the possible lepton field bilinears. Writing e.g. $l_L \sim (\mathbf{2}, -1)$ for an SU(2) doublet with hypercharge $Y = -1$, we have:

$$\bar{l}_L e_R \sim (\mathbf{2}, 1) \otimes (\mathbf{1}, -2) \qquad = (\mathbf{2}, -1) \tag{2.24}$$
$$\bar{l}^C_L l_L \sim (\mathbf{2}, -1) \otimes (\mathbf{2}, -1) \qquad = (\mathbf{1}, -2) \oplus (\mathbf{3}, -2) \tag{2.25}$$
$$\bar{e}^C_R e_R \sim (\mathbf{1}, -2) \otimes (\mathbf{1}, -2) \qquad = (\mathbf{1}, -4) \tag{2.26}$$

The standard model Higgs has structure $(\mathbf{2}, 1)$ with the above notation, and couples to the first kind of bilinear (2.24). The third bilinear does not help us with neutrino mass, but if one is to avoid adding extra leptons, the second bilinear (2.25) only contains left-handed fields. One could then introduce a charged singlet Higgs $\sim (\mathbf{1}, 2)$ and a triplet $\sim (\mathbf{3}, 2)$. This would produce a Majorana neutrino.

Additions to the Higgs sector would also contribute to the ratio of charged to neutral current couplings

$$\rho = \frac{M_W}{M_Z \cos\theta_W} = \frac{v_\phi^2}{v_\phi^2} = 1 \quad \text{in the standard model.} \tag{2.27}$$

A limit on the vacuum expectation of additional Higgs fields is placed by the measurement of $\rho = 0.9998^{+0.0011}_{-0.0006}$ from a global fit in [21].

2.4. Neutrino oscillations

We have seen that there are several ways in which neutrinos may acquire a mass term; we now consider the consequences. Recall from section 2.1 that in general we will have a 3×3 mass matrix linking the three neutrino species. Just as the weak quark eigenstates (d', s', b') are related to the mass eigenstates by the CKM matrix, so the weak neutrino eigenstates $(\nu_\mathrm{e}, \nu_\mu, \nu_\tau)$ are related to the neutrino mass eigenstates (ν_1, ν_2, ν_3) by a unitary matrix U. An N-dimensional unitary matrix has N^2 free parameters—$N(N-1)/2$ Euler angles and $N(N+1)/2$ phases. For Dirac neutrinos, we can without loss of generality absorb $(2N-1)$ phases into the definitions of the fields, leaving $N(N-1)/2$ angles and

$(N-1)(N-2)/2$ physical phases. We consider here only neutrinos with Dirac mass. A full treatment of Majorana neutrinos requires the use of a 6×6 matrix coupling neutrino and anti-neutrino states.

2.4.1. Vacuum oscillations

Two neutrino model

Although in general all three neutrino species will mix together, many theoretical models favour a scenario where one mixing is dominant; it is therefore of interest to consider the limited scheme where only two neutrinos are involved. For two neutrino flavours ν_x and ν_y, related to mass states ν_1, ν_2 with masses m_1 and m_2:

$$\begin{pmatrix} |\nu_x\rangle \\ |\nu_y\rangle \end{pmatrix} = \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix} \begin{pmatrix} |\nu_1\rangle \\ |\nu_2\rangle \end{pmatrix}.$$

A neutrino oscillation experiment has a source of neutrinos of known flavour some distance L away from a detector, which has some flavour measurement capability. So we generate a flavour x neutrino of energy E at our source:

$$|\nu_0\rangle = |\nu_x\rangle = \cos\theta|\nu_1\rangle + \sin\theta|\nu_2\rangle. \tag{2.28}$$

The neutrino then travels a distance L to our detector, so the wavefunction evolves:

$$|\nu_L\rangle = \mathrm{e}^{ip_1L}\cos\theta|\nu_1\rangle + \mathrm{e}^{ip_2L}\sin\theta|\nu_2\rangle. \tag{2.29}$$

So the oscillation probability, the probability of finding our neutrino in state $|\nu_y\rangle$ is

$$|\langle\nu_y|\nu_L\rangle|^2 = |-\mathrm{e}^{ip_1L}\sin\theta\cos\theta|\nu_1\rangle + \mathrm{e}^{ip_2L}\sin\theta\cos\theta|\nu_2\rangle|^2 \tag{2.30}$$

$$= \sin^2\theta\cos^2\theta(2-2\cos((p_1-p_2)L)) \tag{2.31}$$

$$= \sin^2 2\theta \sin^2\frac{(p_1-p_2)L}{2} \tag{2.32}$$

Now, we have ${p_1}^2 = E^2 - {m_1}^2$ and similarly for p_2. A massive neutrino will still have $m \ll E$, so we can expand $p = E - \frac{m^2}{2E}$ and so, writing $\Delta m^2_{21} = {m_2}^2 - {m_1}^2$,

$$|\langle\nu_y|\nu_L\rangle|^2 = \sin^2 2\theta\sin^2\frac{\Delta {m^2}_{21}L}{4E} \tag{2.33}$$

where natural units are used. Writing energy in GeV, length in km, and mass in eV, we obtain the oft-quoted formula

$$P(\nu_x \longrightarrow \nu_y) = \sin^2 2\theta \sin^2 \frac{1.27\Delta m^2{}_{21} L}{E} \tag{2.34}$$

Figure 2.3 shows the probability for a ν_μ to change flavour in a two-generation model over the MINOS baseline.

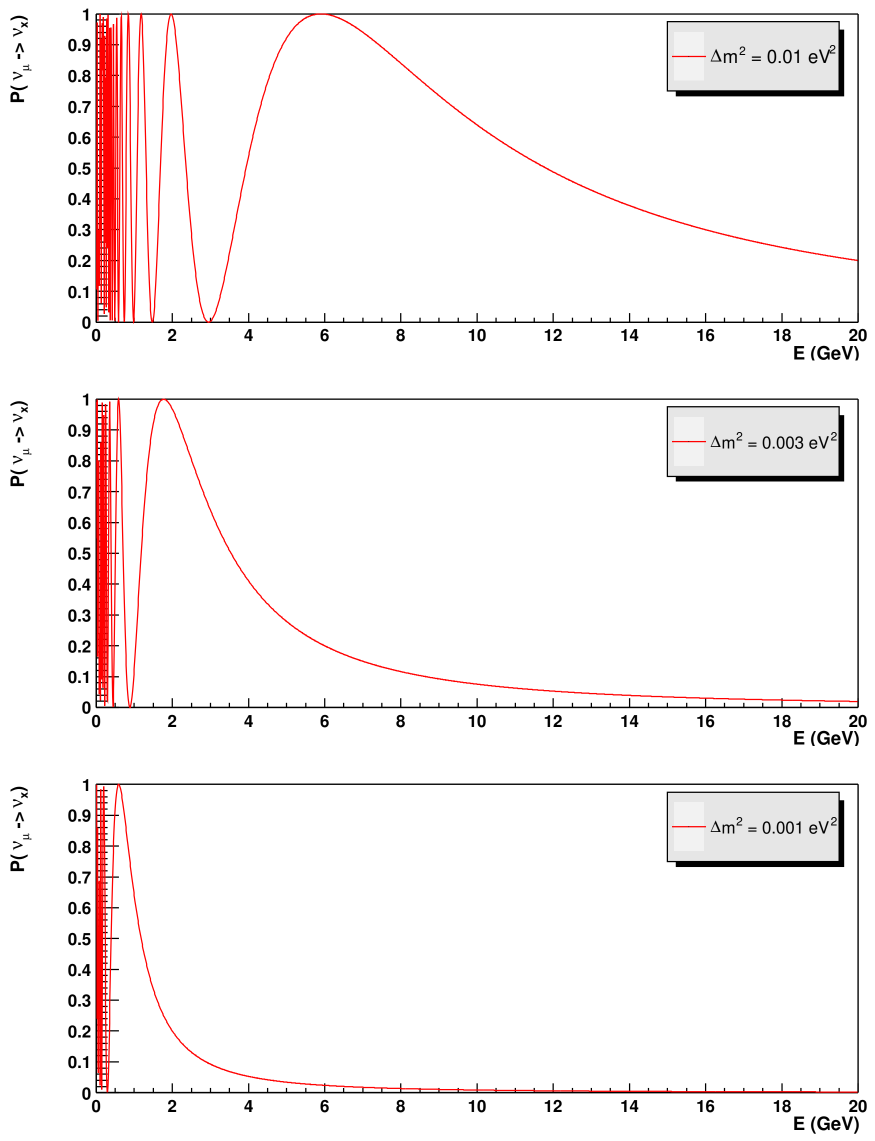

Figure 2.3.: Oscillation probability as a function of energy for the two generation model over the MINOS baseline and $\sin^2 2\theta = 1$.

Three generation oscillations

Now consider three neutrino flavour states $|\nu_a\rangle, a \in \{x, y, z\}$ related to the mass states $|\nu_i\rangle, i \in \{1, 2, 3\}$ by

$$|\nu_a\rangle = \sum_i U_{ai}|\nu_i\rangle.$$

Our initial state is then

$$|\nu_0\rangle = |\nu_x\rangle = U_{xi}|\nu_i\rangle$$

where the Einstein summation convention is implied. At the detector, we have

$$|\nu_L\rangle = \mathrm{e}^{ip_iL}U_{xi}|\nu_i\rangle$$

and so

$$\begin{aligned}
|\langle\nu_a|\nu_x\rangle|^2 &= |\langle\nu_j|U^\dagger_{aj}\mathrm{e}^{ip_iL}U_{xi}|\nu_i\rangle|^2 && (2.35)\\
&= \left|\sum_{ij} U^*_{aj}\mathrm{e}^{ip_iL}U_{xi}\delta_{ij}\right|^2 && (2.36)\\
&= |U^*_{a1}U_{x1} + \mathrm{e}^{i(p_2-p_1)L}U^*_{a2}U_{x2} + \mathrm{e}^{i(p_3-p_1)L}U^*_{a3}U_{x3}|^2 && (2.37)\\
&\approx |U^*_{a1}U_{x1} + \mathrm{e}^{i\Delta m^2{}_{21}L/2E}U^*_{a2}U_{x2} + \mathrm{e}^{i\Delta m^2{}_{13}L/2E}U^*_{a3}U_{x3}|^2 && (2.38)
\end{aligned}$$

For Dirac neutrinos, we may parameterise U:

$$U = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}\mathrm{e}^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}\mathrm{e}^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}\mathrm{e}^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}\mathrm{e}^{i\delta} & c_{12}s_{23} - s_{12}c_{23}s_{13}\mathrm{e}^{i\delta} & c_{23}c_{13} \end{pmatrix}$$

with δ the CP-violating phase, $c_{ij} = \cos\theta_{ij}$ and $s_{ij} = \sin\theta_{ij}$. This parameterisation has the advantage that we may easily recover the two generation case by setting two angles and the phase to zero.

2.4.2. Matter effects—the MSW effect

The effect of matter on neutrino oscillations was first studied by Wolfenstein [22] and Mikheyev and Smirnov [23], and so is known as the MSW effect. When neutrinos

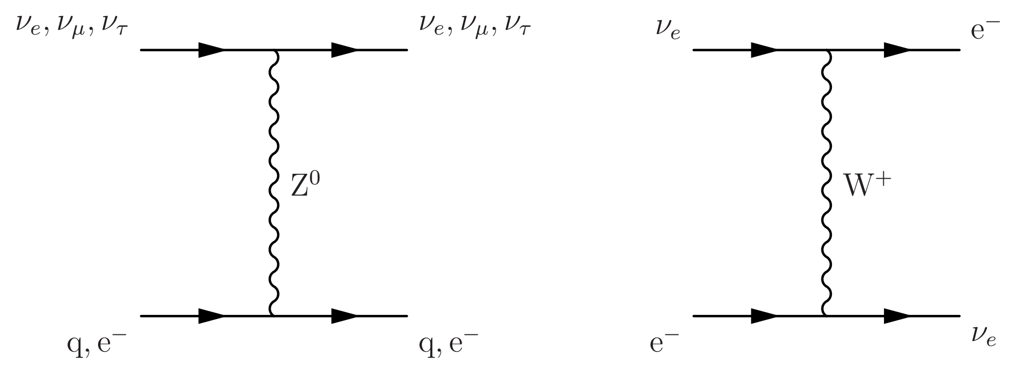

Figure 2.4.: Forward scattering of neutrinos in matter. All three flavours of neutrino experience the neutral current interactions to the left; only ν_e have the additional channel to the right available

propagate in matter, they interact via the interactions in figure 2.4. For electron anti-neutrinos, the additional diagram appears in figure 2.5. The presence of these matter interactions produces a refractive index

$$n = 1 + \frac{2\pi N_{\mathrm{e}}}{p^2} f_{\nu \mathrm{e}} + \frac{2\pi N_{\mathrm{p}}}{p^2} f_{\nu \mathrm{p}} + \frac{2\pi N_{\mathrm{n}}}{p^2} f_{\nu \mathrm{n}} \tag{2.39}$$

where $f_{\nu \mathrm{x}}$ is the forward-scattering amplitude for the neutrino from x, given by the diagrams above. The contributions from the neutral current interactions are the same for all neutrinos, this will merely produce an overall phase, which can be ignored.

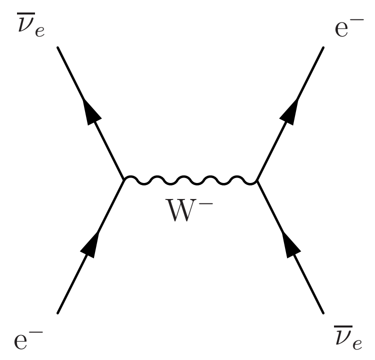

Figure 2.5.: Additional diagram for $\overline{\nu}_e$ in matter.

The charged current interactions are only present for electron neutrinos, and contribute

$$f_{\nu_{\mathrm{e}}\mathrm{e}} = \pm \frac{G_F p}{\sqrt{2}\pi} \tag{2.40}$$

where the plus sign is taken for neutrinos and the minus sign for antineutrinos. Recalling that the refractive index introduces an extra phase factor $\mathrm{e}^{\mathrm{i}(n-1)p.x}$, we see that matter interactions introduce an additional phase $\pm\sqrt{2}G_F N_{\mathrm{e}} x$.

Now we consider the effect of this additional phase on neutrino oscillations. Again we consider here the two-generation model for simplicity and clarity. Recall the evolution of the mass eigenstates

$$-\mathrm{i}\frac{\mathrm{d}}{\mathrm{d}x}\begin{pmatrix}|\nu_1\rangle \\ |\nu_2\rangle\end{pmatrix} = \begin{pmatrix}p_1 & 0 \\ 0 & p_2\end{pmatrix}\begin{pmatrix}|\nu_1\rangle \\ |\nu_2\rangle\end{pmatrix}. \tag{2.41}$$

Writing this in the flavour basis, we have

$$-\mathrm{i}\frac{\mathrm{d}}{\mathrm{d}x}\begin{pmatrix}|\nu_{\mathrm{e}}\rangle \\ |\nu_x\rangle\end{pmatrix} = \begin{pmatrix}p_1\cos^2\theta + p_2\sin^2\theta & (p_2-p_1)\sin\theta\cos\theta \\ (p_2-p_1)\sin\theta\cos\theta & p_1\sin^2\theta + p_2\cos^2\theta\end{pmatrix}\begin{pmatrix}|\nu_{\mathrm{e}}\rangle \\ |\nu_x\rangle\end{pmatrix} \tag{2.42}$$

Now, we add in the matter effects, and note that subtracting p_1 times the unit matrix is the same as removing an overall phase, and obtain the equation for the evolution of the flavour states in matter:

$$-\mathrm{i}\frac{\mathrm{d}}{\mathrm{d}x}\begin{pmatrix}|\nu_{\mathrm{e}}\rangle \\ |\nu_x\rangle\end{pmatrix} = \begin{pmatrix}(p_2-p_1)\sin^2\theta \pm \sqrt{2}G_F N_{\mathrm{e}} & (p_2-p_1)\sin\theta\cos\theta \\ (p_2-p_1)\sin\theta\cos\theta & (p_2-p_1)\cos^2\theta\end{pmatrix}\begin{pmatrix}|\nu_{\mathrm{e}}\rangle \\ |\nu_x\rangle\end{pmatrix} \tag{2.43}$$

This equation has eigenstates

$$|\nu_1'\rangle = \cos\theta_m|\nu_{\mathrm{e}}\rangle + \sin\theta_m|\nu_x\rangle \tag{2.44}$$

$$|\nu_2'\rangle = -\sin\theta_m|\nu_{\mathrm{e}}\rangle + \cos\theta_m|\nu_x\rangle \tag{2.45}$$

with

$$\tan 2\theta_m = \frac{(p_2-p_1)\sin 2\theta}{(p_2-p_1)\cos 2\theta \mp \sqrt{2}G_F N_{\mathrm{e}}}. \tag{2.46}$$

We can now follow the argument that we gave for vacuum oscillations, with our new primed eigenstates as a basis, and immediately write down from equation 2.32 the probability of an electron neutrino changing flavour

$$|\langle\nu_x|\nu_L\rangle|^2 = \sin^2 2\theta_m \sin^2\frac{(\lambda_1-\lambda_2)L}{2} \tag{2.47}$$

where the λ_i are the eigenvalues of equation 2.43. One can then show that

$$\lambda_1 - \lambda_2 = \frac{\Delta m^2_{21}}{2E}\sqrt{\sin^2 2\theta + \left(\frac{\pm\sqrt{2}G_F N_e}{\Delta m^2_{21}/2E} - \cos 2\theta\right)^2}, \tag{2.48}$$

and so, writing

$$\Delta M^2 = \Delta m^2_{21}\sqrt{\sin^2 2\theta + \left(\frac{\pm\sqrt{2}G_F N_e}{\Delta m^2_{21}/2E} - \cos 2\theta\right)^2}, \tag{2.49}$$

and recalling

$$\sin^2 2\theta_m = \frac{\sin^2 2\theta}{\sin^2 2\theta + \left(\cos 2\theta \mp \frac{\sqrt{2}G_F N_e}{\Delta m^2_{21}/2E}\right)^2} \tag{2.50}$$

we have

$$P(\nu_e \longrightarrow \nu_x) = \sin^2 2\theta_m \sin^2 \frac{1.27\Delta M^2 L}{E}. \tag{2.51}$$

At low matter densities ($N_e \to 0$), we recover the vacuum solution. At very high densities, the oscillation amplitude is suppressed as $1/N_e^2$. The key feature of MSW oscillations is that one can achieve resonant enhancement of the oscillations. The enhancement of the mixing angle for different matter densities is shown in figure 2.6, and the effect on the effective ΔM^2 is shown in figure 2.7. Even if the mixing between two states is small (and so $\sin^2 2\theta$ is small), we have maximal oscillations with $\sin^2 2\theta_m = 1$ for

$$\frac{\sqrt{2}G_F N_e}{\Delta m^2_{21}/2E} = \pm\cos 2\theta \tag{2.52}$$

where again the plus sign is taken for neutrinos and the minus for antineutrinos. This allows resonant enhancement for neutrinos if $m_2 > m_1$ and antineutrinos if $m_2 < m_1$. The confirmation of MSW enhancement in solar neutrinos would then demonstrate that the neutrinos follow a conventional mass hierarchy (i.e. $m_{\nu_e} < m_{\nu_\mu} < m_{\nu_\tau}$.)

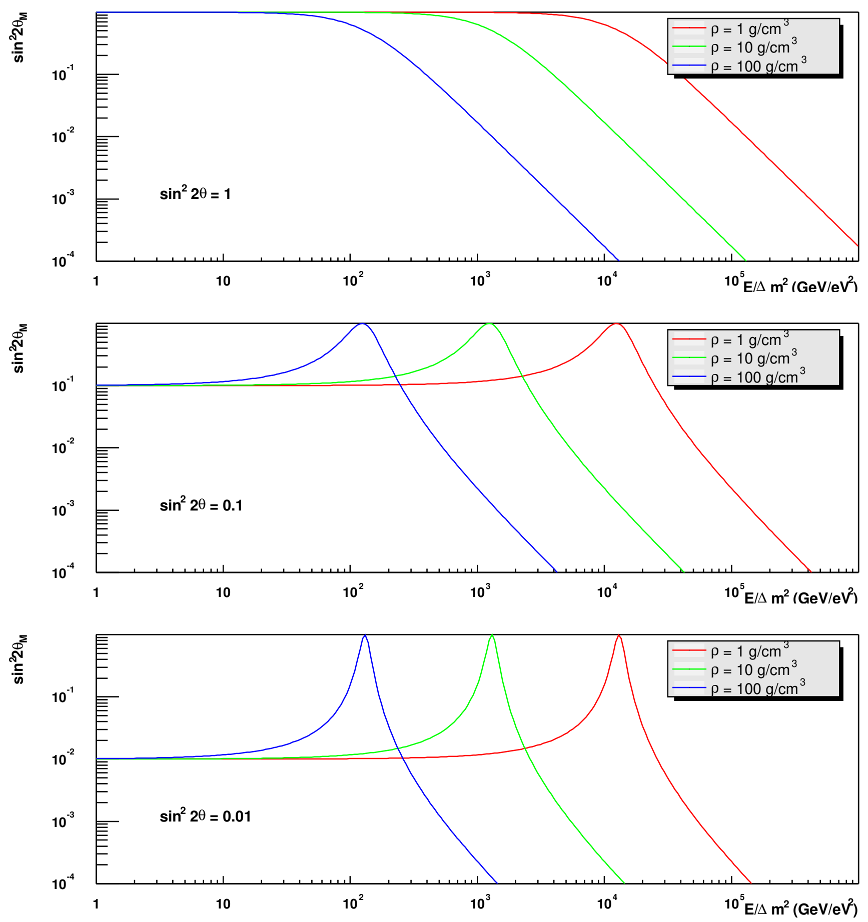

Figure 2.6.: MSW resonant enhancement of neutrino oscillations. For MINOS neutrinos, $E/\Delta m^2 \approx 10^3$ is the region suggested by Super-Kamiokande. We see that resonant enhancement of oscillations is relevant if $\nu_\mu \longrightarrow \nu_{\mathrm{e}}$ oscillations are the cause of the atmospheric neutrino anomaly—the density of the Earth's crust is a few $\mathrm{g/cm}^3$

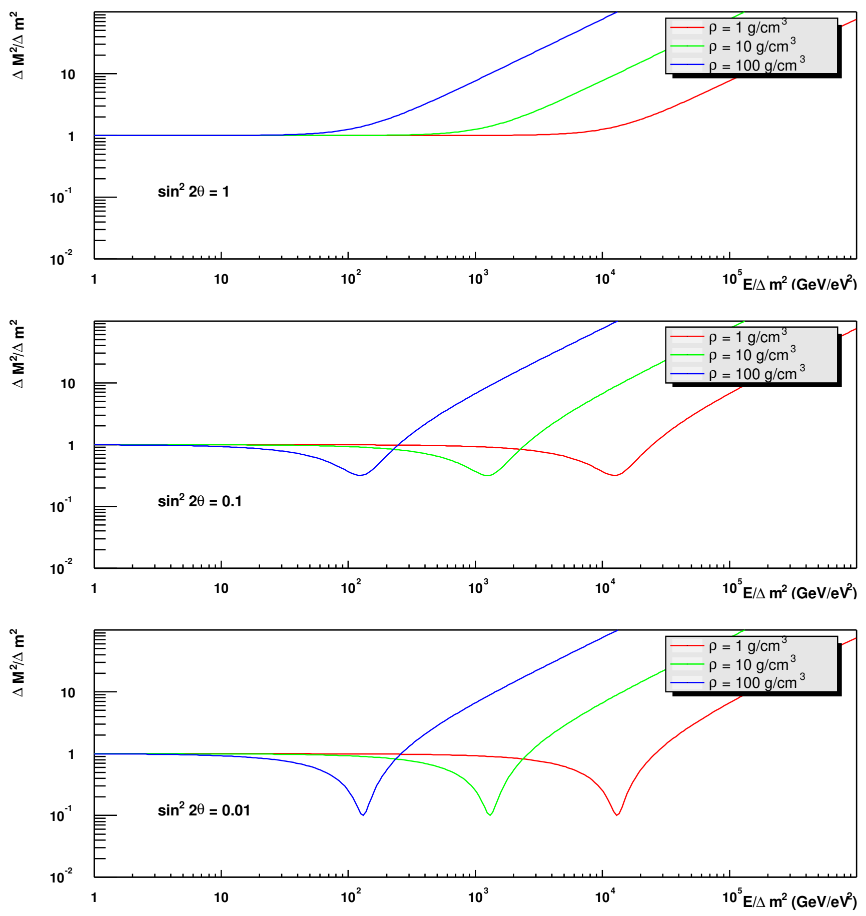

Figure 2.7.: Effective Δm^2 for the oscillations of figure 2.6.

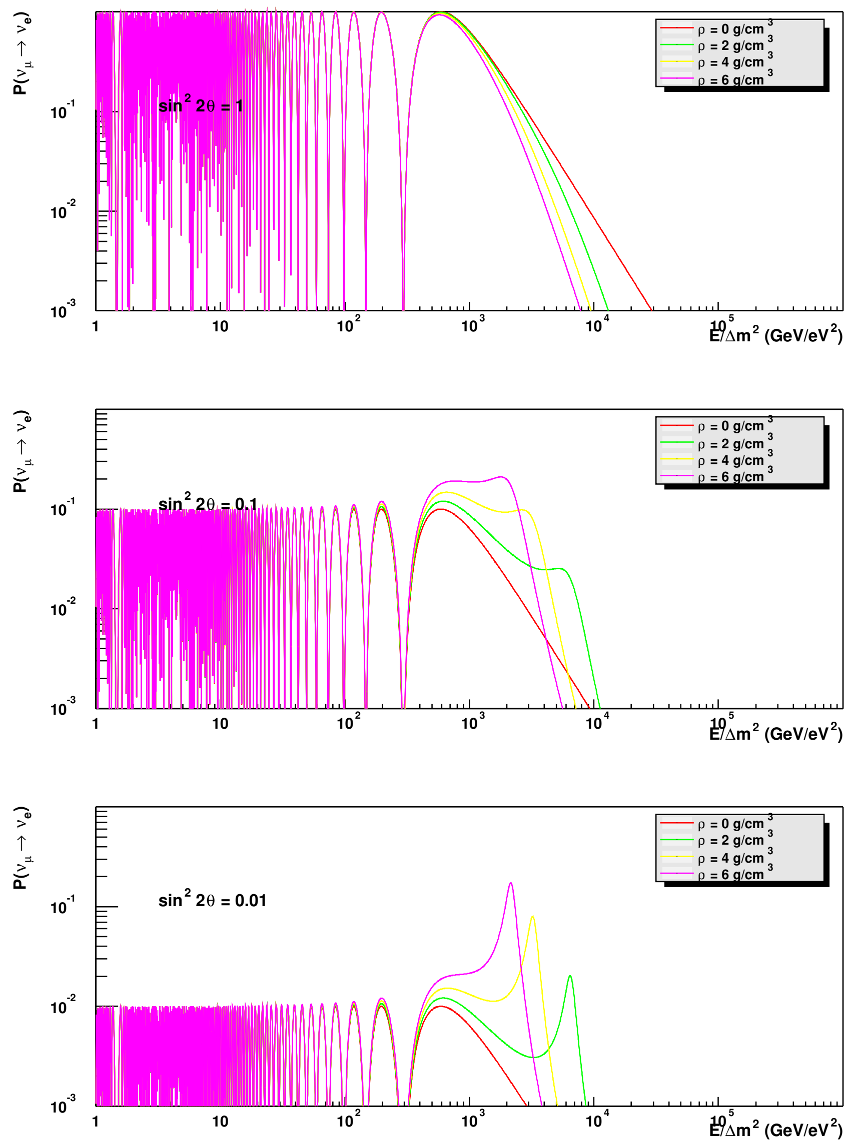

Figure 2.8.: Probability of two-flavour $\nu_\mu \longrightarrow \nu_e$ oscillation in MINOS. Results for different densities of crust are shown.

3. Searches for neutrino oscillations

Frustra fit per plura quod potest fieri per pauciora.

— William of Ockham

Since the first observation [11] of neutrinos from the sun, by Ray Davis in the Homestake mine, experiments around the world have searched for conclusive proof that the solution of the problem of missing neutrinos was oscillations. This chapter presents a summary of the main experiments that have been involved in the oscillation search.

Recall equation 2.34 from chapter 2, the expression for a neutrino to change species in the two generation model:

$$p(\nu_x \longrightarrow \nu_y) = \sin^2 2\theta \sin^2 \frac{1.27\Delta\mathrm{m}^2 L}{E},$$

with L in km and E in GeV. To make a measurement of neutrino oscillations, one therefore needs a source of neutrinos of which one understands the flavour content, a detector (large, because of the small cross-section for neutrino interactions) which is capable of distinguishing between neutrino flavours, and a suitable distance L between them. The combination of L and E used will determine the region of $\Delta\mathrm{m}^2$ to which the experiment will be sensitive.

There are several available sources of neutrinos:

The Sun. The Sun is powered by nuclear fusion in the core. Many of the reactions produce electron neutrinos, as a result of the conversion of protons to neutrons.

Energies are of the order of 1 MeV, and the distance of the sun from the earth is 8 light minutes, giving $L/E \sim 10^{11}$ km/GeV.

The Atmosphere. Cosmic rays incident on the upper atmosphere produce a shower of particles, including many pions. These decay, as in figure 3.1 and will produce ν_μ and ν_e in a ratio of approximately 2:1. Typical energies range from less than 1 GeV to 100 GeV and more.

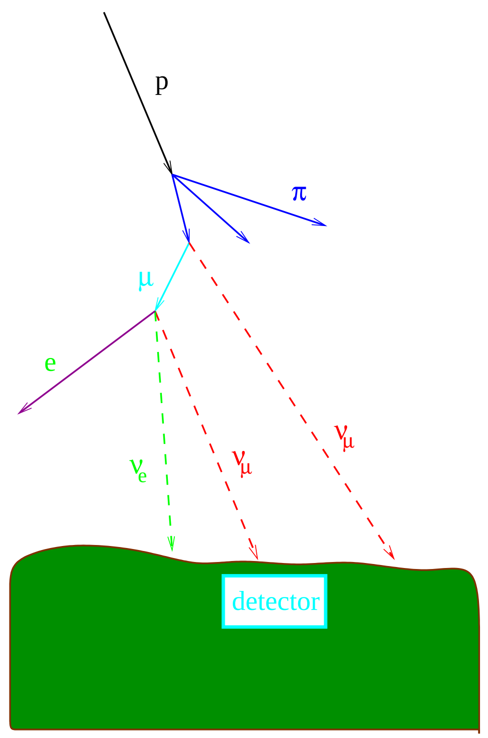

Figure 3.1.: Cosmic rays incident on the upper atmosphere lead eventually to neutrinos. In the limit that all muons decay before reaching Earth, the expected ratio of muon to electron type neutrinos is 2:1.

Nuclear reactors. These produce $\overline{\nu}_\mathrm{e}$ from the β-decay of fission daughter products. Energies are typically a few MeV.

Accelerators. High energy particle beams incident on targets produce large numbers of pions and kaons. These decay principally to muons and muon neutrinos.

Muon rings. Support has been growing in recent years for the construction of a muon storage ring, which would provide a very intense source of muon and electron neutrinos with a well-understood energy spectrum. The neutrino flux from such a device may be sufficiently high as to pose a radiation safety risk due to charged current interactions re-forming muons. This places some constraints on the construction, as neutrinos will pass through any shielding. The only way to ensure safety is to go sufficiently far away from the ring that the divergence of the neutrino beam reduces the flux to a 'safe' level.

3.1. Solar neutrino experiments

Neutrinos from the sun have energies of the same order as nuclear binding energies, typically a few MeV. Three methods have been used to detect them:

Chlorine The method used by the pioneering experiment of Ray Davis in the Homestake mine. Davis had a 100,000 gallon tank of perchloroethane, and used the reaction $\nu_e + {}^{37}\mathrm{Cl} \longrightarrow {}^{37}\mathrm{Ar} + e^-$ which has a threshold energy of 0.86MeV, and so is sensitive to neutrinos from the ${}^{8}\mathrm{B}$ and ${}^{7}\mathrm{Be}$ reactions in the sun (see figure 3.2). After several weeks exposure, the argon is filtered out, and counted by being allowed to decay in a proportional counter. The flux of $\nu_e$was then inferred.

Gallium Another radiochemical method, used by the SAGE [24], [25] and GALLEX/GNO [26], [27] experiments. They use the reaction $\nu_e + {}^{71}\mathrm{Ga} \longrightarrow {}^{71}\mathrm{Ge} + e^-$. Again, after a couple of weeks exposure, the germanium is removed and counted. This experiment has a lower threshold, and so is sensitive so neutrinos from the pp reaction $p + p \longrightarrow {}^{2}H + e^+ + \nu_e$.

Water Čerenkov The method used by the Kamiokande [28], [29] experiment, and its larger successor Super-Kamiokande [30], [31]. Kamiokande consisted of a 3 kT tank of purified water situated 1 km underground in the Kamioka mine in Japan. Neutrinos scatter elastically from electrons in the water, and the recoil electron produces Čerenkov light, which is detected by phototubes surrounding the tank. This method allows one to measure both the direction of the recoil electron, so showing that the neutrinos are indeed coming from the sun, and also its energy, giving one a handle on the energy spectrum rather than just its integral as in the radiochemical experiments. Super-Kamiokande is much larger, holding 50 kT of water, but operates on the same principle.

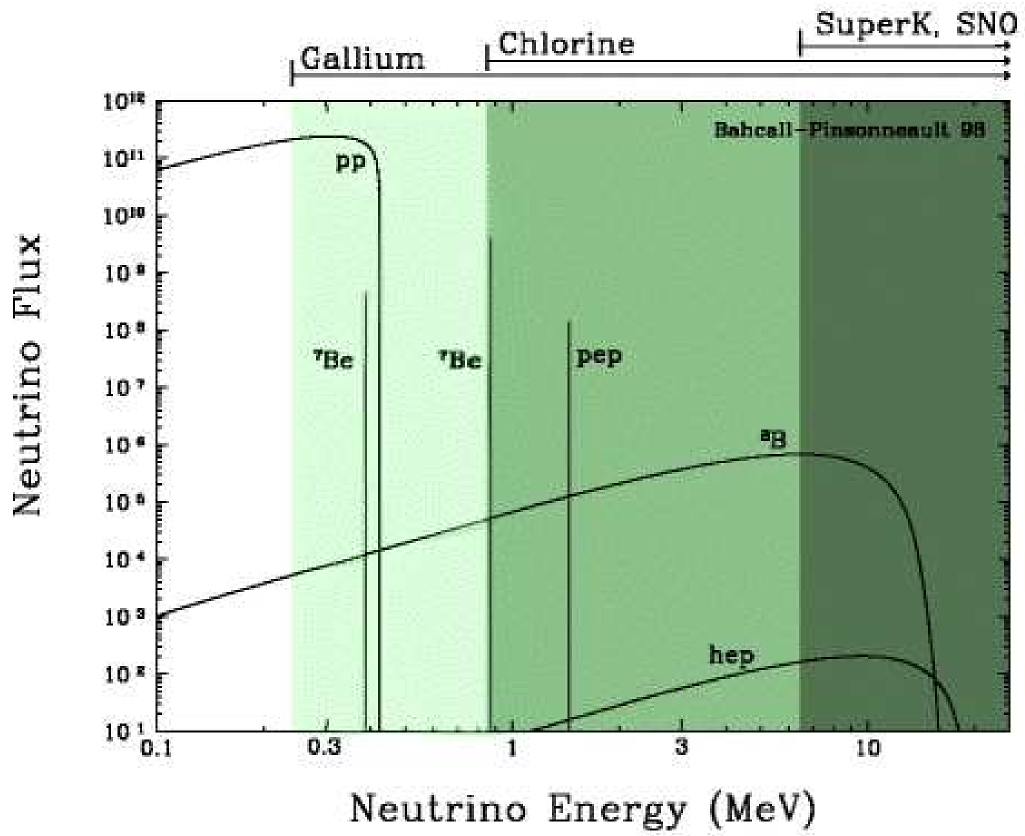

Figure 3.2.: Energy spectrum of solar neutrinos, with threshold energies of the different experiments. From [32].

