Detector Techniques for the Application of Topological and Kinematical Criteria for a $\nu_\mu \rightarrow \nu_\tau$ Oscillation Search

Ioannis M. Papadopoulos

Ph.D. thesis

Physics Department
University of Athens
Greece

20 March 2002

This thesis was supported by the Hellenic State Scholarship Foundation and the CERN Doctoral Student Programme.

Main Supervisor	Dr. Christine Kourkoumelis[1].
CERN Supervisor	Dr. Jaap Panman[2].
Members of advisory committee	
	Dr. Pavlos Ioannou[1],
	Dr. Athanassios Lahanas[1].
Members of examination jury	
	Dr. Emmanuel Dris[3],
	Dr. Dimitrios Fassouliotis[1],
	Dr. Paris Sphicas[1],
	Dr. George Tzanakos[1].

[1] Nuclear & Particle Physics Section, Physics Department, National and Kapodistrian University of Athens, Greece.
[2] EP/CHE Division, CERN, Geneva, Switzerland.
[3] Department of Physics, National Technical University of Athens, Greece.

CONTENTS

1. INTRODUCTION

The question whether neutrinos have mass, is one of the most intriguing questions in particle physics. Experimental searches for neutrino oscillation provide at present the best means to probe the mass of these elusive fermions down to the sub-eV range. Potential experimental evidence for oscillations, which would directly imply the existence of non-zero neutrino masses, would open a window to new physics beyond the Standard Model.

The CHORUS experiment at the CERN-SPS searches for $\nu_\mu \rightarrow \nu_\tau$ oscillations making use of the nuclear emulsion technique. A potential signal is marked with the observation of the τ decay topology in a ν_τ charged current event. Such a topology is searched for in the emulsion sheets of the CHORUS target using fully automatic microscopes. This thesis discusses the development and performance of the microscopes, built at the scanning facility at CERN.

Nuclear emulsion has been the basis of new compact and high resolution trackers, whose performance was measured in a test experiment at the CERN-PS. Within the framework of this test experiment, a study of pion elastic scattering in emulsion was performed. The study aimed at the estimation of the so-called "white kink" background to the CHORUS oscillation signal. The nuclear emulsion trackers were used to measure the transverse momentum spectrum of this background process. The appropriate kinematical criteria were then determined, so that this background process can be separated from potential signal in the $\nu_\mu \rightarrow \nu_\tau$ oscillation search in CHORUS.

The thesis begins with an overview of the neutrino oscillation phenomenology. The CHORUS experiment is described in chapter 3. The nuclear emulsion technique, as it has been developed within the CHORUS experiment, and with the implementation in Nagoya as the working example, is presented in chapter 4.

The scanning facility at CERN is described in chapter 5. Emphasis is given to the design and implementation of the data acquisition system of the microscopes and to the description of the automatic scanning procedures.

The test experiment at the CERN-PS and the "white kink" measurement are described in chapter 6. The thesis concludes with the sensitivity to $\nu_\mu \rightarrow \nu_\tau$ oscillation using the current accumulated statistics by CHORUS, and the background estimation from the "white kink" measurement.

2. NEUTRINO OSCILLATION PHENOMENOLOGY

In this chapter the current experimental status on neutrino oscillation searches is presented. It is also illustrated how different types of experiments, using different types of neutrino sources, explore different ranges of mass differences and mixing angles.

The neutrino oscillation phenomenology is discussed within the context of the simplest case of two-flavour mixing. This choice is made because under such a scheme experiments can express their results, either positive or negative, as allowed regions or exclusion contours in a two-parameter space. It is shown how experimental constants, such as the baseline, the neutrino energy, systematical errors, background rates, and the search method itself affect the final result. To illustrate the current experimental status it is shown how more general mixing schemes affect the interpretation of the results.

Since the existence of non-zero neutrino masses is essential for oscillations to occur, its importance in physics is described, quoting the latest experimental results from direct mass measurements. For completeness, this chapter starts with a brief review of the basic neutrino properties in the Standard Model.

2.1 Neutrino properties and masses

Neutrinos have served often as a tool to study the intrinsic properties of the elementary fermions and their interactions. The theoretical foundation of the electroweak interaction, namely of the gauge invariance of the $SU(2)_L \otimes U(1)_Y$ symmetry, was mainly inspired by experimental results involving neutrinos. However, still very little is known about their intrinsic properties.

2.1.1 Neutrinos in the Standard Model

In the theory of the Standard Model (SM) developed by Glashow, Salam, Weinberg, Veltman, 't Hooft, and others, neutrinos appear as the weak isospin partners of the left handed charged leptons e, μ and τ, coupled to the $W^{\pm}$ bosons via the weak charged current:

$$J_{\mu}^{CC} = \overline{\psi}_{\nu_l} \gamma_{\mu} \frac{1}{2}(1 - \gamma^5)\psi_l \tag{2.1}$$

Therefore, their flavour can be uniquely identified by the production of the corresponding charged lepton l in charged current (CC) interactions.

The first experiment to detect electron neutrinos was set up in 1953 by Reines and Cowan [1], a long time after its original theoretical postulation by Pauli in his famous

letter to the "radioactives" in 1930. Neutrinos were generated at a nuclear reactor and detected via the inverse beta decay process

$$\overline{\nu}_e + p \rightarrow e^+ + n.$$

The detector consisted of hydrogeneous liquid scintillator with dissolved cadmium. A neutrino interaction was triggered by the delayed coincidence of two pulses: The first one originated from the energy deposition of the two γs from the e^+ annihilation. Hence, the deposited energy had to be equal to $2 \cdot m_e$ for the pulse to be valid. The second pulse, which was required to arrive about 5 μs later, was a signal from the neutron capture by a cadmium nucleus. This technique is still being used at reactor and other low energy neutrino experiments.

The existence of the electron neutrino was established though with a more sophisticated experiment in 1956 [2]. The cross section for the above reaction was measured to be $6.3 \cdot 10^{-44}\,\mathrm{cm}^2$, in good agreement with the predicted value.

The absence of an electromagnetic decay of the muon

$$\mu \rightarrow e\gamma$$

was a strong indication for the existence of a second neutrino species associated with the muon. The muon neutrino has first been observed in 1962 using a 10-ton spark chamber detector at Brookhaven in the first accelerator neutrino beam ever built [3]. Muon neutrinos were identified as the neutral products of pion decays which would interact in the detector producing single long tracks, identified as muon trajectories.

The DONUT experiment at Fermilab has made the first direct observation of tau neutrino CC interactions [4]. The delay of the experimental confirmation of the existence of the last fermion in the SM, has been due to the difficulty in identifying the τ lepton production in CC interactions using electronic detectors. Thanks to recent advances in emulsion technology, which is adopted by DONUT, it is feasible to reconstruct the τ track, as well as its production and decay vertices. The ν_τ CC interactions can then be identified applying topological criteria at the region around the neutrino interaction vertex.

The V-A structure of the weak charged current is an immediate consequence of maximal parity violation in weak decays [5]. This *left-right symmetry* violation was experimentally demonstrated by measuring the orientation of the electrons emitted in decays of polarized beta unstable nuclei [6] and by studying the kinematics of the pion and muon decays [7, 8].

The neutral electric charge of neutrinos makes them distinctive among the rest of the fermions in the SM in that their weak neutral current couplings manifest a pure V-A structure. Recall that a weak neutral current has the general form

$$J_\mu^{NC} = \overline{\psi}_f \gamma_\mu \frac{1}{2}(c_V^f - c_A^f \gamma^5)\psi_f \tag{2.2}$$

where,

Q_f the electric charge of the fermion, and in the SM:

$c_V^f = T_f^3 - 2\,Q_f\,\sin^2\theta_W$,
$c_A^f = T_f^3$,
T_f^3 is the third component of the fermion's weak isospin, and
θ_W the electroweak mixing angle.

In the case of neutrinos since $Q = 0$, $c_V = c_A = T^3 = \frac{1}{2}$ the neutral current in (2.2) becomes

$$J_\mu^{NC} = \overline{\psi}_\nu \gamma_\mu \frac{1}{4}(1-\gamma^5)\psi_\nu \tag{2.3}$$

This simple form of the neutral current (NC) interaction had various useful consequences. First of all neutrino beams could be used to reveal the structure of the Z^0 coupling with the rest of the fermions. A typical example is the CHARM II experiment at CERN. Measurements of the c_V and c_A constants for the electron have been made using intense ν_μ and $\overline{\nu}_\mu$ beams, contributing to the calculation of $\sin^2\theta_W$ [9].

Moreover, the simple form of (2.3) allows a precise calculation of the Z^0 decay width as a function of the number of light neutrino generations. The experiments at LEP have measured exactly three light left handed neutrino flavours [10].

Right-handed neutrinos are $SU(2)_L$ weak isospin singlets ($T^3 = 0$) and electrically neutral ($Q = 0$). Therefore, they do not interact neither electromagnetically, nor via Z^0 exchange ($c_V = c_A = T^3 = 0$).

In the simplest form of the SM, all neutrino species are assumed to be massless. This has the following immediate consequences:

- The existence of right-handed neutrinos becomes irrelevant to the SM, since there is no way to interact with other particles. Therefore, there are only left-handed neutrinos and their CP counterparts, right-handed antineutrinos.
- Muon and tau neutrinos are stable, contrary to the corresponding charged leptons.
- Neutrinos have zero magnetic dipole moment, consistent with the current experimental data [11, 12], while if the CPT theorem holds, their electric dipole moment vanishes as well.

2.1.2 Neutrino masses and their importance

Massless neutrinos are theoretically challenging to explain. Supersymmetry was originally invoked to resolve this issue, even though it failed. There is no justification in the SM why neutrinos should be massless. Massive neutrinos would have a major impact on cosmology and particle physics.

The astrophysical and cosmological arguments for the existence of massive neutrinos are related to the dark matter problem. The visible mass of the universe accounts only for a density $\rho_v = 0.1 \cdot \rho_c$, where ρ_c is the critical energy density which leads to a flat universe. Present cosmological knowledge [18] favours a combination of Hot (HDM) and Cold Dark

Matter (CDM) with $\Omega_{Cold} \approx 0.7$ and $\Omega_{Hot} \approx 0.2$, where $\Omega = \frac{\rho}{\rho_c}$. If HDM would be fully attributed to neutrinos, this would imply that $m_\nu \approx 5\,\mathrm{eV}$.

Massive neutrinos are interesting for particle physics, because they require extensions of the SM. The Lagrangian of the electroweak interaction has to include an additional term which has to be hermitian and Lorentz invariant. It can be shown that there are two linear combinations of the neutrino field operator ν obeying these requirements:

- Dirac mass terms, conserving the lepton number, which appear in the Lagrangian as

  $$-m_D(\overline{\nu}_L\nu_R + \overline{\nu}_R\nu_L)$$

  where $\nu_L = \frac{1}{2}(1+\gamma_5)\nu$ and $\nu_L = \frac{1}{2}(1-\gamma_5)\nu$. In this case Lorentz invariance requires the existence of right handed neutrinos as well.

- Majorana mass terms, which violate lepton number conservation, of the form

  $$-\frac{1}{2}m_M^L(\overline{\nu}_L\nu_L^c + \overline{\nu^c}_L\nu_L) - \frac{1}{2}m_M^R(\overline{\nu}_R\nu_R^c + \overline{\nu^c}_R\nu_R)$$

  where, $\nu^c = S_c\nu^{\dagger T}$ is the charge-conjugate field operator, with S_c defined by the relation $S_c^{-1}\gamma^\mu S_c = -\gamma^{\mu *}$.

Majorana particles, i.e. particles with such mass terms, are identical to their antiparticles [13]. From all the elementary fermions of the SM, only neutrinos can be Majorana particles. For charged fermions charge conservation would be violated.

For both Dirac and Majorana particles, their physical state can be identified by the helicity quantum number. Even though it is widely believed that for massless neutrinos it is practically impossible to determine whether they are Majorana or Dirac particles, it can been shown that the structure of the weak neutral current is different for the two cases. In particular, Majorana neutrinos miss the vector coupling in neutral current interactions, which is not the case for Dirac neutrinos [14]. According to this reference, the data from the CHARM II experiment show a discrepancy in the value of the vector coupling constant of more than 2 σ with respect to the zero value which is predicted by the theory under the assumption that neutrinos are pure Majorana particles.

However, most Grand Unified Theories, require that the ν_R exists and lepton number is not conserved, such that both Dirac and Majorana mass terms are naturally present in the Lagrangian. It has been pointed out that the fact that neutrino masses are so small, could be caused by the extremely large value of the Grand Unification scale M_{GUT} [15, 16]. Combining then both Dirac and Majorana mass terms, and considering one flavour for the moment, the Lagrangian can be written in the following matrix notation:

$$L_{mass} = -\frac{1}{2}\begin{pmatrix} \overline{\nu}_L & \overline{\nu^c}_R \end{pmatrix}\mathcal{M}\begin{pmatrix} \nu_L^c \\ \nu_R \end{pmatrix} + h.c. \tag{2.4}$$

where $\mathcal{M}$ is the *see-saw* matrix

$$\mathcal{M} = \begin{pmatrix} m_M^L & m_D \\ m_D & m_M^R \end{pmatrix} \tag{2.5}$$

The relationship among the elements of $\mathcal{M}$ define a particular physics scenario.

A popular choice is to assume hierarchy in the elements of $\mathcal{M}$:

$$m_M^R = M \gg m_D \gg m_M^L = \mu \tag{2.6}$$

with μ negligible or zero. Then one ends up with a heavy particle of mass

$$m_N \approx M \tag{2.7}$$

consisting mainly of the sterile (non-interacting) ν_R and a light one with mass

$$m_\nu \approx \left| \mu - \frac{m_D^2}{M} \right| \tag{2.8}$$

which consists mainly of the active (interacting) ν_L. When one generalizes these arguments to three flavours, then m_M^L, m_M^R and m_D become 3×3 complex matrices. Assuming that m_D is of the order of the mass of the charged lepton, or a quark of the same family, small neutrino masses are generated satisfying the hierarchy

$$m_{\nu_e} < m_{\nu_\mu} < m_{\nu_\tau} \tag{2.9}$$

with the absolute value of the masses being mainly determined by the value of M.

Theory dependent limits on an effective Majorana neutrino mass can be obtained from experimental data on the lifetime limits on neutrinoless double beta nuclear decay. This process is forbidden for Dirac neutrinos because lepton number conservation is violated. It requires helicity flip of the intermediate neutrino, which may happen only for a massive neutrino with a probability proportional to the square of its mass [17].

2.1.3 Direct mass measurements

When attempting to directly measure neutrino masses, it is always assumed that the mass eigenstates ν_1, ν_2 and ν_3 are composed mainly of the weak interaction (flavour) eigenstates ν_e, ν_μ and ν_τ respectively.

The square of the electron neutrino mass $m_{\nu_1}^2$ is measured in tritium beta decay experiments by fitting the shape of the electron spectrum near the endpoint. The current best limits comes from the Mainz experiment. They report an upper limit on the neutrino mass of $\mu_{\nu_e} < 2.8\, eV$ at 95% CL [20]. The same group intend to perform an experiment in the future based on a larger spectrometer which will probe electron neutrino masses below 1 eV.

The ν_2 mass is constrained by measuring the muon momentum in pion decays:

$$\pi^+ \rightarrow \mu^+ + \nu_\mu$$

Combined with the best knowledge of the muon and pion masses as well as four momentum conservation, this yields the square of the ν_2 mass. The present best limit is $m_{\nu_\mu} < 170\, keV$ at 90% CL [21].

Finally, upper limits on m_{ν_3} are obtained studying the kinematics of the hadronic τ-decays. The analysis, which has been last optimized by the OPAL Collaboration, involves the calculation of the hadronic invariant mass and fitting the energy spectrum of the hadronic system [22]. The best result comes from ALEPH Collaboration quoting an upper limit for the tau neutrino mass $m_{\nu_3} < 18.2\,MeV$ at 95% CL [23].

For the moment it seems difficult to explore absolute mass values significantly below the current limits performing direct measurements. The upper limits on the mass differences that can be established from the above results practically coincide with the actual mass values. This happens because the masses of the corresponding leptons follow the pattern $m_e \ll m_\mu \ll m_\tau$. As it is shown in the following sections, neutrino oscillations provide the means to probe mass differences which may be even smaller than the lightest neutrino mass. Potential positive experimental results on neutrino oscillation searches may therefore establish upper limits on the absolute mass values of the heavy neutrinos very close to the upper limit on the mass of the lightest neutrino.

2.2 Lepton mixing angles: the two-flavour case

If neutrinos are massive, the weak eigenstates ν_e, ν_μ, and ν_τ do not necessarily coincide with the physical (mass) eigenstates ν_1, ν_2, and ν_3. It can be shown [24] that the most general introduction of a neutrino mass term in the Lagrangian leads to neutrino mixing, regardless whether Dirac or Majorana states are involved. Flavour is no longer a conserved property of the neutrino, hence flavour transitions may take place in vacuum. The notion of neutrino oscillations, was originally proposed by Pontecorvo in 1957 [25]. In this reference he is discussing neutrino – antineutrino transitions, since the muon and tau neutrino had not been discovered then. His ideas apply though similarly to flavour oscillations as well [26].

To illustrate the effect of mixing in the propagation of a massive neutrino in spacetime, we consider for the moment the simplest case assuming two flavours and two mass eigenstates. Despite its simplicity, this is the working assumption underlying the data analysis carried out at most oscillation experiments. One of the reasons is that the results of such an analysis can be presented in a single plot of a two-dimensional oscillation parameter space.

In the more general situation with three flavours and three masses, this could still provide a good description in the limiting case where the third flavour consisting only of the third mass eigenstate, is completely decoupled from the other two.

2.2.1 Vacuum oscillations

The flavour eigenstates ν_l and $\nu_{l'}$ are a superposition of the mass eigenstates ν_1 and ν_2 according to a unitary mixing matrix:

$$\begin{pmatrix} \nu_l \\ \nu_{l'} \end{pmatrix} = \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \end{pmatrix} \tag{2.10}$$

where θ is the mixing angle. In principle, there is an additional phase in the mixing matrix. It has been omitted because in the two-flavour case it does not appear in the resulting transition probability.

At production (for example in a weak decay) neutrinos are left in a definite flavour eigenstate. Assuming that a neutrino is produced in a weak decay as ν_l with momentum p, its wavefunction at production ($t = 0$) is:

$$|\nu_l\rangle = \cos\theta\,|\nu_1\rangle + \sin\theta\,|\nu_2\rangle \tag{2.11}$$

Working with the *natural* system of units, where $c = \hbar = 1$, the evolution of the wavefunction in time will be

$$|\nu_l(t)\rangle = \cos\theta e^{-iE_1 t}\,|\nu_1\rangle + \sin\theta e^{-iE_2 t}\,|\nu_2\rangle \tag{2.12}$$

where

$$E_i = \sqrt{p^2 + m_i^2} \approx p + \frac{m_i^2}{2p} \tag{2.13}$$

assuming that neutrinos are highly relativistic. From (2.13) it follows that

$$E_1 - E_2 \approx \frac{\Delta m^2}{2p} \approx \frac{\Delta m^2}{2E} \tag{2.14}$$

where $\Delta m^2 = m_1^2 - m_2^2$.

The probability that the neutrino will still be found as ν_l after travelling some distance $L = ct$, is given by:

$$P_{\nu_l \to \nu_l}(L) = |\langle \nu_l(0)|\nu_l(t)\rangle|^2 \tag{2.15}$$

Substituting the expressions from (2.11), (2.12), and (2.13) in (2.15), and using (2.14), one obtains

$$P_{\nu_l \to \nu_l}(L) \approx 1 - \sin^2 2\theta \sin^2\left(\frac{\Delta m^2 L}{4E}\right) \tag{2.16}$$

Similarly, the transition probability $P_{\nu_l \to \nu_{l'}}$ for the neutrino to be detected as $\nu_{l'}$ rather than ν_l, which is commonly quoted as the *oscillation probability*, is given by:

$$P_{\nu_l \to \nu_{l'}}(L) = \sin^2 2\theta \sin^2\left(\frac{\Delta m^2 L}{4E}\right) \tag{2.17}$$

From the above we can define the *oscillation length*

$$L_{osc} \equiv 4\pi \frac{E}{\Delta m^2} \tag{2.18}$$

which is the length needed to go from ν_l to $\nu_{l'}$ and back to ν_l. It is customary to express Δm^2 in eV^2, L in km and E in GeV. The oscillation probability and oscillation length are given then respectively by

$$P_{\nu_l \to \nu_{l'}}(L) = \sin^2 2\theta \sin^2 \left(1.267 \frac{\Delta m^2 L}{E}\right) \tag{2.19}$$

and

$$L_{osc} = \frac{\pi E}{1.267 \Delta m^2}. \tag{2.20}$$

It is clear that in order to observe the neutrino oscillation pattern, one has to place a detector preferably half an oscillation length away from the neutrino source. On the other hand, if neutrinos are detected many oscillation lengths away from their production point, it becomes experimentally very difficult to determine precisely the exact value of the argument in the oscillatory sine term, mainly due to the finite energy resolution of a typical detector. In this case the oscillation pattern is lost, hence the probabilities in (2.16) and (2.17) become equal to their time averaged values

$$P_{\nu_l \to \nu_l} = 1 - \frac{1}{2} \sin^2 2\theta \tag{2.21}$$

and

$$P_{\nu_l \to \nu_{l'}} = \frac{1}{2} \sin^2 2\theta \tag{2.22}$$

respectively.