Figure 3.3 shows a comparison of standard model theory and experimental results for the three kinds of detector. Recall that the water Čerenkov detectors have the highest threshold, with Chlorine in the middle, and Gallium the lowest. The Kamioka neutrinos are primarily from ^{8}B, with a small number from the hep reaction. Explaining the deficit by changing the solar model requires the frequency of the ^{8}B reaction to be halved. But even half the standard model ^{8}B neutrino flux provides all the measured flux for the Homestake experiment, so there is no room for any ^{7}Be neutrinos. These

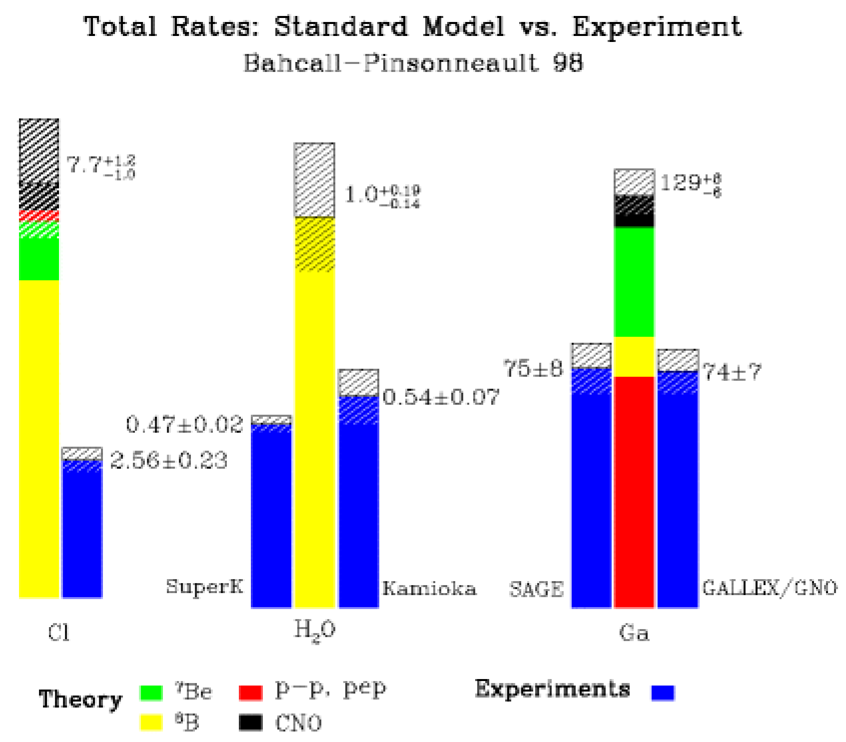

Figure 3.3.: A comparison of measured neutrino event rates and theoretical values for the standard solar model for the different experiments. Units are SNU for the radiochemical experiments, and fraction of expected rate for the water Čerenkov detectors. From [32]

neutrinos must be there though, as ^{8}B is formed from ^{7}Be. It therefore seems impossible to account for the observed solar neutrino flux by an alteration of the solar model. It should also be noted that we can account for all the flux observed by the Gallium experiments with the pp flux, and that it is difficult to alter the expected pp neutrino flux, as it is the rate of the pp reaction that determines the luminosity of the Sun.

3.1.1. SNO

SNO [33] is a new experiment which has recently started taking data, situated in a Nickel mine in Sudbury, Ontario. Like the experiments at Kamioka, SNO is a water Čerenkov detector, but it uses heavy water rather than normal water. SNO's strength is that it offers the ability to separately measure the ν_e charged current rate and the neutral current interaction rate (i.e. all non-sterile neutrinos). In a normal water Čerenkov detector, the interactions observed are primarily charged current, with a small admixture of neutral current scattering. These two processes cannot be distinguished (figure 3.4).

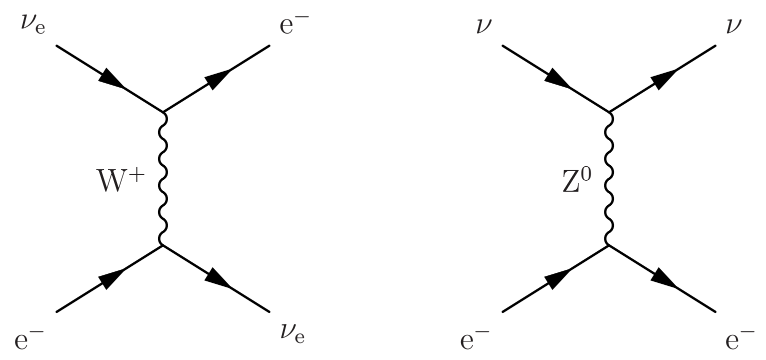

Figure 3.4.: Interactions in a light water Čerenkov detector. Charged current (left) and neutral current (right) are indistinguishable.

In SNO, the much larger cross-section for neutrino interactions with the deuteron dominates (figure 3.5). The charged current interaction $\nu_e + {}^2H \longrightarrow p + p + e^-$ produces an electron which emits Čerenkov radiation. The neutral current interaction $\nu + {}^2H \longrightarrow p + n + \nu$ splits the deuteron if the energy transferred exceeds the 2.2 MeV deuteron binding energy; the detection of the free neutron is the signal for this interaction. SNO will

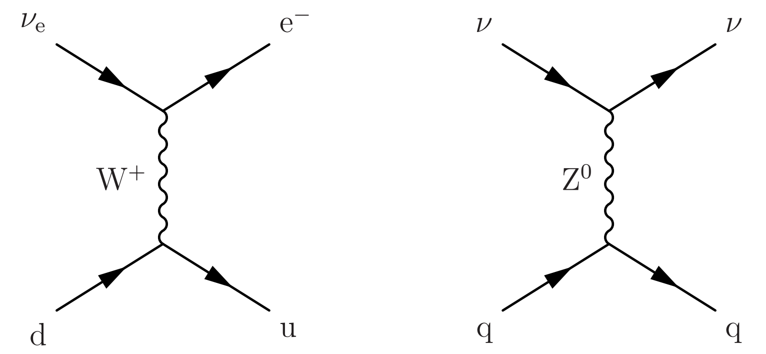

Figure 3.5.: Interactions with the deuteron in SNO—CC (left) and NC (right).

detect these neutrons either by introducing Magnesium Chloride into the heavy water, and detecting the 8.6MeV gamma produced by neutron capture on ${}^{35}Cl$ or by introducing

^{3}He proportional counters, which detect neutrons by the reaction $^3\mathrm{He} + \mathrm{n} \longrightarrow {}^3\mathrm{H} + \mathrm{p}$. With these two measurements, SNO will be able to produce measurements of both total neutrino flux and ν_e flux. If SNO were to measure the total neutrino flux from the sun to be the expected value, but the ν_e flux was lower, it would provide cast-iron proof that neutrinos are changing flavour between the Sun and Earth.

3.2. Atmospheric neutrino experiments

Atmospheric neutrino detectors, operating with neutrino energies ranging from about 1 GeV to many hundreds of GeV can be divided into two groups - water Čerenkov detectors, represented by the IMB [34], Kamiokande [35] and Super-Kamiokande [36], [37] detectors, and iron sampling calorimiter detectors, represented by NUSEX [38], Frejus [39] and Soudan 2 [40].

The experiments usually report their results in terms of the ratio of ratios R, with $R \equiv (N_\mu/N_e)_{data}/(N_\mu/N_e)_{MC}$, with N_μ being the number of muon-like events (events containing a long track) and N_e the number of electron shower-like events. Figure 3.6 shows the values of R measured by the various experiments. For a time it appeared as though the water experiments were producing results of around 0.6, whereas the iron experiments were consistent with 1. Fears that the effect could be due to an insufficient understanding of the water Čerenkov detectors were put to rest by the Soudan 2 experiment, which has higher statistics than the other iron experiments, and does see missing ν_μ.

3.2.1. Super-Kamiokande

Super-Kamiokande is a 50 kT water Čerenkov detector situated in the Kamioka zinc mine in Japan. A larger successor to the Kamiokande experiment, it began taking data in April 1996, and has to date analysed in excess of 1000 days' data.

Super-K divides contained events (events in which the charged lepton does not leave the detector) into Sub-GeV (energies typically a few hundred MeV) and multi-GeV (ener-

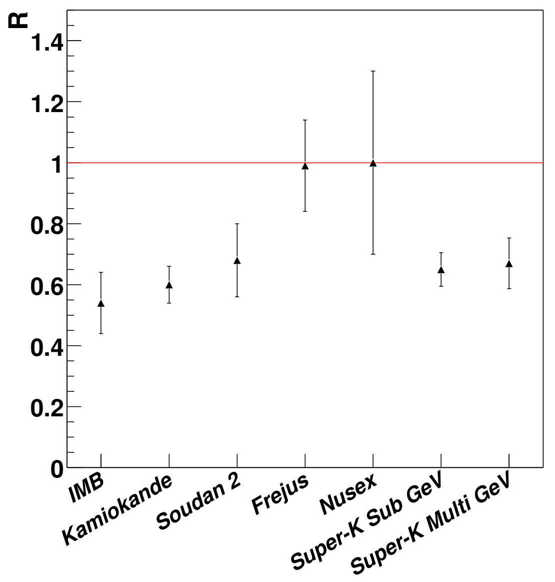

Figure 3.6.: Summary of atmospheric neutrino measurements. For no oscillations, $R = 1$.

gies typically a few GeV). They also analyse partially contained events (muon is formed in the detector and leaves) and upward through-going muons (muons formed in the rock under the detector that pass through it). They achieve their parameter measurement capabilities by measuring the zenith angle of the incident neutrino, and hence the distance it has travelled. This distance ranges from 10 km to in excess of 10,000 km (see figure 3.7).

They observe a deficit of upward-going ν_μ, which varies with zenith angle in a manner consistent with ν_μ oscillations with $\Delta\mathrm{m}^2 \approx 3.2 \times 10^{-3}\,\mathrm{eV}^2$ and $\sin^2 2\theta = 1$. In 1998, they claimed [36], [37] discovery of neutrino oscillations, with a zenith angle effect that differed from their no-oscillation Monte Carlo by 5 σ. They saw no evidence for a distortion in the ν_e spectrum, and so claim $\nu_\mu \longrightarrow \nu_\tau$ or $\nu_\mu \longrightarrow \nu_\mathrm{sterile}$ oscillations. Recently, with improved statistics [41] they exclude $\nu_\mu \longrightarrow \nu_\mathrm{sterile}$ oscillations at the 99% level. As sterile neutrinos do not interact with matter, $\nu_\mu \longrightarrow \nu_\mathrm{sterile}$ oscillations are supressed at energies above about 15 GeV due to the MSW effect in the Earth's core. They do not observe this supression (in their high-energy partially contained and upward muon

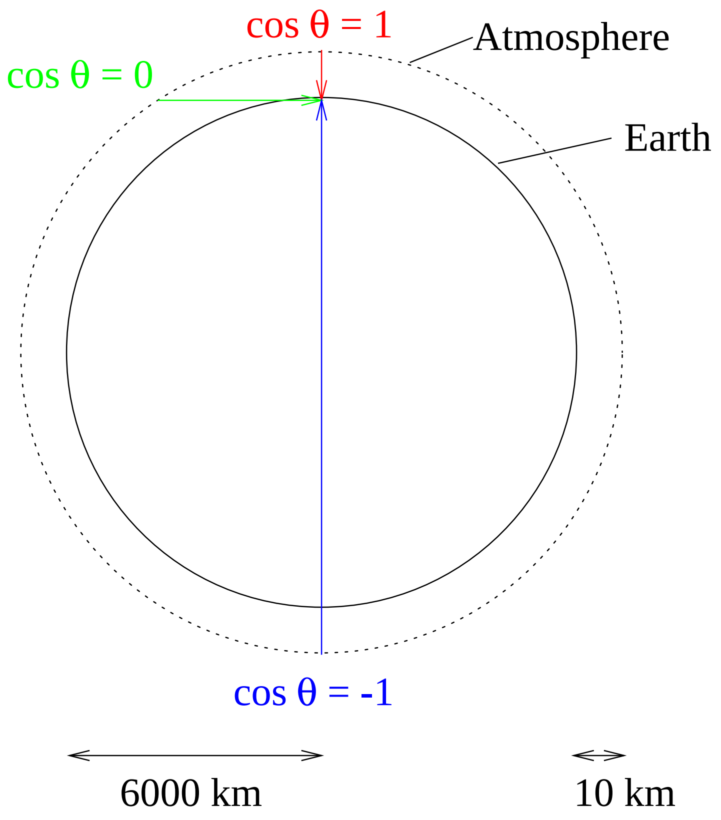

Figure 3.7.: Super-Kamiokande measures neutrino flux as a function of distance travelled.

datasets) and so disfavour $\nu_\mu \longrightarrow \nu_{\rm sterile}$ oscillations.

3.3. Reactor neutrino experiments

Nuclear reactors produce $\overline{\nu}_{\rm e}$ as a byproduct of the fission of ^{235}U and ^{239}Pu. Experiments detect the antineutrinos vis the reaction $\overline{\nu}_{\rm e} + {\rm p} \longrightarrow {\rm e}^+ + {\rm n}$, and typically measure the neutrino energy spectrum at one or more distances of order tens or hundreds of metres from the source. Recalling again equation 2.34, these numbers give a sensitivity down to $\Delta{\rm m}^2 \sim 10^{-2}\,{\rm eV}^2$. The most sensitive reactor experiment is the CHOOZ [42] experiment, completed in 1998. They used a detector 1km from their reactor, and found no distortion in the energy spectrum of $\overline{\nu}_{\rm e}$ when compared to the expected value. They therefore exclude a large region of parameter space down to $\Delta{\rm m}^2 \approx 0.9 \times 10^{-3}\,{\rm eV}^2$ for large

$\sin^2 2\theta$, thus effectively ruling out $\nu_\mu \longrightarrow \nu_\mathrm{e}$ oscillations as a solution to the atmospheric neutrino problem.

3.4. Accelerator neutrino experiments

3.4.1. Short baseline experiments

Short baseline experiments have a detector less than about 1 km from the accelerator source, which produces neutrinos of energies ranging from $\sim 10^{-2}$ MeV to ~ 10 GeV. This gives them sensitiviy down to $\Delta\mathrm{m}^2 \sim 0.1\,\mathrm{eV}^2$. Recent experiments include the CHORUS and NOMAD experiments at CERN, which looked for τ production from a beam of ν_μ, and failed to find any, thus excluding large regions of parameter space for $\nu_\mu \longrightarrow \nu_\tau$ oscillations with $\Delta\mathrm{m}^2 > 1\,\mathrm{eV}^2$.

The only terrestrial experiment to report a positive signal to date is the LSND experiment at Los Alamos. They produce a beam of pions from a proton beam incident on a water target, and use neutrinos both from pion decay in flight (DIF) and μ^+ decay at rest (DAR).

$$\begin{aligned} \text{DIF: } \pi^+ &\longrightarrow \mu^+ + \nu_\mu \\ \pi^- &\longrightarrow \mu^- + \overline{\nu}_\mu \\ \text{DAR: } \mu^+ &\longrightarrow \mathrm{e}^+ + \nu_\mathrm{e} + \overline{\nu}_\mu \end{aligned} \tag{3.1}$$

They detect $\overline{\nu}_\mathrm{e}$ via $\overline{\nu}_\mathrm{e} + \mathrm{p} \longrightarrow \mathrm{e}^+ + \mathrm{n}$—signal = the positron, plus a delayed 2.2 MeV gamma from the capture of the neutron by a proton, and ν_e via $\nu_\mathrm{e} + \mathrm{C} \longrightarrow \mathrm{e}^- + \mathrm{N}$ in the organic scintillator. They report an excess of $\overline{\nu}_\mathrm{e}$ of $51^{+18.8}_{-16.9}$ from the DAR neutrinos, and 21.9 ± 2.1 from the DIF source. Taken together with the excluded regions from the other experiments, these results imply $\nu_\mu \longrightarrow \nu_\mathrm{e}$ oscillations with parameters in a sausage-shaped region around $\Delta\mathrm{m}^2 \sim 1\,\mathrm{eV}^2$ and $\sin^2 2\theta \sim 10^{-2}$. Figure 3.8 shows this region, together with limits from the experiments mentioned below. The KARMEN experiment at the Rutherford Lab is a similar experiment, looking for ν_e appearance in a ν_μ beam from muon decay at rest. KARMEN's results are consistent with expected

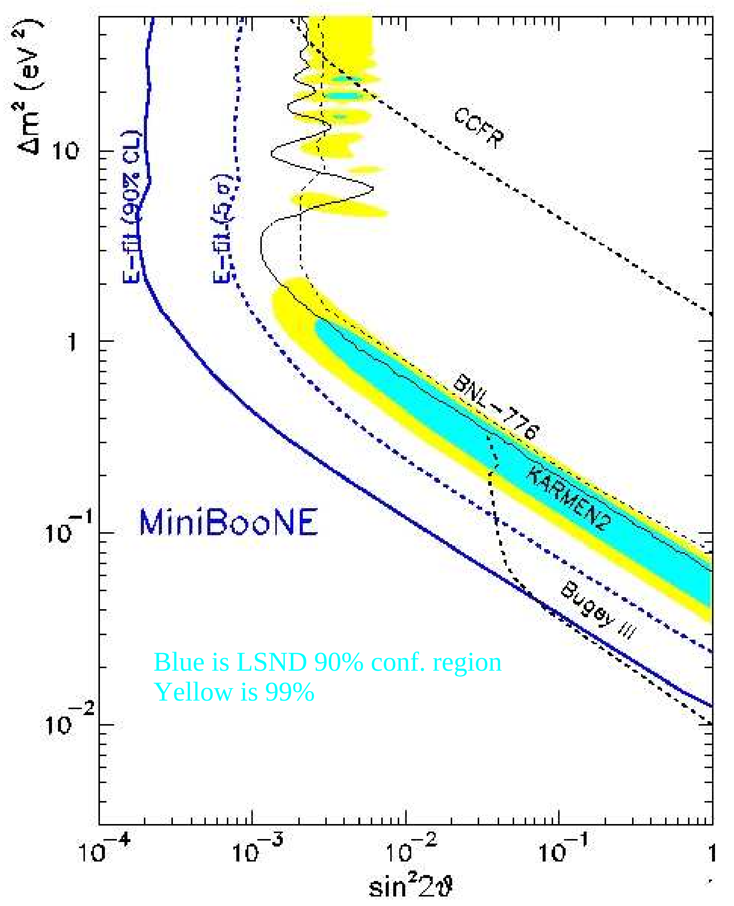

Figure 3.8.: Results of the LSND experiment, and the predicted 2-year limits for MiniBooNE, assuming no oscillations. Also shown are the limits from some existing experiments, including KARMEN.

background, and exclude much of the LSND parameter space; the experiment is not sufficiently sensitive to cover the whole of the LSND allowed region. The forthcoming MiniBooNE experiment at Fermilab is intended to address this result.

3.4.2. Long baseline experiments

The K2K experiment has recently started collecting data. This consists of a beam of ν_μ from the KEK accelerator facility to the Super-Kamiokande detector. Put in baseline, energy. The low energy of the neutrino beam means that K2K can only measure

ν_μ disappearence. Early results are supportive of ν_μ-ν_τ oscillations in the region of parameter space suggested by the Super-Kamiokande experiment, but can make no clear statement.

3.5. Summary of neutrino oscillation searches

If one takes all the experiments at face value, and assumes that neutrino oscillations are the cause of all the discrepancies in the neutrino sector, we have the following position, summarised in figure 3.9:

- The solar neutrino problem can be explained by $\nu_e \longrightarrow \nu_x$ oscillations, with Δm^2 of order either 10^{-5} eV^2 for oscillations enhanced by the MSW effect in the sun, or 10^{-10} eV^2 for vacuum oscillations only.

- The atmospheric neutrino problem requires $\nu_\mu \longrightarrow \nu_x$ oscillations, with Δm^2 of order 10^{-3} eV^2. The Super-Kamiokande zenith angle distribution rules out $\nu_\mu \longrightarrow \nu_e$, and the CHOOZ result also rules out most of the $\nu_\mu \longrightarrow \nu_e$ parameter space. Oscillations from ν_μ to ν_{sterile} are disfavoured by recent Super-K results.

- LSND requires $\nu_\mu \longrightarrow \nu_e$ oscillations with $\Delta m^2 \sim 1$ eV^2.

We see that on the surface, we require three different Δm^2 in order to explain all the data. The fact that we must have $\Delta m^2_{12} + \Delta m^2_{23} - \Delta m^2_{13} = 0$, together with the differing orders of magnitude required by the data seems to preclude the explanation of all these effects in a three-generation neutrino mixing model. Some attempts have been made to do so, for example Cardall and Fuller [43], but all models only really seem to work by ignoring some of the data—Cardall and Fuller chose to ignore the atmospheric neutrino zenith angle dependance. It is fair to say that more recent results from Super-Kamiokande have since ruled out their model. At best, the present situation is unclear. Four generation models including a sterile neutrino are popular, and more exotic models abound. The outlook is good, however. MINOS has the ability to make a definitive statement about whether the atmospheric neutrino anomaly is caused by $\nu_\mu \longrightarrow \nu_\tau$ or

$\nu_\mu \longrightarrow \nu_{\text{sterile}}$ oscillations, or indeed whether we must look elsewhere for the solution to this problem, and to measure the parameters to generally better than 10%. MiniBooNE will be able to confirm or exclude the LSND result. SNO is able to measure both total neutrino flux and ν_e flux from the sun, and so demonstrate whether neutrino oscillations are the cause of the solar neutrino problem, and KamLAND offers a complementary measurement of the solar neutrino oscillation parameters. It seems likely that in 5 years time, physics in the neutrino sector will become significantly clearer.

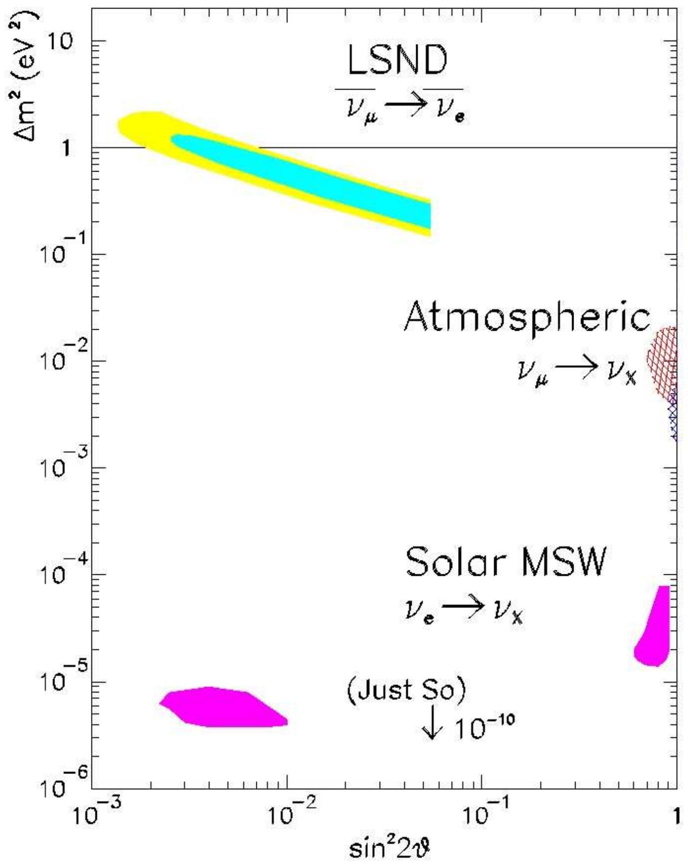

Figure 3.9.: Allowed regions for the various neutrino oscillation experiments. The red hatched region represents Kamiokande and the blue Super-K.

4. The MINOS experiment

The chances of a neutrino actually hitting something as it travels through all this howling emptiness are roughly comparable to that of dropping a ball bearing at random from a cruising 747 and hitting, say, an egg sandwich.

— Douglas Adams, "Mostly Harmless"

MINOS is a long baseline neutrino oscillation experiment. Protons from the Main Injector accelerator at Fermilab are incident on a carbon target, producing pions and kaons. These are focussed by an arrangement of focussing horns, and decay to muons and muon neutrinos.

The muons are absorbed by several hundred metres of rock, leaving a beam of muon neutrinos. These pass through a small detector (the *near detector*) at the edge of the Fermilab site, then travel north under Wisconsin, emerging some 730 km away in the Soudan mine in northern Minnesota. There, the neutrino beam, which has now diverged to ~ 1 km across, passes through the 5000 tonne *far detector*, shielded from cosmic rays by 2700 feet of greenstone rock, before leaving the Earth in southern Canada and heading off into space.

Figure 4.1.: The path of the neutrino beam

The use of two detectors allows one to sample the neutrino beam at widely separated points along their journey, so one may make measurements of neutrino oscillations with-

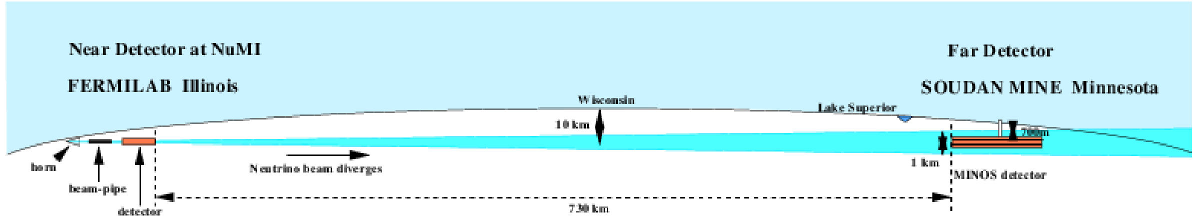

Figure 4.2.: The MINOS experiment

out any of the flux uncertainties inherent in an atmospheric experiment. To reduce any systematic error, one would like the two detectors to be identical; we shall see that due to the widely different environments of the near and far detectors, some differences are necessary, but the detectors have been designed so that their responses are as similar as possible.

The experiment was initially designed with beam energy and baseline to match the parameters suggested by the Kamiokande experiment. Subsequent results from Super-Kamiokande suggest a lower Δm^2, so the beam energy has been lowered to increase our sensitivity to oscillations with Δm^2as low as 10^{-3} eV^2.

The experiment is described in depth in the Technical Design Reports for the beam facility [44] and detectors [45], although some details have since been changed.

4.1. The beam

The NuMI beam uses 120 GeV protons from the Main Injector facility at Fermilab. When the NuMI beam begins operation in 2003, the Main Injector is expected to deliver 4×10^{13} protons on target per spill. The spill is delivered to a 3 m long graphite pencil target via single-turn extraction, producing a spill lasting 10 μs. The π and K produced are focussed by a pair of parabolic magnetic horns and travel down a 675 m evacuated decay pipe, where they decay and produce neutrinos. A water-cooled hadron absorber at the end of the decay pipe absorbs any undecayed hadrons; a further 290 m of rock absorbs the muons. The neutrinos produced are mostly ν_μ, principally from the decay $\pi^+ \longrightarrow \mu^+ + \nu_\mu$. There is approximately 1% ν_{e} contamination.

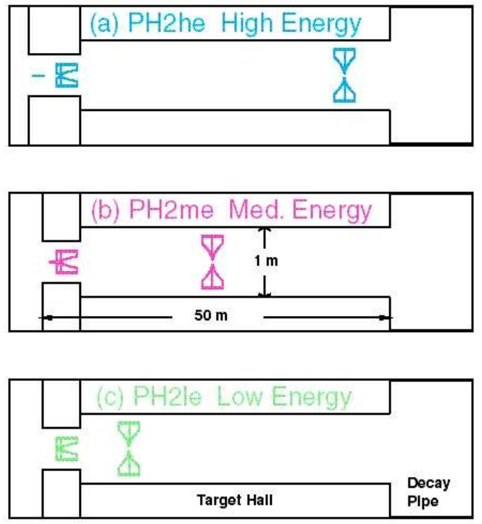

Figure 4.3.: Three beam options with parabolic focussing horns

The use of parabolic horns allows a certain amount of flexibility in the energies of the neutrinos that are produced—by moving the target and second horn, one can choose to focus secondaries with different energies. Three horn configurations were evaluated during the design stages, and are shown in figure 4.3. The neutrino energy spectra that these configurations produce are shown in figure 4.4.

The low energy beam option has been selected for initial MINOS running. This beam has the largest neutrino flux below 3 GeV, which extends the sensitivity of MINOS to oscillations to cover the whole of the region of parameter space preferred by Super-K at 90%.

4.1.1. Possible upgrades

The Hadronic Hose is a proposal to run a current-carrying wire down the centre of the decay pipe. This would focus the pions as they travel down the decay pipe, and so reduce the systematic uncertainties in the neutrino flux (the pion divergence, and hence the production angle of neutrinos in the two detectors, is the major uncertainty in the neutrino beam flux). The hose will be pulsed with a current of 1 kA in phase with the beam spill. MINOS has recently decided to proceed with the hadronic hose in principle,

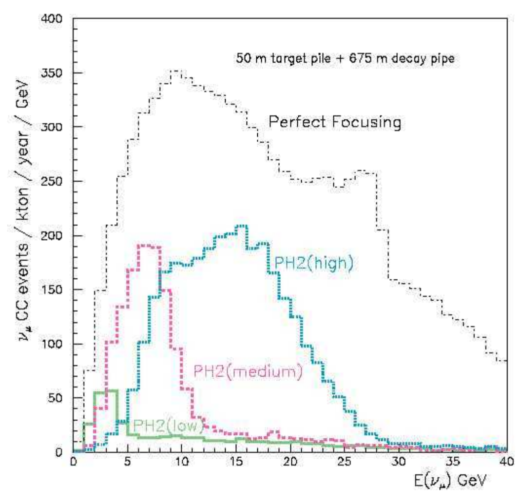

Figure 4.4.: Energy spectra of neutrinos in the far detector for the beams in figure 4.3 together with the spectrum that would be obtained if all the pions produced were focussed.

however extra funds will have to be obtained. Provision will be made in the beam pipe construction so that the hose can be installed at a later date as an upgrade.

MINOS also has the possibility to run with a "narrow band" beam—a neutrino beam with a much narrower range of energies (and a much reduced flux). Parameter measurements obtained with such a beam will have different systematics from those obtained with the wide band beam; running in such a mode will be considered after a couple of years of MINOS running if we have observed and measured neutrino oscillations.

4.2. Detector structure

The MINOS detectors are coarse-grained tracking calorimeters. They consist of 1" thick planes of steel backed with plastic scintillator strips. The strips have a transverse width

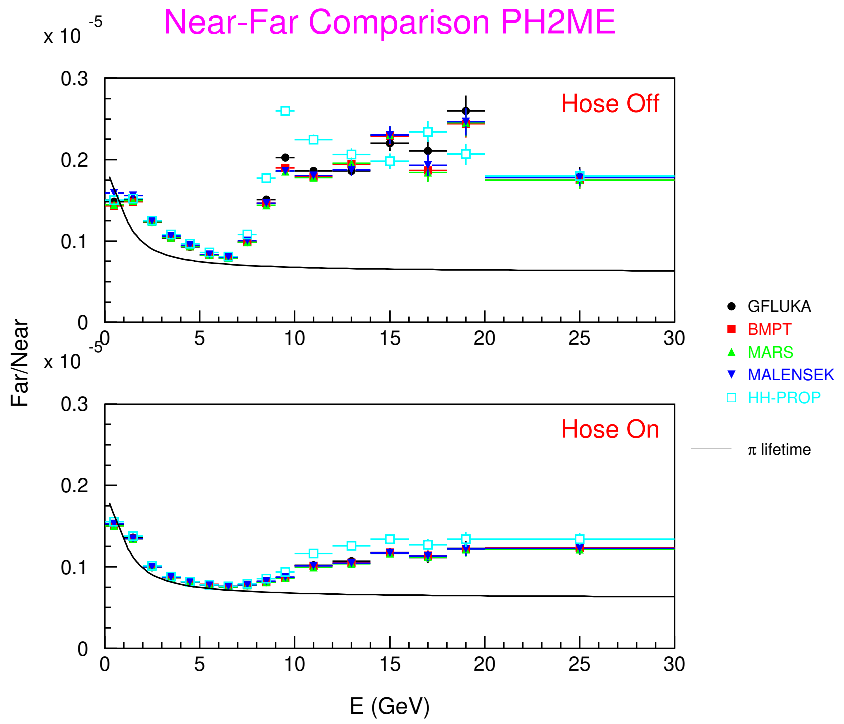

Figure 4.5.: Ratio of far to near flux as a function of energy, with and without the hadronic hose. Different colours represent different beam models. The solid curve is that expected for perfectly focussed pions. This plot is for the medium energy beam, but the effect is similar for all choices of beam.

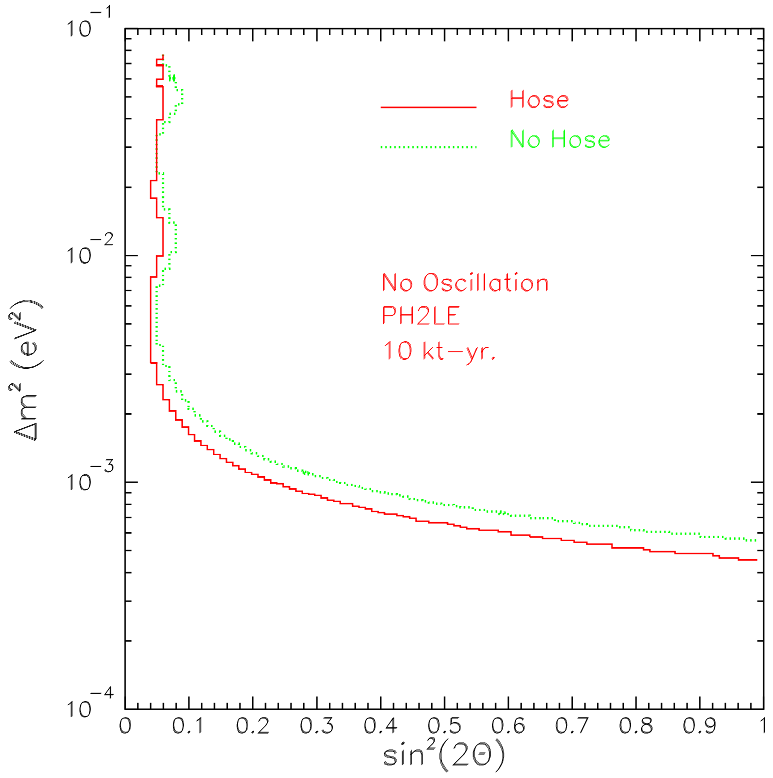

Figure 4.6.: Two year sensitivity with and without the hadronic hose. The region of parameter space above the curves would be excluded if MINOS observed no signal.

of 4.1cm and a thickness of 1cm. The orientation of the strips changes by 90^o in successive planes, providing 3D readout. This gives a detector resolution of approximately $23\%/\sqrt{E}$ for electromagnetic showers, and $55\%/\sqrt{E}$ for hadron showers, where E is measured in GeV.