2.2.2 Neutrino oscillations and the uncertainty principle

If both the energy and the momentum of the neutrino are measured sufficiently accurately that the error in $m_\nu = \sqrt{E_\nu^2 - p_\nu^2}$ is smaller than the mass difference of the two neutrinos, the virtual transition into the other neutrino state cannot take place. In this case the detector will be observing mass eigenstates with probability of occurrence given by the incoherent sum of the two weak eigenstates, namely:

$$P_{\nu_l \to \nu_l} = \cos^4 \theta + \sin^4 \theta \tag{2.23}$$

which is independent of L. It may be noticed that the above value is identical to the time averaged probability given by (2.21). The same applies for $P_{\nu_l \to \nu_{l'}}$ as well.

Although this statement sounds paradoxical, this *puzzle* can be resolved if one considers the uncertainty principle of quantum mechanics. To keep the error of m_ν^2 ($= \delta m_\nu^2$) smaller than Δm^2, one must measure p_ν with a precision δp_ν that satisfies:

$$\Delta m^2 \gg \delta m_\nu^2 = 2 p_\nu \delta p_\nu \tag{2.24}$$

This means that the position becomes undetermined to the order of one oscillation length

$$\delta L > \frac{1}{\delta p_\nu} = \frac{2p_\nu}{\Delta m^2} = \frac{1}{2\pi} L_{osc} \tag{2.25}$$

and hence the oscillation pattern is washed out by this uncertainty.

The problems related to the uncertainty principle are usually treated using a wavepacket formalism [27]. The main points are summarized below.

The transition probability is derived assuming that the wavefunctions of the neutrino which is produced and propagated in time, and the one which is actually detected, are described by gaussian wavepackets of the same central momentum for the corresponding mass eigenstates.

The overall uncertainty in position with respect to the source, σ_x, is a quadratic sum of the coherence widths at production and detection. It is dominated by the maximum between the spatial or temporal coherence widths at the production or the detection process. The coherence length is then defined as

$$L_{coh} \equiv \frac{4\sqrt{2}\sigma_x E^2}{|\Delta m^2|} = \frac{L_{osc} E \sqrt{2}}{\pi} \tag{2.26}$$

It can be shown that neutrino oscillations can occur as long as

$$\sigma_x \ll L_{osc} \tag{2.27}$$

and

$$L \ll L_{coh} \tag{2.28}$$

The first and more stringent condition requires that a detector should not have an infinite energy resolution, the neutrino beam should not be monochromatic, or in a more realistic situation, the dimension of a neutrino source should not be comparable to the oscillation length.

On the other hand, the second condition becomes important, when σ_x, and hence the coherence length, is small. If a neutrino beam is originally localized then after some time the wavepackets of the two mass eigenstates, which move with different group velocities, will be separated in space by more than a coherence width σ_x and will not interfere any more. Since time information is lost a detector will pick up the mass eigenstates incoherently. Therefore the neutrino flavours will appear with the time averaged probabilities (2.21) and (2.22).

However, in most of the cases the coherence length is constrained by the energy spread of the neutrino beam. The oscillation pattern is visible as long as

$$L \lesssim \frac{L_{osc}}{2} \left(\frac{\delta E}{E} \right)^{-1} \tag{2.29}$$

This has the consequence that for any real neutrino beam the oscillation pattern will disappear just after a few oscillation lengths.

2.2.3 Matter enhanced oscillations

When neutrinos undergo elastic scattering with charged leptons, they interact differently according to the flavour of the charged lepton. In particular the process

$$\nu_{l'} + l^- \rightarrow \nu_{l'} + l^-$$

takes place only via a neutral current (NC) interaction if $l' \neq l$, while if $l' = l$ the process can be realized by both Z^0 and W^+ exchange, so that the two transition amplitudes M^{CC} and M^{NC} have to be added, when calculating the scattering cross-section.

When neutrinos propagate in dense matter, they may undergo elastic interactions at forward angles. For macroscopic regions, this leads to coherent effects in their propagation, which can be described in terms of the potential energy

$$V = V_Z + V_W \tag{2.30}$$

where V_Z and V_W are the potential energies due to Z and W exchange respectively.

Since there are no muons or taus in ordinary matter, the V_W term exists only for electron neutrinos. Assuming that electrons in matter are non-relativistic and isotropic, for constant electron density N_e, the potential V_W is

$$V_W = \sqrt{2} G_F N_e \tag{2.31}$$

Wolfenstein, Mikheyev, and Smirnov [28, 29] pointed out that the effect of coherent forward scattering can change the oscillation pattern of neutrinos travelling through matter. With these studies they originally intended to explain the observed deficit of solar neutrino interactions at the earth (described in 2.5.1) in terms of $\nu_e \rightarrow \nu_x$ oscillations, where $x = \mu, \tau$.

For such a mixing, with given Δm^2 and $\sin^2 2\theta$, the observed oscillation pattern in matter could be described by the effective values of

$$\Delta m^2_{\text{eff}} = \Delta m^2 \sqrt{(a - \cos 2\theta)^2 + \sin^2 2\theta} \tag{2.32}$$

and

$$\sin^2 2\theta_{\text{eff}} = \frac{\sin^2 2\theta}{(\cos 2\theta - a)^2 + \sin^2 2\theta} \tag{2.33}$$

where

$$a = 2\sqrt{2} G_F N_e \frac{E}{\Delta m^2} \tag{2.34}$$

It is obvious that even small mixing angles can lead to an effective maximal mixing in matter if the condition $a = \cos 2\theta$ is satisfied. This resonance behaviour, the so-called MSW effect, may take place if N_e has the appropriate value. In this case the observed Δm^2_{eff} is smaller than the vacuum Δm^2.

2.3 Exploring the parameter space and visualizing experimental results

Any neutrino experiment searching for oscillations of the type $\nu_l \rightarrow \nu_{l'}$, establishes constraints on the possible allowed values of the oscillation probability $P_{\nu_l \rightarrow \nu_{l'}}$. The interpretation of this experimental observable in terms of physics parameters such as neutrino masses and mixing angles, depends on the underlying theoretical model and predefined experimental constants, namely the neutrino energy and the flight length.

For the simplest case of two-flavour mixing, experimental results are presented as excluded or favoured regions in the $\sin^2 2\theta - \Delta m^2$ plane. According to the oscillation channel and the region in the parameter space which is to be explored, different neutrino sources and detection techniques have to be adopted for an efficient oscillation search.

2.3.1 Neutrino sources

The parameter space explorable by an oscillation experiment is always predefined mainly by the mass of the detector and the neutrino source. Recalling equation (2.19), the oscillation probability is maximized at mass squared differences around the optimum value

$$\Delta m^2 \approx \frac{E}{L} \tag{2.35}$$

Therefore, given the mass and the position of the detector, the neutrino energy constrains the region of Δm^2 values which can be probed.

Through the energy spectrum and the flavour composition, the neutrino source also determines whether specific oscillation channels are detectable and which detection techniques can be applied. For example, let us consider a $\nu_l \rightarrow \nu_{l'}$ oscillation search, through the direct observation of the l' lepton in $\nu_{l'}$ CC interactions. This search is only feasible as long as the neutrino energy is sufficient to produce the massive charged lepton l' and the intrinsic contamination of $\nu_{l'}$s in the beam is relatively low.

Neutrinos used for oscillation experiments are produced in the sun, the atmosphere, in nuclear reactors and at accelerators.

Solar neutrinos are produced in the various nuclear reaction chains in the sun. The theoretical Standard Solar Model (SSM) [30] predicts an energy spectrum shown in Fig 2.1. From all the contributions to the spectrum, the $p\ p$ flux is estimated to be known best ($\sim 2\%$), because these neutrinos are produced in the same reactions as the photons forming the main luminosity flux. The ^8B flux, is extremely temperature dependent ($\phi(^8B) \propto T^a$, where $a \approx 18$), but recent data on helio-seismology have excluded large variations of the temperature, and validated to a large extent the SSM.

In Fig 2.1 the threshold energies are also shown for various types of solar neutrino experiments (discussed in 2.5.1). Solar neutrinos, being all of the electron type, can be used for $\nu_e \rightarrow \nu_x$ oscillation searches, where x can be any other lepton type. However, their low energy does not allow the production of the corresponding charged lepton in charged current interactions. Therefore, oscillations manifest themselves only through the deficit of the ν_e CC and the relative enhancement of the NC interaction rate.

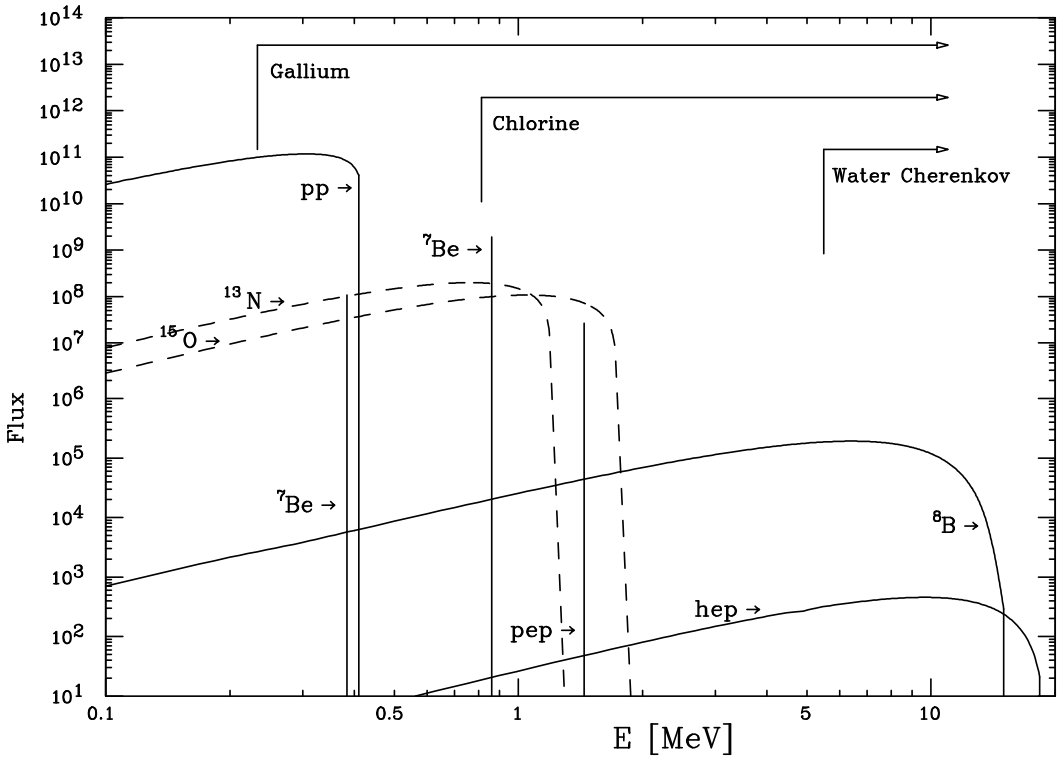

Fig. 2.1: Predicted spectrum of solar neutrinos. Neutrino fluxes from continuum sources (such as $p\ p$ and ^8B) are given in $\mathrm{cm}^{-2}\mathrm{s}^{-1}\mathrm{MeV}^{-1}$, while those from discrete sources (such as ^7Be) are given in $\mathrm{cm}^{-2}\mathrm{s}^{-1}$.

On the other hand, the long baseline defined by the distance between the earth and the sun and the low energy of the solar neutrinos, allows experiments to probe mass squared differences down to $10^{-11}\,\mathrm{eV}^2$.

Electron antineutrinos of high intensity are produced at nuclear reactors. Like the solar neutrinos, reactor neutrinos have energy spectra extending up to several MeV. Due to the much shorter baseline, the lowest mass squared difference values that reactor neutrino experiments have probed so far are of order $10^{-3}\,\mathrm{eV}^2$. The neutrino flux is known with high accuracy ($\sim 2\%$) [31], allowing for precise calculations of the oscillation probability.

Higher energy neutrinos are generated in the atmosphere. Cosmic rays, mainly protons and light nuclei, interacting with atmospheric nuclei produce a large number of secondary mesons:

$$\mathrm{CR} + \text{“Air”} \rightarrow \pi/K + X.$$

Because of the low atmospheric density, most mesons decay before interacting:

$$\pi/K^{\pm} \rightarrow \mu^{\pm} + \nu_{\mu}(\overline{\nu}_{\mu})$$

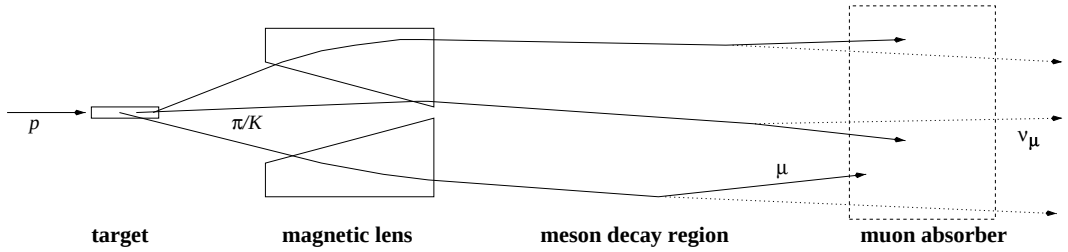

Fig. 2.2: Production of a neutrino beam at a proton accelerator

and most muons, mainly at low energies (few GeV), decay before reaching the ground:

$$\mu^{\pm} \rightarrow e^{\pm} + \nu_e(\overline{\nu}_e) + \overline{\nu}_\mu(\nu_\mu).$$

The absolute flux of atmospheric neutrinos is known with an error $\sim 20 - 30\%$. The uncertainties originate mainly from the poor knowledge of the primary cosmic ray flux at the boundaries of the atmosphere[32], and the hadron production cross-sections. It should be noted that neither the neutrino interaction cross-section is known well at low energies ($E_\nu < 10\,\mathrm{GeV}$). On the contrary, the ratio $\phi(\nu_\mu)/\phi(\nu_e)$ is better known. Ignoring the distinction between neutrinos and antineutrinos and assuming that all muons have decayed before reaching the altitude level of the neutrino detector, one would expect that $\phi(\nu_\mu)/\phi(\nu_e) \approx 2$. It should be noted that this is only the case for low neutrino energies. More generally this ratio is energy dependent. All oscillation searches using atmospheric neutrinos are performed looking for possible deviations of $\phi(\nu_\mu)/\phi(\nu_e)$ from the expected values.

The energy spectrum of the interacting neutrinos in a detector at the earth's surface peaks around 1 GeV. Very few neutrinos have sufficient energy to produce taus. Therefore, it is difficult to search for oscillations including the ν_τ in a direct way.

Proton accelerators are used to create controlled, and intense neutrino beams, with a tunable mean energy. They provide the best means to perform oscillation searches of the type $\nu_\mu \rightarrow \nu_x$, where x can be any flavour: e or τ.

The principle underlying the production of neutrino beams at proton accelerators is illustrated in Fig. 2.2. High energy protons are directed towards a target where charged pions and kaons are produced. These are focused by a system of magnetic lenses to form parallel beams which are left to decay in a long decay pipe. The beam of particles following a muon absorber consists almost exclusively of muon neutrinos.

A small contamination of electron neutrinos ($\sim 1 - 2\%$) is always present in such a neutrino beam. They are produced mainly by kaons decaying into electrons and by muon decay. If toroidal magnets are used for the focusing system, the secondary mesons of a specific charge are focused, while the others are diverted. The relative contamination of antineutrinos in the neutrino beam (or vice versa) is then highly suppressed.

The energy spectrum of the neutrinos produced at accelerators is determined by the proton energy, the target configuration, and the magnetic lenses. The spectrum of the

neutrinos which are actually detected depends also on the position and the size of the detector. For a $\nu_\mu \to \nu_\tau$ oscillation search the mean energy is usually chosen to be higher than $\sim 10\,\mathrm{GeV}$ to overcome the tau production threshold, while for the $\nu_\mu \to \nu_e$ channel the mean energy is usually lower.

Neutrinos of lower energy, up to ∼50 MeV, are produced at proton dumps, which are thick enough to absorb all produced pions and muons. In this case the proton energy does not exceed 1 GeV and the secondary mesons usually decay at rest inside the target. Since negatively charged mesons and muons are captured by the target nuclei there is only a small number of $\overline{\nu}_e$s in the resulting beam. Positively charged pions and muons decay before they are absorbed. The low $\overline{\nu}_e$ contamination allows for $\overline{\nu}_\mu \to \overline{\nu}_e$ oscillation searches to be performed at such facilities.

The values of Δm^2 reachable by accelerator neutrino experiments, depending on the baseline and the mean energy, can be as high as $\sim 10\,\mathrm{eV}^2$ or as low as $10^{-3}\,\mathrm{eV}^2$. According to the mean value of E/L, which is proportional to Δm^2, neutrino oscillation experiments are classified as *Short Baseline* (SBL), *Medium Baseline* (MBL), and *Long Baseline* (LBL), where $E/L \gg 1$, $E/L \approx 1$, and $E/L \ll 1$ respectively.

Accelerator experiments searching for $\nu_\mu \to \nu_\tau$ oscillations, the mean neutrino energy is typically within the range 10-20 GeV. The mean flight length L for SBL and LBL experiments is therefore in the order of 10^3 m and 10^6 m respectively.

2.3.2 Oscillation searches and parameter space plots

appearance experiments

An experiment searching for $\nu_l \to \nu_{l'}$ oscillations which utilizes a detector sensitive to $\nu_{l'}$ interactions is called an *appearance* experiment.

Let us consider the simplest case where a detector is exposed in a neutrino beam consisting exclusively of ν_ls, and the expected number of background events is $N_b \ll 1$. This is usually the aim of experiments performing $\nu_\mu \to \nu_\tau$ searches at accelerators.

The ratio of the observed $\nu_{l'}$ CC interactions over those of ν_l can be calculated as

$$\frac{N_{l'}}{N_l} = BR \frac{\langle \epsilon_{l'} \rangle}{\langle \epsilon_l \rangle} \frac{\langle \sigma_{l'} \rangle}{\langle \sigma_l \rangle} \frac{\langle P_{\nu_l \to \nu_{l'}} \rangle}{1 - \langle P_{\nu_l \to \nu_{l'}} \rangle} \tag{2.36}$$

where BR is the branching ratio of the l' decay channel to which the experiment is sensitive. If the l' lepton is stable, then $BR = 1$. The quantities $\langle \epsilon_i \rangle$ and $\langle \sigma_i \rangle$, $i = l, l'$ are the overall detection efficiencies and production cross-sections respectively, integrated over the energy spectrum of the beam. If more than one l' decay channels are analysed, then the ratio $\frac{N_{l'}}{N_l}$ is a sum of terms as the one in (2.36) each of them with different branching ratio and detection efficiency for the decaying l'.

The mean oscillation probability, which in the case of two-flavour mixing is

$$\langle P_{l \to l'} \rangle = sin^2 2\theta \langle sin^2 \frac{1.267 \Delta m^2 L}{E} \rangle \tag{2.37}$$

is averaged over all possible L/E values of the incoming neutrinos. Solving (2.36) for $\langle P_{\nu_l \to \nu_{l'}} \rangle$, one obtains for small $\langle P_{\nu_l \to \nu_{l'}} \rangle$

$$\langle P_{\nu_l \to \nu_{l'}} \rangle = \frac{N_{l'}}{N_{l',\,max}} \tag{2.38}$$

where

$$N_{l',\,max} = N_l \cdot BR \cdot \frac{\langle \epsilon_{l'} \rangle}{\langle \epsilon_l \rangle} \frac{\langle \sigma_{l'} \rangle}{\langle \sigma_l \rangle} \tag{2.39}$$

is the maximum possible number of $\nu_{l'}$ charged current events that could be detected.

In general, appearance experiments search for a small number of *signal* $\nu_{l'}$ events among the accumulated statistics of ν_l interactions. Then $N_{l',\,max}$ is a random variable normally distributed with a mean value proportional to the number of accumulated events and a standard deviation determined mainly by the overall systematics of the experiment. On the other hand, the background subtracted signal $N_{l'}$ is Poisson distributed with a mean value determined by the product of the mean oscillation probability $\langle P_{\nu_l \to \nu_{l'}} \rangle$ and the normalization factor $N_{l',\,max}$.

Given any experimental outcome, a confidence interval for $N_{l'}$ can be constructed taking into account the expected background rate and the overall uncertainty in $N_{l',\,max}$ [33]. The bounds on $N_{l'}$ are translated then through (2.38) as a confidence interval for $\langle P_{\nu_l \to \nu_{l'}} \rangle$. Until recently, an upper limit on $N_{l'}$ would be derived using the *Bayesian* interpretation of the probability and a uniform prior in the true values of the physics parameters. Applying this statistics in the simplest case of a background free experiment with zero observed events yields an upper limit on $N_{l'} = 2.3$ at 90% Confidence Level (C.L.).