Figure 4.7.: A short piece of MINOS scintillator, illuminated with a blue LED. The coextruded TiO_2 reflective coating and fibre groove can be seen.

The scintillator is read out by 1.2 mm Kuraray Y11 wavelength shifting fibre lying in a groove along its length (see figure 4.7). Between 16 and 28 strips of scintillator are enclosed in a light-tight aluminium case. The fibres run to a connector at the end of the case; clear fibres then run to boxes containing Hamamatsu multi-anode photomultiplier tubes. The light yield for a minimum ionizing particle in the MINOS detectors is shown in figure 4.8.

A current-carrying coil running through the centre of the detector provides a toroidal magnetic field ranging from approximately 1 to 1.5 T, enabling us to measure the momentum and charge of muons from track curvature. The resolution obtained is limited by multiple scattering in the iron planes; a 10 GeV muon can be measured to 14%. Low energy muons which stop in the detector can be better measured by range—this gives a

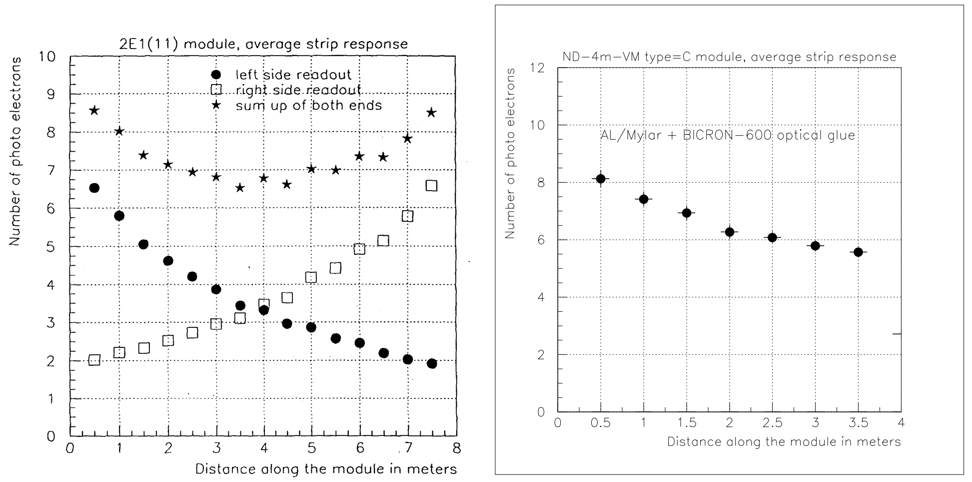

Figure 4.8.: Light yield for a minimum ionizing particle in the far (left) and near (right) detectors. The near detector is read out at one end of the strip; the other end is painted with reflective paint.

resolution of approximately 6%, due to straggling.

4.2.1. The far detector

The far detector is situated 2341 feet under ground in the Soudan mine in northern Minnesota. It consists of two “supermodules” each of 240 8 m diameter octagonal planes. Each plane is backed by a plane of scintillator in alternating ‘U’ and ‘V’ orientations; there are 192 strips of scintillator in each plane. Figure 4.10 shows an image of one supermodule, with the racks containing the electronics. Hamamatsu M16 photomultipliers (figure 4.9) are used, with fibres from 8 scintillator strips (at 1 m intervals across a plane) feeding each PMT pixel. This “optical multiplexing” is possible as the event rate at the far detector is low, so there will never be more than one event in the detector at once, and the typical width of a shower event in MINOS is 50 cm. The cost saving in reducing the number of PMTs and electronics channels required by a factor of 8 is significant!

The electronics used is a modified version of the Viking chip made by IDE AS of Norway. This is a single chip device, containing 32 channels of shaping amplifiers and track and hold circuitry; each channel is then digitised in turn. This electronics, together

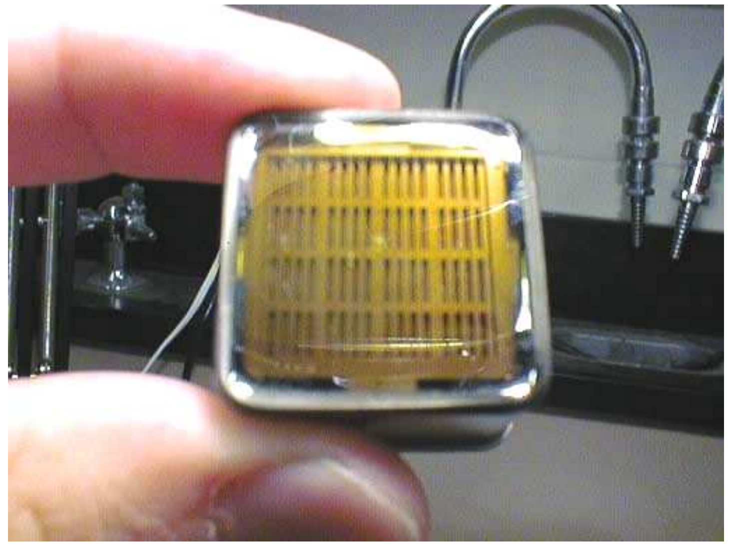

Figure 4.9.: A Hamamatsu R5600-M16 multi-anode photomultiplier as used in the far detector.

with the PMTs and "mux boxes" that map readout fibre cables from the scintillator modules on to PMT pixels, are located in the racks on the top and bottom levels in figure 4.10. The mid-level racks hold the pulser boxes, LeCroy 1440 high voltage supplies for the PMTs and VME readout control crates, with single-board computers running the VxWorks real-time operating system.

4.2.2. The near detector

The near detector serves as a zero-distance reference for the far detector. The MINOS approach is to compare the flux and flavour of neutrinos at Fermilab, and compare them with those measured at Soudan. Such a direct comparison renders us comparatively insensitive to the details of beam Monte Carlos. As well as being technically superior, the ongoing scepticism regarding the LSND result leads us to believe that this approach will produce an easily understandable result that will gain rapid acceptance by the community.

For the near detector to serve as such a reference, we obviously want it to be as similar as possible to the far detector. Unfortunately, the widely differing environments

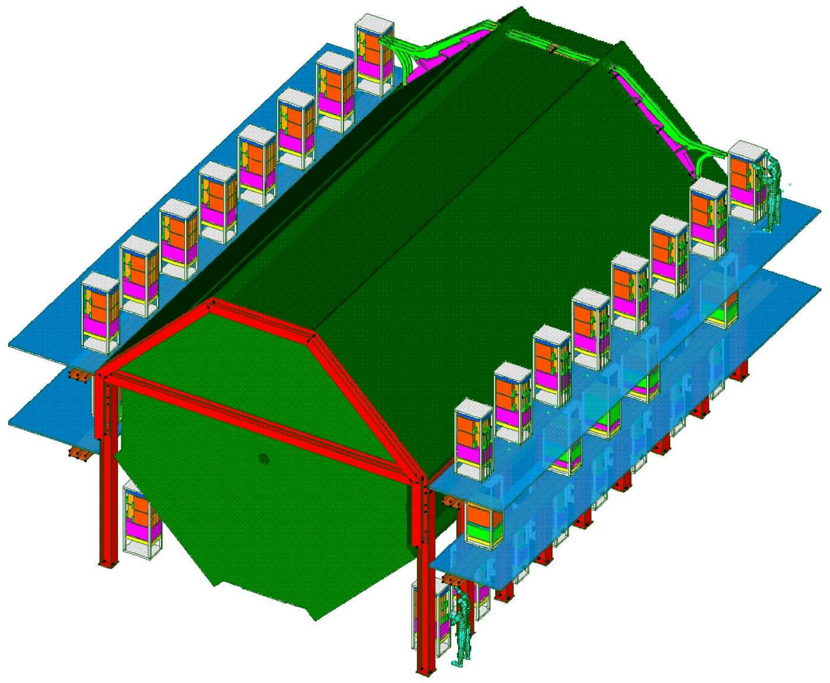

Figure 4.10.: One supermodule of the far detector

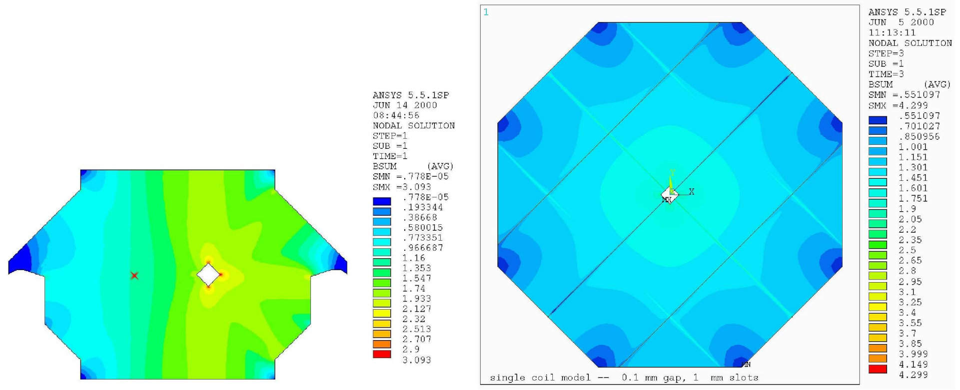

Figure 4.11.: A comparison of the magnetic fields in the near and far detectors. The red cross in the near detector (left) marks the position of the beam centre. The beam is 25 cm in diameter at the near detector; the near detector face is 4.8 m across. The laminate structure of the far detector planes can be seen in the field map—access to the Soudan mine is via an 8 m × 2 m × 2 m lift shaft, so the planes must be welded together underground.

at the two locations necessitate some differences:

- The neutrino spectrum at Fermilab and Soudan will be slightly different. The neutrino source (the decay pipe) is a line source. Lower energy neutrinos tend to be produced at large production angles and upstream in the decay pipe; simple consideration of the geometry indicates that the two detectors will be sampling different spectra. The effects can be minimised by using only the central 25 cm of the neutrino beam. The Hadronic Hose proposal mentioned above focusses pions in the decay pipe, and so further reduces this difference. The predicted near-far ratio as a function of energy is shown in figure 4.12.

- The far detector expects a few tens of neutrino interactions per day; the near detector will see so many in one $10\,\mu$s spill. This necessitates electronics which can cope with these increased rates. The near detector therefore uses QIE electronics. The QIE system is a deadtimeless multi-ranging ADC system used in the KTEV experiment, and has also been adapted for use in CDF. It's principal difference from the far detector electronics is that it divides input signals into 19 ns time buckets, and digitises each separately. These separate digitisations must be recombined during the event building process. Whilst this does offer a potential systematic difference, it is believed that the electronics is well understood. The LED pulser does have an external trigger, which could be used to pulse the LED at different starting positions within the 19 ns bucket, in order to investigate any residual systematics at the bucket boundaries.

Because of the high neutrino flux at the near detector, maximising the fiducial volume is not important. To obtain the cleanest data without extra cost for superfluous hardware, the detector is logically divided into 4 distinct sections, as shown in figure 4.13.

The first 20 planes (the 'veto region') are used to identify muons formed in the rock in front of the detector. The next 40 planes ('target region') must contain the vertex of a neutrino event to be used in near-far comparisons. A further 60 planes contain the

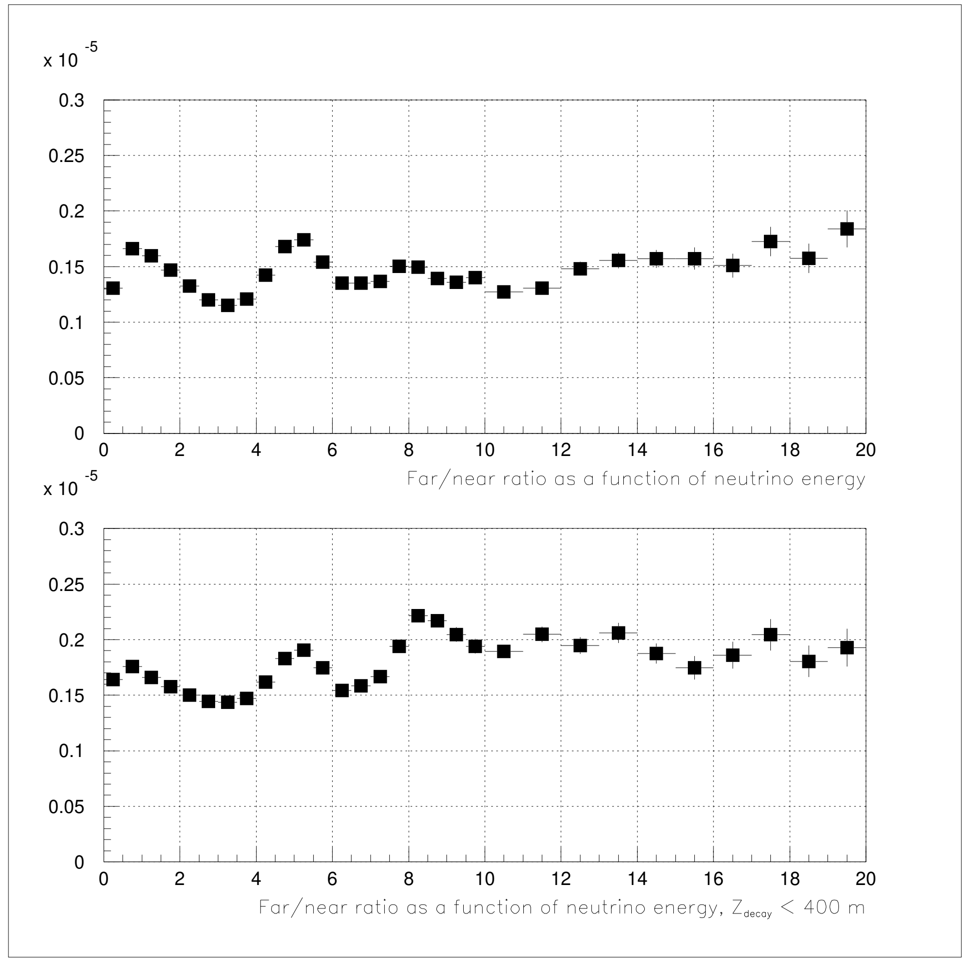

Figure 4.12.: Ratio of neutrino flux in the near and far detectors as a function of energy. Top plot—all neutrinos, bottom plot—neutrinos formed in the upstream half of the beam pipe.

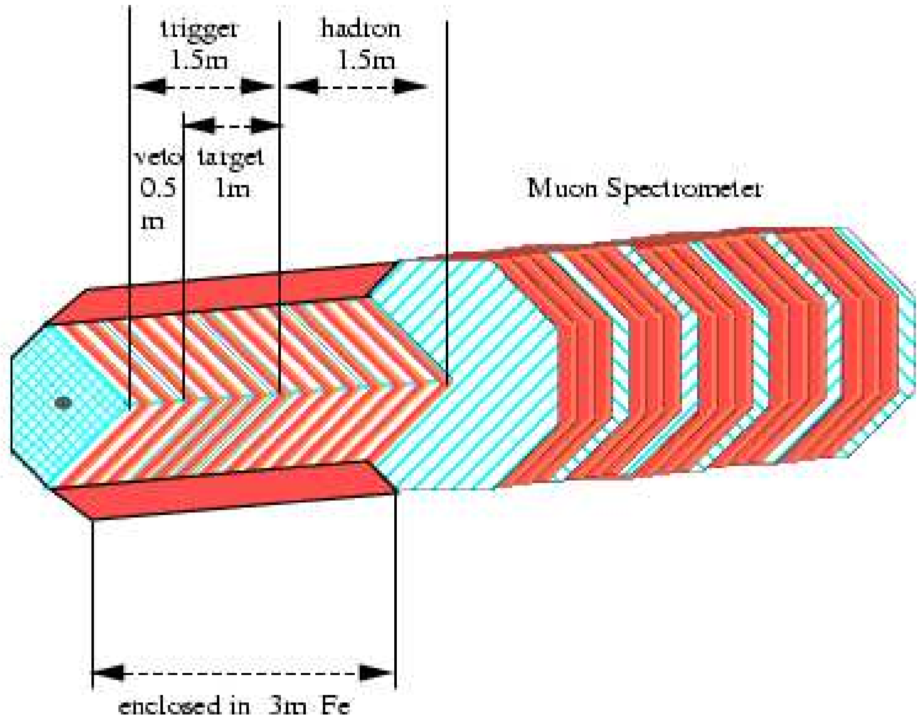

Figure 4.13.: Schematic diagram of the near detector

hadron shower, and the final 160 planes form the muon spectrometer, which measures the momentum of muons produced in the target region. Due to the small size of the beam spot, and hence the small area that can be covered by a shower, only one quadrant of the detector is instrumented with scintillator (although due to the asymmetric nature of the detector this forms rather more than a quarter of the area of the plane) in the calorimeter region (the first three logical sections). Every fifth plane is a full plane of scintillator, as Monte Carlo studies showed that low energy muons tend to wander out of the instrumented quadrant. As the spectrometer section's only function is to measure the muon momentum, by either range or curvature, it does not need to be as fine-grained as the calorimeter. Only every fifth plane is instrumented.

4.3. Measuring neutrino oscillations

There are two tasks facing MINOS. The first is to determine whether or not neutrinos oscillate. Whilst the Super-Kamiokande result is a strong indication that they do, any experiment that does not measure both the 'before' and 'after' states in some way will always leave a certain amount of room for doubt. The second task is to measure the oscillation parameters. The principal oscillation tests in MINOS are described here.

4.3.1. NC/CC

This is the most robust oscillation test available to us. We define the ratio

$$T = \frac{(NC/CC)_{\mathrm{near}}}{(NC/CC)_{\mathrm{far}}} \tag{4.1}$$

where CC is the number of ν_μ charged current-like events (events with a track) and NC is the number of short events. In the presence of neutrino oscillations, this ratio will be reduced from unity by a decrease in ν_μ events. The number of short, showering events will increase if ν_{e} are produced.

Using the ratio NC/CC removes most of the beam-dependent systematics, as one is not sensitive to the details of a Monte Carlo projection of the flux from the near to far detector. The ratio of ν_μ CC to NC cross-sections varies only slowly with energy. The major systematic error is the number of neutral current events recorded—the cross section for neutral current interactions rises sharply at small visible energy. It is therefore important that our lower energy cutoff is well understood; the better our calibration, the lower we can push this bound in energy, and so the more statistics we can accumulate.

4.3.2. CC energy spectrum

Figure 4.15 shows the expected ν_μ energy spectrum at the far detector with and without oscillations for different $\Delta \mathrm{m}^2$. The position of the most depleted region gives a measurement of $\Delta \mathrm{m}^2$, the fractional amount of depletion is given by $\sin^2 2\theta$. Here, absolute energy calibration is obviously important for the measurement of $\Delta \mathrm{m}^2$, but relative energy calibration is also important at low $\Delta \mathrm{m}^2$ for determining whether oscillations occur.

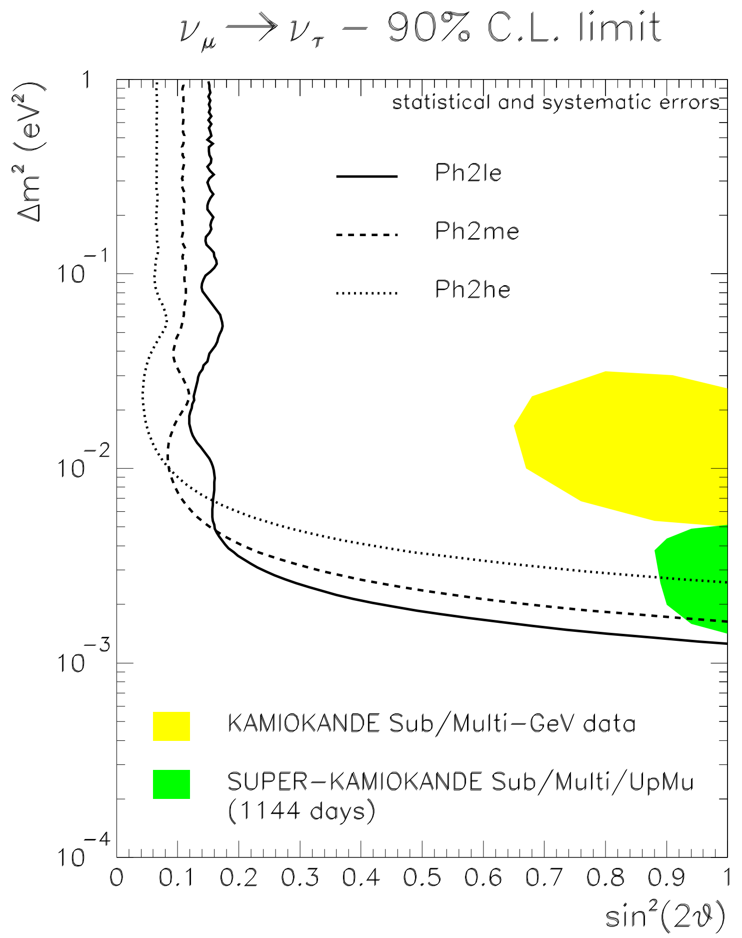

Figure 4.14.: Two year 90% exclusion limits for the NC/CC test.

Consider the lower plot in figure 4.15. The red curve does not look unlike the black one with a small energy shift. A miscalibration here can cause a false positive signal, and uncertain calibration reduces the certainty with which we can claim that oscillations occur.

Despite the more troublesome systematics, and its dependence on the Monte Carlo extrapolation of the neutrino flux from the near to far detector, this method is potentially the most sensitive to low Δm^2 oscillations. Figure 4.16 shows a comparison of the sensitivity of this measurement with the Super-K preferred region. Assuming $\nu_\mu \longrightarrow \nu_x$ oscillations are observed, figure 4.17 shows the precision with which the parameters can be measured.

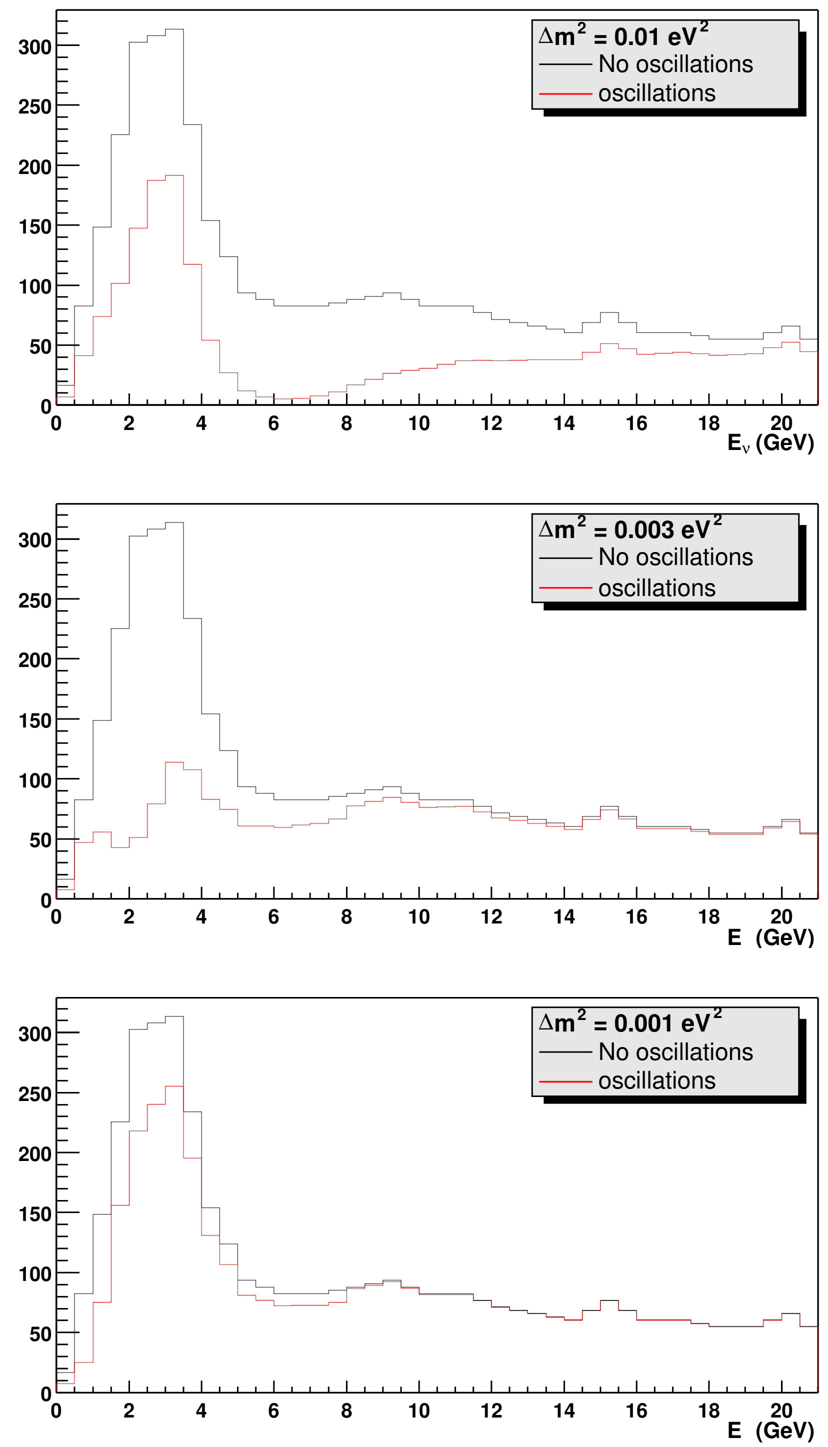

Figure 4.15.: Expected CC energy spectrum for 1 year's running with the low energy beam, for different Δm^2 and $\sin^2 2\theta = 1$. Expected detector energy resolution has been included.

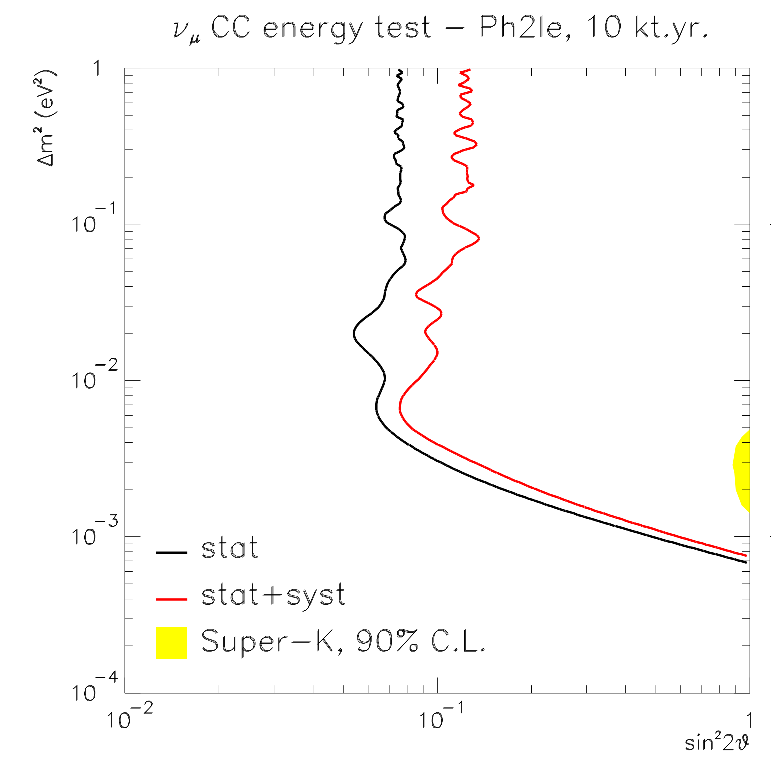

Figure 4.16.: Two year 90% exclusion limits for the CC energy test and the low energy beam, assuming no signal is observed. 2% beam systematics have been assumed.

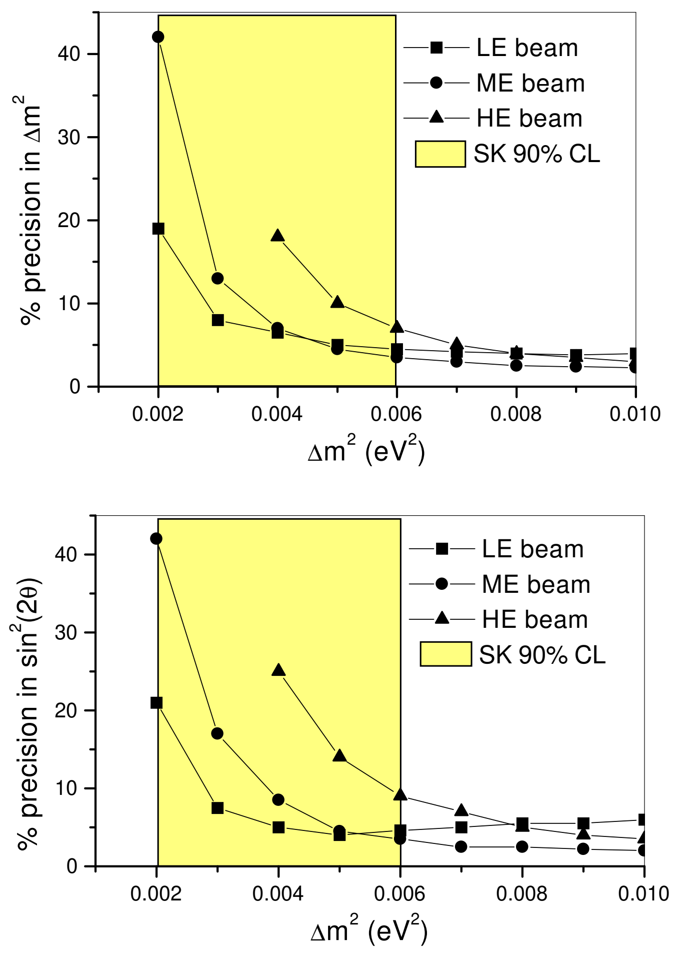

Figure 4.17.: Parameter measurement capabilities of MINOS (assuming 2-generation mixing is dominant) for the different beam options, compared with the Super-K preferred region. The low energy beam is the optimum choice.

5. Calibration overview

Science is built up of facts, as a house is with stones. But a collection of facts is no more a science than a heap of stones is a house.

— Jules Henri Poincare

When MINOS begins to take data, we will be presented with a number of events, each of which will consist of ADC values for many scintillator strips. It is the aim of the calibration systems to convert these sets of ADC values ultimately into energies of muons, hadronic and electronic showers.

5.1. Physical processes in data readout

Consider the readout chain in figure 5.1. The following processes occur to produce a signal in the MINOS data:

1. A charged particle loses energy by ionization in the scintillator, causing it to emit light. The amount of light produced will be a function of the quantity of fluors in the plastic, the type and energy of the particle involved, and its path length in the scintillator.

2. The light is absorbed by the wavelength shifting fibre, re-emitted as green light in the fibre, and travels to the end of the scintillator module, where it transfers to clear fibre cables and on to a photomultiplier pixel. This will contribute a multiplicative

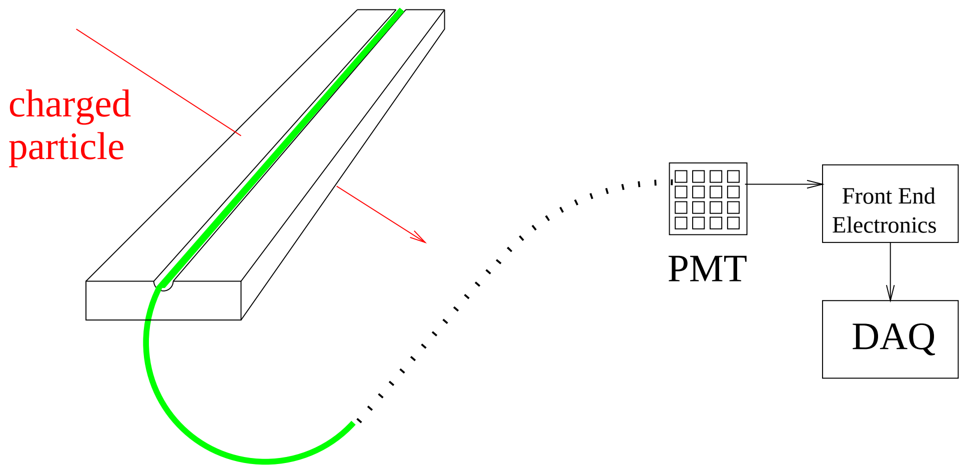

Figure 5.1.: Signal readout path from scintillator to electronics in MINOS

transmission factor, comprising of the transmission at the optical couplings, the conversion efficiency of the wavelength shifting fibre, and attenuation in the green and clear fibres.

3. Light incident on the phototube produces a charge pulse at the anode. Here we have a multiplicative factor for the quantum efficiency of the photocathode and a function for the gain of the PMT. This function is expected to be linear at low light levels, but to deviate from linearity at higher light levels ($\sim$ 100 photoelectrons) due to space charge effects in the later dynode stages.

4. The charge pulse is then digitized by the front-end electronics. The number of ADC counts produced will then be a function of the input charge. The details of the function will depend on the kind of electronics used (the QIE electronics in the near detector has quite different characteristics from the IDE Viking chips used in the far detector) and the electronic characteristics of each channel.

5.2. Calibration methodology

The calibration process consists of reversing the above processes. We begin with the ADC counts from the DAQ, and by applying the calibrations in order, produce particle energies. An overview of the calibration process is given in figure 5.2. The processes outlined in green operate on a per-channel basis; the conversion to energies of muons and showers occurs on a global basis.

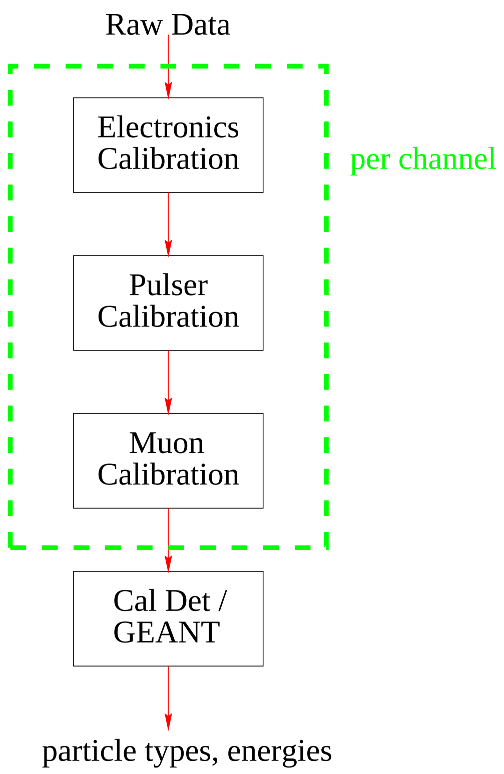

Figure 5.2.: Calibration overview

5.2.1. Electronics

The electronics is calibrated by a charge injection process. Known quantities of charge are injected into the front end electronics, and digitized. This then produces a function enabling us to convert ADC counts to charge.

The charge is set using a digital potentiometer, which is linear to significantly better than 1% [46] over the range of the ADC. The far detector contains a simple single-range ADC, so integral nonlinearities can be measured with only a few points. Differential nonlinearities do not concern us.

In the near detector, the multi-ranging QIE electronics has 32 integrating capacitors, each of which will provide its own slope and offset. Many more points must be measured here.

Pedestals

The subtraction of pedestals from the data forms part of the calibration. The MINOS electronics continuously monitors pedestals in the front end (by digitising channels when there isn't an event in that part of the detector). Summaries are produced for monitoring purposes, and pedestal subtraction is carried out online, in the front end, so that by the time the data reaches the trigger processors and the datastream, the pedestal for each channel is known to be zero. We shall see that the use of muon calibration places a stringent requirement on the measurement of the pedestals.

5.2.2. Light injection

Just as the electronics calibration functions by injecting known quantities of charge and observing the response, so we may calibrate the optical readout chain by injecting known quantities of light. It may seem at first as though this measurement will calibrate the whole optical and electronic chain, so removing the necessity for the charge injection calibration, but we shall see that this is not the case—the electronics calibration will be needed to linearise the PIN diode readout. In any case, the nonlinearities present in the optical and electronic systems have quite different causes, and quite different characteristics, so it makes sense to factor the calibration into a product of these two independent parts.

The light injection system is described in detail in the next chapter. This section presents a summary of the principle involved.

The light injection calibration system is based around a number of blue LED pulsers.

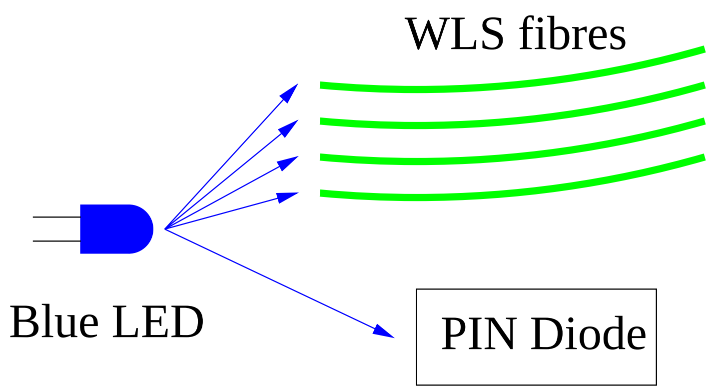

Figure 5.3.: Schematic diagram of the LED calibration system

Light from an LED is fanned out between many channels, and injected into the wavelength shifting fibre at the module end. Given the many optical couplings involved, it would be difficult to arrange for a known amount of light to enter each channel—instead the LED is also monitored by a PIN diode (see figure 5.3). The PIN diode is a reliable solid state device that is stable over time, and produces an output which is a linear function of the intensity of incident light. Plotting the signal in the PMT against the signal in the PIN diode then allows us to measure the nonlinearity of the PMT. This is the reason why we still need the electronics calibration—the PIN diode and its amplifier have been shown to be linear, however the PIN is read out and digitised by the same electronics that reads out the PMTs. If our PIN scale is to be a linear function of light intensity, we therefore require the linearised electronics scale which the charge injection calibration produces.

By pulsing light at about 10 intensities over the range in which we are interested, we can map out the PMT response curve. To do this for the whole detector will take a large fraction of a day, and we would like the ability to monitor PMT gain variations on an hourly timescale. We make the reasonable assumption that whilst the gain might fluctuate on a short-term basis, the nonlinearities, which are caused by space-charge effects in the later dynode stages, will stay constant in charge terms. Therefore we will measure the linearity curve every month as a check, and with the expectation that the nonlinearities will not change over this timescale, but measure the response to one single point, chosen to be in the linear part of the PMT response curve and of a similar size to the typical signals we will be measuring, every hour. This point will enable us to scale the gain curve up and down as in the somewhat exaggerated figure 5.4.

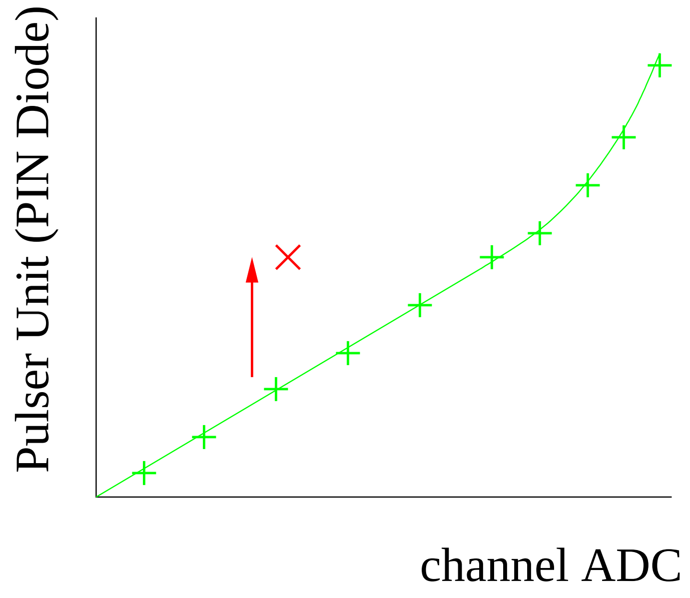

Figure 5.4.: The LED pulser calibration principle

5.2.3. Muons

With the pulser calibration, we have converted ADC counts to a scale which is linearly proportional to the amount of light arriving at the phototube. We still need to account for strip to strip differences in scintillator yield and coupling, and also the channel to channel variations in the coupling of the LED pulser to the PMTs. To do this, one needs to identify a standard source of charged particles that the bulk detector can be exposed to over the lifetime of the experiment, measure the light output of each scintillator strip when exposed to this 'standard candle', and so convert the arbitrary pulser units to a portable unit which is dependent only on the type and energy of the particle in the scintillator.

One common method used in particle physics experiments is to have one or more radioactive sources, and to periodically expose the detector to the source and examine

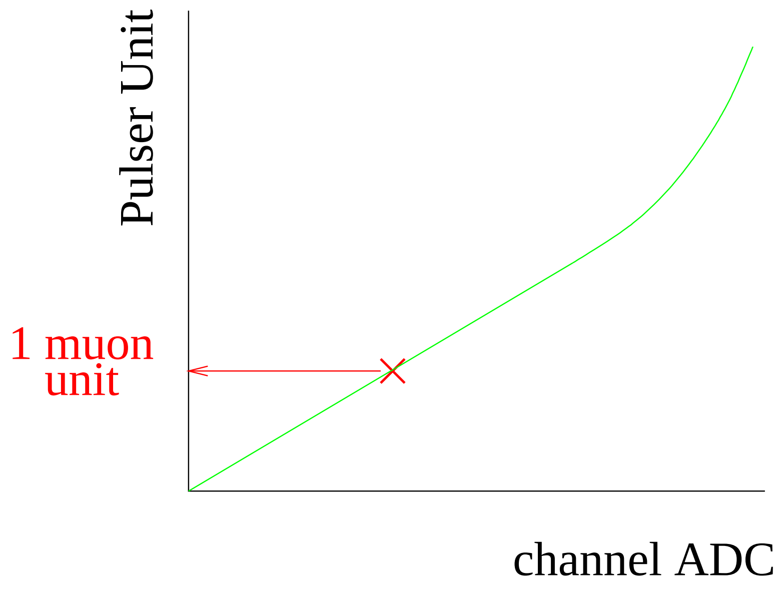

Figure 5.5.: Muon calibration converts arbitrary pulser units to multiples of light produced by a 'standard muon'.

the response. The area of scintillator that would need to be covered, and the engineering tolerances required on the distance from source to scintillator render this method impractical.

Instead, MINOS will exploit the ready availability of muons from cosmic rays. The muon calibration is a two-step process:

Strip to strip calibration

The cosmic ray muon energy spectrum at the surface rises sharply towards low (< 1GeV) energies. At the far detector, the low energy muons stop in the rock above the detector, so the mean energy of a muon that survives as far as the far detector cavern is of order 100 GeV. The near detector is much closer to the surface, and muons here have an average energy of a few GeV. These two sets of muons therefore lose different amounts of energy in scintillator, and so cannot be directly compared. However, cosmic muons in one part of the far detector have the same energy as cosmic muons in another part of the far detector, so we can use the 'far detector cosmic muon' as a standard candle to measure the light production of the far detector scintillator strips against. Similarly, ADC counts in the near detector can be expressed in terms of the 'near detector muon'. Simulations suggest that over a period of 1 month, MINOS will observe sufficient muons to produce a strip to strip calibration to a precision of 2.4% in the far detector, and 0.3% in the near detector. As the scatter on these strip-by-strip calibrations will be statistical and not systematic, it has been shown by Litchfield [47] this much precision is not necessary on a strip-by-strip basis.

Relative calibration between detectors

In order to carry the calibration from one detector to another using muons, one must identify an equivalent set of muons in each detector. For this purpose, we will use stopping muons. The energy of a stopping muon can be measured to 6% by looking at its range. At the far detector, 700 stopping muons are expected per day; the near detector expects a figure two orders of magnitude higher. This yields a detector-detector calibration precision of better than 1% per day.

5.2.4. Calibration Detector

Thus far we have managed to convert ADC counts into muon units, where a muon unit is the amount of light produced in a scintillator strip in the MINOS geometry by a 'standard' stopping muon. The final stage is to convert a pattern of hits measured in muon units into visible energy. To this end we will construct the Calibration Detector—a $1\,\mathrm{m} \times 1\,\mathrm{m} \times 60$ plane mini-MINOS, which we will expose to beams of electrons, protons, pions and muons of known energy in a test beam at CERN. The muons will allow us to carry a calibration from this detector to the near and far detectors; the other measurements will enable us to measure the light production of, for example, a hadronic shower caused by a pion as a function of energy, in terms of the muon unit.

A combination of these measurements, and simulations of the neutrino interactions in MINOS will enable us to convert a pattern of hits in MINOS expressed in muon units to a visible energy from the neutrino interaction.

5.3. Summary

We can now revisit figure 5.2 and mark in more detail the conversions that will be made in the calibration process.

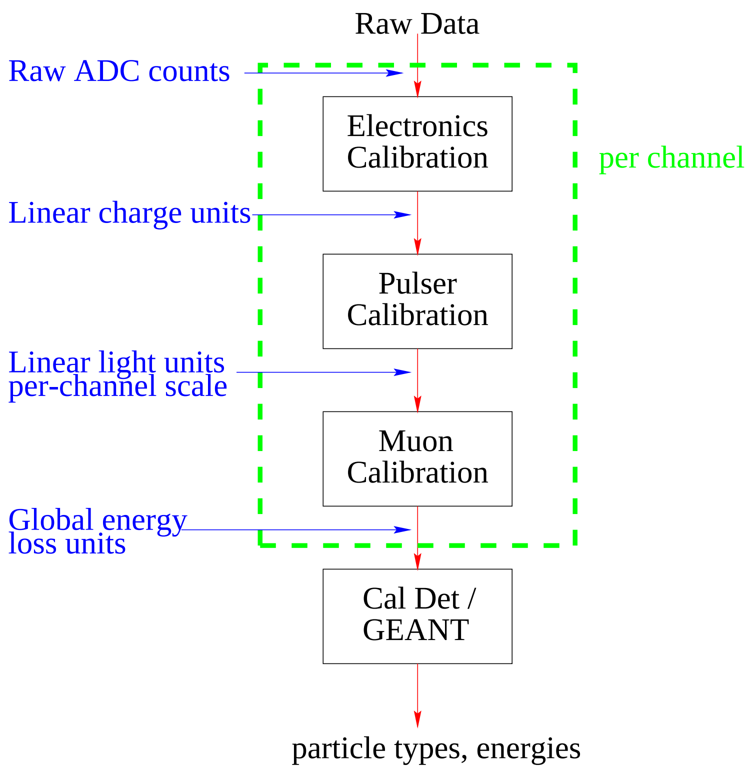

Figure 5.6.: Calibration overview redux

Raw ADC counts are converted by the electronics calibration into units which are linear with charge. The LED pulser calibration then converts these to units which are linear in light incident on the PMT pixel, although the constant of proportionality is different and arbitrary for each channel.

The muon calibration removes this arbitrariness, and the light yield of the scintillator to produce units which depend only on particle energy and type, and the results from the calibration detector feed in to the global event reconstruction to produce energies.

6. The LED light injection system

Thus month after month his papers accumulated, until every corner of the room was stacked with bundles of manuscript which were on no account to be burned, and which could not be put away save by their owner.

— Sir Arthur Conan Doyle, "The Musgrave Ritual"

As described in chapter 4, the MINOS detectors produce blue light in plastic scintillator, which is captured in green wavelength shifting fibre, and transported to Hamamatsu multi-anode photomultipliers. To calibrate the readout chain, we would therefore like to inject light into the system. Experiments in the past have often done this by using a bright light source (typically some sort of laser), a filter wheel with a number of neutral density filters to attenuate the light by known factors, and a series of optical fanouts to distribute the light around the experiment. The recent availability of ultra-bright blue LEDs has enabled us to consider a different approach—rather than employing potentially delicate lasers, and mechanical filter wheels, we use blue LEDs, and vary the light output by varying the current with which we pulse them.

In this chapter, we cover in detail the design of the LED pulser system for MINOS, and our baseline approach to calibration. Wherever possible, we have designed the system to be flexible enough to accommodate debugging tests and uses of the calibration system that we have not yet considered in detail. This chapter is arranged for clarity in sections for the different parts of the calibration system; of course, in reality the different parts were developed in parallel. An unavoidable consequence of this is that some of the designs in earlier sections will only make complete sense when a later section is encountered. I have tried to avoid this as much as possible my my choice of section ordering, but it is not possible to completely remove the forward references.

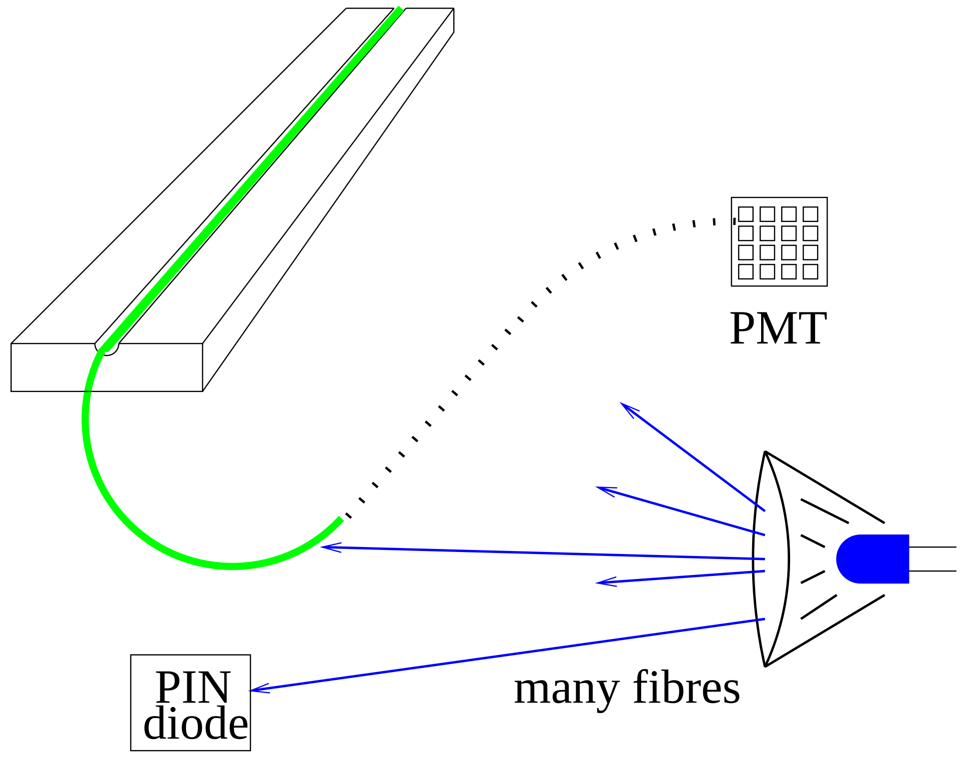

Figure 6.1.: LED calibration scheme. A blue LED fans out into many fibres carrying light, which is injected into the wavelength shifting fibre in the scintillator module. The output of the LED is monitored by a PIN diode, which serves as a constant reference.

6.1. Blue LEDs

The LEDs used in the pulser system are Hewlett Packard HLMP-CB15. Initial development work had been done with the LEDs made by Nichia Chemical Industries; the change to the HP LED was forced as Nichia "upgraded" their LED to one which whilst being 50% brighter, also had a slower response. This made the pulses produced unacceptably long. Some manufacturer's specifications of the HP LED are shown in table 6.1.

Viewing Angle (FWHM)	15^o
Intensity at 20 mA	1575 mcd
Max. mean forward current	30 mA
Max. peak forward current	100 mA

Table 6.1.: LED specifications

Figure 6.2 shows a comparison of the spectra of light emitted from the blue LED with light from an excited piece of MINOS scintillator. The spectra are quite different.

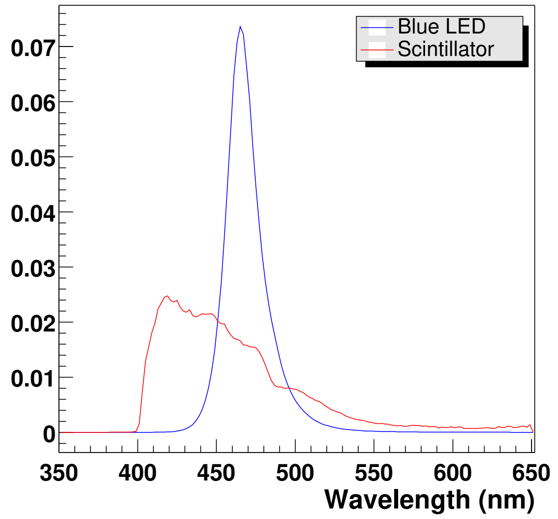

Figure 6.2.: A comparison of the wavelengths of light emitted by the blue LED and stimulated MINOS scintillator.

When the light is absorbed and re-emitted by the wavelength shifting fibre, however, the spectra become far more similar (figure 6.3).

The match is not exact, but it seems sufficiently good that we will not fall foul of any wavelength-dependent aging of the photocathode. Any residual wavelength-dependent aging effects will show up in the monthly muon data, and will look like scintillator aging. The calibration will not be degraded.

Figure 6.4 shows that there is no gain in light for applied current pulses greater than 100 mA. We therefore design the pulser box to deliver a maximum current pulse of

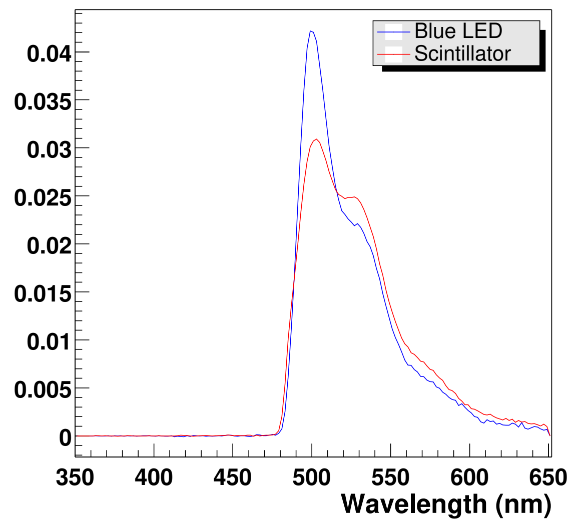

Figure 6.3.: Spectrum of light emitted from wavelength shifting fibre when excited with the light in figure 6.2

100 mA to the LED.

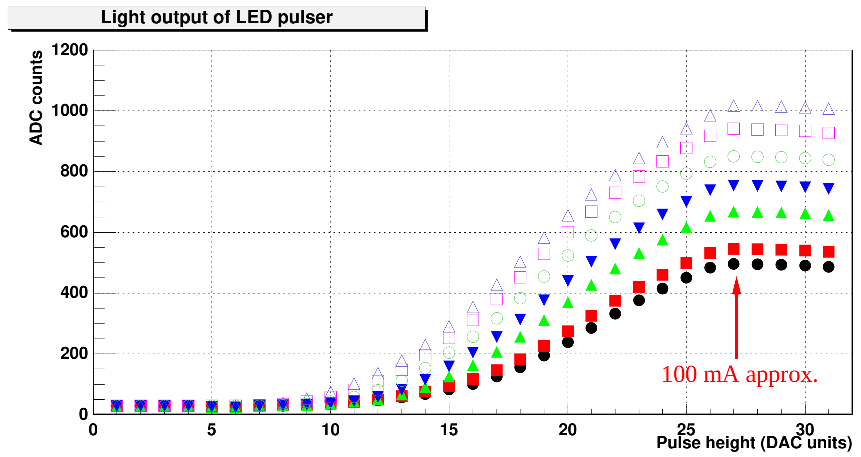

Figure 6.4.: Light output from an early prototype pulser as a function of applied pulse height. An increase in pulse current above 100 mA does not buy you more light. The different data series represent different pulse widths.

6.2. Light injection modules

We generate blue light in the LEDs, but must consider the optimum way to inject this light into the wavelength shifting fibres.

6.2.1. Concepts

Blue light is absorbed by the fluor in the wavelength shifting fibre, and is re-emitted isotropically as green. Most of the green light is trapped in the fibre, and propagates in both directions along it. To achieve maximum absorption, one needs to avoid light being reflected from the fibre, and so would expect the maximum absorption to occur when the incident light made an angle of 90^o with the fibre. A measurement was made to confirm this; the results are shown in figure 6.5.

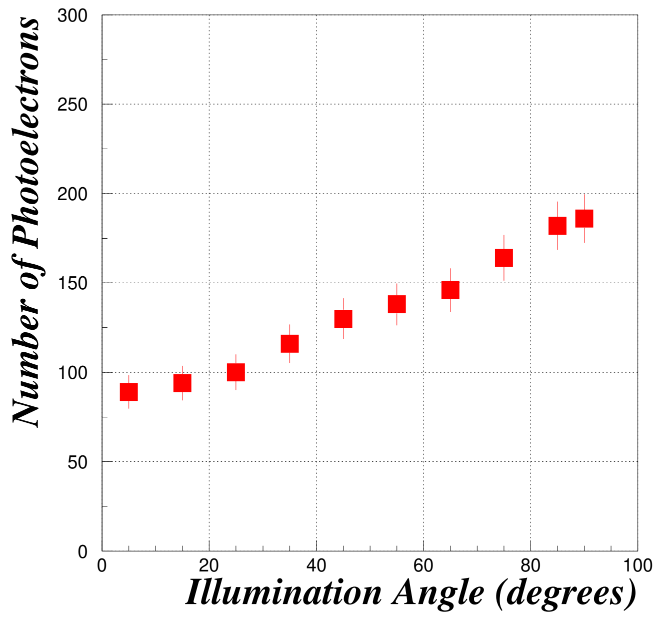

Figure 6.5.: Absorption of blue light from optic fibre at different angles to wavelength shifting fibre.

The scintillator in MINOS is packaged into light-tight modules containing 16, 20, 24, or 28 scintillator strips. The green fibres run to a 28-wide connector at the end of the module, from where the signal is carried by clear fibre cables—the short attenuation length of the wavelength shifting fibre precludes its use for carrying signals all the way to the photomultiplier. If we are to inject light into the green fibres, we must do it at the ends of the modules. Here we are constrained for space—there is a plane of steel adjacent to the scintillator module on one side, and 1 cm away on the other. This does not allow us room to have fibres carrying blue light arriving at 90^o to the planes, so we imagine using a mirror at 45^o instead. We can then run fibres carrying blue light parallel to the readout fibres, as in figure 6.6.

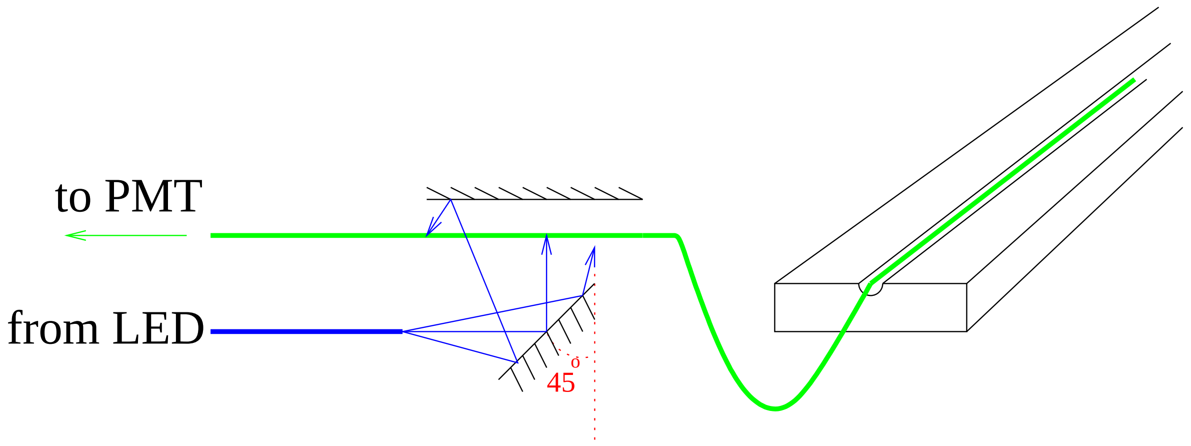

Figure 6.6.: Light injection module concept

6.2.2. Design

At the end of the scintillator modules, there is a manifold (figures 6.7, 6.8) which leads the green fibres from the scintillator to a 28-wide connector; clear fibre cables run from this connector to the PMTs.

The light injection module ("LIM") is made from injection moulded plastic, like the manifold. It forms part of the manifold cover, and fits over the connector end of the manifold, where the green fibres run close together (see figure 6.8).

The LIM is divided into "ashtrays" that hold 10 fibres—the 16 and 20-wide modules use two ashtrays, the wider modules use all three. Each ashtray is supplied with light

Figure 6.7.: A scintillator module with the cover off, with manifold. This is a prototype; the real MINOS manifolds will be black.

via a clear plastic fibre, which connects to the LIM with a push connector. We will see in the next section that the ashtray divisions are required to reduce the effects of cross-talk.

With our design, one LED will illuminate many clear fibres, each of which will in turn illuminate 8 or 10 wavelength shifting fibres. Recall from chapter 5 that we want to pulse at approximately 10 pulse heights to map out the response of the PMT. So that we can map out the full range of signals likely in MINOS in all these channels, without having to pulse at many extra heights, we want the light injected into the wavelength shifting fibres to be as uniform as possible. Our initial design had a flat "mirror" surface, as in figure 6.6; this gave a light output that was strongly peaked towards the centre of the ashtray. A modified ashtray with a surface that curves out towards the edges of the ashtray gives a much more uniform response. The comparison is shown in figure 6.9, and the curved back design LIM in figure 6.10.

Figure 6.8.: Manifold parts. The light injection module is to the lower left of the picture. Ringed in red are the “ashtrays” that reflect the light on to the fibre, the manifold surface behind the fibres is also covered with a reflective coating (ringed green). This is an early prototype using aluminised mylar; the production version will have reflective paint instead.

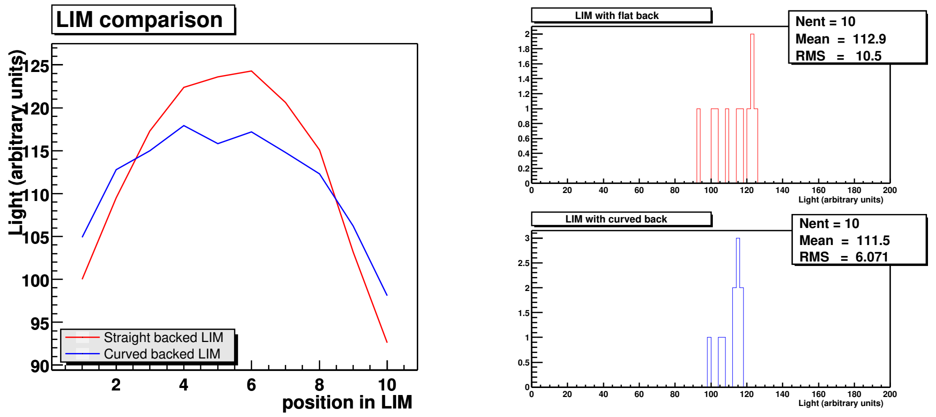

Figure 6.9.: Comparison of signal in fibres across the LIM with two different designs. The curved back design is clearly more uniform.

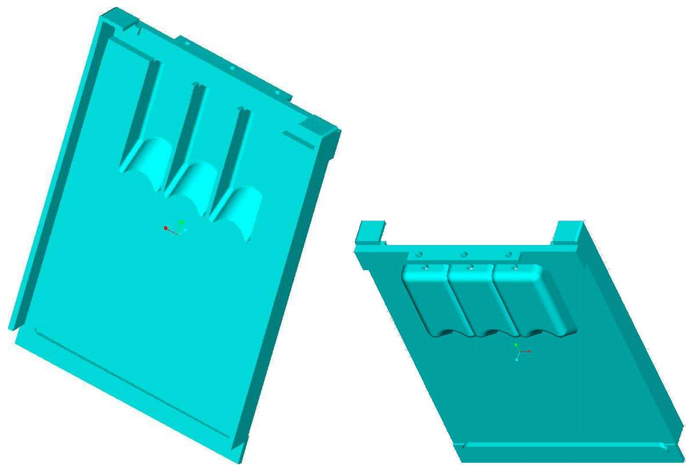

Figure 6.10.: The Light Injection Module (solid model)

The LIM is made from black plastic, and so must be coated with a reflective material. The 45^o mirror design is based upon specular reflection, so our early prototypes used aluminised mylar. It is difficult to cut and stick the mylar in to the ashtray without scratching the surface, so we tested the use of a TiO_2 reflective paint (Bicron BC625) instead. As we see from figure 6.11, using the reflective paint produces 30% more light in the fibres than the mylar. This is likely to be because the mylar used was about 10% transmissive; the middle bar, using the mylar over a painted surface, bears this out to some extent. At any rate, we are able to use the reflective paint, which is easier for module construction, without having less light than our initial design.

6.3. Optical routing constraints

The far detector has more than 180,000 optical readout channels, which are 8-fold multiplexed onto 23,000 PMT pixels. Calibrating these channels one at a time is obviously impractical; this section discusses the constraints that we would like to put on which channels are illuminated at the same time. This will have some impact on the light injection modules described in section 6.2.

We have established that we must illuminate many channels with each LED. We

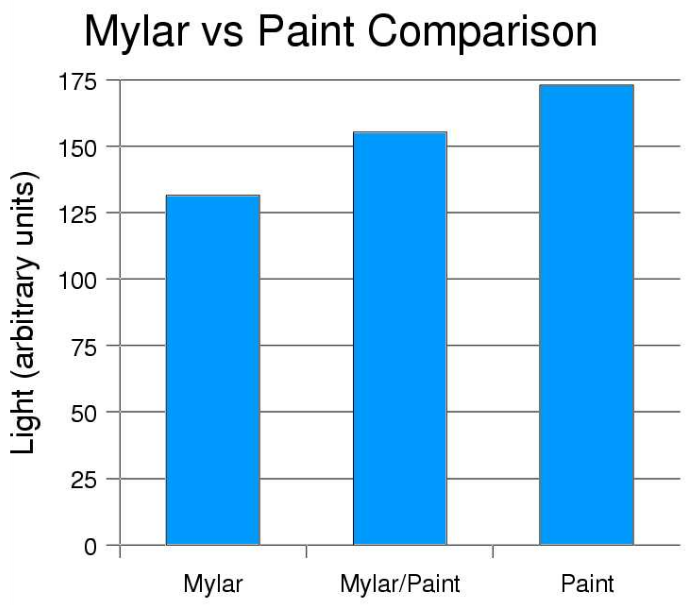

Figure 6.11.: Comparison of light transmitted to wavelength shifting fibre with aluminised mylar and reflective paint-lined ashtrays.

would like to avoid having to run individual light injection fibres to each wavelength shifting fibre, as the fibre cost, and engineering cost to arrange separate light injection for each fibre, would become too much.

We want to inject light into the green fibres, so our only option is to have a fibre running to the light injection module at the end of a scintillator module that illuminates more than one fibre. Recall that the far detector is 8-fold optically multiplexed. In order that the signal in a pixel reflects the response of the pixel spot at which a fibre is positioned, we want to ensure that only one fibre is illuminated per pixel at any one time. This places a constraint on the multiplexing pattern used—any scheme where adjacent pixels are mapped to the same pixel (i.e. to provide a coarse position at one end of the plane and a fine readout at the other) is not permitted, as it would involve 8 adjacent fibres, some of which must surely be illuminated simultaneously, mapping to the same PMT pixel.

We would also like to avoid pulsing adjacent PMT pixels, in order to mitigate the effects of cross-talk. Cross-talk is further discussed in section 6.3.1. We can accomplish

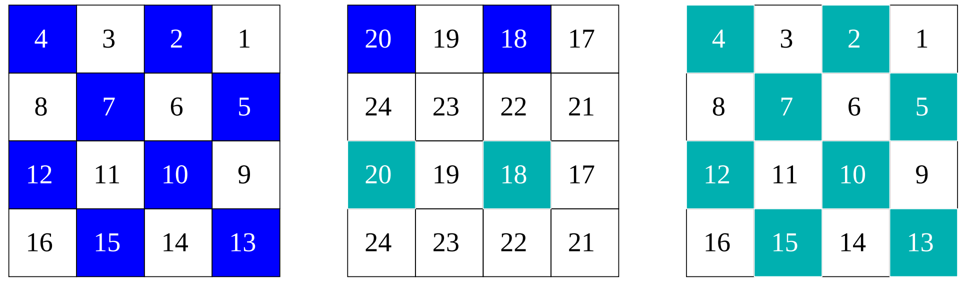

Figure 6.12.: Chessboard illumination pattern for M16s in the far detector. Three PMTs read out two planes. The pixels shown in blue are illuminated by the first ashtray. Numbers are pixel numbers.

this by dividing the light injection module into three compartments, containing 10, 10 and 8 fibres in the widest module. We then arrange the fibres in a chessboard pattern on the face of the M16, as in figure 6.12. The wiring pattern for a far detector plane is shown in figure 6.13.

In the near detector, there is no optical multiplexing, so the wiring is somewhat simpler. We aim to illuminate pixels as far apart as possible on the face of the tube, and produce the mapping in figure 6.14.

6.3.1. Cross-talk

We can identify three different kinds of cross-talk that may be present in the phototubes:

Optical cross-talk Multiple internal reflections inside the glass window of the PMT cause some photons to hit the photocathode on a pixel adjacent to the one that is being illuminated (figure 6.15).

Dynode cross-talk Electrons in the later dynode stages may leak into adjacent pixels.

Capacitive coupling The PMT pixels are read out by a number of traces on the PCB that forms the PMT base. If these traces are close together, we may see signals in dark channels that are induced by a current pulse in a neighbouring trace.