Since 1998 the Particle Data Group (P.D.G.) [34] has adopted the Feldman-Cousins (F.C.) *Unified Approach* to the classical statistical analysis of small signals [35] which sticks to the strict definition of confidence intervals by Neyman [36]. The confidence intervals are constructed using an ordering principle based on likelihood ratios. The power of the F.C. approach is on the fact that experimental data can be presented in an objective way. It should be noted that since then there has been an increasing interest in the problem of establishing confidence limits in searches of new physics phenomena or particles which even led to a dedicated workshop on this subject [37].

In the F.C. approach the 90% C.L. upper limit for the signal corresponds to $N_{l'} = 2.44$ observed events. The corresponding result for the upper limit on the oscillation probability is then:

$$\langle P_{\nu_l \to \nu_{l'}} \rangle_{\text{U.L.}} = \frac{2.44}{N_{l',\,max}} \propto \frac{1}{N_l}\,. \tag{2.40}$$

In terms of the physical parameters Δm^2 and $\sin^2 2\theta$ this result means that a region in the $\sin^2 2\theta - \Delta \mathrm{m}^2$ plane is excluded. A typical *exclusion plot* is shown in Fig. 2.3. The exclusion curve (solid line) which separates the excluded from the allowed region at some C.L., is obtained by solving (2.37) for $\sin^2 2\theta$ at various values of Δm^2. The calculation

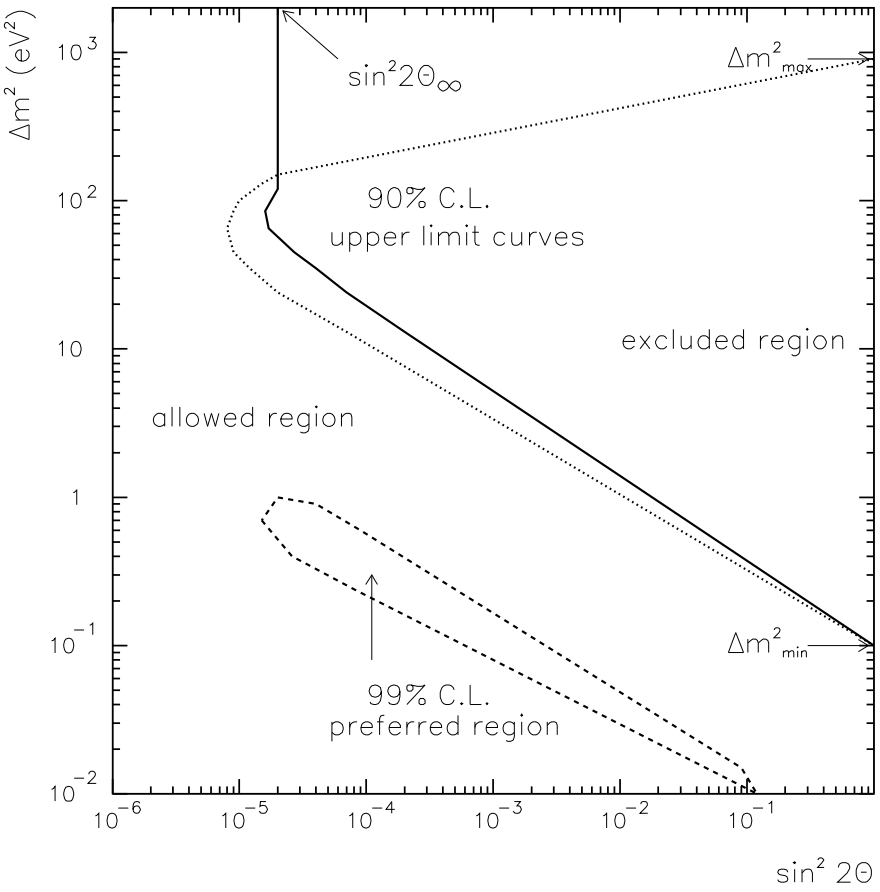

Fig. 2.3: Typical parameter plots from oscillation experiments. The solid and dotted curves correspond to 90% C.L. upper limit exclusion curves from single- and two-detector experiments respectively. The area closed by the dashed contour line is an example of the allowed region at 99% C.L. for an experiment claiming signal for oscillations.

takes into account the energy spread of the beam and the detector resolution as well: If the quantity $1.267 \cdot L/E$ is distributed normally around the central value b_0 with a standard deviation σ_b, then (2.37) yields:

$$\langle P_{l\to l'}\rangle = \frac{1}{2} sin^2 2\theta \left[1 - \cos(2b_0\Delta m^2)\cdot e^{-2\sigma_b^2(\Delta m^2)^2}\right] \tag{2.41}$$

As a consequence, the oscillatory pattern practically disappears after the first oscillation length. Therefore, for high Δm^2 values the mixing angle converges to the asymptotic value of

$$\sin^2 2\theta_\infty = 2\cdot\langle P_{\nu_l\to\nu_{l'}}\rangle_{\text{U.L.}} \tag{2.42}$$

On the other hand, at full mixing ($\sin^2 2\theta = 1$) the parameter Δm^2 reaches a minimum:

$$\Delta m^2_{min} = \left\langle \frac{E}{L} \right\rangle \cdot \frac{\sqrt{\langle P_{\nu_l \to \nu_{l'}} \rangle_{\text{U.L.}}}}{1.267}. \tag{2.43}$$

The existence of systematical uncertainty and background events forces the number of signal events $N_{l'}$ fluctuate stronger, so that its larger variance expands the confidence interval. Consequently, the upper limit on the oscillation probability $\langle P_{\nu_l \to \nu_{l'}} \rangle_{\text{U.L.}}$ is increased. The exclusion curve shown as a solid line in Fig. 2.3 is then displaced along the horizontal direction towards the right hand side of the plot.

If a small number of background events is expected, it is good practice together with the upper limit curve to present the *sensitivity curve* as well [34, 35]. This is the curve that an experiment would quote as its upper limit curve in case the expected number of background events would have been observed. The quotation of the sensitivity curve allows proper comparison of various experimental results. By definition, for background-free experiments the two curves coincide.

If the neutrino beam lacks a $\nu_{l'}$ component, the main source of background for appearance experiments is the misidentification of ν_l CC or ν NC as $\nu_{l'}$ CC interactions. On the other hand, the presence of a $\nu_{l'}$ component leads to an irreducible background rate. An experiment can search then for $\nu_l \to \nu_{l'}$ oscillations on a statistical basis: A signal is established if there is a statistically significant excess of observed $\nu_{l'}$ CC events above the expected background rate. This is a typical case for experiments performing $\nu_\mu(\overline{\nu}_\mu) \to \nu_e(\overline{\nu}_e)$ oscillation searches at accelerators.

The relative excess of $\nu_{l'}$ events is calculated as

$$\frac{N_{l'} - N^b_{l'}}{N_l} = \frac{\langle P_{\nu_l \to \nu_{l'}} \rangle \left(1 - \langle f(E) \rangle\right)}{1 - \left(1 - \langle f(E) \rangle\right) \langle P_{\nu_l \to \nu_{l'}} \rangle} \tag{2.44}$$

where $N^b_{l'}$ is the number estimated background events, and $\langle f(E) \rangle$ is the fraction of $\nu_{l'}$s in the beam with respect to the ν_ls, integrated over the energy spectrum of the latter.

The number of excess events, $N_{l'} - N^b_{l'}$, the *signal* in this case, is a random variable following a normal distribution with a mean value determined by the mean oscillation probability and a standard deviation $\sigma = \sqrt{N^b_{l'}}$, in case the probability is small. If the quantity $N_{l'} - N^b_{l'}$ is consistent with zero, a 90% C.L. upper limit on the mean probability can be set:

$$\langle P_{\nu_l \to \nu_{l'}} \rangle_{\text{U.L.}} \propto \frac{\sqrt{N^b_{l'}}}{N_l} \propto \frac{1}{\sqrt{N_l}} \tag{2.45}$$

where the last relation holds because the background rate is a constant fraction of the ν_l interaction rate.

The exclusion curve is then calculated in a way identical to the one for the background free appearance experiments. However, comparing the relations (2.40) and (2.45) it becomes clear why the presence of background decreases the sensitivity of an oscillation search. If these relations are combined with (2.42) and (2.43), it can be shown that

for background-free experiments the sensitivity in $\sin^2 2\theta$ and Δm^2 improves with N_l and $\sqrt{N_l}$, respectively, while in the presence of background it only improves with $\sqrt{N_l}$ and $\sqrt[4]{N_l}$ respectively.

A more interesting case appears when the lower boundary of the confidence interval for the number of signal events is greater than zero, i.e. an oscillation signal is observed. From a fit to the energy spectrum of the signal events the favoured oscillation parameters are then obtained. The preferred region in the $\sin^2 2\theta - \Delta\mathrm{m}^2$ plane is calculated as a two-dimensional confidence interval, usually at 99% C.L. (as shown with the dashed line in Fig. 2.3). The calculation of the preferred region involves the integration of the χ^2 probability around the point of the best fit. The contour of the confidence interval is determined by the requirement that the integrated χ^2 probability within this region has to be equal to the desired confidence level.

disappearance experiments

An experiment sensitive to CC interactions only of ν_l neutrinos, can search for oscillations of the type $\nu_l \to \nu_x$, where ν_x stands for any neutrino different from ν_l. This is achieved through the detection of possible deficit of the ν_l flux after some distance from the source, hence such experiments are defined as *disappearance experiments*.

Most commonly, it is attempted to measure the effective oscillation parameters of the generic channel $\nu_l \to \nu_x$, from the modulation of the ν_l energy spectrum, with respect to the one predicted in the absence of oscillations. An alternative way to establish oscillations is through the observation of relative enhancement of the *NC/CC ratio*. This type of search has the advantage that it does not rely on the accurate knowledge of the absolute neutrino flux.

Disappearance experiments can be used to favour or exclude the participation of ν_ls in an oscillation mode at specific regions of the parameter space. In case oscillations occur, they can measure the Δm^2 parameter from the modulation of the ν_l spectrum. In principle, such a measurement could be done by an appearance experiment as well, but the production threshold energy E_t of the l' lepton should be below the energy where the oscillation pattern would appear:

$$E_t \ll \frac{\Delta m^2}{L}$$

two-detector experiments

For both appearance and disappearance experiments, the best way to control background and systematics is by introducing a second detector, preferably close to the neutrino source. Oscillations can be probed with the measurement of the ratio of observed interaction rates for the *far* over the *near* detector. Some experiments present their results normalizing the

interaction rates to the ones expected from simulation programs:

$$r_{f/n} = \frac{\left(\frac{N_l^{obs}}{N_l^{exp}}\right)_{\text{far}}}{\left(\frac{N_l^{obs}}{N_l^{exp}}\right)_{\text{near}}} \tag{2.46}$$

which in absence of oscillations should be statistically consistent with unity.

If the detectors can efficiently discriminate between NC and CC interactions, the following *double ratio* can be used instead:

$$r_{f/n} = \frac{\left(\frac{NC}{CC}\right)_{\text{far}}}{\left(\frac{NC}{CC}\right)_{\text{near}}} \tag{2.47}$$

This double ratio can be obtained as function of the neutrino energy, once the missing energy distribution in NC events is unfolded. A deviation from a flat distribution is interpreted as manifestation of oscillations. The oscillation parameters can then be measured from the shape of the observed distribution.

One can also exploit the two-detector method in appearance experiments. The near detector serves for a precise determination of $\nu_{l'}$s in the beam. The quantity which is checked for deviation from zero due to oscillations is the difference

$$d_{f/n} = \left(\frac{N_{l'}}{N_l}\right)_{\text{far}} - \left(\frac{N_{l'}}{N_l}\right)_{\text{near}}. \tag{2.48}$$

When designing a two-detector experiment, care should be taken that both detectors are exposed to a neutrino beam of similar energy spectrum. Otherwise deviations of the relevant quantities from the expected values can not be fully attributed to oscillations. At an accelerator experiment, one of the ways to achieve this, is to install the far detector slightly off-centered from the beam axis. This prevents the spectrum from being dominated by the central beam component, which is usually harder because of the way that secondary hadrons are produced in the proton target.

The normalization of the results of the far detector with respect to the near one, may on one hand increase the sensitivity of oscillations around a region of Δm^2 values determined by the energy spectrum at the far detector and its position from the source, on the other hand the non-zero distance of the near detector from the source poses an upper limit on the highest explorable Δm^2 value. A typical exclusion curve obtained performing such an experiment is shown as a dotted line in Fig. 2.3, which is to be compared with the dashed one obtained by single detector experiments. The boundary Δm^2 values are inversely proportional to the distances from the source:

$$\Delta m^2_{min} \propto \frac{E}{L_{far}}, \qquad \Delta m^2_{max} \propto \frac{E}{L_{near}}. \tag{2.49}$$

In a two-detector experiment one could also treat one of the detectors as a standalone experiment, analysing only the data obtained with it. The explorable Δm^2 range can be

extended then to higher values. In this case though, due to the increased systematical uncertainties the sensitivity in small $\sin^2 2\theta$ values is worse.

Since two-detector experiments can be very selective for the Δm^2 parameter, they can be used to measure precisely the oscillation parameters indicated by another experiment that has already claimed a signal, as long as the Δm^2 is indicated.

2.3.3 Design criteria for an oscillation experiment

It has been mentioned already, that the explorable region of the parameter space is always predefined for every experiment by its baseline length and the neutrino energy spectrum. This holds especially for Δm^2, since the *nose* of the exclusion curve shown in Fig. 2.3 only moves along the $\sin^2 2\theta$ axis towards lower values of $\sin^2 2\theta$ with increasing statistics, as long as background is well under control. Its vertical position though always remains the same, at the optimum value of $\Delta m^2 \approx E/L$.

Solar and atmospheric neutrinos have a fixed, but broad energy spectrum. However, a single experiment is efficient in detecting neutrinos of energy within a relatively narrow band of the overall spectrum. Given the fact that the neutrino flight length is fixed, the Δm^2 region which can be explored in an oscillation experiment is fully determined by the detector.

There is a higher flexibility from the side of the neutrino source with artificially made neutrinos, produced at reactors and accelerators. The energy spectrum of reactor neutrinos is well known but not tunable. However, there is only little constraint on the baseline length. On the other hand, accelerator experiments obtain their flexibility from the energy spectrum, which is tunable as well.

As it has already been mentioned (p. 2.3.1) experiments using artificial neutrinos are traditionally classified as SBL or LBL according to the distance L of the detector from the source. The distinction of the two classes is made due because probing low values of one of the parameters $\sin^2 2\theta$ and Δm^2 is feasible only at experiments where the complementary parameter is kept at moderate values. This happens for the following reasons.

For any source, the neutrino flux and hence the accumulated statistics drops with the distance as L^{-2}, while the sensitivity in low Δm^2 improves with L. Recalling the behaviour of the sensitivity in the two parameters as a function of the statistics (p.19), one can summarize the dependence of the oscillation parameters on the experimental constants L and M_d (the detector target mass) in Table 2.1. For background-free experiments the

	No Background	Background
$\sin^2 2\theta$	$\propto L^2$, $\propto 1/M_d$	$\propto L$, $\propto 1/\sqrt{M_d}$
Δm^2	No L dependence, $\propto 1/\sqrt{M_d}$	$\propto 1/\sqrt{L}$, $\propto 1/\sqrt[4]{M_d}$

Tab. 2.1: Dependence of the oscillation parameters on the experimental constants.

sensitivity in Δm^2 is eventually independent of the baseline length L and the number of detected oscillated neutrinos depends only on the mass of the detector. For the more

realistic situations where some background rate is present, there is a clear gain with L, and very little with the total active mass M_d. On the contrary, $\sin^2 2\theta$ is in either case highly affected by the loss of flux due to the longer distance from the source.

As a consequence, short baseline experiments are best suited to probe low mixing angles at moderate mass differences. They are based on detectors with relatively small mass in order to easily achieve high background rejection power. The intensity of the neutrino flux compensates for the limited detector mass, and the search for small mixing angles is only limited by the background rate and the accumulated statistics.

On the other hand, long baseline experiments are set up to probe low mass differences. The achieved mixing angles are moderate, because the neutrino flux at large distances is not sufficient to generate a very large number of events. This can be compensated by the use of massive detectors, which need not have excellent background rejection capabilities. For constant detector mass the background rate itself is lower in this case since neutrinos at moderate mixing angles oscillate faster and therefore the signal-to-background ratio is higher.

If an experiment aims for discovery or confirmation of oscillations at a predefined narrow region of Δm^2, the relevant quantity to be maximized is the number of potentially detected oscillated neutrinos. For an accelerator experiment with a given baseline and financial resources this quantity is a convolution of the following factors:

- the accelerator duty cycle as a function of the peak proton energy,
- the maximum number of protons that the target can sustain during an accelerator spill as a function of the proton energy,
- the pion yield per proton on target and its angular dispersion as a function of the proton energy and the target structure,
- the focusing efficiency of the magnetic lenses as a function of pion energy and angular dispersion,
- the absorption rate of the pions when they traverse the material of the beam line elements as a function their energy and angular dispersion,
- the dimensions of the pion decay region. An extended region leads to more useful pion decays but also to an increase of the ν_e contamination from muon decays, and hence beam related irreducible background rate at the detector,
- the neutrino angular dispersion as a function of the parent pion energy,
- the $\sigma_{\nu_{l'}}$ cross-section as a function of the mean neutrino energy,
- the detector efficiency as a function of the $\nu_{l'}$ energy,
- the amount of target mass which can be increased only at the cost of background rejection power, and of course

- the oscillation probability as a function of the neutrino energy.

From the above it becomes clear that the optimization of an experiment for a parameter space region that one wishes to explore, depends on the intrinsic capabilities of the neutrino beam line.

If the aim of an experiment is the precise measurement of the oscillation parameters indicated by a previous experiment, one should first of all choose the E/L that matches exactly the proposed Δm^2. The effect of oscillations would then be maximally enhanced and clearly manifest in the neutrino spectrum. Therefore the experiment should be such that the detector is located half an oscillation length away from the source. Placing the detector further only deteriorates the measurement: The flux drops rapidly as $1/L^2$, while the oscillation probability $P_{\nu_l \to \nu_{l'}} \propto \sin^2 2\theta$ does no longer increase as L^2. Hence the danger of background events from accidental coincidences and cosmic rays increases. Depending on the detector resolution and the neutrino energy spread, the oscillation pattern may also start fading due to decoherence effects.

Finally, it should be noted that it is very difficult to establish from a single experiment precise knowledge about the physics parameters themselves. It should be remembered that the only experimental observable is the transition probability $P_{\nu_l \to \nu_{l'}}$ as function of the neutrino energy.

2.4 General treatment of lepton mixing

Two-flavour oscillation is only a specialization of the general case of lepton mixing. It has been considered so far for various reasons. It illustrates in the simplest way the oscillatory behaviour of neutrino propagation in spacetime. Moreover, experimental results are easy to represent in two-dimensional plots, as explained in 2.3.2.

Like in the quark sector, weak mixing includes in general all flavours. The mixing angle of the two-flavour case resembles the Cabbibo angle, which is used to describe the mixing between d and s weak quark flavours. The three known active lepton weak-interaction eigenstates $(\nu_e, \nu_\mu, \nu_\tau)$ are mixed with the physical eigenstates (ν_1, ν_2, ν_3) through a unitary matrix U, corresponding to the CKM matrix of the quark sector:

$$\nu_l = \sum_i U_{li} \nu_i \qquad \text{and} \qquad \nu_j = \sum_{l'} U^\dagger_{jl'} \nu_{l'} \tag{2.50}$$

where $l, l' = e, \mu, \tau$ and $i, j = 1, 2, 3$.

The calculation of the transition probability is performed similar to the two-flavour case. One obtains:

$$P_{\nu_l \to \nu_{l'}}(L) = \delta_{ll'} + 2Re \sum_{i>j} \left| U_{li} U^\dagger_{il'} U_{l'j} U^\dagger_{jl} \right| \left[\cos\left(\frac{m_j^2 - m_i^2}{2E} L - \phi_{ll'ij} \right) - \cos \phi_{ll'ij} \right] \tag{2.51}$$

where $\phi_{ll'ij}$ is the phase of the product $U_{li} U^\dagger_{il'} U_{l'j} U^\dagger_{jl}$.

As a consequence of the CPT theorem $P_{\nu_l \to \nu_{l'}} = P_{\overline{\nu}_{l'} \to \overline{\nu}_l}$, while under the assumption of CP invariance $P_{\nu_l \to \nu_{l'}} = P_{\overline{\nu}_l \to \overline{\nu}_{l'}}$. For the latter case, there is no CP-violating phase present in the mixing matrix U, so that all its elements are real. The oscillation probability then becomes

$$P_{\nu_l \to \nu_{l'}}(L) = \delta_{ll'} - 4 \sum_{i>j} \left| U_{li} U^{\dagger}_{il'} U_{l'j} U^{\dagger}_{jl} \right| \sin^2 \left[\frac{m_j^2 - m_i^2}{4E} L \right] \tag{2.52}$$

where one can immediately recognize the oscillatory terms, which appear as an immediate consequence of the non-degeneracy of the physical eigenstates. The partial oscillation lengths

$$L_{osc,ij} = \frac{\pi E}{1.267 \Delta m^2_{ji}} \tag{2.53}$$

have a form identical to the two-flavour case (Equation.2.20). The strength of each oscillatory term is determined by the elements of the mixing matrix.