Optical cross-talk is not a huge problem from the point of view of the LED calibration. We are looking for gain changes as a function of light level and time. Optical cross-talk

Strip	Pix. Side A	Side B(0)	B(1)	B(2)	B(3)	B(4)	B(5)	B(6)	B(7)
1+24n	1	1	3	6	8	9	11	14	16
2+24n	3	3	6	8	9	11	14	16	17
3+24n	6	6	8	9	11	14	16	17	19
4+24n	8	8	9	11	14	16	17	19	22
5+24n	9	9	11	14	16	17	19	22	24
6+24n	11	11	14	16	17	19	22	24	2
7+24n	14	14	16	17	19	22	24	2	4
8+24n	16	16	17	19	22	24	2	4	5
9+24n	17	17	19	22	24	2	4	5	7
10+24n	19	19	22	24	2	4	5	7	10
11+24n	22	22	24	2	4	5	7	10	12
12+24n	24	24	2	4	5	7	10	12	13
13+24n	2	2	4	5	7	10	12	13	15
14+24n	4	4	5	7	10	12	13	15	18
15+24n	5	5	7	10	12	13	15	18	20
16+24n	7	7	10	12	13	15	18	20	21
17+24n	10	10	12	13	15	18	20	21	23
18+24n	12	12	13	15	18	20	21	23	1
19+24n	13	13	15	18	20	21	23	1	3
20+24n	15	15	18	20	21	23	1	3	6
21+24n	18	18	20	21	23	1	3	6	8
22+24n	20	20	21	23	1	3	6	8	9
23+24n	21	21	23	1	3	6	8	9	11
24+24n	23	23	1	3	6	8	9	11	14

Figure 6.13.: 8-fold multiplex chessboard wiring pattern for far detector. The mapping is staggered on side B so a hit can be located in the plane by using both ends of the readout. From [48]. Far detector scintillator modules are 24 planes wide.

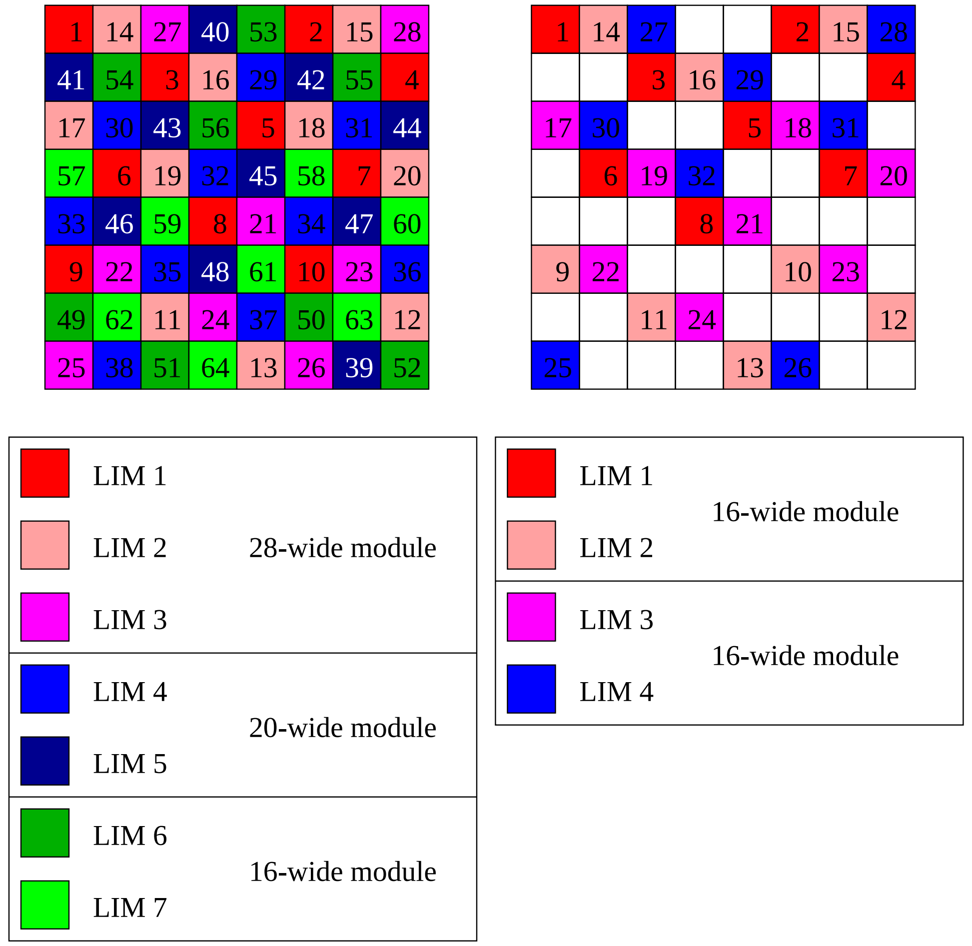

Figure 6.14.: Wiring pattern for M64s. No pair of illuminated pixels are closer than a knight's move from each other. The pattern to the right is for the additional M64 for the fully instrumented planes. Numbers correspond to strip numbers

is a random process that merely changes the effective coupling between PMT pixel and LED, and hence PMT pixel and PIN diode. The light incident on the photocathode is still proportional to the PIN diode signal; the reflectivity of the glass does not change as a function of light level or quickly with time. Large amounts of optical cross-talk in only some channels would be something of an inconvenience, as it would increase the variation in light supplied to a set of pixels by an LED, and so increase the number of pulse heights we need to use to map out a gain curve in every channel.

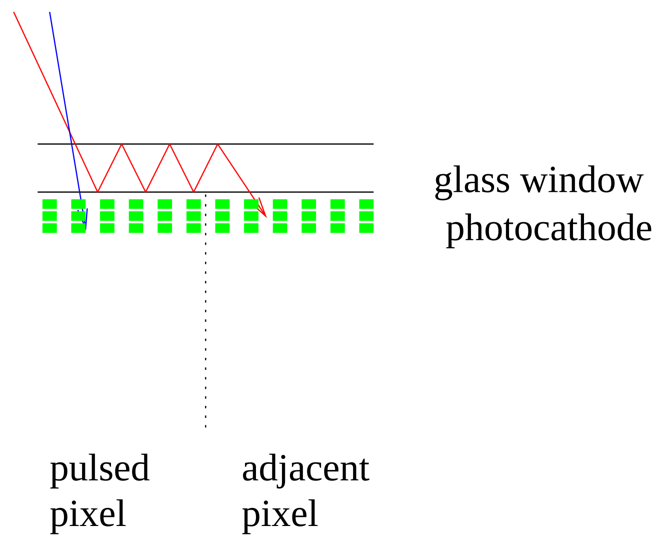

Figure 6.15.: Optical cross-talk

Capacitive coupling must be avoided by careful design of the base—studies at UCL [49] and Oxford [50] have observed capacitive coupling in an M64 with the (small) base supplied by Hamamatsu, but observed no effect with a larger base designed at Oxford for use with M64s in MINOS.

Dynode cross-talk has the potential to cause a miscalibration—we have recently observed indications that the fraction of dynode cross-talk occurring varies with applied light. If this effect is significant, it will cause us to measure a fake nonlinearity caused by dynode cross-talk from the surrounding pixels.

Figure 6.17 shows the variation of fractional cross-talk with light for an M16. We constructed a test rig at Sussex consisting of an M16 cookie (figure 6.16) filled with wavelength shifting fibres running to a patch panel. Single fibre connectors can be plugged in to the patch panel, thus allowing us to choose to pulse the phototube in any pattern we choose.

Most of the cross-talk in this plot is optical cross-talk, there could be an element of dynode cross-talk causing the increasing fraction of cross-talk in some plots. Prior to installation in MINOS, all the PMTs will be scanned, and their cross-talk measured, which should allow us to correct for small effects caused by cross-talk.

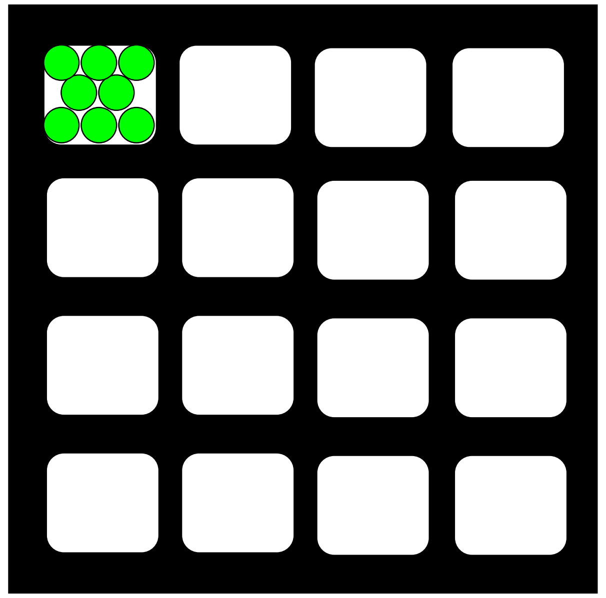

Figure 6.16.: 8-fold multiplex "cookie" for M16, with fibre positions indicated.

6.4. LED Pulser

The LED pulser offers the ability to change both the height and the width of the pulse electronically, by varying the charge pulse that is applied to the LED. It contains no moving parts, so should prove to be robust. Each pulser box can contain 20 LEDs; in the near detector only 11 LEDs are used, so the PCB will not be fully populated. There is an additional "trigger" LED, which pulses at a fixed intensity, and is used to flag the light injection data in the DAQ.

6.4.1. Electronics

The pulser box is controlled by a PIC16F877 microprocessor. It communicates with a control PC via ASCII commands over a serial RS485 link. To avoid a proliferation of different types of data-carrying cable running around the detector, the signals will be carried over the detector ethernet, and converted for the pulser box by an ethernet-RS485 converter. In the standard mode of operation, the PC will instruct the pulser to pulse so many times at a given height and width, and with a given frequency. The

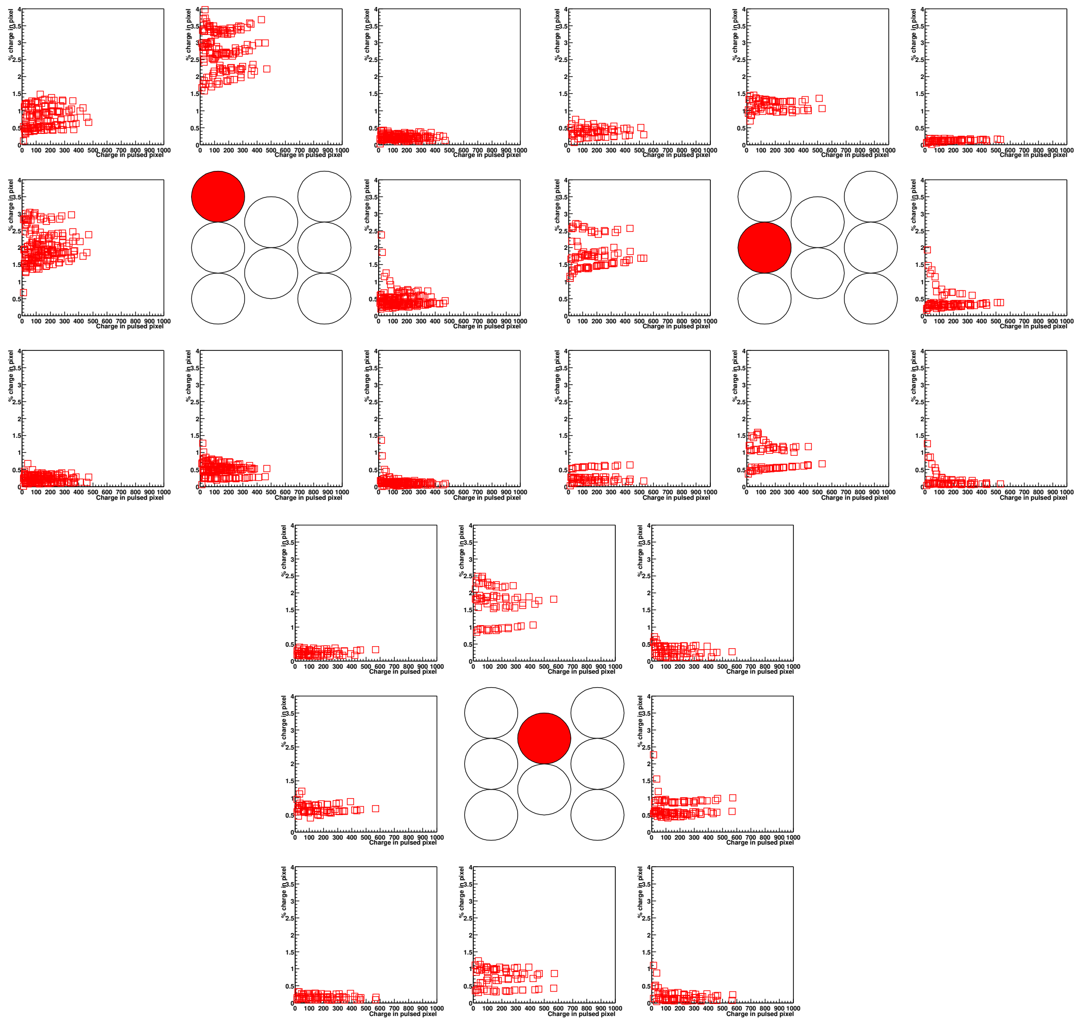

Figure 6.17.: Cross-talk in neighbouring pixels as a function of light level for different pulsed positions within a pixel. Full scale on the y axis is 4% cross-talk in all cases. The x axis is charge in ADC units; full scale is approximately 500 photoelectrons. The fibre spot highlighted in red has been pulsed. Several pixels are plotted on each graph.

microprocessor then controls the actual pulsing. The pulser box is supplied with an external trigger input, which will force a pulse of its current setup in response to a NIM signal. This can be used in the near detector to investigate the effects of the timing of a pulse with respect to the 19 ns time buckets in the QIE electronics.

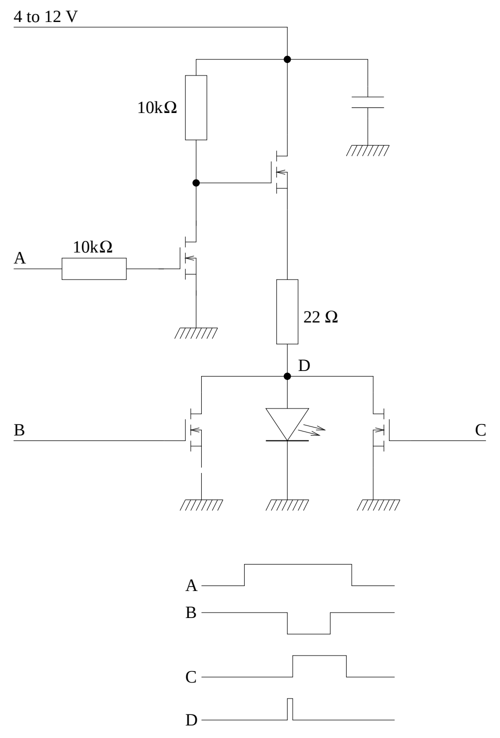

Figure 6.18.: LED pulser drive circuitry.

Figure 6.18 shows the LED driver stage from the pulser box circuit. The applied voltage is controlled by a 10 bit DAC, and controls the pulse height. When the signal at A goes high, the driver stage is turned on—voltage is applied to the 22 Ω resistor, and the current is shunted to ground via the left-hand FET. When the signal at B goes low, current flows through the LED, which emits light. A short time later, signal C goes high, and the current is shunted to ground via the right-hand FET; LED is turned off. The pulse width is determined by the delay between B turning off and C turning on; this is controlled by a programmable delay line, giving a current pulse width of 15 to 36 ns in 3 ns steps. The light pulses delivered to the phototube produce signals ranging from a fraction of a photoelectron to in excess of 300 photoelectrons with the narrowest pulse.

Figure 6.19 shows a comparison of a pulse from the LED pulser and a data pulse due to a 3 GeV pion at the CERN mini-module run. The pulses are seen to be of comparable shape. This is of interest when measuring nonlinearities—space charge effects depend on the instantaneous current, so it is important that the pulse shapes are similar.

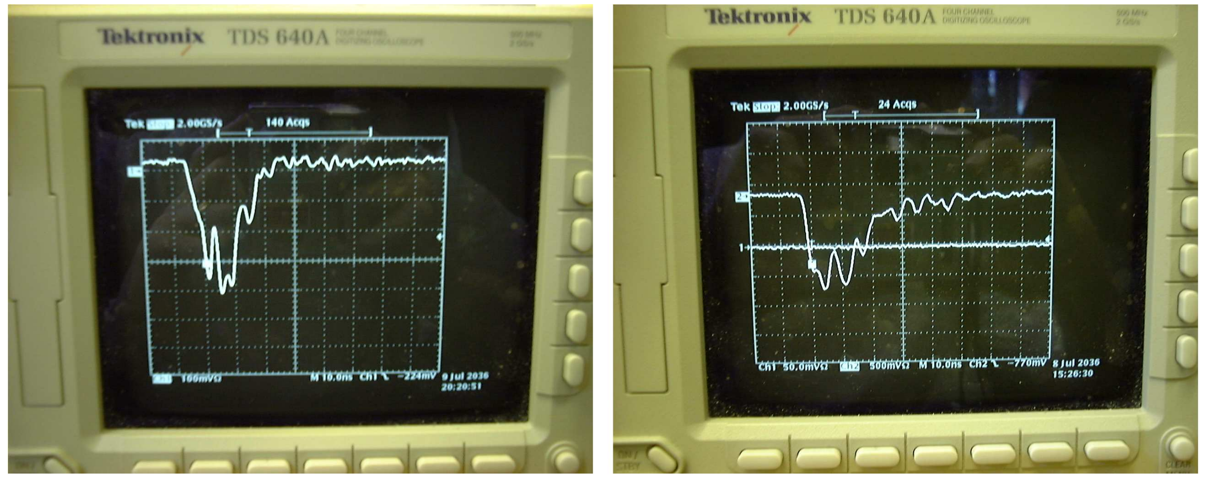

Figure 6.19.: Comparison of pulse from LED (left) and data (right) taken at the mini-module run at CERN. The pulses are of comparable length in time. The scale on the oscilloscope is 10 ns per division (100 ns full scale).

Figure 6.20 shows the short term pulse-to-pulse variation of the LED pulser, as measured over a few minutes by the PIN diode and a digital oscilloscope. We see that 1.1% is an upper limit on the variation of light from the LED. Most of this width is in fact electronics noise (the amplitude was measured with the 'max' function on a digital oscilloscope). This good pulse to pulse stability allows us to monitor the LED statistically, using only the mean of 1000 pulses in the PIN diode to compare to the mean of those 1000 pulses in the PMT channels, without any loss of accuracy.

6.4.2. Cones and collars

Each LED must illuminate up to 64 light injection fibres, plus two for the PIN diodes. The PIN diodes will benefit from as much light as possible—more light implies a better signal to noise ratio in the PIN diode amplifier. We again require the light in our 64

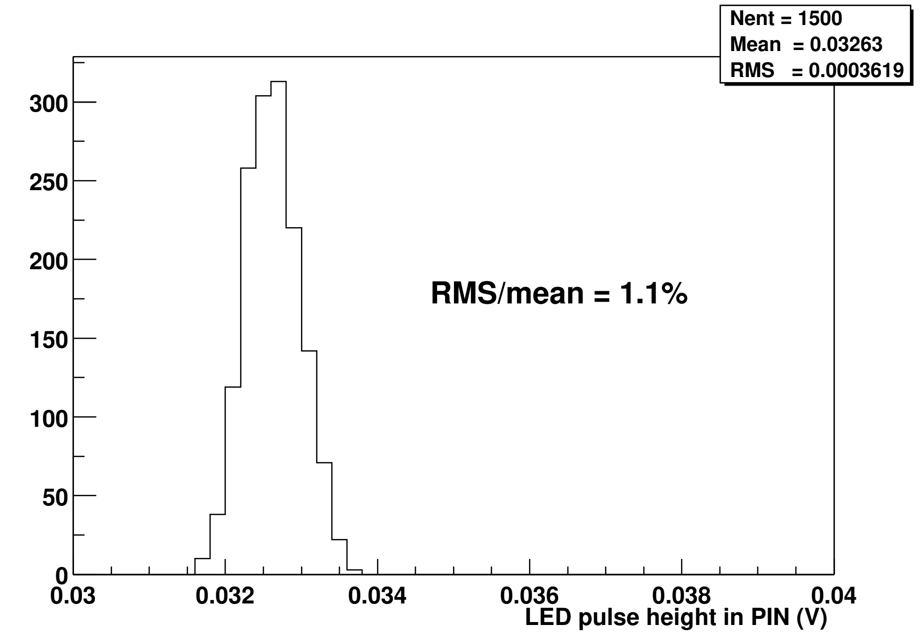

Figure 6.20.: Pulse-to-pulse stability of the LED pulser.

fibres to be as uniform as possible.

The basic design is a machined aluminium cone, lined with aluminised mylar, with a hole at the apex to admit the LED. The base of the cone is of a sufficient size to fit a bundle of 71 fibres (allowing some spares). The fibre bundle is glued in a plastic “collar” for mechanical strength, and to provide for fixing the bundle to the cone. The fibres are placed in a “teastrainer”—a plastic disc with holes for the fibres (see figure 6.21). This allows the assembly of the fibre collar without twisting of the fibre bundle.

Some modifications have been made to improve this basic concept:

- We want as much light as possible to go to the PIN diodes. Three special fibres (including one spare) are designated to serve the PINs, and are placed in the centre of the fibre bundle. To provide more light, these fibres are extended towards the LED, and encased in a rigid plastic tube for support (see figure 6.23).

- In the near field, the pattern of light in the fibre teastrainer is highly structured (figure 6.22). This produces large nonuniformities in the output of the fibres in the

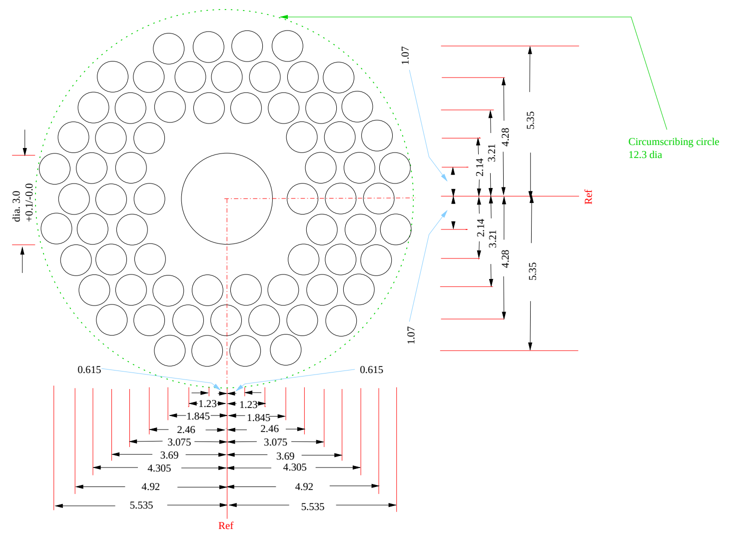

Figure 6.21.: Fibre hole pattern in the "teastrainer" plate. The larger central hole accommodates the tube containing the 3 PIN diode fibres.

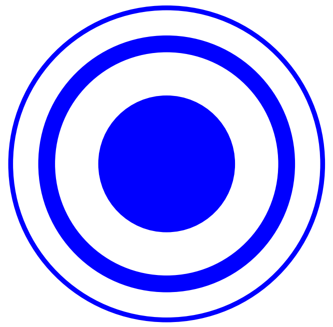

Figure 6.22.: Schematic diagram of structure of light emitted from LEDs.

bundle. To overcome this, the tip of the LED is sandblasted to act as a diffuser. Figures 6.24 and 6.25 show a comparison of the uniformity obtained with and without the sandblasted LED for the same cone and collar assembly. It is clear that the sandblasting is a huge improvement.

A prototype cone and collar assembly is shown in figure 6.26.

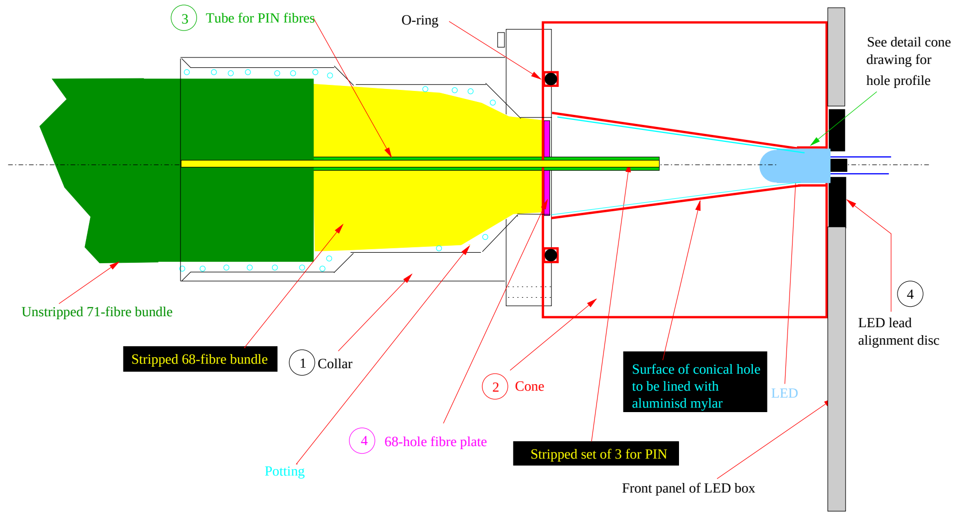

Figure 6.23.: Cone and collar assembly for pulser box.

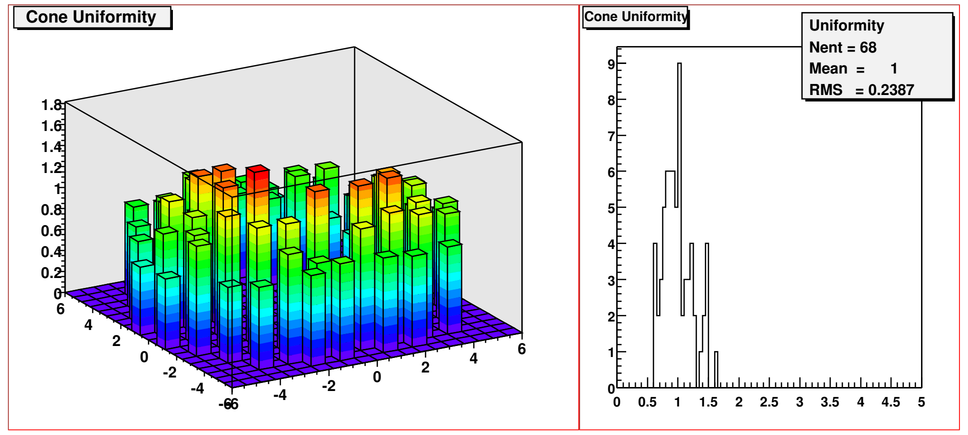

Figure 6.24.: Uniformity of cone and collar with sandblasted LED.

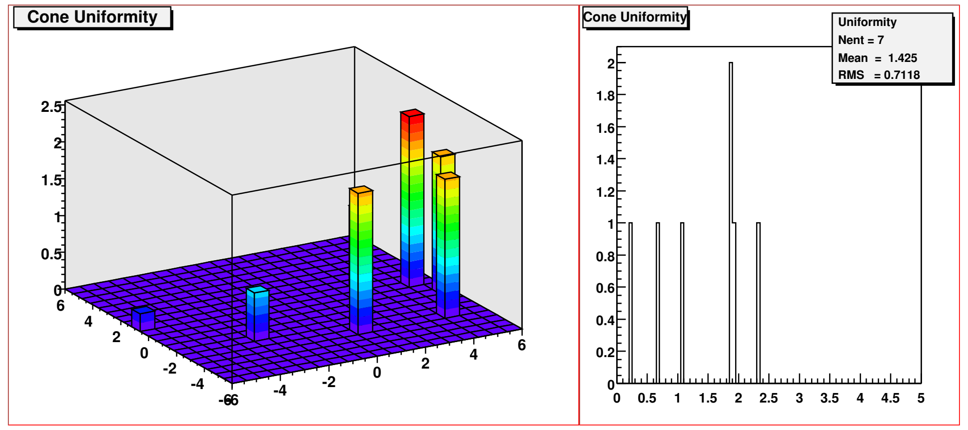

Figure 6.25.: A few channels only, uniformity of cone and collar with clear LED. Normalised to the same scale as figure 6.24.

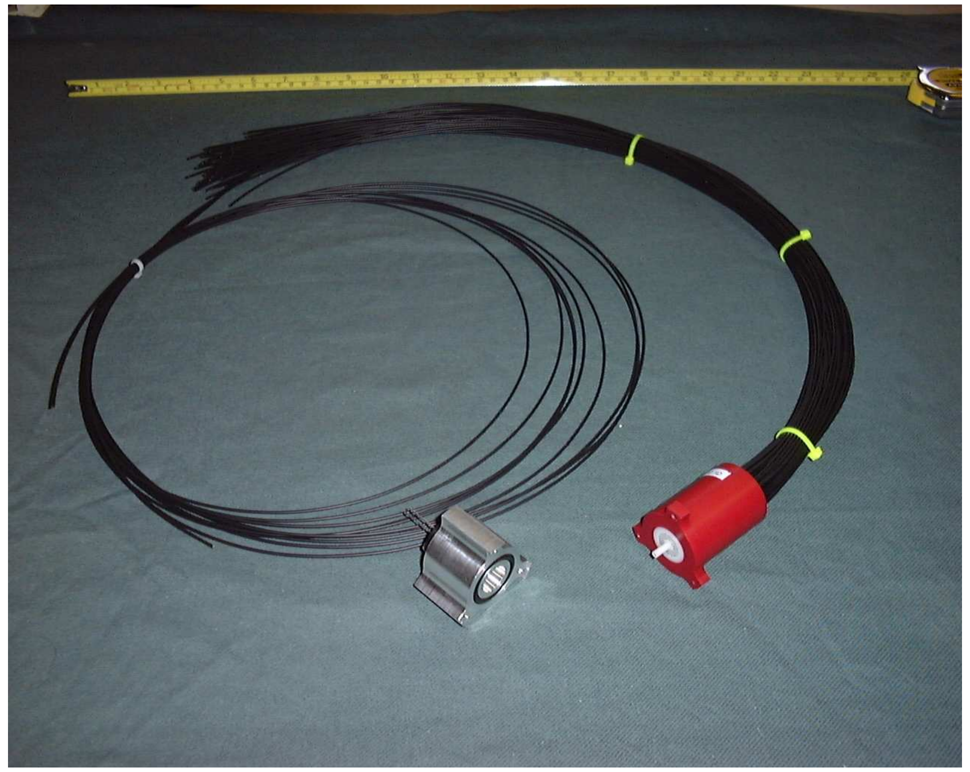

Figure 6.26.: A prototype cone and collar. The legs of the LED can be seen protruding from the rear of the cone.

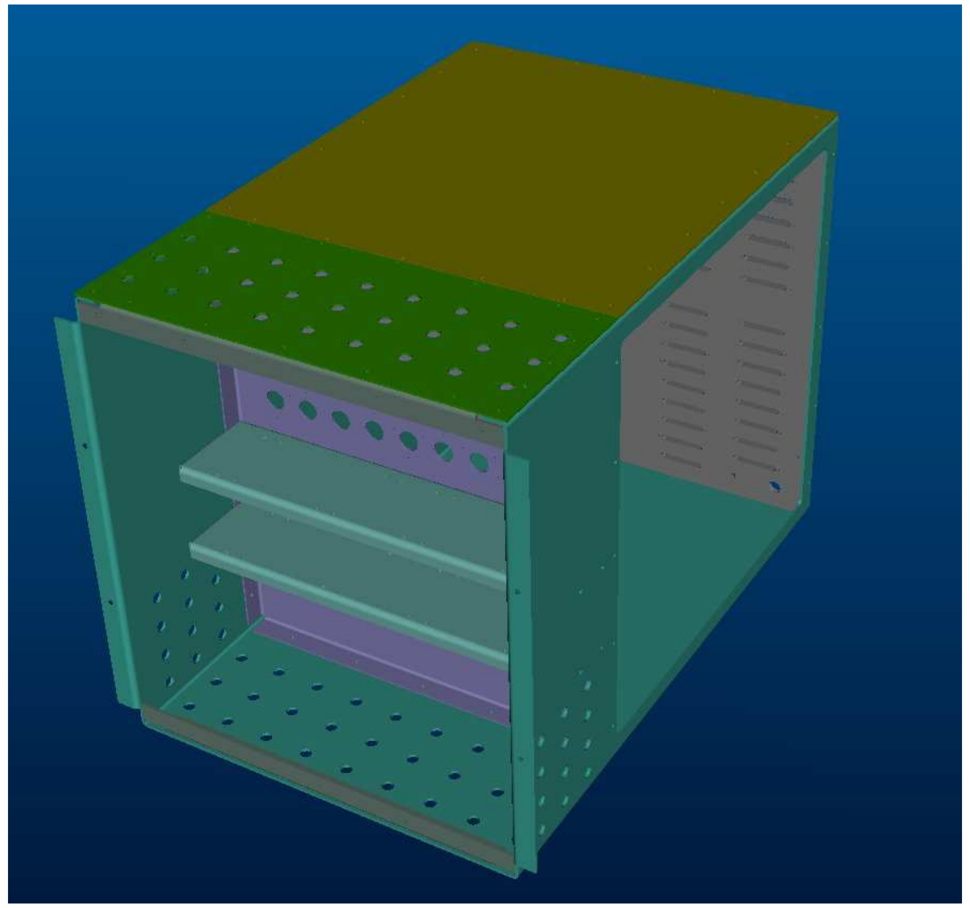

Figure 6.27.: The pulser box. The trays at the front hold the pulser circuit boards. Cones are fixed to the bulkhead in the centre of the box and the LEDs fit in from behind.

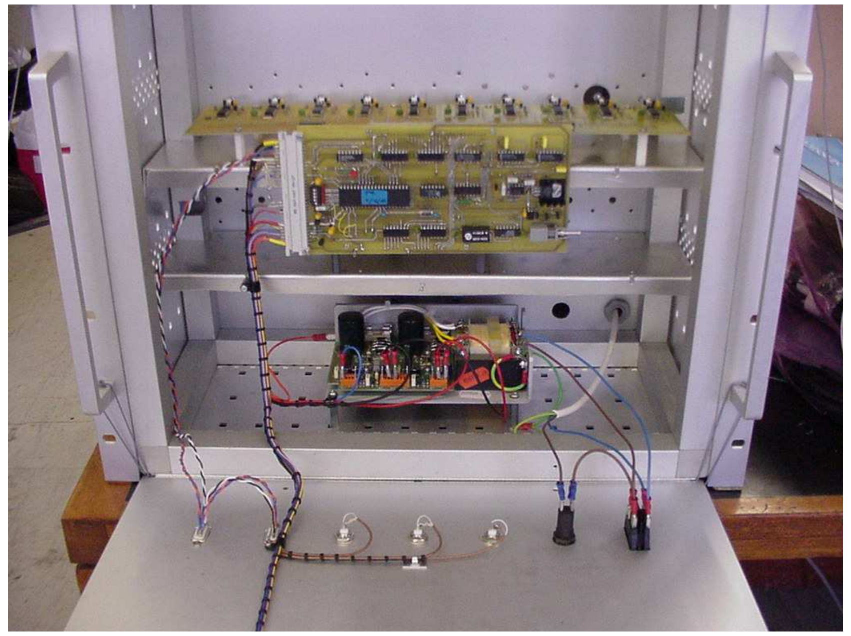

Figure 6.28.: Electronics compartment in a prototype pulser box.

6.4.3. Pulser box construction

The pulser box is 21” high, and fits in a standard 19” rack. The pulser electronics is housed on two trays in a compartment at the front of the box (see figure 6.27). A controller board, containing the microprocessor and RS485 interface connects the two pulser boards. Figure 6.28 shows the controller board and one of the two pulser boards in the prototype pulser box. The LEDs are mounted to overhang the end of the circuit boards, and poke through holes in the bulkhead wall into the rear of the cones (see figure 6.29). Light is distributed to the detector via 4 columns of 16 20-wide connectors on the back panel (figure 6.30) of the box; a fibre from each cone runs to each connector. There are an additional 3 20-wide connectors—two provide light to the PIN diode (and so are fed by the special PIN diode fibres), the third is a spare, containing a fibre from each cone, installed against the possibility that we might want to introduce a radioactive source-monitored PMT as an additional monitor at some future date. There are also a pair of single-fibre connectors carrying fibres to the “trigger PMT”.

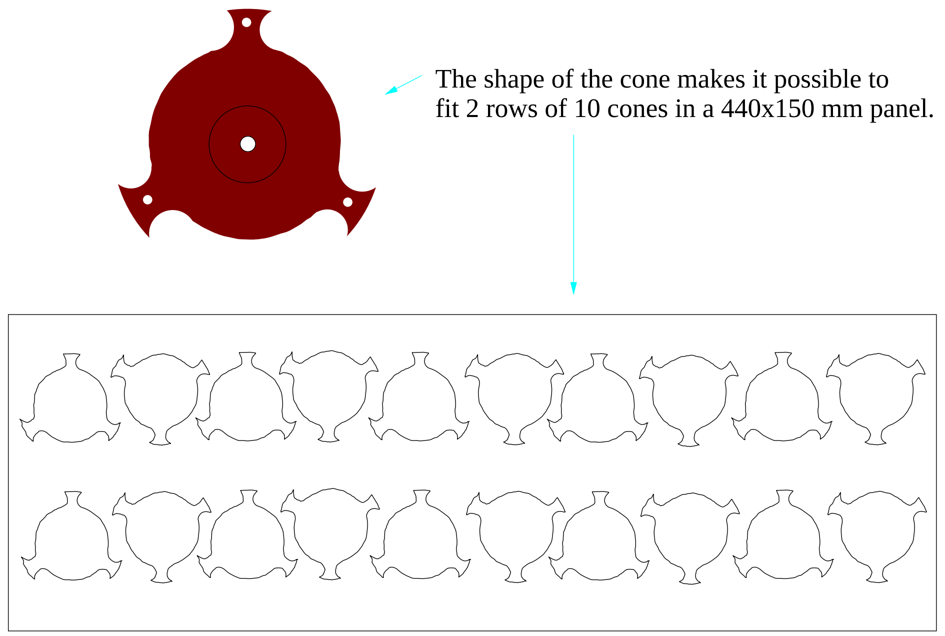

Figure 6.29.: Arrangement of cones on the bulkhead in the pulser box.

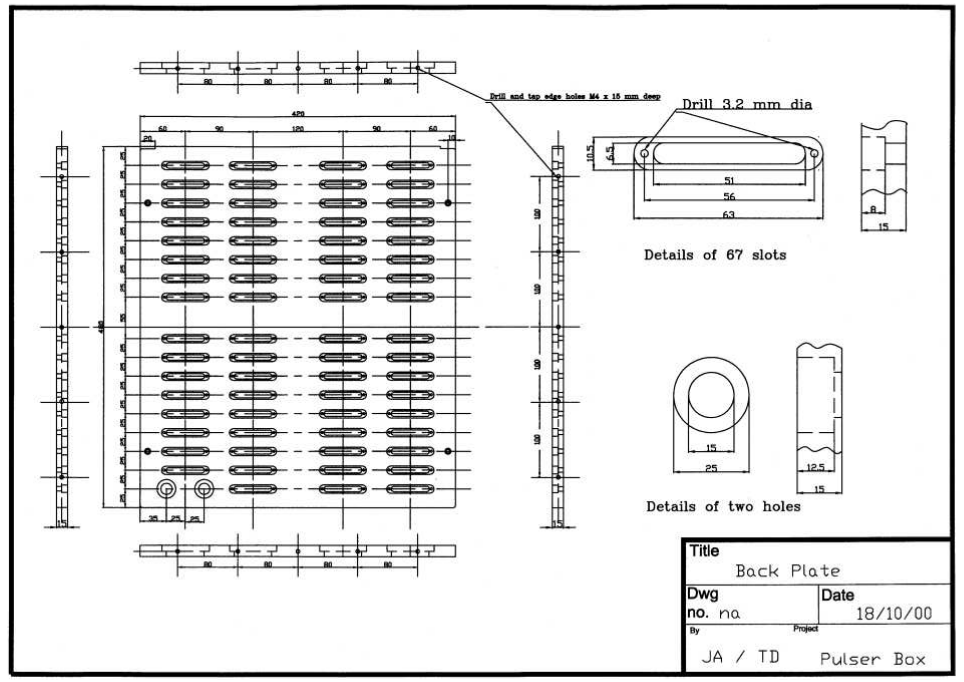

Figure 6.30.: Back of pulser box, showing slots for connectors. A fibre from each cone runs to each 20-wide connector.

6.5. PIN diode

The PIN diode used is the Hamamatsu S5971. This is chosen primarily as it offers a fast response (cutoff frequency 100 Hz) whilst having a large enough surface area to accommodate our 1.2 mm diameter fibres.

Active area size	1.2 mm diameter
Spectral response	320 nm to 1020 nm
Dark current	0.07 nA typ.
Temperature coefficient of dark current	15%/^{o}C
Photosensitivity @ 470 nm	0.24 A/W
Temperature coefficient of photosensitivity	$<$ 0.01%/^{o}C at lab temperatures

Table 6.2.: PIN diode specifications

The PIN diode is amplified by a two-stage amplifier, which amplifies and shapes the current pulse produced by the PIN. The shaping time of the output stage is matched to the far detector readout VA chip which is used to read out the PIN.

The VA chip requires a PMT dynode signal to trigger its readout. In the far detector, this will be accomplished by mounting the PIN diodes on the front end readout cards (VFBs). Two PIN diodes will be mounted on each front end card, the wiring diagram in table 6.3 ensures that the PMT that "owns" the PIN will always be illuminated when the PIN should be read out. It also ensures that the two PIN diodes that serve each LED are on different front end cards; this allows one front end card to be replaced (to change a photomultiplier or to replace an amplifier) whilst the LED calibration is held by the other PIN diode, which remains undisturbed. In the near detector, the phototubes are

Upper Walkway		Plane	1,3	5,7	9,11	13,15	17,19	21,23	25,27	29,31	33,35	39,39	41 - 63
	LED	High Gain PIN	1	3	5	7	9	11	13	15	17	19	PINs unused
		Low Gain PIN	2	4	6	8	10	12	14	16	18	20	PINs unused
Lower Walkway		Plane	2,4	6,8	10,12	14,16	18,20	22,24	26,28	30,32	34,36	38,40	42 - 64
	LED	High Gain PIN	2	4	6	8	10	12	14	16	18	20	PINs unused
		Low Gain PIN	1	3	5	7	9	11	13	15	17	19	PINs unused

Table 6.3.: PIN diode to LED mapping in the far detector. Planes are numbered within the 64-plane block owned by a mid-level VME readout rack, LEDs are numbered within the pulser box in that rack. The pattern is duplicated on each side of the detector.

read out by the QIE electronics, so cannot be used to trigger the PIN readout. Instead, we mount all the PIN diodes in the same rack, and use the dynode signal from the trigger PMT as a trigger.

The PIN diode is the lynch-pin of the LED calibration, so it is important that its properties are understood.

6.5.1. Linearity

The PIN diode is used to provide a scale which varies linearly with applied light, in order to calibrate the nonlinearities of the PMT. It is therefore important that the PIN itself is linear. The linearity of our PIN diode/amplifier system was measured by comparing the signal due to a pulsed LED with the signal due to the same pulsed LED, but with a neutral density filter that blocked 90% of the light. The PIN diode was read out by a digital oscilloscope in peak-finding mode. Repeating this measurement for several different pulse heights produces the linear response in figure 6.31. We also require the

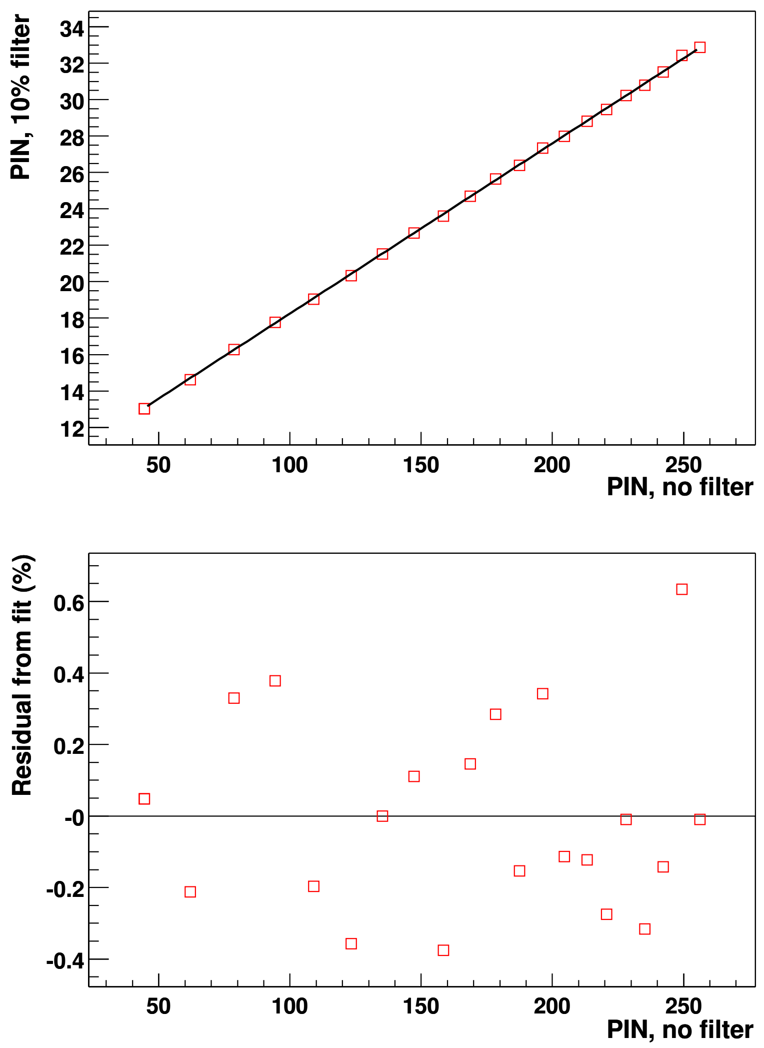

Figure 6.31.: PIN diode linearity. The top plot shows the signal in the PIN due to light from the LED pulser passing through a 10% neutral density filter, plotted against the signal with no filter, for several different pulse heights. The bottom plot shows the residuals from the straight line fit to the top plot.

ADC conversion to be linear. The charge injection system will calibrate the PIN readout channels to an accuracy of 1%, thus ensuring that the PIN scale is indeed proportional to the light input.

6.5.2. Stability

The light injection system is used to track gain changes over a short timescale ($\sim$ 1 month). It is therefore important that the system itself is stable over this timescale. Figure 6.33 shows the result of a test carried out in the lab over a period of 2 months. An LED pulser was set pulsing at a fixed height, and monitored several times a second by a pair of PIN diodes, which were read out by a digital oscilloscope in peak-finding mode (figure 6.32).

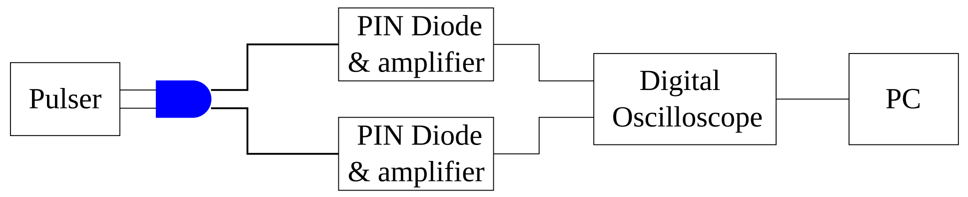

Figure 6.32.: Schematic diagram of the apparatus used to measure the PIN diode stability

We observe that, although the light output of the LED changed (probably due to instabilities in the LED pulser power supply), the ratio between the individual PIN readings remained constant to about 0.5% over a period of two months. Each point in figure 6.33 is an average of 1000 consecutive PIN readings.

Some features of this plot are worth pointing out: over the first 10 days, the signal from PIN 1 gradually decreased; this is thought to be caused by the fibre settling slightly in the LED connector. The contribution from this first 10 days is filled in red in the histogram in figure 6.33. The feature at 46 days occurred when the apparatus was accidentally disturbed—we see again a sharp deviation in one PIN's signal, followed by a slow return to roughly the previous level.

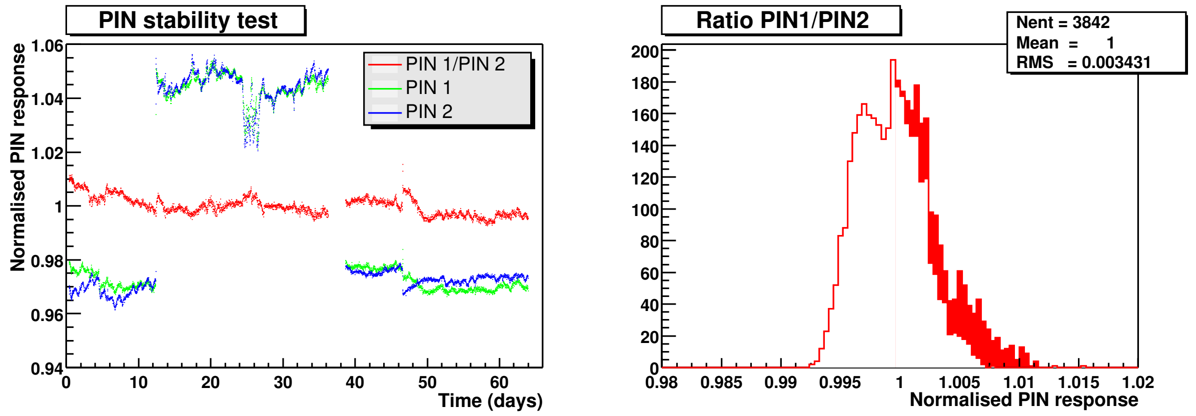

Figure 6.33.: Test of pin diode stability. The first 10 days' contribution is filled red in the histogram. The PIN is stable to better than 0.5% over two months.

6.6. Data acquisition & offline software

6.6.1. Volume of data required for calibration

Signals of interest in MINOS range from the photoelectron level up to about 300 photoelectrons. Assuming poisson statistics dominates the spread of the pulser data, the fractional error on the mean is given by

$$\frac{\sigma_\mu}{\mu} = \frac{1}{\sqrt{Nn}}, \tag{6.1}$$

where n is the expected number of photoelectrons, and N is the number of pulses. To achieve 1% accuracy, we therefore require $Nn = 10000$, and so need 34 pulses at the 300 photoelectron level, and 1000 pulses at the 10 photoelectron level. The logical extension of this requires taking 10,000 pulses at 1 photoelectron, however the PMT is expected to be highly linear at this low light level, so one loses nothing by extrapolating linearly from 10 photoelectrons down to the well-measured pedestal at zero. Ten different pulse heights should be sufficient to map out the nonlinearity of the PMT.

Our hourly point will be taken at a level of around 50 photoelectrons. In principle, 200 flashes are sufficient to establish the mean to 1%; in practice, we would like to take 1000, to allow for the variation in the light supplied to each channel and a safety margin.

6.6.2. Data rates

Far Detector

Consider first the amount of data produced by one flash of an LED. A pulser box serves up to 64 planes (32 of each orientation). In each plane, up to 10 fibres (one ashtray) are illuminated by one LED. Each flash therefore illuminates up to 640 channels in each of two front end crates. Each hit requires 8 bytes of data in the front end, so we generate 5 kB of data in each of two FE crates per flash. We have ignored the effects of cross-talk here—if there is sufficient cross-talk in a channel to raise it above the 0.3 photoelectron threshold, it will be read out. Potentially all 24 PMT pixels per plane could be read out, producing 12 kB per flash per FE crate.

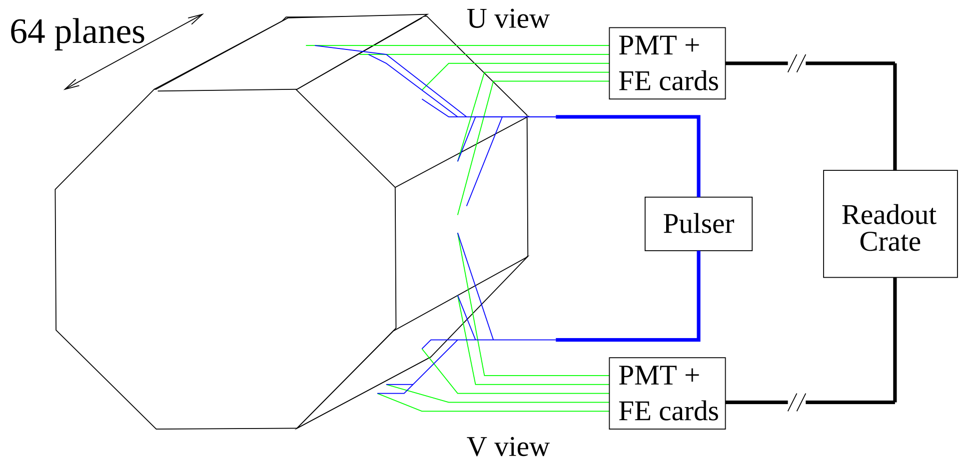

Figure 6.34.: Schematic diagram of one block of the far detector. The readout and pulser shown are mirrored on the other side.

Each front end crate can handle an instantaneous rate of 2.5 MB/s. We should aim to stay well below this level. Pulsing at 100 Hz will produce a maximum data rate of 1.2 MB/s per FE crate, which is as large as one would want to go.

To accumulate 1000 pulses for our hourly point would then require 10 s per LED. There are 16 pulser boxes in the far detector, each with 20 LEDs. To pulse all of them every hour at this rate would take 53 minutes. This would involve pulsing each of the 8 fibres on a pixel separately. We expect that any gain changes that occur on an hourly timescale will occur across a whole pixel; three LEDs are sufficient to pulse each pixel once, which would take 8 minutes per hour.

To map out the linearity, we want to pulse at 10 different heights, and need to illuminate all fibres on a pixel individually. This would take a total of 530 minutes at the above rates, which could be achieved by pulsing for 22 s per minute for 24 hours.

Near detector

There are two 11-LED pulser boxes in the calorimeter section, each illuminating one orientation of plane, and one 11-LED box in the spectrometer section. Each LED illuminates up to 10 strips in up to 60 planes in the calorimeter section, and up to 10 strips in all 33 planes of the spectrometer section. Digitised data is buffered in the readout crates, which is where the bottleneck occurs; there are 2 readout crates for the spectrometer section and 6 for the calorimeter section. The bottleneck is therefore a spectrometer readout crate, which receives up to 170 channels of data.

Large signal pulses may have tails of up to 50 ns caused by the time constant of the wavelength shifting fibre. Light will be reflected by the reflector at the far end of the scintillator module, adding an extra 40 ns to the length of the event. Each hit may then occupy 5 19 ns QIE time buckets. Each of these buckets is digitised to produce 8 bytes of data, so we generate 6.8 kB of data in one readout crate per flash. A limit of 1 MB/s per crate allows us to flash at up to 140 Hz, so completing the hourly point for the detector in 4 minutes. Cross-talk in the M64 may cause a signal in all channels in the worst case; this multiplies the data rate by 6.4, and so the hourly point would take 25 minutes.

The gain curve could be mapped out over 24 hours by pulsing for 1.7 s per minute if no crosstalk is important, or for 10 s per minute if it is.

6.6.3. Deadtime issues

A far detector channel goes dead for 5 μs after a hit. If we flash every LED 1000 times per hour (our worst case 53 minutes per hour scenario above) we have a total deadtime of 1.6 s per hour. This figure is negligible, and even overestimates the problem slightly, as it assumes that the whole detector is dead after a pulse, when in reality only those channels that have been flashed (which are localised in space) are dead.

In the near detector, the high-speed QIE electronics does not suffer from deadtime, as long as we don't fill all its buffers. As a worst case, a flash might last 200 ns in the whole detector (reflection from far end of scintillator adds 40 ns, light injection fibres are different lengths which adds more spread). We flash 33 LEDs 1000 times per hour, giving a total length of time for which there is pulser data in the detector of 6.6 ms per hour. This too is negligible.

The calibration detector is a small version of either the near or far detector, depending on which electronics is used, but without multiplexing. It is clear that no scheme which is possible in the "real" detectors can cause a problem for the calibration detector.

6.6.4. Online averaging

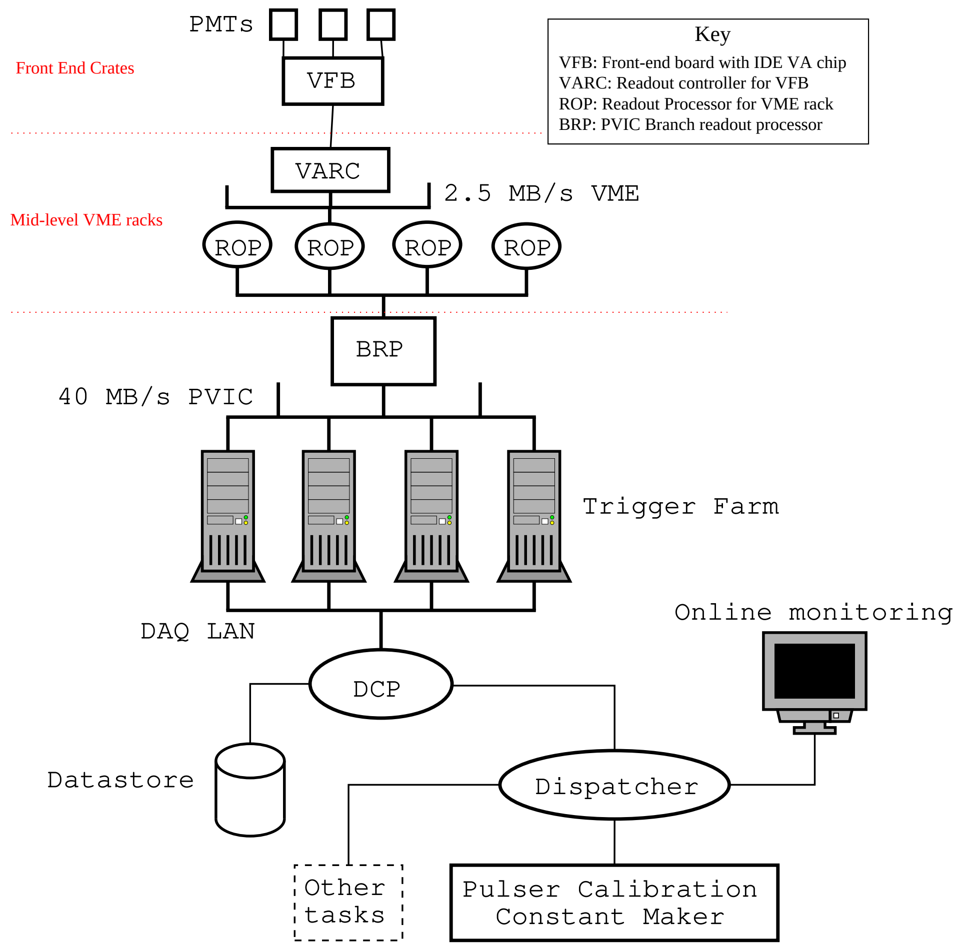

Figure 6.35.: Far detector DAQ. The near and calibration detectors have similar architectures.

Whilst these data rates are manageable in the front end, the amount of data involved (in the region of 200 GB/day in the far detector) is too large to be handled by the back-end, without needless extra expense and effort. We therefore perform preliminary averaging online, in the trigger farm.

The MINOS DAQ (figure 6.35) divides the data into 1 s timeframes, and each timeframe is passed to a different trigger processor. The trigger processors will recognise an LED pulser event by the presence of a hit in the trigger PMT, and pass the event to the online averaging process. This tracks the mean and mean square of the LED signals—the running average of n data is given by

$$\overline{x_n} = \overline{x_{n-1}} + \frac{x - \overline{x_{n-1}}}{n} \tag{6.2}$$

At our maximum pulsing rate of 100 Hz, each timeframe then contains up to 100 pulses for each channel served by a given LED. The averages of x and x^2 for each channel are written out in the datastream with the number of pulses n, and header information obtained from the pulser control PC indicating which LED was flashing, the total number of pulses expected and so on.

6.6.5. Offline calibration software

The offline calibration software can be divided into two distinct tasks—the constant maker, which must determine calibration constants from the partial averages produced by the online code and write them to the database, and the calibrator, which runs at the start of a reconstruction job, and uses the constants in the database to calibrate the data. The latter is essentially a trivial task, once the constants have been determined.

The pulser calibration constant maker task runs as a quasi-online task as a client of the data dispatcher (see figure 6.35). This enables it to provide real-time feedback on the performance of the light injection system. It receives in the datastream the encoded partial averages of PMT and PIN signals produced in the trigger farm, and must provide a calibration to convert linearised charge to time-independent PIN-based units.

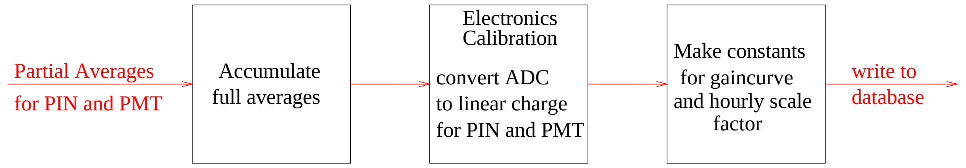

Figure 6.36.: Data flow in pulser calibration constant maker task.

The partial summaries will arrive at a rate of about one summary block per second; each summary will contain 50-100 pulses in all the channels served by a given LED. We must first accumulate a weighted average of the sub-averages for each channel. When all the data for a pulse point has been received, the average for each channel is passed on to the electronics calibrator to convert ADC counts to charge. The next stage depends on whether a gaincurve or hourly point is being taken.

Gaincurve

For each channel, we accumulate 10 different pulse heights, plus the corresponding signals in the two pin diodes—one higher gain, one lower gain. This allows us to plot channel charge against PIN diode charge for each PIN diode. It is likely that at the top of the range the high-gain PIN will overflow the ADC, and at the bottom of the range the low-gain PIN will be degraded by electronic noise. We define a pulser unit for each channel such that

$$\text{Light (Pulser Units)} = \alpha p_1 = \beta p_2 \tag{6.3}$$

where p_1 and p_2 are the charge in the first and second PIN diodes respectively. A least squares fit of the compatible data allows us to determine the ratio α/β, and we arbitrarily chose to set the scale of the pulser unit so that $\alpha = 1$ i.e. a pulser unit is equal to a PIN charge unit for the first PIN diode.

This use of two PIN diodes and the pulser unit allows us to easily hold the light injection calibration over the exchange of a PIN diode. In the far detector, the PIN diodes are located on the front end boards which also contain the PMT readout. Should a front end board, and thus its PIN diodes, require replacing, the other PIN can be used to hold the calibration—when the new board is fitted, a gaincurve may be taken, the constant relating the remaining PIN to the pulser unit is held constant, and the least squares fit determines to constant for the new PIN.

One now has a set of points relating charge to pulser units; the simplest way to use these to convert a data hit to pulser units is via interpolation. It may be superior to fit a straight line to the lower part of the graph, where the PMT is expected to be linear, and then a spline in the non-linear region. This will be investigated at the calibration detector run at CERN, where we will be able to see what the light injection data looks like in real MINOS electronics.

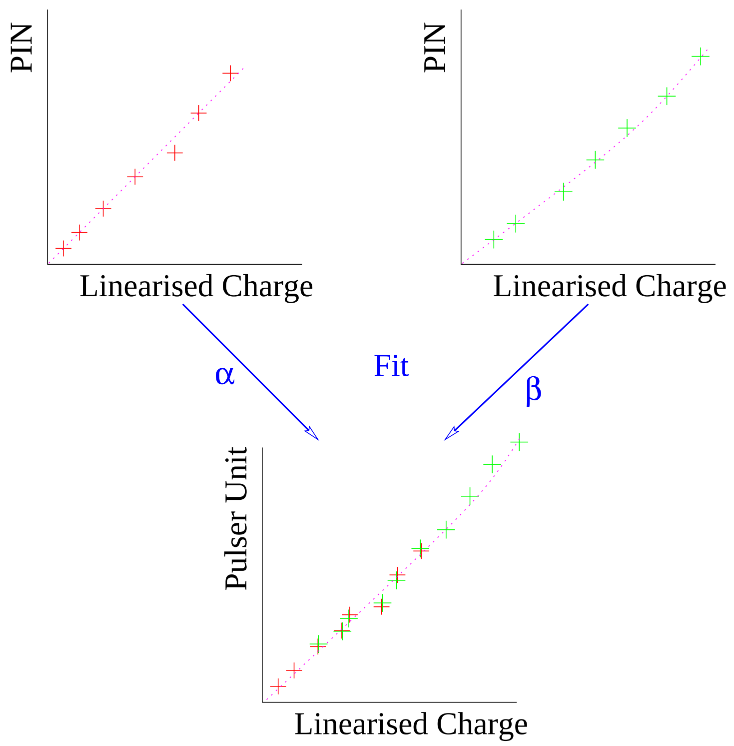

Figure 6.37.: The Pulser Unit

Hourly point

This is completely trivial—the calibration code holds the most recent gaincurve. When an hourly point with charge c arrives, we look up the corresponding pulser unit p_0 with the gain curve. We use the stored values of α and β with the PIN diode readings to give measurements of the current pulser unit p_a and p_b, and write the average of the ratios p_a/p_0 and p_b/p_0 to the database as this hour's scale factor. The code will check that the two PIN diodes are compatible, and will flag a warning if they are not. A step change in one PIN measurement would indicate that someone had disturbed a connection, and no further action would be necessary. A gradual change of one reading or a complete failure would indicate a fault, probably a broken amplifier, which would require manual intervention to fix.

7. Results from prototype

On two occasions I have been asked, "Pray, Mr. Babbage, if you put into the machine wrong figures, will the right answers come out?" I am not able rightly to apprehend the kind of confusion of ideas that could provoke such a question.

— Charles Babbage

We have constructed a small calorimeter of similar design to the MINOS detectors, and exposed it in the T11 test beam at CERN for a period of two months. This chapter presents an analysis of the performance of the light injection system using results from this prototype.

7.1. Experimental setup

The calorimeter (the "mini-module") consists of 16 planes of 15mm thick steel, measuring approximately 18 cm × 18 cm, interleaved with planes of 4 strips of MINOS scintillator (recall from chapter 4 this is 4.1 cm wide and 1 cm thick.) The scintillator used was from an early development batch made by Quick Plastics, and had a somewhat lower light yield than the production scintillator that will be used for MINOS.

The scintillator was read out at both ends using 1.0 mm Kuraray wavelength-shifting fibre; the fibres passed through early prototype light injection modules which were illuminated by an LED pulser. All channels were illuminated simultaneously when an LED was pulsed. One side of the module was read out with 4 M16 PMTs, the other with an M64. The whole calorimeter was encased in a light-tight aluminium box. The inside of

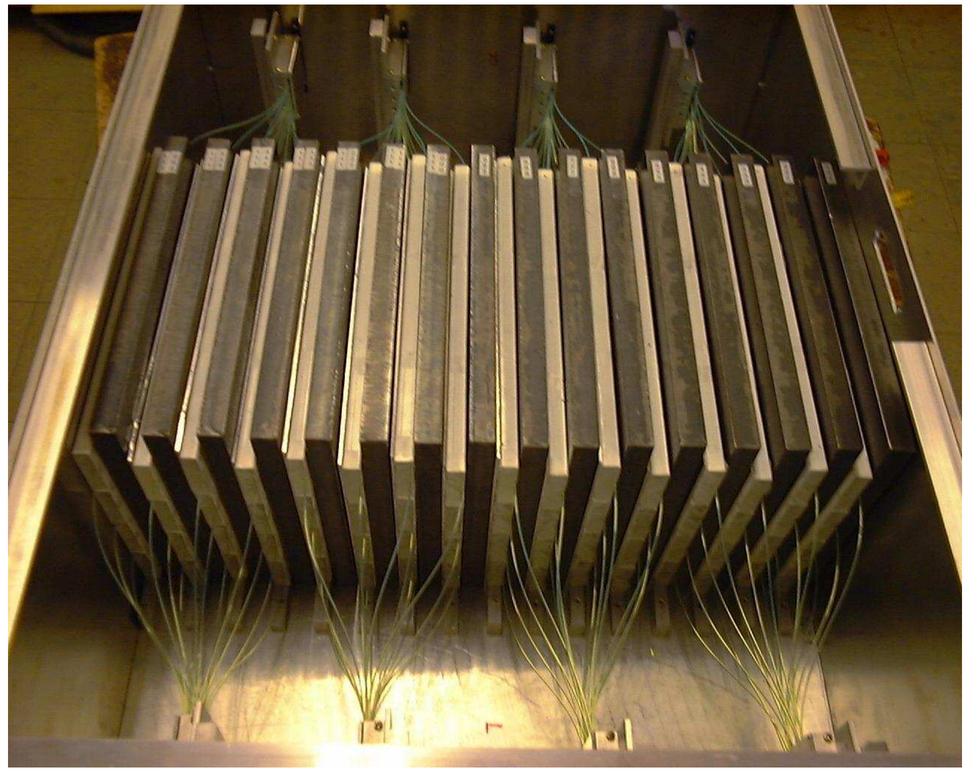

Figure 7.1.: The mini-module.

the mini-module is shown in figure 7.1.

7.1.1. Electronics and DAQ

The PMTs used bases made at the University of Texas for M16s in MINOS; these had a resistor chain for the dynodes in the ratio 2:2:3:1:1:1:1:1:1:1:2:5. This base produces good linearity and a gain of about 10^6 in the M16 tubes, when run at about 800 V. In the M64, space charge effects are much worse, due to the smaller pixel size, so nonlinearities appear at lower light levels.

The PMTs were read out via Phillips Scientific 779 pre-amplifiers and LeCroy 2249A charge-integrating ADCs. These ADCs integrate charge within a supplied gate; the system was triggered by overlapping scintillator paddles[1] placed in front of the module, and a gate of length 100 ns was generated by a CAEN timing unit.

[1] These originally were part of the OPAL forward detector.

The ADCs were read out by a PC running Linux, via a Jorway 73A SCSI CAMAC controller. The system had a maximum readout rate of about 600 Hz.

7.2. Understanding the PMT spectrum

This section specifically describes the Hamamatsu R5600-M16 and -M64 photomultipliers that will be used in MINOS, but with the exception of minor details will apply to any PMT.