The oscillation parameters in a 3-flavour mixing scheme are two independent Δm^2 values, three mixing angles, and a CP-violating phase. A convenient parameterization of the mixing matrix U is the one proposed in Ref. [38]:

$$U = \begin{pmatrix} c_1 c_3 & s_1 c_3 & s_3 e^{i\delta} \\ -s_1 c_2 - c_1 s_2 s_3 e^{i\delta} & c_1 c_2 - s_1 s_2 s_3 e^{i\delta} & s_2 c_3 \\ s_1 s_2 - c_1 c_2 s_3 e^{i\delta} & -c_1 s_2 - s_1 c_2 s_3 e^{i\delta} & c_2 c_3 \end{pmatrix} \tag{2.54}$$

where c_i stands for cosine and s_i for sine of the three mixing angles θ_i, and δ is the CP-violating phase.

If $\Delta m^2_1 \ll \Delta m^2_2$, which is the case for *One-Mass-Scale-Dominance* (OMSD) models [40], an experiment with properly chosen L/E is sensitive to oscillations with contribution only from the dominant mass difference. In this case a two-flavour analysis would result in an effective $\sin^2 2\theta$ value, which is just a combination of all three mixing angles, and a Δm^2 which coincides with the actual dominant mass difference.

Lepton mixing can be generalized further to include more than three physical neutrino eigenstates. This would be the case if one introduced a fourth, *sterile* (weakly inactive) neutrino. If the neutrino is a *pseudo-Dirac particle* [39] there is a fine mass splitting of each neutrino into two mass eigenstates, and $\nu \to \overline{\nu}$ transitions may also occur. A 6×6 mixing matrix would then be required.

2.5 Present experimental status

Recently, there have been experimental results which indicate the existence of oscillations. The most representative experiments that have contributed to the study of the neutrino puzzle are described.

2.5.1 Solar neutrino experiments

The first experimental hint for neutrino oscillations is about 30 years old. It comes from a discrepancy between the measured detection rate in the Homestake solar neutrino experiment[41] and the theoretical calculations [42]: a statistically significant deficit in the ν_e flux was observed. The detector consists of a tank filled with C_2Cl_4 (perchloroethylene). The total Cl mass amounts to 520 tons with 24% consisting of the ${}^{37}_{17}Cl$ isotope. Solar neutrinos (ν_e) induce the inverse β-decay reaction

$$ {}^{37}_{17}Cl + \nu_e \rightarrow {}^{37}_{18}Ar + e^- $$

which has an energy threshold of $E_\nu = 0.813$ MeV, just below the 8.62 MeV ${}^7Be$ line (see Fig. 2.1). Therefore the detected flux is mainly determined by ${}^7Be$ and ${}^8B$ reactions. The radioactive nuclei ${}^{37}_{18}Ar$, whose lifetime is ~ 34 days, decay through the electron capture reaction

$$ {}^{37}_{18}Ar + e^- \rightarrow {}^{37}_{17}Cl + \nu_e $$

Every few months the ${}^{37}_{18}Ar$ nuclei are extracted chemically from the tank and the decays of the nuclei are detected using a proportional counter. From the decay rate the neutrino interaction rate is deduced.

The GALLEX[43] and SAGE[44] experiments make use of gallium instead of chlorine. Solar neutrinos induce a β-decay reaction on ${}^{71}Ga$:

$$ {}^{71}_{31}Ga + \nu_e \rightarrow {}^{71}_{32}Ge + e^- $$

which has a threshold energy of $E_\nu = 0.233$ MeV. The experiments are therefore able to study neutrinos from the $p\ p$ reaction chain of the solar neutrino spectrum. The ${}^{71}_{32}Ge$ nuclei which have a lifetime of ~ 11 days, are chemically extracted and their decays

$$ {}^{71}_{32}Ge + e^- \rightarrow {}^{71}_{31}Ga + \nu_e $$

are detected in a way similar to the method in the chlorine experiment.

Another technique used to detect solar neutrinos is the one originally performed with the Kamiokande experiment[45]. Its 4.5 kton water Cherenkov detector is sensitive to the electron-neutrino scattering process

$$ \nu + e- \rightarrow \nu + e^- $$

by detecting the Cherenkov light of the scattered electron. For this reaction, a threshold of ~ 5 MeV is required to reject background, hence the experiment is only sensitive to ${}^8B$ neutrinos. If $\nu_e \rightarrow \nu_x$ oscillations occur, where $x = \mu, \tau$, there should be a deficit in the observed event rate of this reaction because $\sigma(\nu_e e^-) \approx 6\sigma(\nu_x e^-)$. The experiment's successor, Super-Kamiokande[46], detects solar neutrinos in an identical way but with a much more massive detector.

Experiment	Data / Theory	Reference
Homestake	0.33 ± 0.029	[47]
GALLEX	0.60 ± 0.06	[43]
SAGE	0.52 ± 0.06	[44]
Kamiokande	0.54 ± 0.07	[45]
Super-Kamiokande	0.474 ± 0.020	[46]

Tab. 2.2: Comparison of solar neutrino data to the SSM expectations.

All experiments have measured a solar neutrino flux which is about half of the one expected by the model, except the chlorine experiment which measures a significantly larger reduction. Table 2.2 summarizes the comparisons between the experimental results and the predictions of the Standard Solar Model (SSM).

Despite the already mentioned very large dependence of the ^7Be and ^8B fluxes on the solar core temperature and other theoretical uncertainties of the SSM, one has found no way to accommodate the data with a modification of the model. Even if the temperature is kept as a free parameter the data can not be reproduced. Neutrino oscillation, though, can reproduce the experimental results.

Given the large magnitude of the effect, it is expected that in a two-flavour mixing scheme, the mixing angle is rather large. A combined analysis [48] of all the experimental data available, results in a solution with $\Delta m^2 = 6.5 \times 10^{-11}\, eV^2$ and $\sin^2 2\theta = 0.75$. The χ^2 probability of the fit is 6%. Fig. 2.4(a) shows the preferred region in the parameter space at 99% C.L. Apart from the total interaction rates in the various experiments, the fit takes into account the recoil electron energy spectrum and the measurements of day-night differences at the water Cherenkov experiments.

Another plausible interpretation of the solar neutrino data is matter enhanced oscillation. As described in 2.2.3, the high density of the sun can affect the propagation of the solar neutrinos towards the detectors on earth. The best-fit global MSW solution for active neutrinos, which has a χ^2 probability of 7%, yields $\Delta m^2 = 5 \times 10^{-6}\, eV^2$ and $\sin^2 2\theta = 5.5 \times 10^{-3}$. Fig. 2.4(b) shows the 99% C.L. allowed region in the parameter space. The large mixing angle solution is disfavoured, though not completely excluded, if one includes in the fit zenith angle[1] distribution of solar neutrinos measured by Super-Kamiokande. The latter takes into account the possible matter effect in the earth.

Absence of large mixing angle solutions occurs also if oscillations to sterile neutrino are considered, even without the zenith angle data. In this case the χ^2 probability of the fit is 8% and yields $\Delta m^2 = 4 \times 10^{-6}\, eV^2$ and $\sin^2 2\theta = 7 \times 10^{-3}$. The 99% C.L. allowed region is almost identical to the one for active neutrinos, without the large mixing solution.

Data from new experiments and upgrades of the existing ones are needed to further confirm the interpretation of the solar ν_e deficit as oscillations. Perhaps the most significant contribution will come from the SNO experiment [49] in the Sudbury mine , which started

[1] defined by the direction pointing to the sun

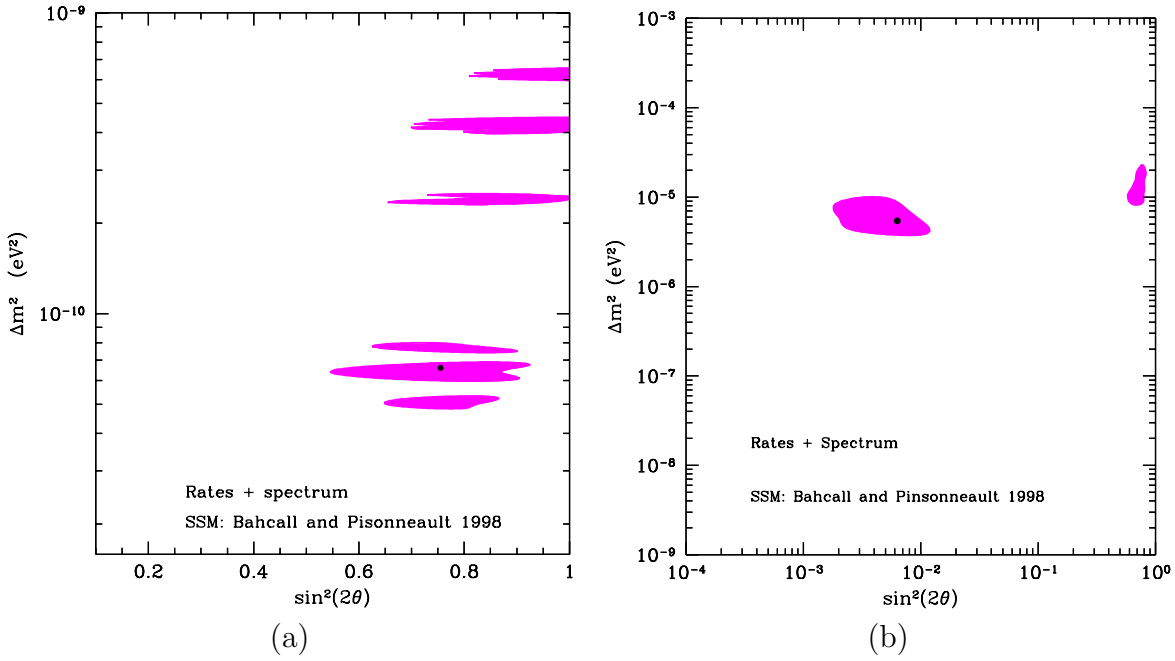

Fig. 2.4: Oscillation fits to the solar neutrino problem. (a): Vacuum solutions. (b): MSW solutions. Contours are drawn at 99% C.L.

taking data in 1999. The detector, which is based on a heavy water target, is not only sensitive to CC interactions of the type

$$\nu_e + d \rightarrow p + p + e^-$$

but will also aim at the detection of NC events of the type

$$\nu + d \rightarrow \nu + n + p$$

with initial detection of the proton, and subsequently of the neutron. Moreover, it is also sensitive to the electron neutrino elastic scattering

$$\nu + e^- \rightarrow \nu + e^-$$

which for the case of electron neutrinos can be either a CC or a NC interaction. If a relative enhancement of the NC/CC ratio is observed, then the existence of oscillations among active flavours can be established for the solar neutrinos.

A first measurement of the CC rate of solar neutrinos with the SNO detector has already been published [50]. The results are compatible with the observations made with all the above experiments. The measured rate of the neutrino - electron elastic scattering is also compatible with the high-statistics measurement performed in Super-Kamiokande, which indicate an enhancement of the NC/CC ratio. The SNO Collaboration claims that there

is evidence at a level of 3.3 standard deviations that there is a non-electron flavour active neutrino component in the solar flux.

Finally, in the Gran Sasso underground laboratory, the BOREXINO experiment [51] will make use of its liquid scintillator target to detect νe scattering, with an energy threshold below the ^7Be monochromatic line. Making use of its good energy resolution, the ^7Be line in BOREXINO will be visible as a step in the electron recoil spectrum. Requiring that events have a total visible energy within a narrow window around the ^7Be line, only the corresponding solar neutrinos will be selected. The experiment will measure then the flux of this specific component from the solar neutrino spectrum, which is especially sensitive to different oscillation hypotheses. In case vacuum oscillations occur, seasonal variations in the event rates will be observed.

2.5.2 Atmospheric neutrino experiments

The strongest evidence so far for the existence of oscillations comes from experiments studying atmospheric neutrinos. As explained in 2.3.1, oscillations may be manifest from the deviation of the ratio $\phi(\nu_\mu)/\phi(\nu_e)$ from the predicted value. Usually the experiments quote the ratio of observed and simulated events

$$R = \frac{\left(\frac{N_\mu}{N_e}\right)_{DATA}}{\left(\frac{N_\mu}{N_e}\right)_{MC}} \tag{2.55}$$

where N_μ and N_e are the numbers of μ-like and e-like events respectively in a detector sensitive to charged current interactions of both ν_μs and ν_es.

So far six experiments, NUSEX[52], Frejus[53], IMB[54], Soudan2[55] Kamiokande[56] and Super-Kamiokande[57] have published results on the measurement of R, summarized in Table 2.3. The double ratios quoted in the table are averaged over all the possible directions and the energy spectrum of the atmospheric neutrinos.

Experiment	R
NUSEX	$0.99^{+0.35}_{-0.25}$
Frejus	$1.00 \pm 0.15 \pm 0.08$
IMB	$0.54 \pm 0.05 \pm 0.11$
Soudan2	$0.64 \pm 0.11^{+0.06}_{-0.05}$
Kamiokande (sub-GeV)	$0.60^{+0.06}_{-0.05} \pm 0.05$
Kamiokande (multi-GeV)	$0.57^{+0.08}_{-0.07} \pm 0.07$
Super-Kamiokande (sub-GeV)	$0.638 \pm 0.017 \pm 0.05$
Super-Kamiokande (multi-GeV)	$0.675^{+0.034}_{-0.032} \pm 0.08$

Tab. 2.3: Results on the *double ratio* from various atmospheric experiments.

The experiments establish a value of R consistent with the scenario of mixing. From the magnitude of the effect (R being very close to 0.5) it can be inferred that rather large

mixing angles are involved. Taking into account the mean neutrino flight length, the data suggest relevant Δm^2 values higher than $10^{-4}\,\mathrm{eV}^2$.

If an experiment can measure with sufficient accuracy the direction of the incoming neutrino, then an estimate of its actual flight length L can be achieved. It is related to the zenith angle θ_z through the formula

$$L = \sqrt{R_\odot^2 + (R_\odot - d)^2(1 - \cos^2\theta_z)} - (R_\odot - d)\cos\theta_z \tag{2.56}$$

where $R_\odot$ is the radius of the earth ($\sim$ 6000 km) and d is the typical height from the earth's surface where neutrinos are produced ($\sim$ 10 km). If neutrinos oscillate with a Δm^2 in a relevant range (i.e. $\frac{E}{L_{max}} < \Delta m^2 < \frac{E}{L_{min}}$), a strong dependence of the observed event rate on θ_z is expected. Such a dependence has actually been observed by the Super-Kamiokande experiment[57], the 50 kton water Cherenkov detector located in the Kamioka mine in Japan. The experiment has studied "single ring events", i.e. quasi-elastic charged current interactions of the type

$$\nu_l + N \to l + N'$$

where N stands for a proton or a neutron. Fig. 2.5 shows the zenith angle dependence of the observed interaction rate for μ-like and e-like events. The fully contained events at Super-Kamiokande have been divided into two classes, according to the total visible energy: *sub-GeV* for events with $E_{vis} < 1.33$ GeV and *multi-GeV* for events with $E_{vis} > 1.33$ GeV. Partially contained events comprise a 98% pure sample of ν_μ CC interactions, with mean energy of about 15 GeV. Assuming no oscillations, e-like events behave as expected, but the zenith angle dependence on the observed μ-like event rate is clearly different from the expected one for all event classes. On the other hand, the data is consistent with the scenario of $\nu_\mu \to \nu_\tau$ oscillations.

The unprecedented statistics collected with the Super-Kamiokande detector, allowed for a measurement of the *Upward-Downward* asymmetry of the produced leptons considering more than the two mentioned energy bins. In Fig. 2.6 the observed asymmetries for μ-like and e-like events are plotted as a function of the neutrino energy. Similar to the zenith angle measurements, the observed asymmetry is consistent with the expectations only for electron neutrinos, while all data are reproducible if $\nu_\mu \to \nu_\tau$ oscillations are considered.

The measurement of the U-D asymmetry simulates in some sense a two detector experiment. The difference here is that a single detector is exposed in two beams of neutrinos whose sources are almost identical, but located at different distances.

Using the statistics of 1289 live-days data, a combined $\nu_\mu \to \nu_\tau$ oscillation fit using all the Super-Kamiokande measurements yields $\Delta\mathrm{m}^2 = 2.5\,\mathrm{eV}^2$ for full mixing ($sin^2 2\theta = 1.0$). The allowed regions in the parameter space for various confidence intervals are plotted in Fig. 2.7.

It should be noted, that transitions to a sterile neutrino ($\nu_\mu \to \nu_s$) at a similar Δm^2 value could explain the muon neutrino deficit as well. The experiment can distinguish between the two cases by studying NC processes with a single π^0 in the final state, or

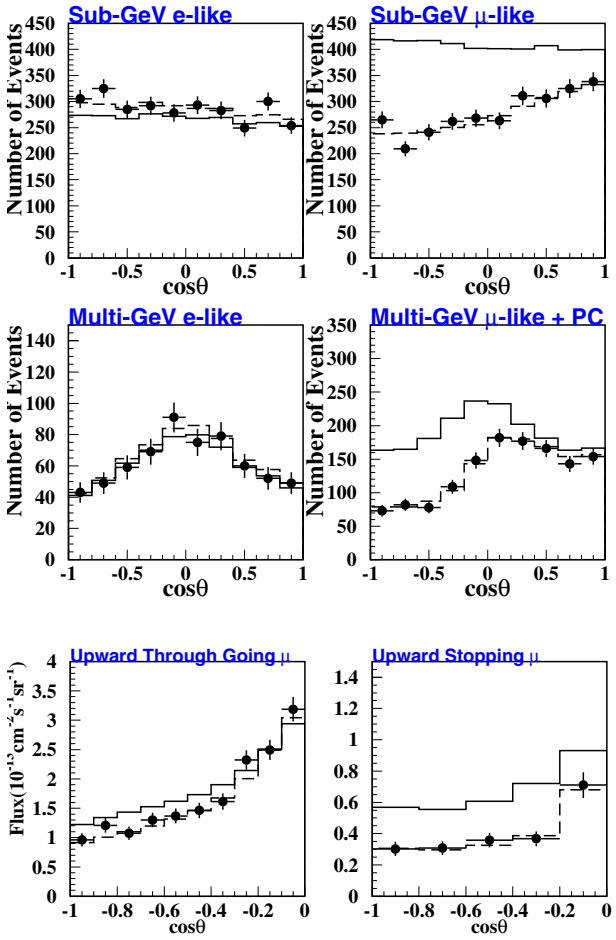

Fig. 2.5: Zenith angle dependence of e-like and μ-like events in the Super-Kamiokande detector. The solid line shows the Monte Carlo expectation in the absence of oscillations normalized to the expected rate with statistical errors. The dashed line is the best-fit expectation for $\nu_\mu \to \nu_\tau$ oscillations with the overall flux normalization fitted as a free parameter. The two bottom plots show the zenith angle dependency of the rate of upward through going and upward stopping muons.

by exploiting the different MSW effect in both cases. If muon neutrinos oscillate into tau neutrinos this event rate should be enhanced. Most of the systematic uncertainties for this measurement come from the poor knowledge of the cross-section for this single π^0 process. Valuable input is expected from the measurement at the near detector of the K2K experiment (discussed in 2.5.5), whose result will be available within the next years. However, already from the existing data pure $\nu_\mu \to \nu_s$ oscillation is ruled out at 99% C.L.

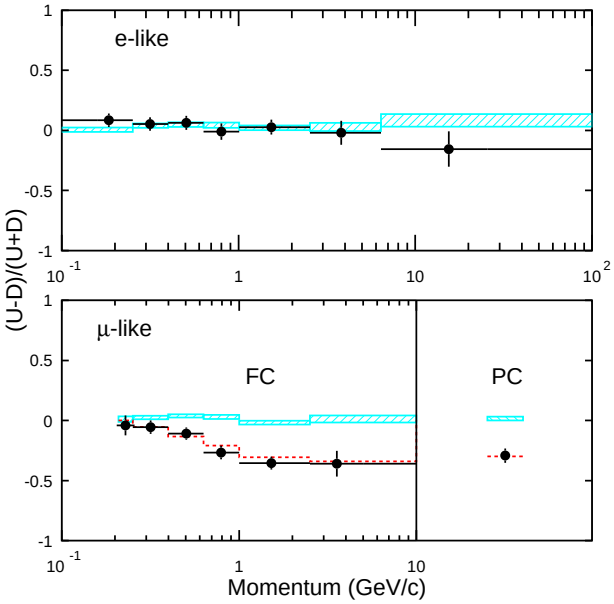

Fig. 2.6: Upward-Downward going lepton asymmetry as a function of momentum. The hatched region shows the Monte Carlo expectation in the absence of oscillations, while the bold line is the best-fit expectation for $\nu_\mu \rightarrow \nu_\tau$ oscillations.

2.5.3 Reactor neutrino experiments

Even though there is no claim so far for oscillations from reactor experiments, the limits they obtained contribute significantly to the solution of the neutrino puzzle by excluding those channels involving the electron neutrino in regions of the parameter space where they are sensitive.

The most recent result has come from the long baseline CHOOZ experiment[58]. The detector has been installed in an underground laboratory about 1 km from two 8.5 GWatt nuclear reactors. Its target consists of hydrogen-rich paraffinic liquid scintillator loaded with 0.09% gadolinium. Hence, the detector identifies an electron antineutrino charged current interactions of the type

$$\overline{\nu}_e + p \rightarrow e^+ + n \qquad \text{with} \qquad E_{e^+} = E_{\overline{\nu}_e} - 1.804\,MeV$$

from the delayed coincidence of the e^+ signal (two 511 keV annihilation photons) and the signal from the neutron capture by the gadolinium nuclei (γs from the subsequent nuclear decay).

Neutrino oscillation including the electron neutrino would manifest itself as a modulation, or deficit of the expected positron energy spectrum. This has not been observed, yielding a null result. The exclusion plot for $\nu_e \rightarrow \nu_x$ oscillations obtained with CHOOZ is shown in Fig. 2.8. The results of other reactor experiments are superimposed, as well as the allowed region of the original Kamiokande result. Comparing Figures 2.8 and 2.7

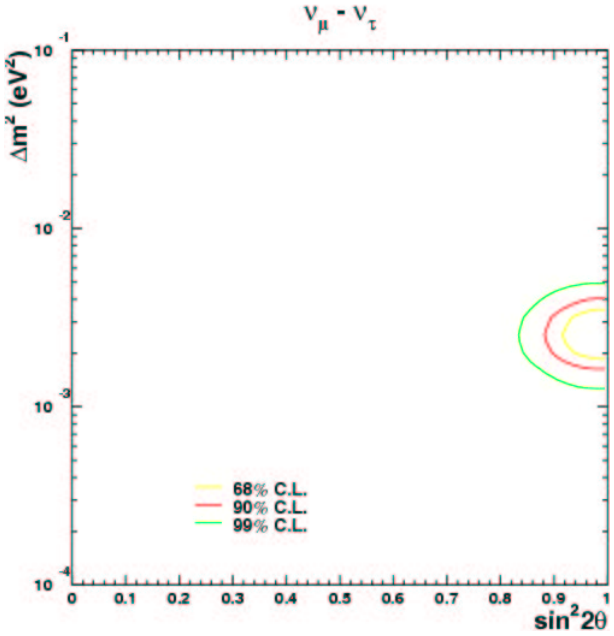

Fig. 2.7: Allowed regions for $\nu_\mu \rightarrow \nu_\tau$ oscillations, from the Super-Kamiokande data.

it can be seen that also the Super-Kamiokande claim is fully covered at 90% C.L., which is an additional hint that the atmospheric muon neutrinos do not oscillate into electron neutrinos.

Nuclear reactors will continue contributing to the solution of the neutrino puzzle. In particular the KamLAND project[59] will attempt to explore the parameter space which covers the large mixing MSW solution of the solar neutrino deficit. Based on similar detection techniques as the CHOOZ experiment, but without the Gd doping and using a more massive target located in the original Kamiokande hall, the experiment will be sensitive to neutrinos produced by nuclear reactors about 200 km from the detector. The envisaged exclusion plot in case of negative result is shown in Fig. 2.9.

2.5.4 Current accelerator neutrino experiments

The only claim for oscillations from an accelerator experiment which has not been excluded by others, comes from LSND[60]. The detector is made of a tank filled with oil, doped with a low concentration of scintillator, which allows both scintillation and Cherenkov light to be detected. It is exposed to a beam of ν_μ, ν_e, and $\overline{\nu}_\mu$. There is practically no $\overline{\nu}_e$ component present (see 2.3.1). However, a significant excess of $\overline{\nu}_e$ events above the background have been observed. The oscillation probability for the $\overline{\nu}_\mu \rightarrow \overline{\nu}_e$ channel has been determined as $P_{osc} \approx 3.1 \cdot 10^{-3}$. This translates into the allowed region in the parameter space shown in Fig. 2.10. An independent, but less significant measurement of the CP conjugate channel $\nu_\mu \rightarrow \nu_e$ confirmed the initial result[61].

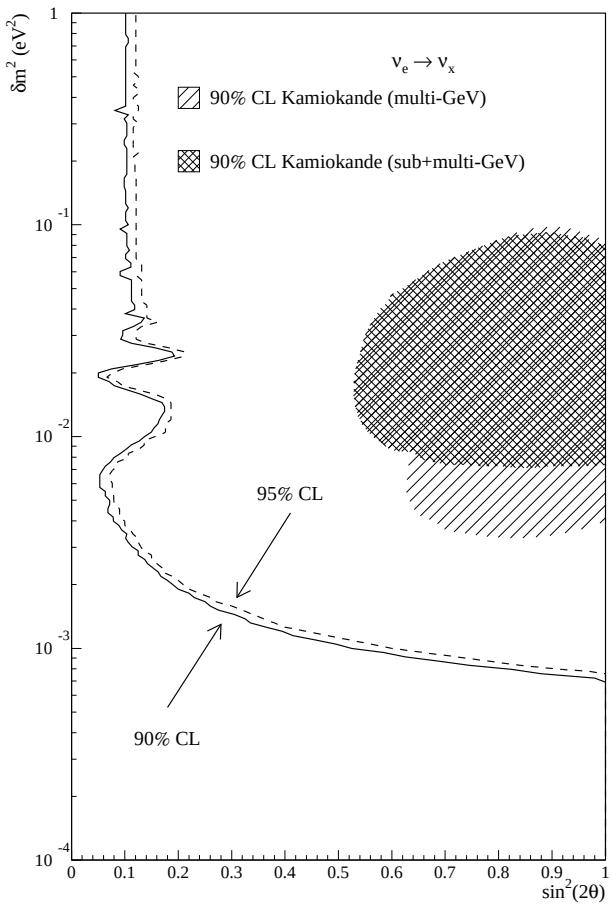

Fig. 2.8: Exclusion plot for the $\nu_e \rightarrow \nu_x$ channel obtained by the CHOOZ experiment. Results from the original Kamiokande claim on atmospheric neutrino oscillations are also overlaid.

The currently running experiment KARMEN[62] has the possibility to confirm the claim from LSND since it is sensitive to $\nu_\mu \rightarrow \nu_e$ oscillations at a similar region of the parameter space. The detector is exposed to a similar neutrino beam. In their case though, the proton beam is pulsed. The relative timing between the proton cycle and the observed signals at the detector helps in suppressing the background.

KARMEN has not confirmed the LSND claim. The results obtained by this experiment are shown in Fig. 2.11, where the favoured LSND region is also presented. From this plot it can be inferred that despite the fact that KARMEN has not observed oscillations, the results of the two experiments are consistent with each other. The reason for that is that KARMEN is not sensitive enough to fully exclude LSND [63]. A definite answer to the LSND claim needs further experimentation (see 2.5.5).

Oscillation searches in the $\nu_\mu \rightarrow \nu_\tau$ channel have been performed in mass ranges interesting for cosmological reasons (see 2.1.2). If the solar neutrino problem is interpreted

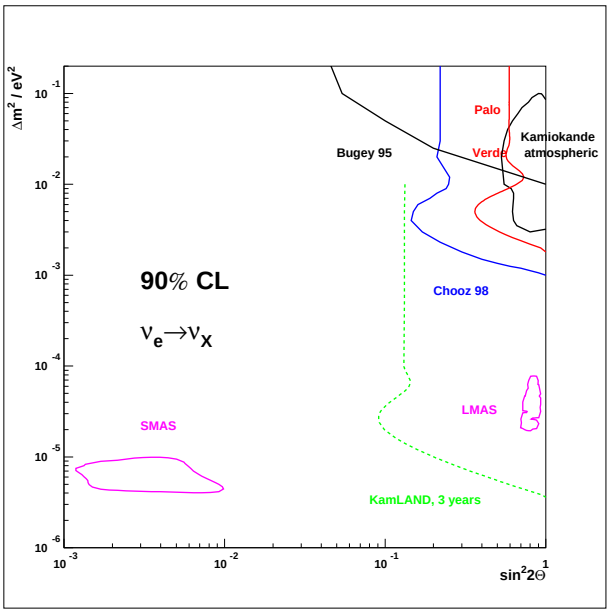

Fig. 2.9: Sensitivity of the KamLAND experiment. The results of other major reactor experiments are superimposed along with the MSW solar neutrino solutions and the atmospheric claim from Kamiokande.

as $\nu_e \rightarrow \nu_\mu$ oscillations, the see-saw mechanism predicts a mass for the ν_τ in the region of a few eV, which is interesting for cosmology. Within this context, the short baseline experiments at CERN, CHORUS[64] and NOMAD[65], have searched for $\nu_\mu \rightarrow \nu_\tau$ oscillations at small mixing angles ($< 10^{-3}$) and Δm^2 values in the range $1 - 100\,\mathrm{eV}^2$. These experiments took data from 1994 to 1998 and the collaborations are currently analysing their data. So far no evidence for oscillations has been claimed by NOMAD[66] or CHORUS (see following chapters). However, if the LSND result is valid, according to some phenomenological models[67], CHORUS could eventually observe some signal events. The beam line at CERN, the CHORUS experiment and its results are described in the following chapter of this thesis.

2.5.5 *Accelerator experiments in the immediate, near and far future*

Summarizing the experimental data, there are three indications for oscillations: The deficit of solar neutrino flux, the atmospheric neutrino anomaly, and the results from the LSND experiment. It is necessary that new experiments are performed to confirm these results. Each of them indicate mixing at different effective Δm^2 values. It is very difficult to accommodate these in a simple model with three neutrinos. If OMSD models are considered, then there are three independent mass differences which lead to four massive neutrinos.

While the hints from the solar and atmospheric neutrinos are supported by more than one experiment, this is not the case for LSND. At Fermilab, the BooNe experiment[68] has

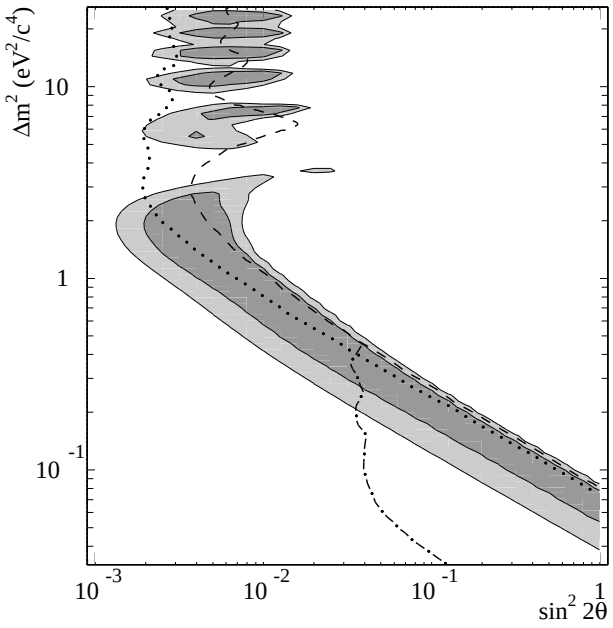

Fig. 2.10: Allowed region in the Δm^2-$\sin^2 2\theta$ plane of the LSND experiment. The shaded regions are the 90% or 99% likelihood regions. Also shown are the 90% C.L. exclusion limits from KARMEN before upgrade at ISIS (dashed curve), E776 at BNL (dotted curve), and the Bugey reactor experiment (dot-dashed curve).

been proposed to run with neutrinos of a mean energy around 1 GeV, produced at the proton booster. The detector is based on a sphere filled with mineral oil doped with scintillator, and aims for excellent electron identification. An oscillation signal would manifest itself as an excess of electron neutrino events, above the expected beam-related background. Data taking is expected to start in 2002 using a single detector. The experiment is referenced to as "miniBooNe" for the one-detector phase. In case of a positive result a second detector will then be installed at a different distance from the source. A better control over the background and the systematics will be achieved then, hence confirming the result and further constraining the favoured region in the parameter space.

The sensitivity of miniBooNe is presented in Fig. 2.12. The experiment will be able to completely cover the region favoured by LSND.

The strongest claim for oscillations, namely the results from the atmospheric neutrinos urge for confirmation by long baseline experiments. K2K[70], the first accelerator long baseline experiment ever built, has been taking data since the spring of 1999. Neutrinos of mean energy around 1 GeV, produced at the KEK proton synchrotron, are directed towards the Super-Kamiokande detector, which is situated 250 km away. It is a ν_μ disappearance and ν_e appearance experiment. At the near location there are two detectors. The first is a fine-grained calorimeter and the other a water Cherenkov, which is a down-scaled copy of the Super-Kamiokande detector. The fine-grained calorimeter will be used to study the energy spectrum, composition and profile of the neutrino beam, minimizing the

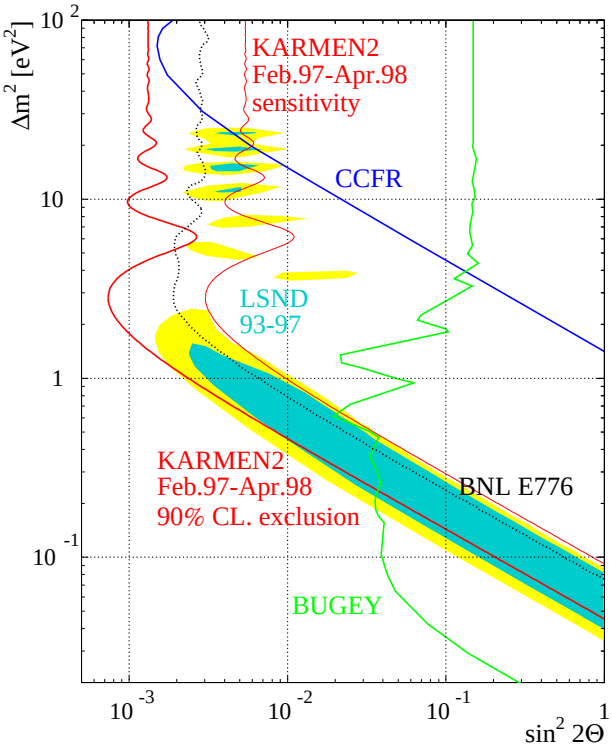

Fig. 2.11: The sensitivity and the 90% C.L. exclusion curves for $\nu_\mu \rightarrow \nu_e$ oscillation from KARMEN. The favoured LSND region as well as results from previous experiments are superimposed.

systematical uncertainties related to the beam itself. The water Cherenkov, will not only provide a good handle on the systematics of the Super-Kamiokande detector, but also perform the cross-section measurement for the process

$$\nu + N \rightarrow \nu + N' + \pi^0$$

under similar conditions as in the Super-Kamiokande experiment. The latter, being a NC interaction, provides the means for Super-Kamiokande to decide whether atmospheric muon neutrinos oscillate into tau or sterile neutrinos, as discussed in 2.5.2.

In case K2K confirms the atmospheric results, i.e. muon neutrino disappearance, a measurement of the effective Δm^2 will be performed, by fitting the observed modulated muon spectrum at the Super-Kamiokande detector.

MINOS[71] is a long baseline experiment which will be performed at the NuMI[72] facility, under construction at Fermilab. Neutrinos produced at Fermilab's Main Injector will be directed towards the Soudan mine, 730 km away, where the far detector will be installed. Both near and far detectors are fine-grained calorimeters with magnetized iron. The focusing system of the beam line will allow the mean neutrino energy to vary within the range 1-20 GeV. As a first step, low energies will be used for the confirmation of

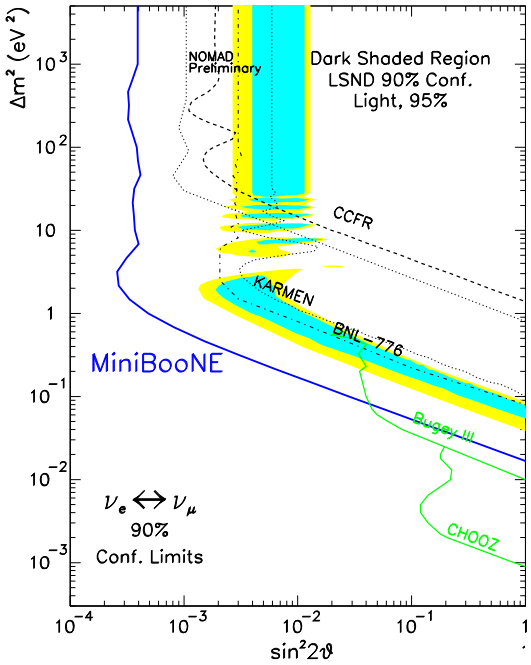

Fig. 2.12: The proposed sensitivity for the miniBooNe experiment.

the atmospheric neutrino effect (muon neutrino disappearance) and the measurement of the relevant Δm^2. The near detector will provide the necessary normalization for the ν_μ spectrum and NC/CC ratio. At the same time, like K2K, a $\nu_\mu \rightarrow \nu_e$ appearance search can be performed.

Searches for ν_τ appearance will also be performed to check whether the dominant oscillation channel at the atmospheric Δm^2 is $\nu_\mu \rightarrow \nu_\tau$. If this is the case an increased NC/CC ratio will be observed. For the run with higher neutrino energies, a detector upgrade will possibly include the installation of nuclear emulsion plates to identify the tau decay topology in ν_τ CC interactions.

MINOS and K2K have adopted complementary detection approaches. The K2K experiment will come first with physics results. However, MINOS will be able to exploit lower Δm^2 values thanks to the longer baseline, and it could possibly cover most of the Super-Kamiokande allowed region. The advantage of the water Cherenkov detectors at K2K is that they do not suffer from background due to events with π^0s misidentified as electron neutrino events. On the other hand calorimetric detectors such as MINOS can measure precisely the total energy of every event, and efficiently distinguish NC from CC events. Therefore, they are better suited for the measurement of Δm^2, from the modulation of the ν_μ energy spectrum and the NC/CC ratio distribution.

These experiments aim at sensitivities shown in Fig. 2.13

At CERN the CNGS [73] facility will be focused to the direct observation of ν_τ ap-

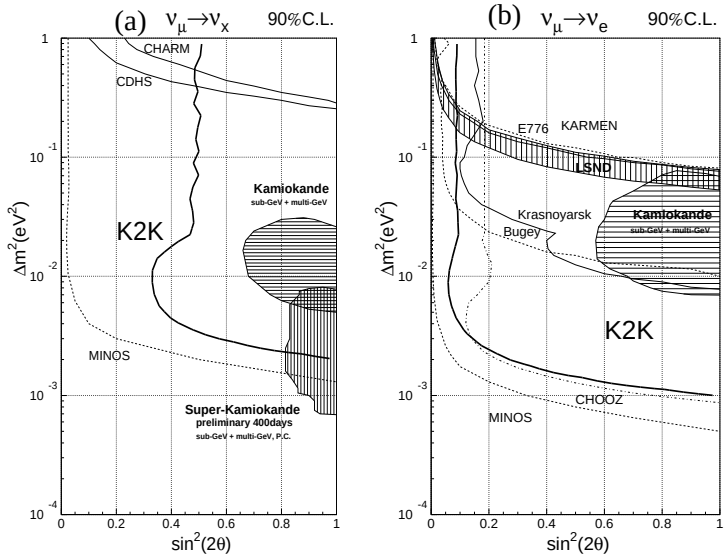

Fig. 2.13: The projected sensitivities for the K2K and MINOS long baseline experiments for (a) $\nu_\mu \to \nu_x$ and (b) $\nu_\mu \to \nu_e$ oscillations.

pearance in a ν_μ wide-band beam which will be sent towards the Gran Sasso underground laboratory in Italy, 730 km from CERN. Since the tau observation requires high energies to overcome its production threshold, the beam has been tuned so that the neutrino flux peaks around 20 GeV. From the various experiments proposed to run at CNGS, OPERA [74] has adopted the emulsion technique for the detection of the ν_τ CC interactions. Lead-emulsion target modules will be followed by electronic hadron and muon spectrometers to suppress the background events, which originate mainly from charmed meson decays.

The ICARUS experiment [75] will be based on a liquid argon TPC. Its excellent electron identification capability allows for efficient oscillation searches of the type $\nu_\mu \to \nu_e$ and $\nu_\mu \to \nu_\tau$ with the τ decaying to e. The separation of ν_τ from ν_e interactions will be achieved by applying kinematical criteria.

The sensitivities that LBL experiments at Gran Sasso are expected to reach, are plotted in Fig. 2.14. The potential sensitivity of a short baseline experiment in the same beam line is plotted as well.

Let us assume for the moment that the atmospheric neutrino anomaly is a consequence of pure $\nu_\mu \to \nu_\tau$ oscillation. An ideal accelerator experiment is able to check this scenario if it can obtain simultaneously (i.e. using the same neutrino beam) the energy spectra from both ν_μ and ν_τ CC events. If the fits to the spectra do not yield the same Δm^2 and $\sin^2 2\theta$ values, then the observed effect can no longer be attributed to a pure $\nu_\mu \to \nu_\tau$ oscillation. A more general mixing scheme should be considered then, which involves at least a third neutrino.