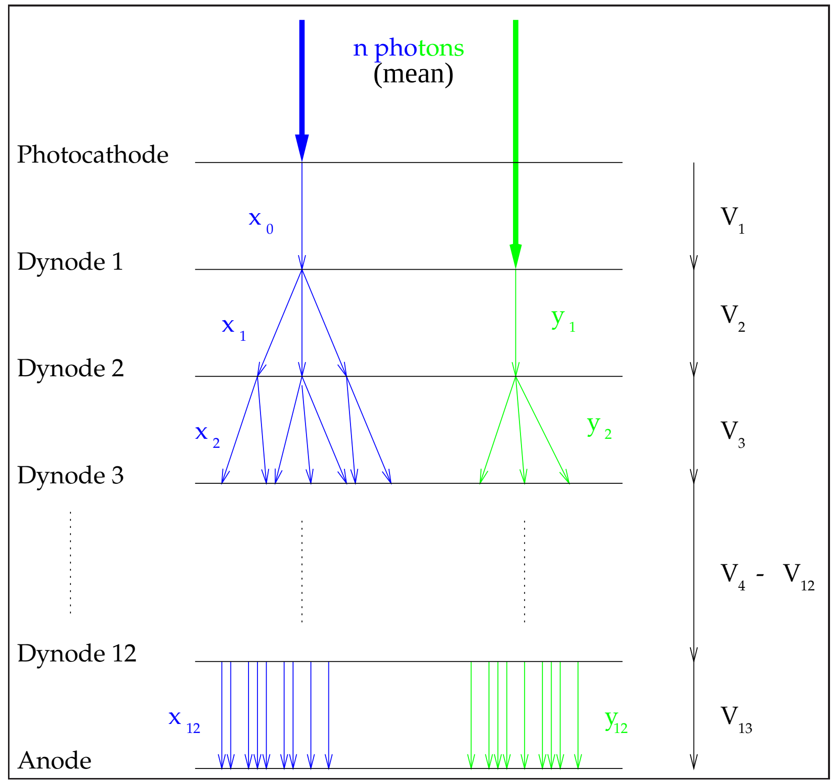

Figure 7.2.: Schematic diagram of a photomultiplier

Consider first the blue arrows in figure 7.2. This is the basic action of a photomultiplier. Our LED pulser produces a large number of photons which, given the random nature of light emission, will be distributed as a poisson distribution. These photons travel through several optical connectors, fanouts and the green fibre on their journey to the PMT pixel; at each stage, many of them are lost. We see that the number of photons arriving at the photocathode is given by a poisson distribution:

$$p(\gamma|n) = \frac{n^\gamma}{\gamma!}e^{-n} \tag{7.1}$$

Each photon has some probability (the quantum efficiency) to produce an electron. This electron then has a further probability (the collection efficiency) to hit the first dynode, rather than being absorbed by the focussing electrodes or part of the mechanical support structure. If we write q for the product of these numbers, then γ photons hit the photocathode, and we see that the number of electrons arriving at the first dynode x_0 is then:

$$p(x_0|n) = \frac{\gamma!}{x_0!(\gamma - x_0)!}q^{x_0}(1-q)^{\gamma - x_0}p(\gamma|n) \tag{7.2}$$

$$= \frac{(qn)^{x_0}}{x_0!}e^{-qn}. \tag{7.3}$$

Our x_0 electrons have been accelerated towards the dynode by the voltage V_1. These electrons collide with the first dynode, and knock out several other electrons. The number of electrons produced by each incident electron is the secondary emission ratio, and depends on the energy of the incoming electron, and hence V_1. Writing $s(V)$ for the secondary emission, we see that, approximately,

$$p(x_1|x_0) = \frac{(x_0 s(V_1))^{x_1}}{x_1!}e^{-x_0 s(V_1)}. \tag{7.4}$$

At each further dynode, we have more amplification, and another poisson distribution until we reach the electrons emitted from the final dynode:

$$p(x_{12}|x_{11}) = \frac{(x_{11} s(V_{12}))^{x_{12}}}{x_{12}!}e^{-x_{11} s(V_{12})}. \tag{7.5}$$

The final stage voltage V_{13} serves to attract the electrons to the anode. It has no multiplication effect.

We may approximate the distributions $p(x_1)\ldots p(x_{12})$ by Gaussians:

$$p(x_i|x_{i-1}) = \frac{1}{\sigma_i\sqrt{2\pi}}e^{\frac{(x_i - x_{i-1}s(V_i))^2}{2\sigma_i^2}} \tag{7.6}$$

with $\sigma_i = \sqrt{x_{i-1} s(V_i)} = \sqrt{x_0 \prod_{j=1}^{i} s(V_j)}$. We then observe that, convolving the gaussians,

$$p(x_{12}|x_0) = \frac{1}{\sigma\sqrt{2\pi}} e^{\frac{(x_{12}-Gx_0)^2}{2\sigma^2}} \tag{7.7}$$

with $\sigma^2 = G^2 x_0 \left(1 + \sum_{i=1}^{12} \frac{1}{\prod_{j=1}^{i} s(V_j)}\right)$ and $G = \prod_{i=1}^{12} s(V_i)$. So then the probability distribution for electrons at the anode is

$$p(x|n) = \sum_{x_0=0}^{\infty} \frac{(qn)^{x_0}}{x_0!} e^{-qn} \frac{1}{G\sqrt{2\pi x_0 (1 + \frac{1}{s(V_1)} + \frac{1}{s(V_1 V_2)} + \cdots)}} e^{\frac{(x-Gx_0)^2}{2G^2 x_0 (1+\frac{1}{s(V_1)}+\frac{1}{s(V_1 V_2)}+\cdots)}}. \tag{7.8}$$

Now consider the green path from figure 7.2. The photocathode is semi-transparent, so some photons may pass through it, and hit the first dynode. Although it is less efficient, the photoelectric effect may also occur in the first dynode, we generate an additional number of electrons y_1 arriving at the second dynode:

$$p(y_1|n) = \frac{\gamma!}{y_1!(\gamma - y_1)!} q_{\prime}^{y_1} (1 - q_{\prime})^{\gamma - y_1} p(\gamma|n) \tag{7.9}$$

$$= \frac{(q_{\prime} n)^{y_1}}{y_1!} e^{-q_{\prime} n} \tag{7.10}$$

with $q_{\prime}$ the efficiency for this process, due to the transmission of the photocathode, the quantum efficiency of photoelectric emission at the first dynode and the collection efficiency of the second dynode. Following the previous argument again, we obtain

$$p(y|n) = \sum_{y_1=0}^{\infty} \frac{(q_{\prime} n)^{y_1}}{y_1!} e^{-q_{\prime} n} \frac{1}{G'\sqrt{2\pi y_1 \left(1 + \frac{1}{s(V_2)} + \cdots\right)}} e^{\frac{(y-G' y_1)^2}{2G'^2 y_1 \left(1+\frac{1}{s(V_2)}+\cdots\right)}}. \tag{7.11}$$

with $G' = \prod_{i=2}^{12} s(V_i)$.

However, things are unlikely to be this simple. In reality, there are likely to be many such contributions, with different $q_{\prime}$ corresponding to different parts of the dynode. We approximate the non-zero element with a single gaussian, and note that as the gain is smaller by a factor of $1/s(V_1)$ than the gain for the normal channel, this is unlikely to cause much of a problem.

We therefore fit the following function to the pedestal subtracted photoelectron spectrum:

$$f(x) = (1-\alpha)\left(\sum_{i=0}^{\infty}\frac{n^i}{i!}e^{-n}\mathrm{Gaus}(Gi, \sigma^2 i + \sigma_{\mathrm{ped}}^2)\right) \tag{7.12}$$

$$+\alpha\left(\sum_{i=0}^{\infty}\frac{n^i}{i!}e^{-n}\mathrm{Gaus}(Gi+\delta, \frac{\delta^2}{G^2}\sigma^2 + \sigma^2 i + \sigma_{\mathrm{ped}}^2)\right) \tag{7.13}$$

where $\mathrm{Gaus}(a, b)$ is a gaussian distribution with mean a and variance b. The mean number of photoelectrons incident on the pixel is n, and the gain of the PMT is G. The intrinsic width of the single photoelectron peak due to poisson statistics at the early dynodes is σ and the pedestal width, due to electronics noise, is σ_{ped}. The fraction α governs the fraction of incident photons that pass through the photocathode and produce a photoelectron in the first dynode, and δ is the effective gain that those electrons see (approximately the gain from the second dynode stage onward). Figures 7.3 and 7.4 show good and bad fits of this function to spectra from the mini-module. In the following plots, the relation between the parameter labels in the fit and equation 7.12 is as follows. Note that the x axis of the plots (ADC counts) is not the same as the x in the functional form (number of electrons) so quantities that are measured in ADC counts differ from those in the equation by a factor (charge on electron / charge per ADC count).

pedpos	position of centre of pedestal peak in ADC counts	(not in eqn)
meanpe	mean no. of photoelectrons at pixel	n
spepos	distance from pedestal to single photoelectron peak in ADC counts—this is the gain multiplied by some constants (charge on electron / charge per ADC count)	G
spewidth	Width of single photoelectron peak due to dynode statistics, in ADC counts	σ
dynodefrac	Fraction of photons that penetrate the photocathode and photoproduce on the first dynode	α
pedsigma	Width of pedestal peak	σ_{ped}
dynscale	Ratio of gains of 'normal' channel in PMT to gain of the channel that photoproduces on the first dynode	G/δ

7.2.1. Improvements to fit

We see from figures 7.3 and 7.4 that the function as fitted has some component below the pedestal. This is due to our use of gaussians to approximate the photoelectron peaks. It

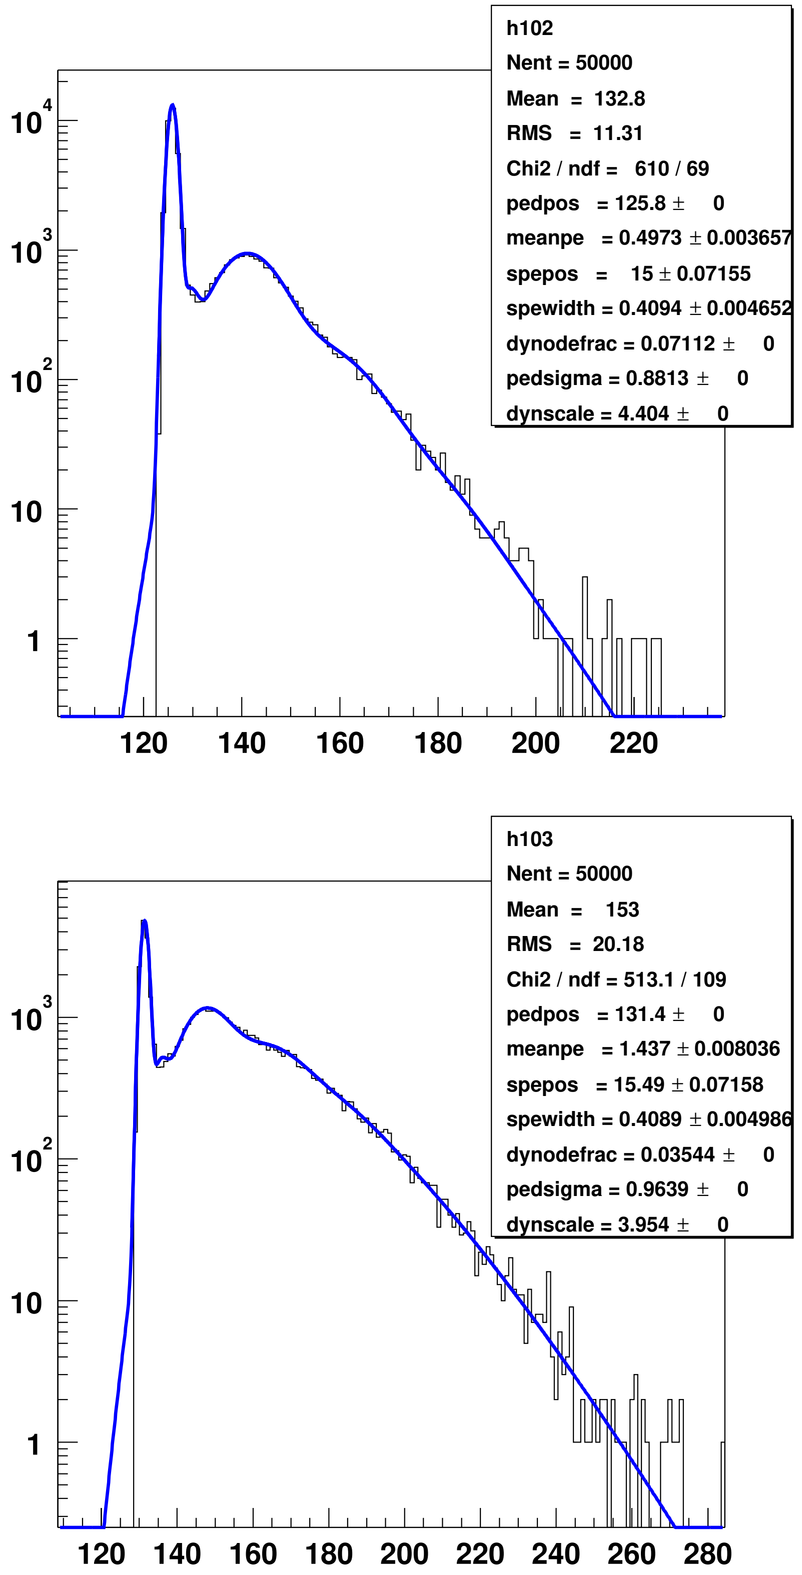

Figure 7.3.: Fits of the expected PMT spectrum to low light level data from the mini-module.

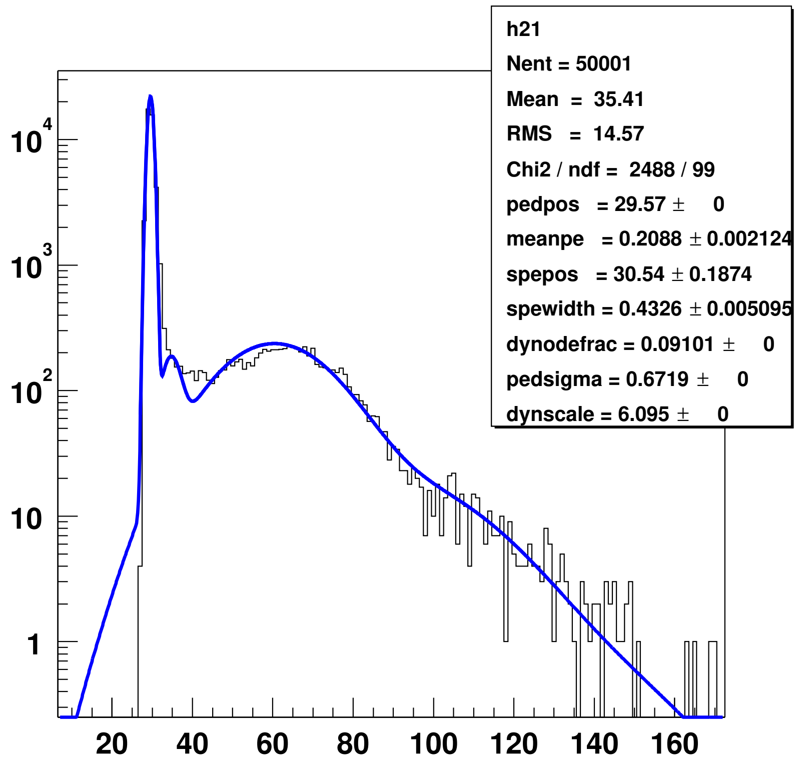

Figure 7.4.: A bad fit to the PMT spectrum, so producing poorly measured gain and number of photoelectrons.

provides too large a contribution under the pedestal, which is probably the reason for the poor χ^2 of the fits (most of the contribution to this comes form the pedestal region), and is probably also the cause of the poor fits such as that in figure 7.4. One can imagine two similar ways of improving this. The problem comes from the fact that the convolution of a number of poisson distributions is being modelled as a gaussian. The first solution is to follow the approach used in [51], and to carry on the use of several gaussians normalised according to a poisson distribution to the number of electrons produced at the first dynode, so instead of writing the single photoelectron peak as a gaussian, we write it as the sum of several gaussians, the relative weightings of which are determined by a poisson distribution with mean the secondary emission of the first dynode.

The second method is to use a gamma distribution in the place of the gaussian. This gives a better description of the true shape of the peaks, but when attempted has produced a fit which is somewhat unstable, and rather sensitive to the initial values of the fit parameters. Further work may produce a more reliable fit.

7.3. Performance of light injection calibration system

The experimental work at CERN was plagued with difficulties. The combination of the pre-amplifiers and the LeCroy ADCs was highly unstable—pedestal shifts of 10 ADC counts in a matter of seconds was not uncommon. In addition, the peak-sensing ADC used to read out the PIN diode proved to be unstable over time. As a consequence, I am not able to present a demonstration of the system working in the fashion in which it will work for MINOS. However, I will present evidence that the pulser system itself can produce the required accuracy given sufficient data, which, taken together with the PIN diode linearity and stability data from chapter 6, will be sufficient to demonstrate the effectiveness of the system.

7.3.1. Short term stability of system

The LED pulser system is designed to calibrate gain changes over a period of between an hour and a month. In order that our strategy of measuring one point per hour will work, we require the PMT gain to be stable over a period of an hour. Figure 7.5 shows the fractional change of gain in each PMT pixel from hour to hour, measured using the low light level spectrum fits described above. If the fit is performing well, it produces independent measurements of the PMT gain and the applied light level. In a poor fit, such as the one in figure 7.4, there tends to be a correlation between the errors on gain and N_{pe} (number of photoelectrons). We see that in the main, the fitted gains are stable

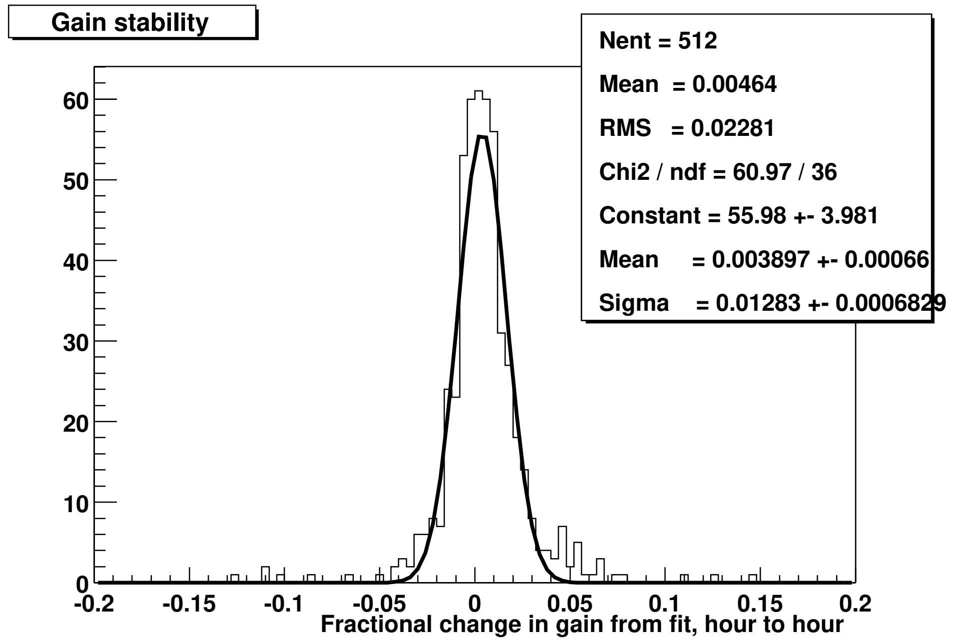

Figure 7.5.: Stability of phototubes over short time periods measured with the PMT spectrum fits.

to a little over 1%; there are some tails to the distribution, which are caused by the poor fits. If the PMT gains remain stable over an hour in the rather hostile environment of the East Hall at CERN, with fairly basic power supplies, there should be no problems in ensuring their stability in the mine, where the environment is rather more stable.

Figure 7.6 shows the number of photoelectrons extracted from the fit from hour to

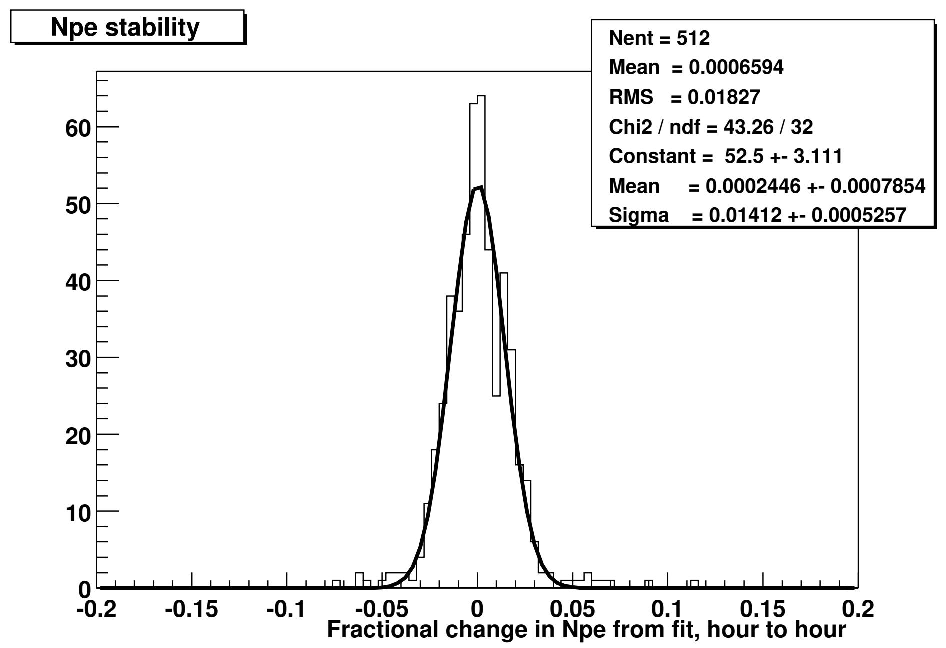

Figure 7.6.: Stability of pulser box "in the field" at CERN.

hour during a run, with the same pulse height setting on the pulser. We see that the pulser is stable to about 1.5%. Again, the distribution has a small tail due to the poor fits. We see that in the field, as well as in the lab, the pulser is stable.

7.3.2. Behaviour of system over longer time periods

Figure 7.7 shows the number of photoelectrons measured with the spectrum fits as a function of time. The plot is for the sum of all M64 pixels. The different colours on the plot represent different runs. The pulser was pulsed at the same light level for all these plots. We see that the LED is dimming at an apparently steady rate, losing 3% of its light over a fortnight. However, this was an early prototype pulser, which flashed the LED at about 3 kHz almost continuously for that time. This is equivalent to more than 2000 years running with the anticipated far detector usage, so we need not worry about the LED becoming too dim to see. Figure 7.8 is the same plot, but with the number of photoelectrons calculated as $(\mu/\sigma)^2$ for a higher light level. Again, we plot the sum of all M64 pixels. If we recall the calculations in section 7.2, the number of photoelectrons

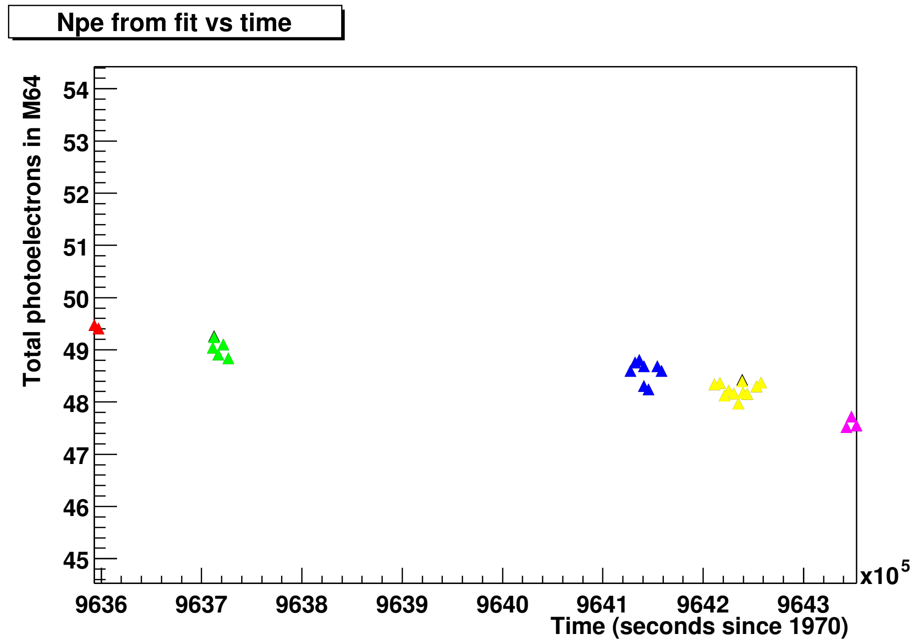

Figure 7.7.: Number of photoelectrons from spectrum fits vs. time. Note that "seconds since 1970" is the standard unix time. The full range of the x axis is about 9 days.

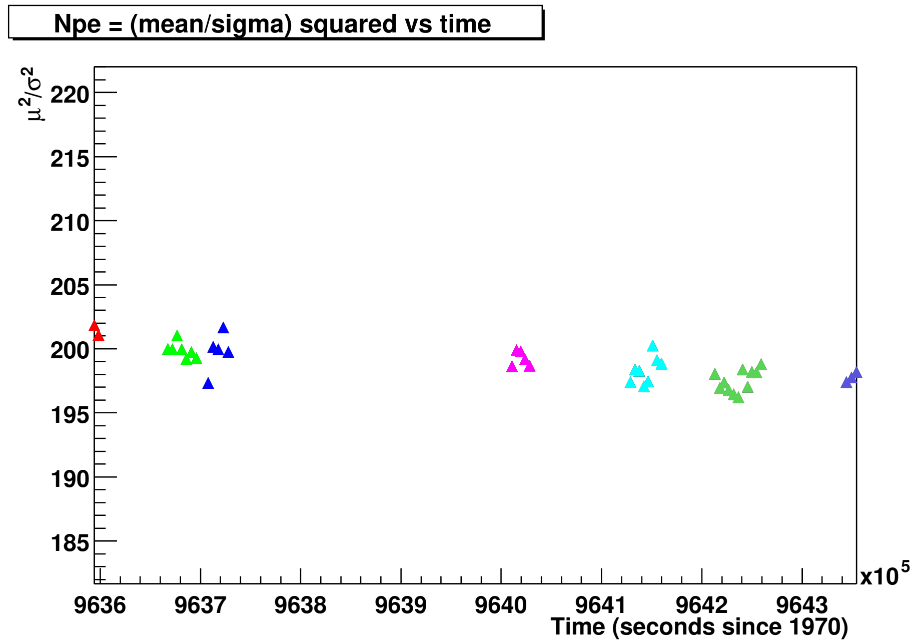

Figure 7.8.: Number of photoelectrons calculated with the $(\mu/\sigma)^2$ method vs. time.

at the first dynode is distributed as a poisson distribution, with mean and variance n. If we then had one very large gain stage, we would see:

$$\text{MeanADC} \quad \mu = Gn \tag{7.14}$$

$$\text{Variance} \quad \sigma^2 = G^2 n. \tag{7.15}$$

Then $\mu^2/\sigma^2 = G^2n^2/(G^2n) = n$. In a real phototube, of course, we don't have an infinite gain stage, we have a stage with a gain of about 7. We can then see that $\mu^2/\sigma^2 = n/(1 + \frac{1}{g_1} + \frac{1}{g_1 g_2} + \cdots)$, where g_i is the gain of the first stage and so on. The quantity $(\mu/\sigma)^2$ is still proportional to the number of photoelectrons produced. Note that the pedestal width must be subtracted from σ in quadrature when doing this calculation.

We see that this plot looks the same as figure 7.7, which gives us confidence that the result is correct.

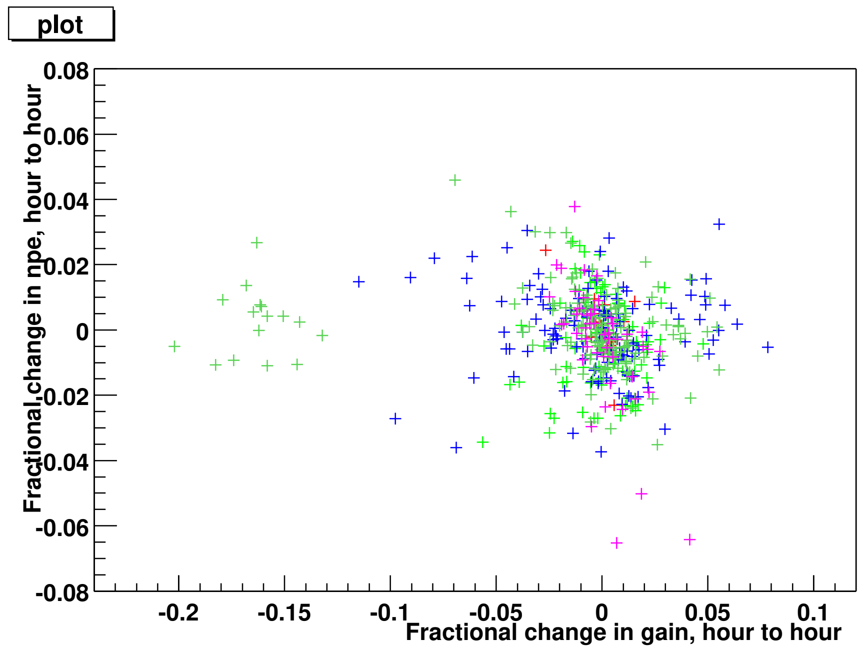

Figure 7.9.: Scatter plot of change in gain vs. change in npe for one M16. There was a 15% gain change between two runs—caused when the high voltage supply was changed. All channels in the M16 are plotted. Different colours represent different runs.

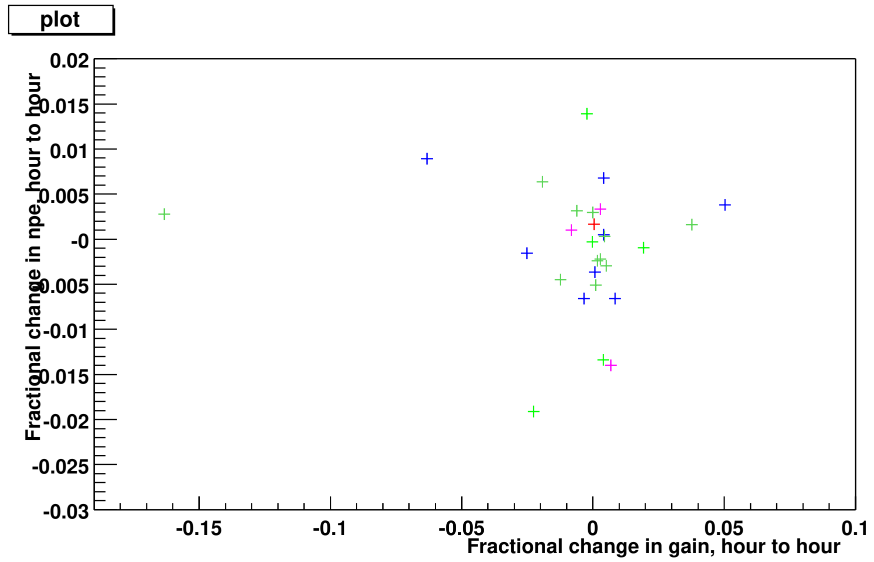

Figure 7.10.: The same plot as figure 7.9 but with the averages of the 16 pixels used. The large scatter in the blue dataset coincided with a severe thunderstorm at CERN.

Figures 7.9 and 7.10 show fractional change in number of photoelectrons compared with fractional change in gain as measured by the spectrum fits. We see that when there was a large gain change (the high voltage supply was changed), the measured number of photoelectrons remains constant to within about $\pm1\%$. This shows that our measurements of gain and number of photoelectrons are independent, and that we could achieve the required calibration, at least at low light levels, by using the spectrum fits. In practice, spectrum fitting cannot be used as the calibration method for MINOS, as a spectrum fit requires at least 50,000 flashes to make a measurement, which would take too long to accumulate on a regular basis. The software will have the ability to produce histograms for fitting, however, and they are expected to be an invaluable debugging tool at first installation.

7.4. Demonstration that the calibration system can calibrate to better than 2%

In order to demonstrate that the calibration system can achieve 2% calibration, we must consider again its scheme of operation. It should be noted that I do not attempt to show the linearity measurement properties here; we have shown in chapter 6 that the PIN is a linear measure of the amount of light produced in the LED. Further progress cannot be made towards demonstrating the linearity of the complete system *in situ* until MINOS electronics is available. This demonstration will be one of the tasks undertaken during the Calibration Detector setup at CERN this spring.

The principle behind the calibration is that, every hour, we pulse the channels with some amount of light, and monitor it with the PIN diode. The ratio of channel ADC to PIN ADC is the calibration constant for that hour for that channel; one then uses that calibration constant to convert data ADC readings into equivalent PIN units, which should be constant.

Unfortunately we don't have any reliable data from the PIN diode, so we are unable to do this. Instead, I have used the total number of photoelectrons in the M64, as measured by the μ/σ method, with LED pulse height[2] 17 as my constant measure of light, equivalent to the PIN diode. The procedure is then:

1. For pulse height 17, measure mean and sigma for all channels.
2. Calculate total number of photoelectrons in M64 as $N = \sum_{i=0}^{63} \mu_i^2/\sigma_i^2$.
3. Calculate calibration constants as a function of time for all channels as $\alpha_i = N/\mu_i$.
4. Look at pulser light for a different light level. I chose height 18. For all channels, use calibration constants α_i to convert ADCs to our "PIN-like" M64 unit.

[2]This prototype pulser had 31 pulse height settings, evenly spaced in current applied to the LED. Pulse heights 17 and 18 correspond to a peak current of around 500 mA in the LED, which due to the poor optical transmission in the prototype parts, corresponded to an average of around 15 photoelectrons per channel.

The results of this procedure are shown in figure 7.11. We see that we have achieved calibration to the 1% level.

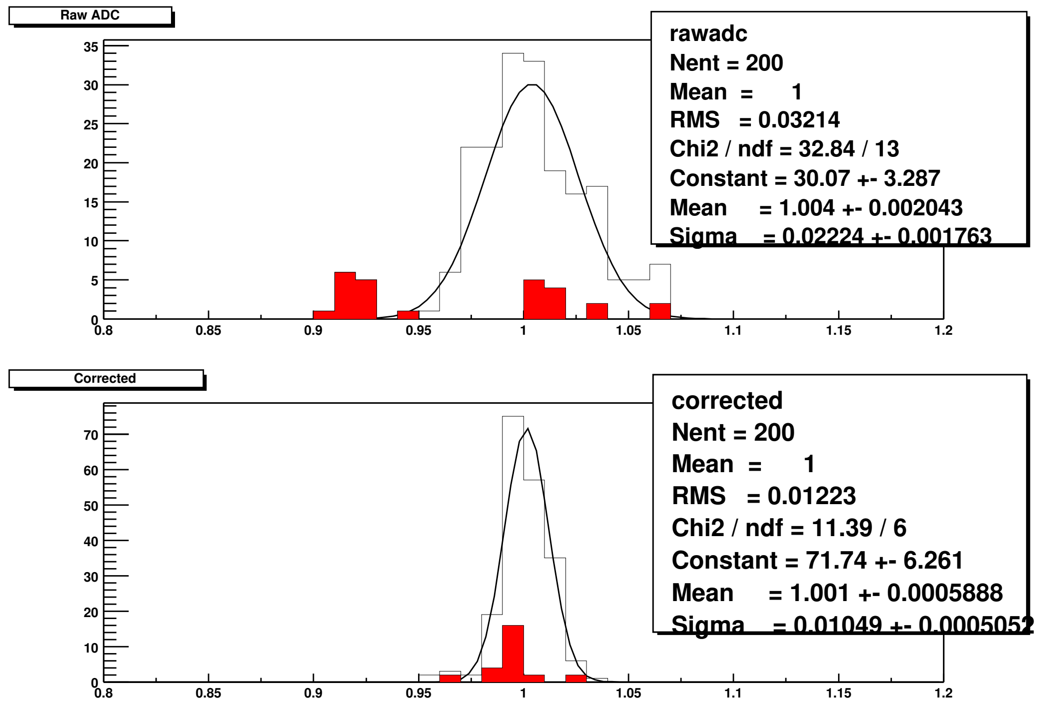

Figure 7.11.: The calibration system works! Upper plot: raw ADC values for pulser height 18. Plotted are averages for each individual M16, normalised to 1. Lower plot: ADC counts from the top plot as calibrated by the "height 17 M64 PIN" method described above. The red shaded areas correspond. Some of the 1% could be jitter on the LED pulser. The histograms include data that cover the same period of two weeks used previously in this chapter.

8. Conclusions

Houston, Tranquility Base here. The Eagle has landed

— Commander Neil Armstrong

An LED-based pulsed light injection calibration system has been designed for the MINOS experiment. It has been demonstrated that the system achieves its design objectives—namely to be able to calibrate PMT channels to an accuracy of 2%. The system has been designed to be as flexible as possible, so as to accommodate any unforeseen calibration needs that MINOS discovers. As well as calibrating the gain and linearity of the PMT readout by means of reference PIN diodes, the system offers the ability to acquire low light-level photoelectron spectra for detailed offline study, and, by knowing the lengths of optic fibre between the pulser box and each channel, allows the relative timing of the readout channels to be measured to better then 1 ns. This timing calibration will facilitate the direction-finding of through-going muons, and so enable measurements of the atmospheric neutrino flux to be made at angles close to the horizontal.

The NuMI beam will turn on in autumn 2003. Within two years of that date, MINOS will have clarified whether or not neutrino oscillations are the cause of the atmospheric neutrino anomaly, and if they are, it will measure the mixing parameters $\sin^2 2\theta$ and $\Delta \mathrm{m}^2$ to better than 10%.

Bibliography

[1] Wolfgang Pauli. Letter, reproduced in translation in September 1978 *Physics Today*.

[2] C. L. Cowan, Jr., F. Reines, F. B. Harrison, H. W. Kruse, and A. D. McGuire. Detection of the free neutrino: A confirmation. *Science*, 124(3212):103, 1956.

[3] C.S. Wu et al. *Phys. Rev.*, 105:1671, 1957.

[4] M. Goldhaber, L. Grodzins, and A.W. Sunyar. *Phys. Rev.*, 109:1015, 1958.

[5] The L3 Collaboration, B. Adeva, et al. *Z. Phys.*, C51:179, 1991.

[6] The OPAL Collaboration, G. Alexander, et al. *Z. Phys.*, C42:175, 1991.

[7] The DELPHI Collaboration, P. Abreu, et al. *Nucl. Phys.*, B367:511, 1991.

[8] The ALEPH Collaboration, D. Decamp, et al. *Z. Phys.*, C53:1, 1992.

[9] The ALEPH, LEP, DELPHI, L3 and OPAL Collaborations, D. Decamp, et al. *Phys. Lett.*, 276B:247, 1991.

[10] The MARK II Collaboration, G.S. Abrams, et al. *Phys. Rev. Lett.*, 63:2173, 1989.

[11] Ray Davis, Jr., D. S. Harmer, and K. C. Hoffmann. *Phys. Rev. Lett.*, 20:1205, 1968.

[12] B. Pontecorvo. *JETP*, 33:549, 1957.

[13] B. Pontecorvo. *JETP*, 34:247, 1958.

[14] B. Pontecorvo. *JETP*, 53:1717, 1967.

[15] V. Lobashev et al. *Phys. Lett. B*, 460:227, 1999.

[16] C. Weinheimer et al. *Phys. Lett. B*, 460:219, 1999.

[17] K. Assamagan et al. *Phys. Rev.*, D53:6067, 1996.

[18] R. Barate et al. *Eur. Phys. J*, C2:395, 1998.

[19] J. Bachall and S. Glashow. *Nature*, 326:476, 1987.

[20] Heidelberg-Moscow Collaboration. *Phys. Rev. Lett.*, 83:41, 1999.

[21] D.E. Groom et al. Review of particle physics. *The European Physical Journal*, C15, 2000.

[22] L. Wolfenstein. *Phys. Rev.*, D17:2369, 1978.

[23] S. P. Mikheyev and A. Y. Smirnov. *Sov. J. Nucl. Phys.*, 42:1441, 1985.

[24] The SAGE Collaboration, A. I. Abazov, et al. *Phys. Rev. Lett.*, 67:3332, 1991.

[25] The SAGE Collaboration, D. N. Abdurashitov, et al. *Phys. Rev. Lett.*, 77:4708, 1996.

[26] The GALLEX Collaboration, P. Anselmann, et al. *Phys. Lett. B*, 285:376, 1992.

[27] The GALLEX Collaboration, W. Hampel, et al. *Phys. Lett. B*, 388:384, 1996.

[28] K. S. Hirata et al. *Phys. Rev. Lett.*, 65:1297, 1990.

[29] Y. Fukuda et al. *Phys. Rev. Lett.*, 77:1683, 1996.

[30] Y. Fukuda et al. *Phys. Rev. Lett.*, 81:1158, 1998.

[31] Y. Fukuda et al. *Phys. Rev. Lett.*, 82:1810, 1999.

[32] J. N. Bachall. `http://www.sns.ias.edu/~jnb/`.

[33] The SNO Collaboration. *NIM A*, 449:172, 2000.

[34] R. Becker-Szendy et al. *Phys. Rev. Lett.*, 69:1010, 1992.

[35] S. Hatakeyama et al. *Phys. Rev. Lett.*, 81:2016, 1998.

[36] Y. Fukuda et al. *Phys. Lett. B*, 433:9, 1998.

[37] Y. Fukuda et al. *Phys. Lett. B*, 436:33, 1998.

[38] M. Aglietta et al. *Europhys. Lett.*, 8:611, 1989.

[39] K. Daum et al. *Z. Phys. C*, 66:417, 1995.

[40] W. W. M. Allison et al. *Phys. Lett. B*, 449:137, 1999.

[41] S. Fukuda et al. *hep-ex/0009011*, 2000.

[42] M. Apollonio et al. *Phys. Lett. B*, 420:397, 1998.

[43] C. Y. Cardall and G. M. Fuller. *Phys. Rev. D*, 53:4421, 1996.

[44] The MINOS Collaboration. NuMI Beam TDR. Technical report, Fermilab, 1999.

[45] The MINOS Collaboration. MINOS TDR. Technical report, Fermilab, 1999.

[46] Roy Lee. Private communication.

[47] Peter Litchfield. Presentation to minos calibration working group.

[48] P. G. Harris. NuMI-L-535: Checkerboard muxing: How to minimise crosstalk problems. Technical report, Fermilab, 1999.

[49] Ruben Saakyan. Private communication.

[50] Alfons Weber. Private communication.

[51] S. Tokár, I. Sýkora, and M. Pikna. Technical Report ATL-TILECAL-99-005, CERN, 1999.

A. Muon fitting

I mentioned in the main text that muons were used to carry the calibration between the MINOS detectors, and to provide an overall energy scale. It is necessary for correct calibration to identify an equivalent set of muons in each detector. This appendix describes work I did on muon reconstruction. The software was written for the old Fortran/ADAMO system (MINOS has since moved to an ROOT-based C++ framework). The algorithm used is based on that used by the CDHS experiment.

A.1. Range Measurement

Where a muon stops in the detector, a measurement of its momentum from range may be made. A lookup table of range vs muon momentum is produced at the start of the run, and this consulted for each contained muon. The results are shown in figure A.1, with a fitted gaussian. The gaussian width is about 5%, with something of a tail towards low reconstructed momenta—these are mostly events with a large discrete energy loss, which is not yet accounted for in the fit. There are a small number of events with reconstructed energies that are very small—these are failures of the current track-finding code to find the complete track. The current track-finding code does not cope well with noise in the same z-region as the track, and needs improving. It is expected that when the "extra"

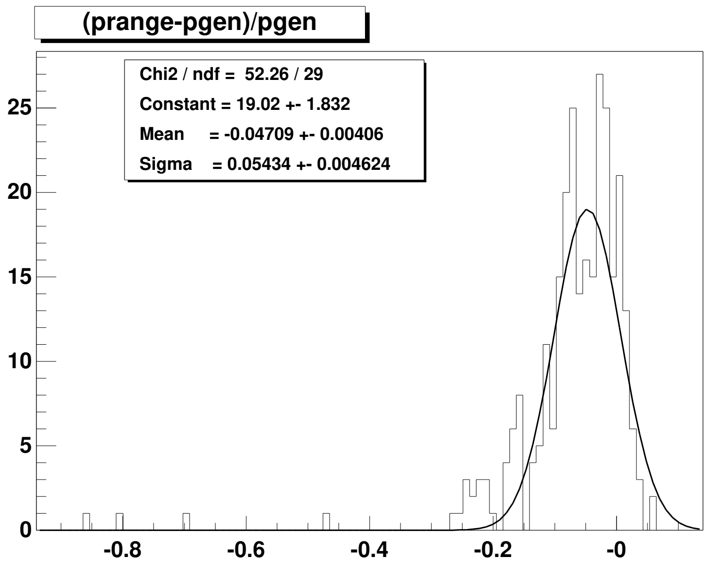

Figure A.1.: $\frac{p_{range}-p_{gen}}{p_{gen}}$ for contained muons in the far detector.

visible energy is added back into the track, much of this tail will disappear. As well as providing a measurement of momentum for some muons, which may or may not be of interest of itself, this measurement can be compared with the curvature measurement for the muon to check that the magnetic field model is correct.

A.2. Curvature measurement

The reconstructed track (x_i, y_i) is described by five parameters: the track origin (x_0, y_0), the slope of the track at its origin $\left(\left.\frac{dx}{dz}\right|_0, \left.\frac{dy}{dz}\right|_0\right)$, and the reciprocal of the initial momentum $\frac{1}{p_0}$. The method is an iterative one, and proceeds as follows:

1. Initial values for the 5 track parameters are chosen.

2. This muon is then propagated through the detector. The points at which the track intersects each active plane are then parameterized as

$$x_i = x_0 + \left.\frac{dx}{dz}\right|_0 (z_i - z_0) + \frac{1}{p_0} b_i^x \tag{A.1}$$

and

$$y_i = y_0 + \left.\frac{dy}{dz}\right|_0 (z_i - z_0) + \frac{1}{p_0} b_i^y \tag{A.2}$$

where $b_i^x = x_i^m - x_0 - \left.\frac{dx}{dz}\right|_0 (z_i - z_0)$, the measured displacement of the track due to bending, and the x_i^m are the measured track positions, so

$$b_i^x = \left(x_i - x_0 - \left.\frac{dx}{dz}\right|_0 (z_i - z_0) \right) p_0 \tag{A.3}$$

and similarly for b_i^y. Note here that it is important that the actual position where the track went (obtained from interpolating the hits) be used to determine the B-field, rather than using the position of the projected track. Due to the way that multiple scattering is accounted for in a correlated covariance matrix, the "best fit" track whilst having the correct initial parameters may actually follow a rather different path in the detector from the real muon.

3. Write $\chi^2 = \sum_{i,j} (x_i^m - x_i) w_{ij} (y_i^m - y_i)$, where the weight matrix w_{ij} contains contributions from the strip width and uncertainties due to multiple scattering and energy loss, and x^m and y^m are the measured values of track position.

4. Minimize χ^2. This gives new values for the track parameters.

5. Repeat until some convergence criterion is satisfied. The present code terminates when successive values of muon momentum differ by less than 200 MeV—this is an empirical number based on the performance of the code, and has no special significance. This choice is sufficient to ensure that the momentum measurement is limited by the physics (i.e. multiple scattering) and not a software artifact.

A.2.1. Measurement of muon momentum

The curvature fit performs qualitatively as one would expect—fully contained high momentum events are measured better than low energy ones, for a given momentum the measurement is better when more of the track is used, and events that exit the detector having travelled through only a small number of planes are not usefully measured at

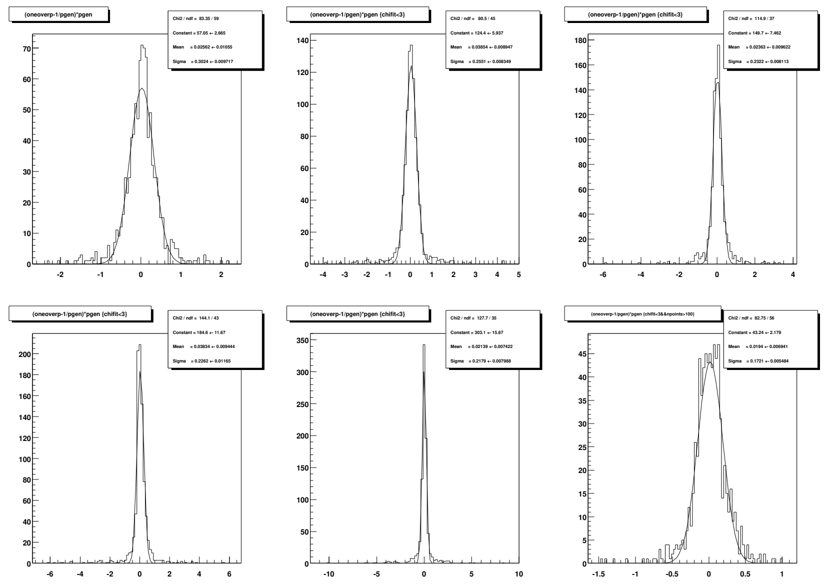

Figure A.2.: $\frac{\frac{1}{p_{reco}}-\frac{1}{p_{gen}}}{\frac{1}{p_{gen}}}$ for 4,6,8,10 and 15 GeV muons, and 15 GeV muons which pass more than 100 planes. The scales are unreadably small; the RMS width ranges from 30% at 4 GeV to 17% at 15 GeV.

all—in the case of 10 GeV muons, the track must pass through about 100 planes before a useful measurement of the momentum can be made. We also note that the muon track is hugely oversampled - selecting only every 5th plane in the far detector (as a mimic of the near detector geometry) makes no significant difference to the performance of the fit, save that it runs faster with a smaller matrix to invert! It is clear that the amount of light produced by a muon in the scintillator will have no effect on the curvature fit unless the light drops to such a low level that other physics measurements have already become impossible. One could imagine losing half the points or more at no cost for tracks of a reasonable length.

Figure A.2 shows the reconstructed $\frac{1}{p}$ for muons with various momenta. We see that our 4 GeV muons are measured to about 30%, and that the measurement gets better with increasing momentum. The set of 15GeV muons with more than 100 datapoints are known to 17%, and the fully or almost-fully contained 15 GeV events to about 12%. Shown in figure A.3 is the reconstructed momentum for our 5 different muon energies

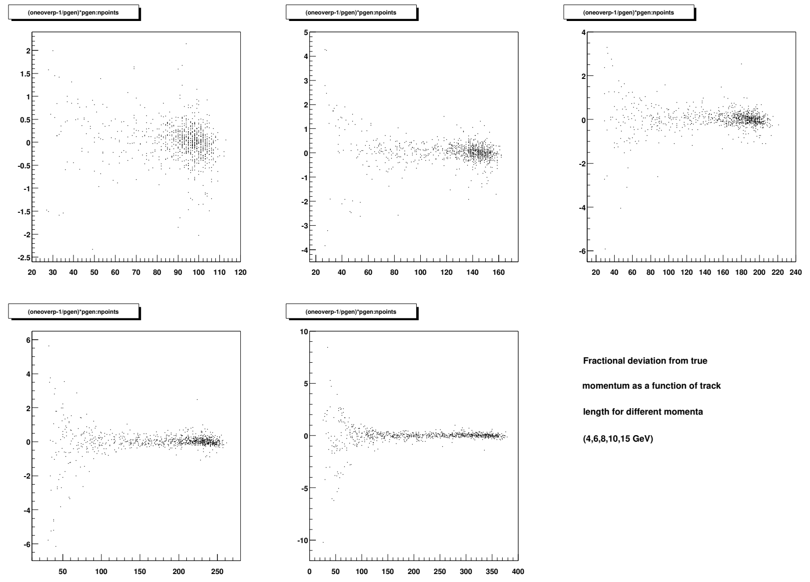

Figure A.3.: $\frac{\frac{1}{p_{reco}}-\frac{1}{p_{gen}}}{\frac{1}{p_{gen}}}$ as a function of number of planes in the event.

plotted against increasing number of planes in the event. We see that any event which is measured over less than the first third of its length is essentially unknown.

Figures A.4 and A.5 show the stretch function $\frac{\frac{1}{p_{reco}}-\frac{1}{p_{gen}}}{\sigma_{\frac{1}{p}}}$ and the χ^2 from the fit for the different muon momenta. If the system is well modelled by the reconstruction code, we expect $\frac{\frac{1}{p_{reco}}-\frac{1}{p_{gen}}}{\sigma_{\frac{1}{p}}}$ to be a gaussian with width 1 and mean 0, and we expect χ^2 to have mean 1. We see that the stretch function does look like this, although it has a long negative tail—these are events which undergo multiple scattering and fit to a straighter track than would be expected from their momenta, and so the error calculation uses the 'wrong' geometry, and calculates too small an error. If it is ever desired to use the error reported by the fit in some calculation, this must be corrected for. The χ^2 distribution is indeed centered on about 1, and has a low tail. On examining the plot of χ^2 against number of planes, it can be seen that the events that leave the detector quickly have a low χ^2. This is due to the position being quantized into 4.1 cm wide bins, so the straightish part of the track at the beginning actually consists of many points with the same x or y coordinate, rather than the gaussian spread that is 'expected' by the model.

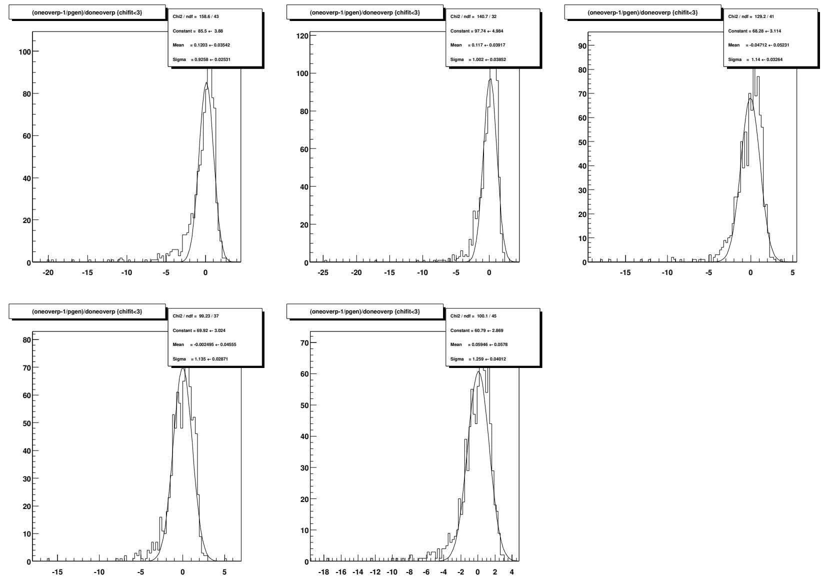

Figure A.4.: $\frac{\frac{1}{p_{reco}}-\frac{1}{p_{gen}}}{\sigma_{\frac{1}{p}}}$ for 4,6,8,10,15GeV muons

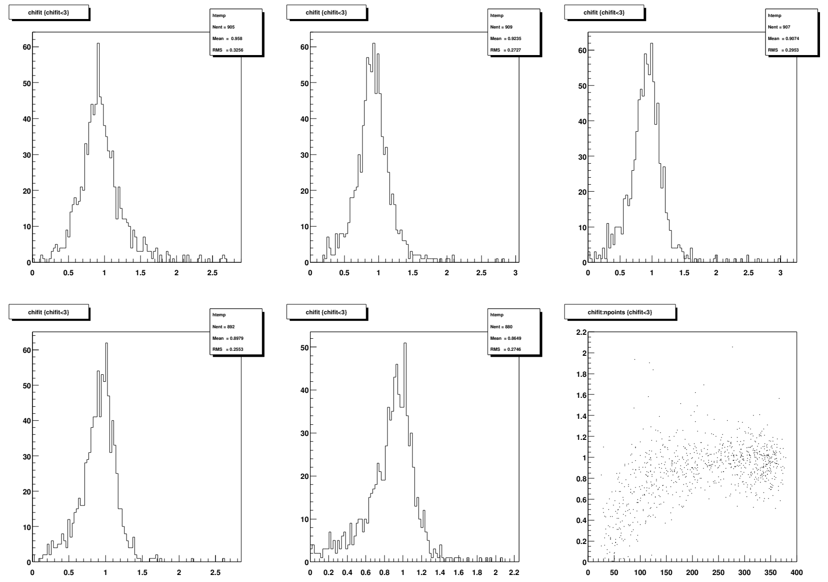

Figure A.5.: χ^2 from curvature fits to 4,6,8,10,15GeV muons, and χ^2 vs number of planes for 15GeV muons.

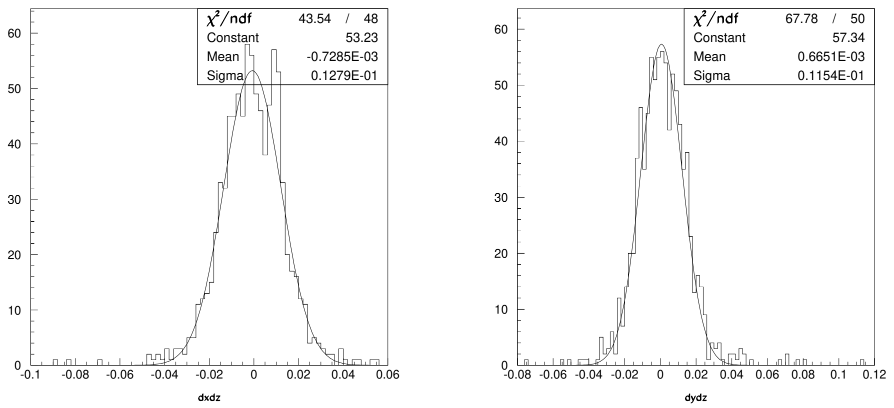

Figure A.6.: Directional resolution of fit.

A.2.2. Measurement of muon direction

A measurement of the initial direction of the muon is provided for free with the curvature fit. This measurement (figure A.6) of the slope at the vertex has a gaussian width of about 0.01. These plots are for muons, not CC ν_μ events, so in the real data there is a little extra uncertainty due to the projection of the muon track back into the shower region, but the angular resolution achievable is still of order a degree.

A.2.3. Sources of error

There are several effects that one might think would introduce error into this measurement; some of these are shown to be small, and others understood and accounted for in the fit. The stretch function $\frac{\frac{1}{p_{reco}}-\frac{1}{p_{gen}}}{\sigma_{\frac{1}{p}}}$ has a mean of about zero, and a gaussian width of about 1, and the distribution of χ^2 is centred close to 1, so we are confident that there are no significant outstanding sources of error.

Multiple scattering

Multiple scattering introduces correlated errors. This is the largest effect for muons in a dense medium such as iron. Considering the geometry in figure A.7 and using the

RMS multiple scattering angle from the Particle Data Book, we see that the variation in displacement due to multiple scattering at plane N and the covariance between the displacements at planes K and N are given by:

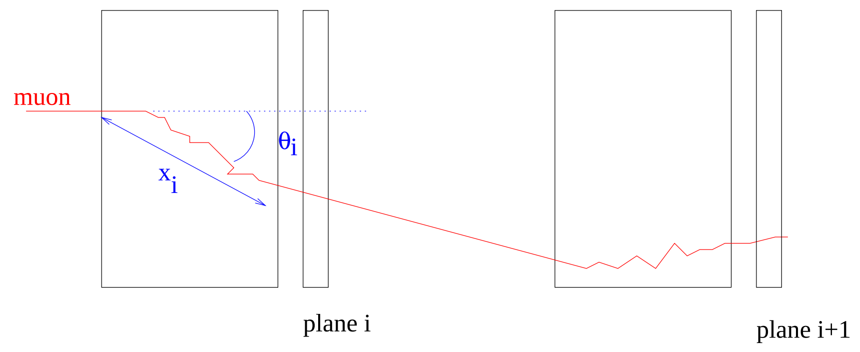

Figure A.7.: Multiple scattering in MINOS.

$$\langle P_N^2 \rangle = \sum_{i=1}^{N} \theta_{0i}^2 \left(\frac{x_i^2}{3} + \mathcal{D}(i,N)(x_i + \mathcal{D}(i,N)) \right) \tag{A.4}$$

$$\langle P_K P_N \rangle = \sum_{i=1}^{K} \theta_{0i}^2 \left(\frac{x_i^2}{3} + \frac{x_i}{2}(\mathcal{D}(i,K) + \mathcal{D}(i,N)) + \mathcal{D}(i,K)\mathcal{D}(i,N) \right) \tag{A.5}$$

where $\mathcal{D}(i,K)$ is the distance along the path of the muon from plane i to plane K. Here the RMS multiple scattering angle from the Particle Data Book is θ_0, and

$$\theta_{0i} = \frac{1}{1 - \frac{1}{p_0}(dE/dx)ix}\theta_0 \tag{A.6}$$

with x the distance along the track to the ith plane. Care must be taken to ensure that the gap between the supermodules is properly accounted for.

For muons in MINOS, the expected RMS deviation due to multiple scattering from the path that the muon would have followed in the absence of scattering for a number of energies is shown in figure A.8. Note that the values for points close to the end of the track are underestimated—no account is taken of the departure from mip-ness at low energies. (The minimum ionizing point occurs for momentum of order 3 times the rest mass. When the muon energy drops below about 200 MeV, energy loss and multiple scattering become large.) In the reconstruction code, no attention is paid to the end

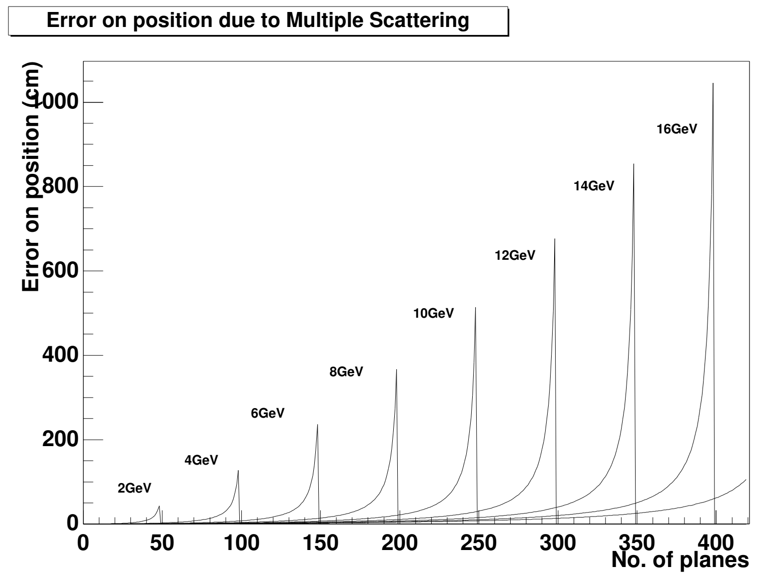

Figure A.8.: RMS deviation from projected track position at nth plane due to multiple scattering.

of the track, as there is essentially no information, and potentially much confusion contained in it if the weight matrix isn't quite right.

Strip size

The MINOS detectors contain plastic scintillator strips of width 4.1 cm. The uncertainty in the measurement of position due to this width is represented in the fit by a gaussian with $\sigma = 4.1/\sqrt{12}$ cm. One might worry that this will tend to pull the reconstructed track toward the centre of the strips, and thus introduce extra error, and possible systematic effects, however this turns out not to be significant. Shown in figure A.9 is the reconstructed momentum for a set of muon data binned into scintillator strips, and that same data with a gaussian error of $\sigma = 4.1/\sqrt{12}$ cm on the true positions. We see that there is no need to worry about any effects due to the real error being a rectangular distribution rather than a gaussian—the effect is a small one. Some effect is seen for the short, straight tracks caused by high momentum muons near the end of the detector, but these events are likely to fall outside the fiducial volume.

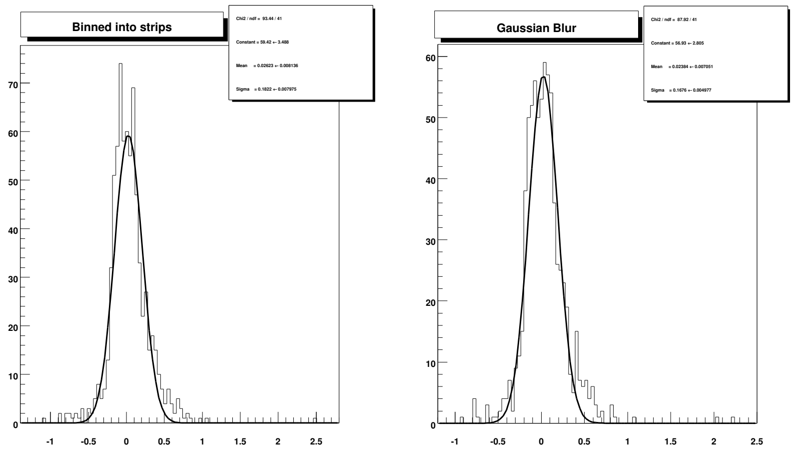

Figure A.9.: Recontructed momentum for a sample of events binned into scintillator strips, and the same events with a gaussian blur on the "true" position.

Variations in energy loss

The loss in energy of a muon travelling through a plane in MINOS is given by the Vavilov distribution. This will introduce a distribution on the muon momentum at any point on the track, and so extra error on the magnetic bending. An exact calculation is rather involved, but approximately, we may proceed as follows:

The energy loss in one plane of the MINOS detectors for a muon of energy of order 1 GeV is given by a Vavilov distribution with mean $\sim$ 40 MeV and RMS width $\sim$ 10 MeV. So the energy distribution of a muon of initial energy E after n planes is given by the convolution of a δ at energy E with n Vavilov distributions with mean -40 MeV. Variance is additive under convolution, so after n planes, we have $\sigma_E = 10\sqrt{n}$ MeV.

Consider a muon of momentum p GeV travelling a distance t m through a magnetic field B T perpendicular to the path of the muon. We define z as being the initial direction of the muon. Approximately, we have $p = 0.3B\rho$, where ρ is the radius of curvature of the track in metres. Then the transverse displacement due to magnetic bending is

$$\Delta x = \frac{0.3Bt^2}{2p} \tag{A.7}$$

and the slope of the track at the end of the region is

The contribution of variations in energy loss to the position of hits in MINOS can then be estimated as follows:

$$\sigma_{b_ib_j} = \left\langle (b_i - \overline{b_i})(b_j - \overline{b_j}) \right\rangle . \tag{A.8}$$

Now, we lose 40 ± 10 MeV per plane, so we may write $p_k = p_0 - 0.04k - \sum_{m=1}^{k} \Delta_m$, with $\langle \Delta_m \rangle = 0$ and $\langle \Delta_m^2 \rangle = (0.01)^2$, where numbers are in GeV. So then

$$\frac{1}{p_k} - \overline{\left(\frac{1}{p_k}\right)} \approx \frac{1}{p_k} - \frac{1}{(\overline{p_k})} \tag{A.9}$$

and

$$\sigma_{b_ib_j} = \left\langle \left(\sum_{k=1}^{i} \left(\frac{0.3B_k t_k^2}{2} + 0.3B_k t_k \Delta_{ik} \right) \frac{\sum_{m=1}^{k} \Delta_m}{(p_0 - 0.04k)^2} \right) \left(\sum_{l=1}^{j} \left(\frac{0.3Bk_l t_l^2}{2} + 0.3B_l t_l \Delta_{jl} \right) \frac{\sum_{n=1}^{l} \Delta_n}{(p_0 - 0.04l)^2} \right) \right\rangle . \tag{A.10}$$

The Δ_i are independant, so we have $\langle \Delta_i \Delta_j \rangle = (0.01)^2 \delta_{ij}$, and so

$$\sigma_{b_ib_j} = (0.01)^2 \sum_{m=1}^{j} \left(\sum_{k=m}^{i} \left(\frac{0.15B_k t_k^2 + 0.3B_k t_k \Delta_{ik}}{(p_0 - 0.04k)^2} \right) \sum_{l=m}^{j} \left(\frac{0.15B_l t_l^2 + 0.3B_l t_l \Delta_{jl}}{(p_0 - 0.04l)^2} \right) \right) \tag{A.11}$$

$$\sigma_{b_ib_j} = (0.01)^2 \sum_{m=1}^{j} \left(\sum_{k=m}^{i} \frac{1}{p_k} (\Delta_k + \left.\frac{dx}{dz}\right|_k \Delta_{ik}) \sum_{l=m}^{j} \frac{1}{p_l} (\Delta_l + \left.\frac{dx}{dz}\right|_l \Delta_{jl} \right) . \tag{A.12}$$

We see from figure A.10 that this error is not important when compared to that from multiple scattering, so we are justified in ignoring it in our treatment.

A.3. Summary

Code exists that will reconstruct muons in the MINOS detectors to about 6% from range, and to between 15 and 30% dependant on momentum from curvature. The code runs as part of the FORTRAN reco_minos framework.

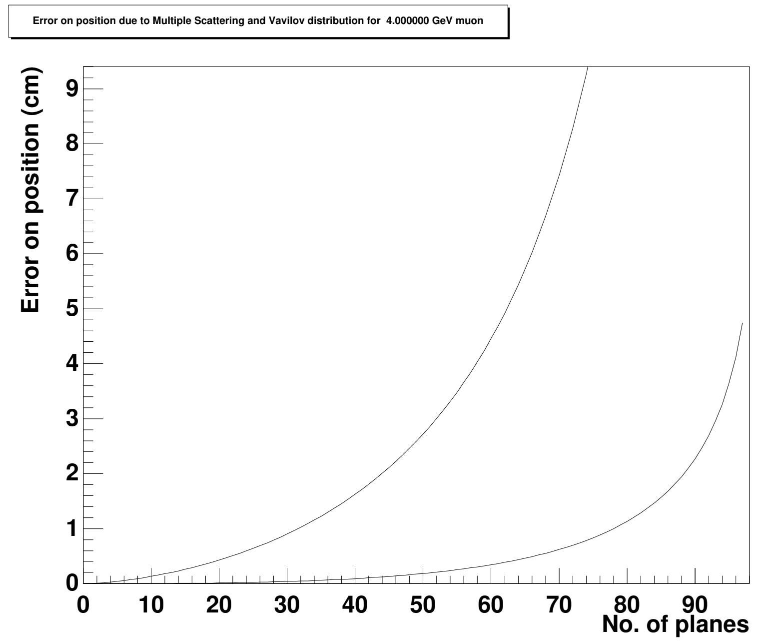

Figure A.10.: Mean error on position for 4GeV muon due to multiple scattering (upper curve) and width of Vavilov distribution. At higher momenta, multiple scattering dominates more.