There are two basic requirements for such an experiment:

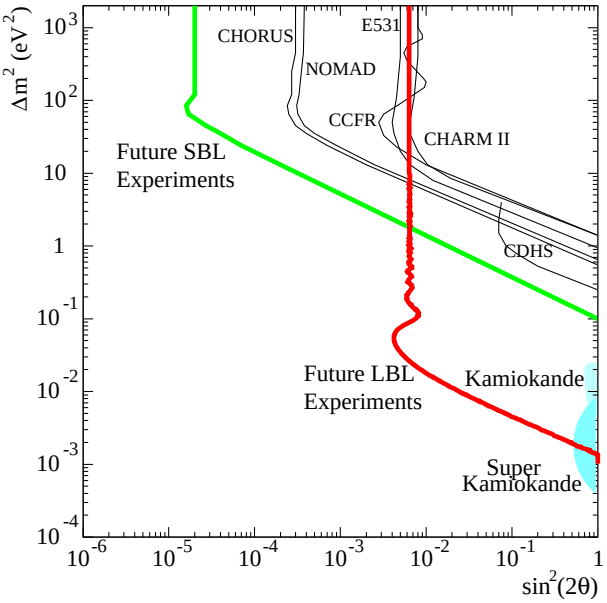

Fig. 2.14: The sensitivity of a generic long baseline $\nu_\mu \to \nu_\tau$ appearance experiment at the CNGS facility. The result of a potential SBL experiment is also shown.

- The mean energy of the neutrino beam should be at least of the order of 10-20 GeV to overcome the tau production threshold.
- The distance of the detector from the source should be such that $E/L \sim \Delta m^2_{\text{atmospheric}}$, so that a precise Δm^2 measurement can be achieved.

None of the proposed experiments meet both requirements. When MINOS will attempt to measure the Δm^2 from the ν_μ energy spectrum, in the best case only few ν_τ CC interactions will be produced, due to the low neutrino energy. At this Δm^2 only the disappearance $\nu_\mu \to \nu_x$ can be measured. On the other hand, the CNGS experiments, which aim at observing ν_τ CC events, will run at higher neutrino energies. Given the CNGS baseline length of 730 km and the actual value of $\Delta m^2_{\text{atmospheric}} \sim 10^{-3}\,\text{eV}^2$, the oscillation probability becomes so low, that the observed ν_τ CC events will be just over the tau production threshold. Therefore, no information will be obtained from the energy spectrum of these events. From their number one could possibly measure the quantity $(\Delta m^2)^2 \cdot \sin^2 2\theta$. However, this requires the underlying assumption of a pure $\nu_\mu \to \nu_\tau$ oscillation, or an OMSD mixing scheme. It also assumes full knowledge of the detector efficiencies and the W_4, W_5, W_6 structure functions at the threshold energy region.

The neutrino puzzle can be fully solved only if very intense neutrino sources are developed, which will allow the setting up of experiments at larger distances. Neutrino factories at muon storage rings are among the most promising candidates. They could provide 2-3 orders of magnitude more intense neutrino flux than conventional neutrino facilities. Among the advantages of the neutrino factories with respect to conventional beams are

the excellent knowledge of the energy spectrum, and the well defined two-flavour and opposite-charge composition of the beams. Such features do not only allow for a very detailed measurement of the mixing matrix elements, but one could also envisage studies of possible CP-violating and matter effects using very long baseline experiments. Neutrino factories are currently entering the R&D phase. Prospective studies have been made both at CERN[76] and the US[77]. The currently performed R&D studies concentrate in the determination of the pion production yield from the primary proton target [78] and the ionization cooling of the produced muons [79].

Another intense source of neutrinos is the so-called β-beam, an idea which has been brought forward quite recently [80] and is gaining increasing interest. Like the neutrino factory it is based on a storage ring where the neutrino parent is a light radioactive ion instead of a muon. Its major advantages with respect to the neutrino factory are that the technology to build such a facility is available, and mostly important there is only one neutrino flavour in the resulting beam. This allows the exploitation of a larger number of experimental signatures during an oscillation search, putting therefore less constraints in the design of a potential detector.

3. THE CHORUS EXPERIMENT

The CERN-WA95/CHORUS (**C**ERN **H**ybrid **O**scillation **R**esearch Apparat**US**) experiment [64] was proposed to search for neutrino oscillations in the $\nu_\mu \rightarrow \nu_\tau$ appearance channel with the CERN SPS wide-band neutrino beam (WBB). It was designed to explore small mixing angles ($\sin^2 2\theta \sim 10^{-4}$) at Δm^2 values of order $10 - 100\,\mathrm{eV}^2$. This region of the parameter space is relevant for the hypothesis of ν_τ as a Dark Matter candidate [81].

The detector itself is fully described in Ref. [82]. In this chapter the main role of each subdetector is summarized. In section 3.5, a list of potential background sources and the design sensitivity are presented.

3.1 Basic concepts in the design of CHORUS

In the SPS WBB, there are practically no ν_τs present in the resulting beam. It is conceivable then to design a background-free $\nu_x \rightarrow \nu_\tau$ appearance experiment probing very small oscillation probabilities. However, this requires that ν_τ CC events are unambiguously identified at the detector and distinguished from any other type of neutrino interactions.

At a typical neutrino detector exposed in a high energy beam ($E_\nu > 1\,\mathrm{GeV}$), neutrino interactions are identified as ν_μ CC, when a single minimum ionizing particle (MIP) traverses a large fraction of the detector before it is absorbed, contrary to the rest of the particles, which release all their energy closer to the interaction vertex. The former is recognized as the primary muon, while the rest constitute the hadronic shower.

A ν_e CC event is recognized by the absence of such a penetrating MIP and the presence of an electromagnetic shower caused by a single charged particle coming from the vertex. The latter is identified as the primary electron.

For ν_τ CC interactions there is no such simple signature. The τ lepton, being short-lived, travels a distance of the order of 1 mm before decaying. The characteristic topology of the τ decay near the vertex of a ν_τ CC interaction is shown in Fig. 3.1. All particle tracks, except the one of the τ charged daughter (a muon in the case shown in Fig. 3.1), originate from the primary interaction vertex. The τ decay is seen as a *kink* on the daughter track. If the interaction was due to a ν_μ, the muon track would be directly attached to primary vertex, without the presence of such a kink.

Generally, for a detector with a moderate three-dimensional spatial resolution at the vertex region, ν_τ CC interactions can be misidentified as ν_e CC, ν_μ CC or ν NC when the τ decays in the electron, muon, or hadron channel respectively. Therefore without knowledge of the vertex topology, all neutrino interactions are potentially background processes. In

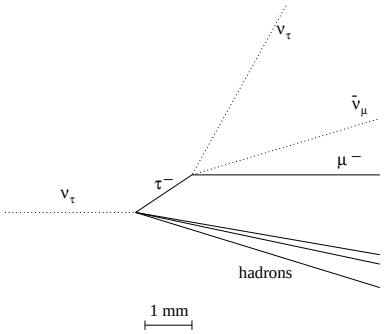

Fig. 3.1: Topology of a ν_τ CC interaction at the vertex region. In this example the τ decays in the muon channel.

principle, a high resolution spectrometer could separate them from the signal ν_τ CC events by applying kinematical criteria. These are based on the fact that neutrinos produced at τ decays carry away a significant fraction of the initial τ momentum, as can be inferred from Fig. 3.1. Signal can be separated from background by posing cuts on the magnitude and direction of the missing transverse momentum. This approach has been adopted by the NOMAD experiment [65] at CERN.

The CHORUS collaboration has chosen to rely mainly on the detection of the τ decay topology. To this end, an fully active target is needed, with a spatial resolution such that allows the identification of this topological signature. The nuclear emulsion has been identified as the most suitable medium meeting this requirement. Its three-dimensional spatial resolution of $\sim 1\,\mu$m and the high hit density along a MIP track (300 hits/mm) allows the reconstruction of the τ track, as well as its production and decay vertices.

However, there are still numerous background processes, where a kink is present along a track close to the vertex in neutrino interactions (see 3.5). These background processes can be suppressed by posing kinematical cuts on the tracks of the τ daughter particle. Therefore, an electronic spectrometer is needed to validate any possible ν_τ candidate event.

Electronic trackers should be used to predict at the emulsion plates the position and slope of tracks originating from neutrino interactions, as they have been electronically recorded during the exposure. This way only the relevant tracks are followed in the emulsion target, with only a small fraction of the total emulsion surface scanned. Good predictions make the scanning also more efficient, since the fraction of background tracks picked up during the scanning procedure is also minimized.

For the above reasons, the design of the CHORUS detector, which is depicted in Fig. 3.2, is driven by the *hybrid* approach of combining emulsion and electronic detection techniques. It is composed of an emulsion target, a scintillating fiber tracker system, scintillator trigger hodoscopes, a hadronic spectrometer based on a hexagonal air core magnet, a lead/scintillator calorimeter and a muon spectrometer.

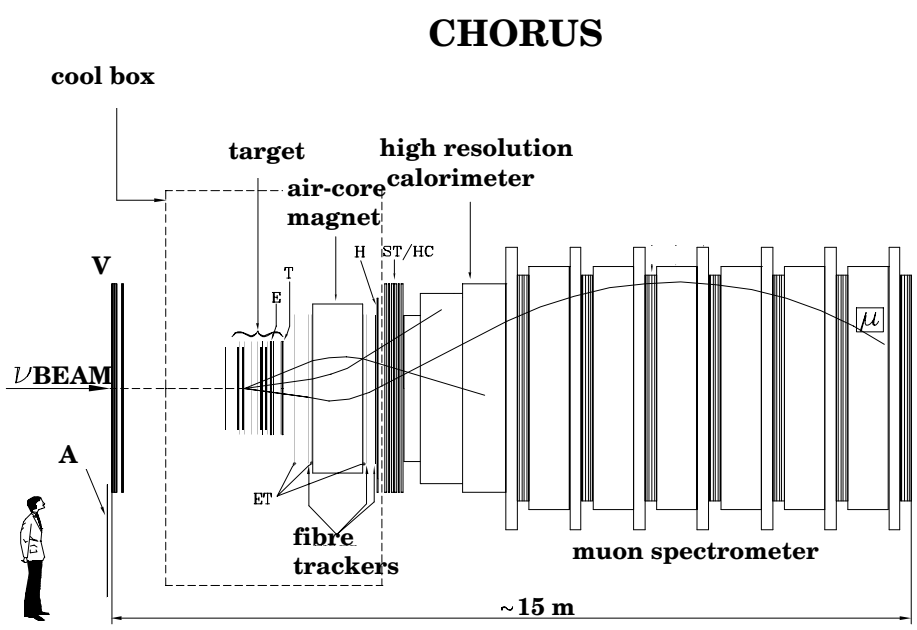

Fig. 3.2: Schematic diagram (side view) of the CHORUS apparatus. The trigger and veto hodoscopes are denoted by V, A, T, E, H. The detector upgrades in 1996 on the emulsion trackers and honeycomb chambers are denoted by ET and HC respectively. In the actual setup, individual modules are mounted vertically while the axis of the apparatus is inclined at 42 mrad to match the neutrino beam axis.

3.2 The neutrino beam

For the needs of a $\nu_\mu \rightarrow \nu_\tau$ appearance experiment there are two basic requirements on the neutrino beam. It has to be as intense as possible to allow the exploration of small mixing angles, and its mean energy should be well above the tau production threshold ($\sim 3.5\,\mathrm{GeV}$).

The most intense neutrino source available at CERN meeting these requirements has been the West Area Neutrino Facility (WANF)[83] at the SPS. This facility has been operating for 20 years serving various neutrino experiments such as CDHS, CHARM, BEBC, and CHARM II.

For the CHORUS and NOMAD run a major reconstruction [84] took place in 1992 and 1993, while in spring 1995 some crucial elements of the line were also re-aligned with respect to the proton beam to provide for a better centered neutrino beam and an increased neutrino flux at the experiments [85]. A schematic layout of the WANF beam line is shown in Fig. 3.3.

Protons are accelerated by the SPS to 450 GeV with a cycle of 14.4 s. They are extracted in two 4 ms long spills separated in time by 2.6 s, one at the beginning and one at the end of the energy flat top (see Fig. 3.4). The proton beam on the target has a base-to-base width of about 2 mm in each projection. The beam line has been operating

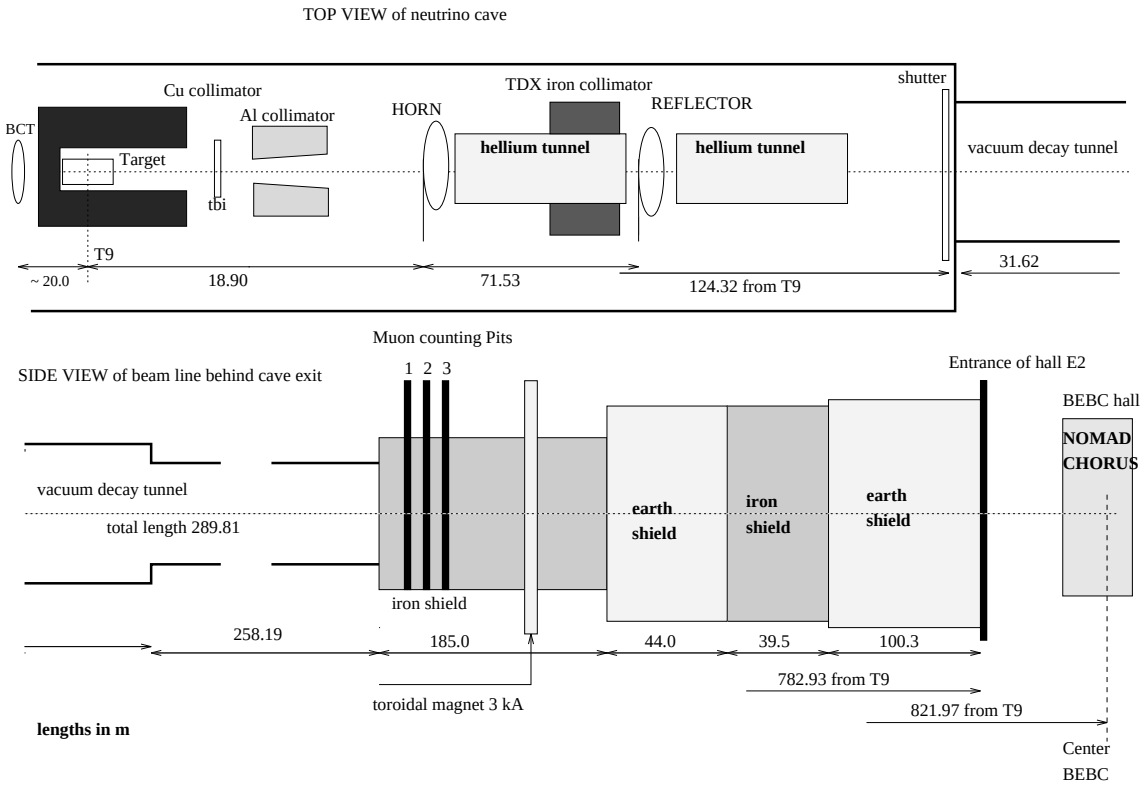

Fig. 3.3: Schematic layout of the WANF beam line

for 5 years (1994-1998) at record proton intensities – up to $\sim 3 \times 10^{13}$ protons on target (pot) per 14.4 s accelerator cycle.

The T9 beryllium target consists of 11 Be rods of 3 mm diameter and 10 cm length, separated by 9 cm gaps. It is situated 35 m underground in the Neutrino Beam Cave (NBC), at 400 m distance from the extraction point in the main SPS ring. The Be target is cooled with a flow of He gas. The target box is contained in a Cu collimator.

For the steering of the proton beam onto the target a number of Secondary Emission Monitors (SEM) are placed in the Target Beam Instrumentation (TBI) boxes upstream and downstream the target. The calibration of the SEMs is described in [86].

Proton intensities are monitored with the aid of two Beam Current Transformers (BCT). The first is installed close to the extraction point ($\sim$ 70 m upstream from T9) and another close to the target ($\sim$ 20 m upstream from T9).

A water cooled, cylindrical, 2.85 m long Al collimator is situated 3.55 m downstream the target center. The geometry of its bored beam hole limits the acceptance for the secondary mesons produced in the target below 25 mrad. This collimator absorbs a large fraction of K_0s, which contribute to the ν_e contamination in the resulting beam.

The focusing of the parent particles over a wide range of momenta is achieved with a system of coaxial magnetic lenses, commonly quoted as the *horn* and the *reflector*. The

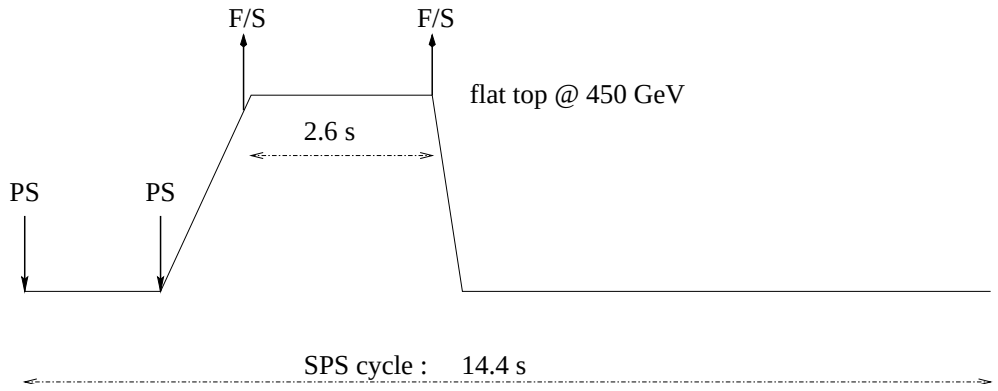

Fig. 3.4: The SPS super-cycle. The solid line is a graph of the current of the accelerator's dipole magnets versus time, during the proton cycle. The flat top corresponds to the maximum deliverable energy of 450 GeV. There are two *Fast/Slow* proton extractions towards the neutrino target, one just before the start and another at the end of the flat top. Protons are delivered by the PS and injected into the SPS in two PS cycles before the acceleration in the SPS starts.

horn operates at 100 kA and focuses positively charged mesons of energy around 40 GeV. The reflector, which is situated downstream, operates at 120 kA correcting the effect of the horn's magnetic field on the over-focused low energy (around 20 GeV) and the under-focused high energy mesons (up to 100 GeV). Negatively charged particles, especially those of low energies, are defocused by the system. With current inversion of the magnetic lenses a pion beam with a different polarity is focused. The polarity of the current can be therefore used to select between a ν and a $\overline{\nu}$ beam.

Between the horn and the reflector, and the latter and the end of the NBC, there are helium pipes to reduce particle absorption in air. An additional iron collimator, just upstream the reflector, protects the walls and the reflector itself from irradiation due to the defocused particles.

The charged mesons are left to decay in a 289.9 m long vacuum tunnel, where muon neutrinos and positive muons are produced. The muons are filtered out in the iron and earth shields which fill the space between the vacuum tunnel and the experimental hall where CHORUS and NOMAD are situated, 822 m from the T9 target.

In the first meters of the iron shield, where the muon flux is not significantly reduced, three muon counting stations utilizing silicon detectors are installed to measure the muon flux associated with the muon neutrinos. By monitoring the resulting muon flux, this system serves for the alignment of the primary proton beam, and for the monitoring of the effects of the horn and reflector currents. They have played an important role during the beam alignment procedure described in [85]. However, from the physics point of view, this system is the most important handle in calculating the absolute neutrino flux reaching the detectors.

The neutrino beam composition and the spectra for all flavours are given in Table 3.1

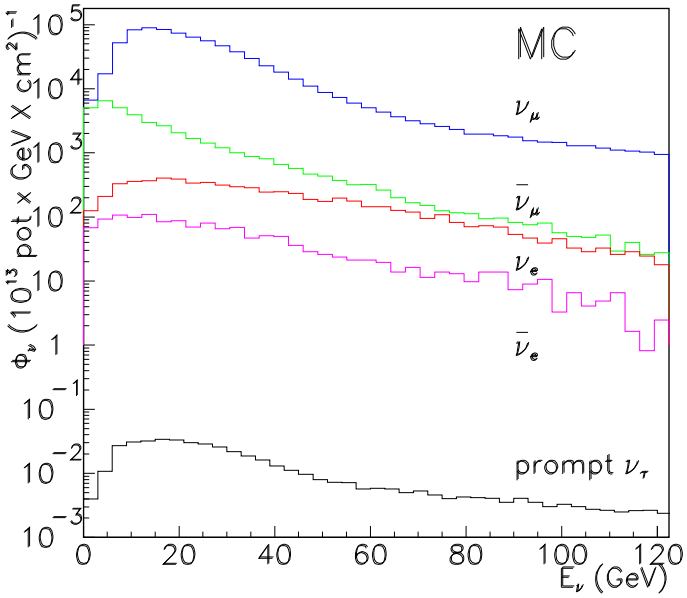

Fig. 3.5: Energy spectra of the different neutrino components in the wide band beam at the location of the CHORUS emulsion target (transverse dimensions $1.44 \times 1.44\,\mathrm{m}^2$), as derived from the neutrino beam simulations.

	$\langle E_\nu \rangle$	relative abundance
ν_μ	26.1	1.000
$\overline{\nu}_\mu$	19.4	0.057
ν_e	39.8	0.008
$\overline{\nu}_e$	28.5	0.002

Tab. 3.1: Neutrino beam composition

and Fig. 3.5 respectively. The fluxes have been calculated at the surface of the CHORUS emulsion target, using the beam simulation [87]. The integrated flux of the main component (ν_μ) is estimated to be $4.6 \cdot 10^{-3}\,\nu/pot$.

The $\overline{\nu}_\mu$ contamination is due to decays of negatively charged mesons, which have not sufficiently defocused by the horn and the reflector, or have decayed before being absorbed by the collimators or the NBC walls. Electron neutrinos and antineutrinos come mainly from kaon decays. A small fraction of them is due to muons decaying before reaching the iron shield.

Tau neutrinos in the beam constitute an irreducible background for the oscillation

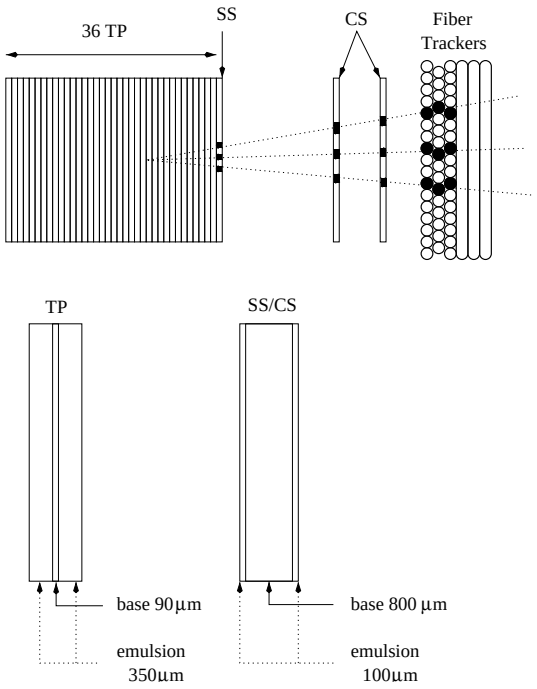

Fig. 3.6: Schematic view of an emulsion target module. The 36 Target Plates (TP), are followed by one Special Sheet (SS) and two Changeable Sheets (CS), which serve as an interface between the target plates and the fiber tracker.

search. They are produced in the Be target and the beam dump from decays of charmed mesons ($D_s^{\pm} \to \tau^{\pm}\nu_\tau$) produced by proton interactions with the Be and Fe nuclei. From calculations on the ν_τ flux [88, 89] about 0.1 background events for the total CHORUS data taking are expected.

3.3 The emulsion target

The heart of the CHORUS experiment is its 770 kg active target [91] consisting of emulsion sheets. It fills a total volume of 206 lt, and covers an area transverse to the neutrino beam of $1.44 \times 1.44\,\mathrm{m}^2$. It is segmented longitudinally in 4 stacks, each of them followed by a scintillating fiber tracking system.

Every stack is divided in 8 modules with transverse dimensions of $72 \times 36\,\mathrm{cm}^2$. A schematic view of target module is depicted in Fig. 3.6. The bulk mass of a target module is formed by stacking 36 plates, each of them having two 350 μm emulsion layers on a 90 μm tri-acetate cellulose foil. During development the emulsion shrinks and gets distorted, so that the angular resolution of the tracks identified in these plates is not better than 5 mrad. Such a resolution does not allow an unambiguous matching of the tracks found in the emulsion to these reconstructed by the fiber tracker.

An improved angular resolution of $\sim 1\,\mathrm{mrad}$ is achieved when thin layers of emulsion

(100 μm) are poured on a thicker acrylic base (800 μm). Such a plate, hereafter quoted as the Special Sheet (SS), is placed downstream the 36 target plates and provides an interface to the electronic tracker.

To further reduce mismatches between fiber tracker and emulsion tracks two Changeable Sheets (CS) are placed between the SS and the fiber tracker. They are identical to the SS, but they are not stacked together with the rest of the planes. Moreover, they are exchanged regularly, so that their track multiplicity is rather low. The relative alignment among the CSs and the SS is realized by the means of X-ray "guns" (brass cylinders filled with ^{55}Fe sources) that produce a pattern of black dots of 1 mm diameter in each pair of sheets. The relative alignment between adjacent plates is performed using tracks of passing through muons.

The emulsion plates used for CHORUS have been produced in a laboratory set up at CERN. After pouring the gel on the plastic base, it was left to dry for 2 days at 20 °C with a relative humidity of 89%. The 36 TPs and the SS are packed together in a thin shielding envelope of aluminized paper, to preserve their initial water content. To avoid emulsion fading, the emulsion stacks in the experiment have been kept at a temperature of (5.0 ± 0.5) °C in a cool box containing the whole target area, the hexagonal magnet, the associated fiber trackers and the trigger planes, as shown in Fig. 3.2.

The TPs and the SSs of all modules have been exposed in the neutrino beam for two years (1994-1995), while for the 1996-1997 run they have been replaced. The CSs used for the first run were replaced every three to six weeks. For the second run only one CS was present per module, instead of two, and it was replaced only once per year. This has been done because the background from secondary muon beams in the SPS West Area has been reduced, and the quality of the emulsion gel was improved (the ET-7B by FUJI, Japan replaced the BR-2M by NII KHIMFOTO PROEKT PLC, Russia).

The emulsion plates were developed also at CERN. Before development, a grid of fiducial marks is printed on every TP and SS using ultra-violet light. They provide a solid coordinate system on each plate, and the effects of transverse shrinkage and distortion can be corrected during the scanning procedure, by measuring the position of the fiducial marks closest to the track being searched for.

3.4 The electronic detectors

3.4.1 Target tracker

The major roles of the target fiber tracker [90] are to reconstruct the position of the neutrino interaction vertices and to make accurate predictions for track locations and directions at the CS interface emulsion sheets, both of which are essential for track scanning and following in the emulsion target. The tracker is built of scintillating fibers because they offer the two-track and spatial resolution necessary for predicting accurately real tracks.

The target tracker (TT) is formed of planes made of 500 μm diameter fibers. Each plane is a ribbon of 7 staggered fiber layers. A tracker module consists of planes oriented in four projections denoted by Y, Z (vertical and horizontal axis normal to the beam direction)

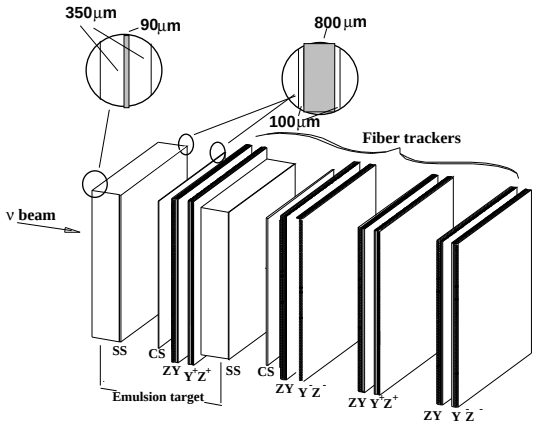

Fig. 3.7: Geometric structure of the CHORUS target.

and $Y^{\pm}$, $Z^{\pm}$. The latter two are rotated by $\pm 8°$ with respect to Y and Z respectively. There are 8 tracker modules which alternate with the emulsion stacks as shown in Fig. 3.7.

At the readout ends the fiber ribbons are bundled together and coupled to opto-electronic readout chains, each of them consisting of four image intensifiers and a CCD camera in series. The CCD camera contains an image and a memory zone allowing the system to record two events during a neutrino spill.

Tracks reconstructed in the TT system are extrapolated upstream to the CS with a position and angular resolution of $\sim 150\,\mu$m and $\sim 2.5\,$mrad respectively.

3.4.2 Hadron Spectrometer

Downstream the target region a magnetic spectrometer of hexagonal shape is used to measure the momenta of charged particles before they reach the calorimeter. It is strongly connected to the $\nu_\mu \rightarrow \nu_\tau$ oscillation analysis, where the tau decays through a hadronic channel; through the momentum measurement of potential tau daughters, background processes can be suppressed posing kinematical criteria.

The spectrometer is based on an air-core hexagonal magnet [92], shown schematically in Fig. 3.8. It is made of six equilateral triangles with 1.5 m sides, 0.75 m in depth, and housed in a cylinder with 3.6 m diameter. Windings of thin aluminum sheets cover all faces of the triangles, producing in each triangle section a homogeneous field parallel to the outer side. The magnet current is pulsed simultaneously with each proton extraction, and develops a magnetic field of 0.12 T with no radial dependence.

The main tracking detector associated to the hadron spectrometer is the diamond-shaped fiber tracker (DT). It is made of the same fibers, and assembled and read out the same way as the TT. There are three DT modules, one upstream and two downstream the magnet. Each module consists of two planes formed by three adjacent *paddles* giving measurements of three coordinates at 120° with respect to each other. The second plane is rotated by 60° with respect to the first (see Fig. 3.9).

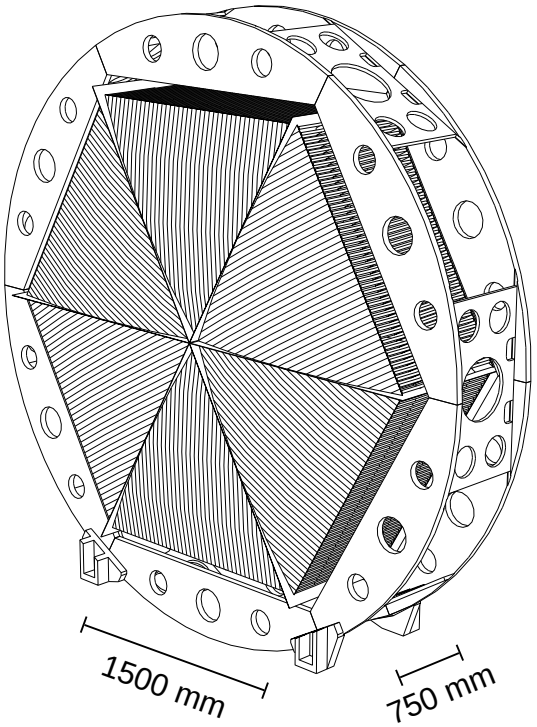

Fig. 3.8: Air-core hexagonal magnet

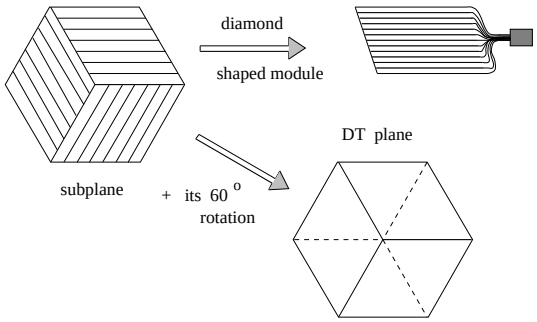

Fig. 3.9: Diamond Tracker paddle geometry

The momentum resolution $\Delta p/p$ of the hexagonal spectrometer using the DT is

$$\left(\frac{\Delta p}{p}\right)_{data} = \sqrt{(0.22)^2 + (0.035p)^2} \tag{3.1}$$

where p is given in GeV. The constant term is due to the multiple scattering in the traversed material, while the one proportional to the actual momentum reflects the measurement accuracy. The sign of the charge for particles with momentum of 5 GeV can then be determined better than the three standard deviation level.

To extend the measurable momentum range of hadrons up to 30 GeV, and hence increase the $\nu_\mu \rightarrow \nu_\tau$ oscillation sensitivity for the $\tau^- \rightarrow h^-$ channel by $\sim 26\%$, the spectrometer was equipped for the 1996-1997 data taking with emulsion trackers (ET). There are three ET planes, as shown in Fig. 3.2. One is placed immediately downstream the TT region, while the other two occupy the space between the magnet and the neighbouring DT trackers. Each plane is composed of juxtaposed emulsion sheets similar to the CS/SS, which were replaced once per year.

For the 1996-1997 run, honeycomb chambers (HC) [93] replaced the existing streamer tube (ST) planes located between the hadronic spectrometer and the calorimeter. They complement the magnet trackers by providing a module with independent and standalone tracking, resolving track ambiguities especially in high multiplicity events. The three honeycomb chambers are mounted at 60°, 0° and 120°. Each chamber is composed of six planes of hexagonal cells with a wire spacing of 11 mm. The HC cover a total area of $3 \times 2.75\text{m}^2$, matching the aperture of the hexagonal magnet.

3.4.3 Calorimeter

One electromagnetic and two hadronic sections compose the CHORUS calorimeter [94]. It is used to measure the total energy and direction of electromagnetic and hadronic showers, and to identify neutral particles. The electromagnetic and the upstream hadronic sections use the *spaghetti technique*, where scintillating optical fibers are embedded in lead. For the downstream hadronic section, the classical sampling approach was adopted, where lead layers alternate with planes of scintillator bars. Both sides of each calorimeter module are read out by photomultipliers. In total the calorimeter corresponds to 144 radiation lengths and 5.2 interaction lengths.

The resolutions obtained in testbeams are

$$\frac{\sigma(E_{had})}{E_{had}} = \frac{32.3 \pm 2.4\%}{\sqrt{E(GeV)}} + (1.4 \pm 0.7)\% \tag{3.2}$$

and

$$\frac{\sigma(E_{em})}{E_{em}} = \frac{13.8 \pm 0.9\%}{\sqrt{E(GeV)}} + (-0.2 \pm 0.4)\% \tag{3.3}$$

where E_{had} and E_{em} refer to the hadronic and electromagnetic energy respectively.

As shown in Fig. 3.10 a group of streamer tube planes have been inserted at regular intervals in the calorimeter to reconstruct trajectories of particles that are not drawn in a shower. Their main purpose is to link muons identified in the muon spectrometer with tracks reconstructed in the upstream trackers.

3.4.4 Muon spectrometer

The main function of the CHORUS muon spectrometer is to identify muons in neutrino interactions and to measure their charge and momentum. It uses the calorimeter modules

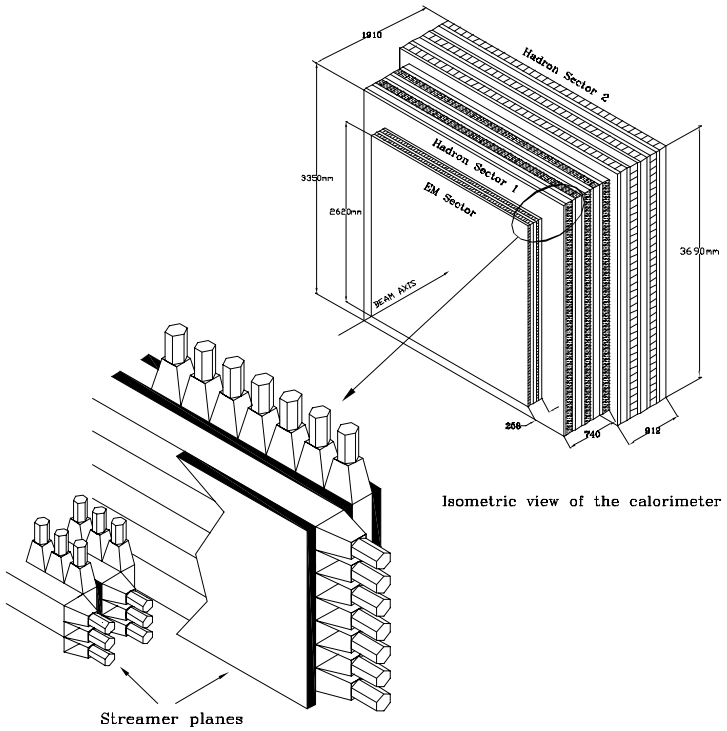

Fig. 3.10: Geometric structure of the CHORUS calorimeter

of the CDHSW detector [95], which were later reused by the CHARM II experiment [96]. It consists of six magnets composed of 20 iron disks of 2.5 cm thickness. A current of 700 A along four water-cooled copper coils generates a toroidal magnetic field with a strength of $0.85\,\mathrm{T}\cdot\mathrm{m}$ per magnet. The polarity of the current is selected such that negative (positive) muons are focused (defocused).

The magnets alternate with seven tracking modules, composed of drift chambers and streamer tubes. As shown in Fig. 3.11, scintillator strips are sandwiched between the iron disks to get fast signals for standalone trigger. They are also used for calorimetry, and pattern recognition purposes.

The spectrometer's resolution is

$$\frac{\Delta p}{p} = \sqrt{(0.165)^2 + (0.0018p)^2} \tag{3.4}$$

where p is given in GeV. For low energy muons (below 5 GeV) the momentum measurement is improved by using the range of the muons in the magnets.

3.4.5 Trigger and Data Acquisition

The CHORUS trigger system [97] was designed to select neutrino interactions in the emulsion target region. The main difficulties in its design came from the need to form a *strobe*

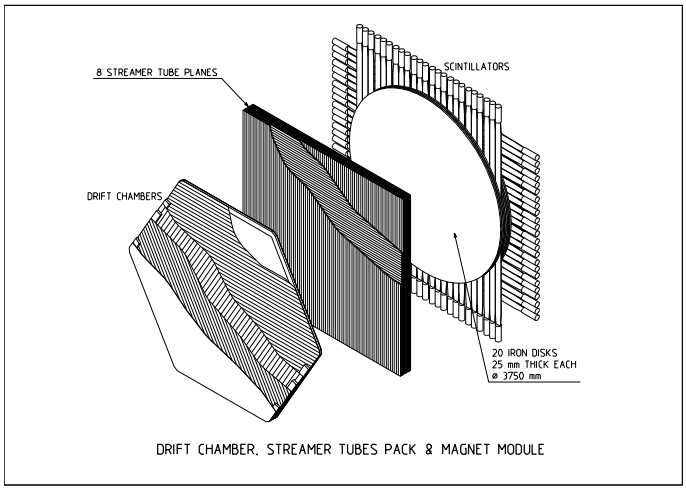

Fig. 3.11: Geometric structure of a spectrometer magnet unit

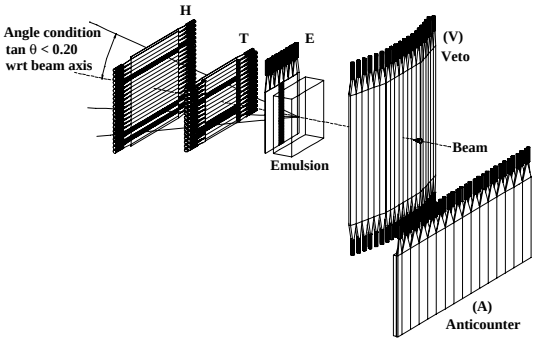

Fig. 3.12: Trigger planes of the CHORUS detector

(i.e. a timing signal to synchronize logic decisions), and the limit of two events per neutrino spill, imposed by the opto-electronic system used to read out the scintillating fiber tracker. As is shown in Fig. 3.12, the trigger system is realized by a set of scintillator planes, denoted by H (hodoscope), T (timing plane), E (emulsion plane), V (veto plane), and A (anti-counter).

The oscillation search trigger is initiated with a hit coincidence in E, T and H planes and no activity in the veto system (V, A), within the time gate of a proton spill. The combination of T and H strips requires consistency with a particle trajectory with $|\tan\theta| < 0.20$, where θ is the polar angle with respect to the neutrino beam direction. The presence of the E plane reduces the trigger rate from neutrino interactions occurring in the iron support structure of the target region and the image intensifier chains used in fiber tracker readout. The anti-counter A vetoes the events of neutrino interactions in the concrete floor upstream of the detector. To avoid rejecting valid neutrino interactions, where a back-scattered particle reaches the veto system the relative timing with respect to the T

plane is used to identify such cases.

The final trigger decision also requires that at least one calorimeter plane or one spectrometer magnet scintillator is hit. All these criteria result in a trigger rate of 0.5 events per 10^{13} protons on target (pot), with a total dead time of 10%. Eventually, about 45% of the triggered events correspond to neutrino interactions inside the emulsion target.

Apart from the main oscillation trigger, up to 16 different triggers may be formed for various physics analyses. Such are the alignment, the calorimeter charged-current and the Neutral Heavy Lepton triggers.

The CHORUS data acquisition (DAQ) system [98] was constructed as much as possible on machines running a standard UNIX operating system. For the low-level *real-time* parts, however, the OS-9 kernel was used, which guarantees sufficiently short interrupt response times.

The essence of the DAQ hardware is formed by a tree of VME CPUs, interconnected via the VME Inter Crate (VIC) bus, as well as the Ethernet. The branches of the tree are formed by the four detector subsystems: opto-electronics (OPTO), trigger (TRIG), calorimeter (CALO), and muon spectrometer (SPEC). The local DAQ of each subsystem is equipped with a local trigger system, such that it can operate in stand-alone mode. In global mode, the data from all subsystems are assembled at the root of the tree by the event builder (EVB), as soon as a trigger condition is met.

The EVB delivers the full events as messages to a dispatcher process on a UNIX machine, that makes the data available to client programs, via shared memory or Ethernet. Programs which listen to the dispatcher are for example the process which stores the actual data on tapes, and the various programs monitoring the performance of each subdetector.

The Dispatcher [99] is a general-purpose message-based data distribution program, which is also used in the DAQ system of the CERN microscopes (see chapter 5). When a client opens a connection to a dispatcher, it subscribes itself to a list of message tags. A tag indicates the type of the data contained in a particular message. Whenever a message is sent to the dispatcher, a copy is made available to those clients that are subscribed to the corresponding tag. Messages are also used to send commands to DAQ programs. Such an architecture allows the user interface to be completely decoupled from the actual DAQ process, which does not need to be running on the same machine.

3.5 Background processes and the (estimated) oscillation sensitivity

The CHORUS proposal as well as the results of the $\nu_\mu \to \nu_\tau$ oscillation search which have been reported in the first two publications [100, 101], are based on the tau decay modes into a muon or a single hadron:

$$\tau^- \to \mu^- \overline{\nu}_\mu \nu_\tau$$

and

$$\tau^- \to h^- n\pi^0 \nu_\tau$$

where h is a pion or a kaon.

The muonic and single-prong hadronic channels represent about 18% and 50% of the total tau decays [105]. However, the overall sensitivities are comparable for both channels, because of the smaller efficiencies of the hadron track reconstruction.

In principle, the oscillation search can be extended to include also the decay modes $\tau^- \to e^- \overline{\nu}_e \nu_\tau$ [103] and $\tau^- \to \pi^-\pi^+\pi^- n\pi^0 \nu_\tau$ [104]. They are not included yet in the present analysis, mainly because of the more complicated scanning procedures they require.

For the muonic channel, decay kinks are searched for along the trajectory of a negatively charged particle which is identified in the muon spectrometer as the primary muon of a ν_μ CC interaction. The main background sources for this channel are:

- The ν_τ contamination of the WANF neutrino beam.
- Short kaon decays

$$K^- \to \mu^- \overline{\nu}_\mu$$

- Large angle muon scattering.
- Production of charmed mesons from the antineutrino components of the beam:

$$\overline{\nu}_\mu(\overline{\nu}_e) N \to \mu^+(e^+) D^- X$$

  followed by

$$D^- \to \mu^- X^0$$

  where the $\mu^+(e^+)$ escapes detection or is misidentified.
- Associated $c\overline{c}$ production in NC interactions, where the negatively charged charmed meson (a D^-) and its subsequent decay is detected, while its associated charmed meson is missed.
- The production of positive charmed mesons in ν_μ CC interactions, if the primary muon is not identified and the charge of the charmed particle daughter is incorrectly measured.

The following cuts are applied during the event selection, in order to suppress the background rates and the scanning load:

- The candidate muon should have a reconstructed momentum of less than 30 GeV.
- In case a decay topology is found, the muon transverse momentum P_t with respect to the τ candidate direction should be larger than 250 MeV. Using this cut the background from kaon decays is eliminated.

- The tau path length should be less than 3.95 mm (corresponding to a maximum of 5 emulsion sheets from the interaction vertex).

For the hadronic channel, the potential tau daughters are searched for among negatively charged hadrons from NC interactions. The main background sources for the hadronic channel are:

- As in the case of the muonic channel, the ν_τ contamination of the neutrino beam, and the single-prong charm decays into a negatively charged hadron.
- The so-called hadron *white kinks*, defined as single-prong nuclear interactions with no heavily ionizing tracks (*black* or *gray* tracks in emulsion terminology) and no evidence for nuclear break-up (evaporation tracks, recoils, blobs, or Auger electrons). These interactions constitute the main background to the hadronic tau decays. The difficulty treating them is due to the large uncertainty in the white kink cross-section. An estimation of this cross-section has been obtained by test beam measurements and Monte Carlo simulations (see chapter 6).

In order to suppress the background sources in the hadronic channel, apart from requirements concerning the quality of the reconstructed tracks, such as the minimum number of hits in the fiber tracker, the kinematical cuts applied also for the $\tau \rightarrow \mu$ channel are used, slightly modified though: The reconstructed momentum should be within the range of 1-20 GeV, and the decay kink should occur within three emulsion sheets downstream from one where the neutrino interaction vertex is located.

Table 3.2 summarizes the background sources and their estimated rates assuming the above analysis. At the end of the oscillation analysis in CHORUS less than two background events are expected. The proposed sensitivity of the experiment is shown in Fig. 3.13.

Channel	Background	Rate / 10^6 ν_μ CC events
$\tau^- \rightarrow \mu^-$	prompt ν_τ	< 0.005
	K decays	< 0.026
	μ scattering	< 0.017
	associated charm in NC	< 0.022
	single charm in antineutrino CC	< 0.46
$\tau^- \rightarrow h^-$	prompt ν_τ	< 0.006
	associated charm in NC	< 0.021
	single charm in antineutrino CC	< 0.48
	white kinks	1.5

Tab. 3.2: Background sources and rates for the $\nu_\mu \rightarrow \nu_\tau$ oscillation search in CHORUS.

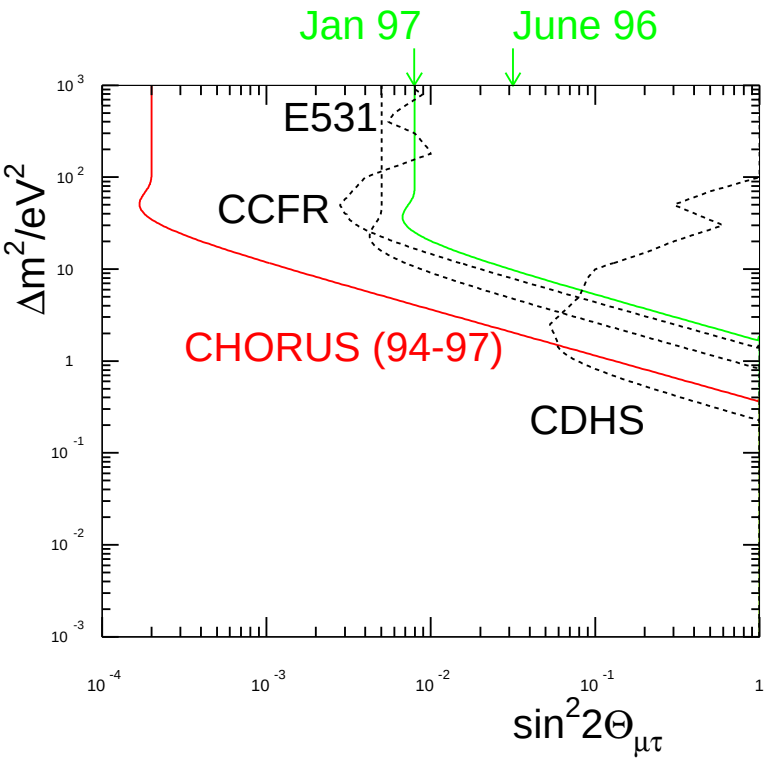

Fig. 3.13: Estimated sensitivity of the CHORUS experiment

4. THE EMULSION TECHNIQUE IN CHORUS

Nuclear emulsion is connected to many major discoveries in the early days of nuclear and particle physics. It still remains the detection technique with the best known three-dimensional spatial resolution, and zero intrinsic dead time. Its use has gradually decreased after the development of electronic detectors. Contrary to emulsion, electronic detectors offer immediate time correlated readout, and digitized output storeable to computer accessible media. They serve therefore better the needs of most present experiments, which require large statistics of accumulated data for prompt and accurate physics results.

Recent developments in automatic emulsion scanning, mainly from R&D efforts within the framework of the CHORUS experiment, have lead to a renaissance to emulsion as a particle detector. Automatic scanning allows for fast extraction of digital information from emulsion sheets, after they have been exposed to particle radiation. Not only does it make handling of large data sets possible, it also ensures that physics results can be produced shortly after the running of the experiment.

This chapter gives a short overview of the emulsion capabilities as a particle detector, and describes the principle of automatic scanning. The working example is the implementation adopted by group of the Nagoya University, Japan [110]. This choice is made because this group pioneered in the field of automatic emulsion scanning, and most of the techniques adopted by other laboratories have underlying principles inherited from Nagoya. The automatic microscopes developed at CERN are extensively described in chapter 5. Finally, it is described how different physics analyses and searches affect the scanning procedures.

4.1 The use of nuclear emulsion in particle physics

4.1.1 Nuclear emulsion as detection and storage device

Nuclear emulsion consists of about equal parts by volume of silver halide crystals, a few tenths of a micron in diameter, and a matrix material, which is chiefly gelatin. An ionizing particle on encountering a crystal may render it "developable": under the action of a chemical reducing agent, conversion of the halide to metallic silver will proceed more rapidly than in an unirradiated crystal. The physical condition in the crystal that makes it developable is called the "latent image".

After development, followed by washing to remove the undeveloped crystals, the gelatin is transparent. Using a microscope, the paths of charged particles which penetrated the emulsion, are visible as trails of minute silver grains. A true three-dimensional image is produced. The paths of particles, outlined by silver, literally exist in space.

The position of the silver crystals in the gelatin remain stable after development. Therefore, the information on the particle trajectories is permanently registered in the emulsion. This feature makes emulsion a high density and long–term data storage device as well.

Developed silver crystals cluster into grains which have a mean diameter of $\sim 0.8\,\mu$m. The grains are dispersed around the actual trajectory of a particle by $\sim 1\,\mu$m. For a MIP the grain density along the track is ∼300 grains/mm. For these reasons, nuclear emulsion has an intrinsic three-dimensional spatial resolution better than $1\,\mu$m, and a two-track resolution of the same order.

Angular measurements in emulsion are limited by the distortion during development. Distortion is a side effect of the shrinkage which the emulsion undergoes during development. Since the developed silver crystals correspond to only a negligible fraction of the total crystals, when the rest are removed from the gelatin during washing, the volume of the emulsion is decreased by a factor of two.

In principle the distortion could be parameterized and corrected for, using reference tracks of high energy particles. In practice, when precise measurement of the orientation of a track segment is required by an emulsion experiment, one resorts to the use of sheets with thin emulsion layers on a relatively thick transparent base. This is the case, for example, with the SS and CS and emulsion sheets in CHORUS (see 3.3). The position of the grains belonging to a track segment are less affected by distortion if they are close to the base. Their transverse position in the sheet reference system is known with a precision of $\sim 1\,\mu$m. The direction of a track segment crossing the sheet, can be measured then with a precision which improves with the thickness of the base. Such a type of slope measurement is quoted as the "base measurement".

For low energy particles emulsion can be used for momentum measurement from multiple scattering. One measures the direction of the track in various depths of the emulsion. The variance of the angular distribution, after unfolding the measurement resolution and distortion effects, can be fully attributed to multiple scattering.

Particle identification is also possible from dE/dx measurement. What is actually measured in this case is the grain density along a particle trajectory.

The most significant limitation in the use of emulsion is the fact that there is no real time readout of its response to particle radiation. One of the consequences is that for all tracks registered in it their time correlation is lost.

The second disadvantage is that registered tracks are continuously accumulated during exposure to radiation. The lifetime of an experiment is therefore limited by the track density tolerable by the scanning and analysis procedures afterwards. Since there is no way to veto response to any radiation, there is also background accumulation from sources other than the nominal one. For example, if an emulsion stack is exposed in a beam of particles produced by an accelerator, there is an irreducible background from cosmic rays and radioactivity from sources in the material surrounding the stack.

For the above reasons, the use of nuclear emulsion is restricted to a small class of experiments, which still require an excellent spatial resolution without having the need for high beam luminosities. Emulsion is the ideal detector for searches for rare topological signatures in an environment of low interaction rates. This is the case for experiments studying

the production of short–lived particles (eg. tau leptons, charmed or beauty mesons) in neutrino interactions.

4.1.2 Hybrid detectors

During the last two decades, nuclear emulsions have been used often together with fully electronic detectors. This *hybrid technique* allows experiments benefit from the use of emulsion, moderating though the impact of the above mentioned disadvantages.

A hybrid apparatus consists of an emulsion target followed by an electronic tracker. Magnetic spectrometers and calorimeters may follow as well. The target is segmented in sheets, which are usually mounted vertical to the beam. Such a configuration is preferable if automatic scanning is used to analyse the emulsions (see following section).

The electronic tracker is used to reconstruct the trajectories of particles originating from interactions in the emulsion target. For the tracks found in the emulsion which are matched to reconstructed ones in the tracker, their time stamp which is lost in the emulsion can be recovered from the electronically recorded data.

In practice, only those tracks predicted by the tracker are searched for in the emulsion. Tracks are extrapolated from the tracker onto the most downstream emulsion sheet. The scanning is performed around the predicted position searching for tracks with a direction similar to that of the predicted track. The total area and angular space covered are determined by the extrapolation errors, which are determined by the resolution of the tracker. This practice leads to a significant reduction of the scanning load and to high background suppression.

The rest of the electronic detectors are used for the reconstruction of the kinematics of the events, and for particle identification. Using event and track pre-selection criteria for the scanning, the background rates and the scanning load can be further reduced.

Some of the past experiments that adopted the hybrid technique are WA17 [106], WA75 [107] at CERN, and E531 [108], E653 [109] at Fermilab. All of them were aiming at detecting short lived particles produced in their emulsion targets by pion or neutrino beams.

4.2 Automatic scanning systems

4.2.1 The Principle of Automatic Emulsion Scanning

Conceptually, a microscope for automatic emulsion scanning consists of a computer driven mechanical stage, the appropriate optical system, a photodetector –typically a CCD[1] camera– and its associated readout. Digital images are analysed for the recognition of the silver grains inside the developed emulsion gel.

A *microscope view* is a series of successive images taken along the beam direction. Grains belonging to different planes can be combined to form a straight line, of which the parameters (slope and position) are compared to those of a track reconstructed by the electronic trackers.

[1] Charge Coupled Device

In practice, the microscope scans some area around a nominal position at the emulsion sheet, searching for straight tracks segments of a nominal slope. For the most downstream sheet these nominal values are given by the parameters of the tracks reconstructed in the electronic detectors. For all other sheets, the updated parameters from a segment found in the previously scanned sheet are used. The number of views required to complete the scanning of a prediction is determined not only by the prediction precision, which defines the area of emulsion surface to be covered, but also by the size of the camera sensitive area, and by the size of the field of view given by the optics.

Since the emulsion sheets are exposed perpendicular to the neutrino beam at the actual experiment, the tracks from neutrino interactions are oriented around the right angle to the sheet plane. Most reconstructed tracks in the CHORUS electronic detectors which are interesting for the physics analysis have a polar angle of less than 400 mrad. Segments of such tracks in the emulsion sheets are fully contained within a single microscope view.

Data taking of a view is performed with the stage driving the objective at a constant speed along the perpendicular to the emulsion sheet. At the same time the camera captures emulsion images at the objective's current focal position. The motor speed and the camera frame rate determine the spacing between successive planes. This spacing should be such that the information extracted from every plane is independent. Thus it should not be smaller than the optical depth of focus.

Track reconstruction is performed while the stage moves to the central position of the next view, and completed before a new view can be acquired.

4.2.2 The Implementation in Nagoya

The microscopes developed in Nagoya [110] are based on a NIKON stage capable of handling emulsion sheets with dimension up to 36×35 cm^2. The three axes of the stage are driven by micro-stepper motors, and their position is read out by linear encoders. The position accuracy of the horizontal X,Y axes and the vertical Z axis is 1 μm and 0.25 μm respectively. The computer operating the microscope drives the stage through an ISA[2]-based controller.

The NIKON optical system is fixed-focus, as the depth-focusing is done by moving the Z-stage, where the objective is mounted. The condenser height is tuned so that the light from the source is focused at a position in the middle of an emulsion sheet mounted on the stage. The optics houses a Tiyoda 50$\times$ oil immersion objective. It delivers a flat field with diameter of 200 μm, and depth of focus of about 3 μm.

The light source is a 100 Watt Halogen lamp with a green filter. The light intensity is controlled by an analog signal fed to the lamp power supply, generated by an ISA-based DAC[3] card. The purpose of the light control is to keep the same light conditions at both sides of an emulsion sheet, independent of the darkness of the emulsion.

The image collected by the objective is focused on a SONY CCD camera with 485 lines of 640 pixels each. It operates at 120 Hz and has an electronic shutter which leaves the

[2] Integrated System Architecture
[3] Digital–to–Analog Converter

active elements of the CCD under light exposure for 1 ms. The readout is done via two separate ports, one for the even and one for the odd lines of the image. During the readout each port supplies 35 Mpixels/s as analog pulse-height information.

The emulsion readout is based on the *Hardware Track Selector* (TS) [111], a bus system equipped with a number of electronic modules responsible for both image analysis and track reconstruction. The TS is connected to the computer driving the microscope with its ISA bus.

A microscope view consists always of sixteen frames. The signal from each output port of the CCD camera is first amplified and then driven to an 8-bit FADC[4]. In the digital image thus obtained, the grains in-focus are seen as small objects with high contrast to the background, typically occupying not more than three pixels along a line. To enhance the grains in the emulsion image, a high-pass digital FIR[5] filter is applied along every line. After a threshold cut of the 4-bit filter output, the pixels belonging to grains are identified and assigned a gray value of 1, while the rest are reset to zero; hence a binarized image is formed. In the binarized image the grains are expanded by one or more pixels around their circumferences. This is done for the track finding to be less influenced by emulsion distortion and the physical grain dispersion along a particle trajectory. The binarized images are zero-suppressed and the positions (column, row) of the signal pixels are stored in FIFO[6] buffers to be used in the next TS operation, the track finding. The treatment of a single emulsion image is completed within the 8 ms of the camera cycle.

The track finding procedure starts as soon as the output from the last image is available. If the binarized images are summed, after applying to each of them a shift proportional to the opposite of the nominal track angle, candidate tracks can be identified as pronounced clusters on the sum image. The summing is performed using processor elements based on FPGA[7] chips. The clustering on the sum image is done on a Pentium processor integrated in the same board as the rest of the processor elements.

For normal prediction scanning, covering 25 mrad around the nominal slope, relatively few processor elements can try out all shifts in less than 300 ms. If 40-50 processor elements are used in parallel, up to 400 mrad can be fully covered within the same time. This enables general scanning of all potentially interesting tracks within every microscope view, at a frequency limited only by the camera frame rate and the stage speed. Such a system is operational since the beginning of 1999. It should be noted that the scanning speed of these system is almost 10 and 100 times larger than the ones available two and four years ago respectively. In the Nagoya emulsion laboratory all scanning stations are gradually upgraded to implement the latest version of the TS. These scanning stations are expected to offer the most significant contribution to the next phase of the CHORUS data analysis.

In summary, using 140×160 μm^2 views, the TS technology and the light mechanics of the microscope stage and the optical system allow for real time general scanning of tracks with slopes up to 400 mrad at a rate of 3 Hz. Despite the hardware limit of the

[4] Flash Analog–to–Digital Converter
[5] Finite Impulse Response
[6] First In First Out
[7] Field Programmable Gate Array

sixteen planes per view, the achievable measurement accuracy ensures high track finding and vertex reconstruction efficiencies.

4.3 Scanning strategies in CHORUS

The scanning strategy for the $\nu_\mu \rightarrow \nu_\tau$ oscillation search which has been followed by CHORUS so far (*Phase I analysis*) is described in Refs. [100, 101, 102]. The main steps are summarized here.

4.3.1 Reaching a vertex

For every neutrino event which has been reconstructed in the electronic detectors, one identifies the tracks of potential tau daughter particles. The selection criteria of these tracks are described in 3.5. The selected tracks have to be followed in the emulsion target, until a decay topology or the primary interaction vertex is reached. This constitutes the so–called *scan-back* procedure.

The scanning of a CS starts by covering an area around the predicted position which is larger than the one defined by the extrapolation errors. This procedure which is followed for a few hundred predictions aims at finding the proper calibration and alignment constants of the sheet with respect to the TT detector. Once a good calibration has been achieved, for the rest of the predictions a smaller area ($\sim 1\,\mathrm{mm}^2$) is analysed. This reduces the possibility of picking up fake or background tracks, which will also have to be followed in the emulsion target. The local reference frame of the CS is defined by the set of X-ray marks (see 3.3). These are always measured before the scanning starts. The coordinates of all tracks found are transformed in this local reference frame before storing them for subsequent use.

A track segment in the CS is accepted whenever a pair of segments found in the upstream and downstream emulsion layer of the sheet, define a segment of which the slope is within some acceptance cone around the predicted direction. This cone has usually an opening angle of $\sim$30 mrad. Before passing through this selection criterion, the slope of the segments found in the emulsion layers are corrected by a factor (≈ 0.5) which accounts for the shrinkage of the emulsion during development. For every prediction, the downstream surface is initially scanned for segments. For every segment found there, one extrapolates its measured position to the upstream surface along the predicted angle. If a segment corresponds to the actual track which is searched for, then its counterpart in the upstream side should be found close to the extrapolated position. In that case, the slope which is eventually stored, is a result of a *base measurement*.

Track segments that have been found on the CS are extrapolated onto the SS surface. The predictions are more accurate than the ones for the CS, thanks to the higher precision of the measurements in the emulsion sheets. Therefore, after calibration and alignment with respect to the CS, the total area to be scanned for each prediction is reduced down to $\sim 0.2\,\mathrm{mm}^2$. The scanning procedure itself is otherwise identical to the CS case.

When track segments found in the SS are extrapolated to the first TP, the predictions are very accurate in position, because the plates are touching each other. This also holds for segments extrapolated from a TP to the subsequent TP upstream. After calibration and alignment with respect to the downstream sheet, track segments are found close to their predicted position with a variance of less than $3\,\mu$m. When scanning a TP, a thin emulsion layer in the most upstream region of the sheet is analysed. Angular and position measurements in the TPs are sensitive to distortion, which is not only the result of the emulsion development. Environmental conditions during the scanning may change, and in any case they are different from the ones during the exposure in the neutrino beam: the temperatures are $\sim 25°$ and $\sim 4°$ respectively. For this reason, before searching for a predicted track, one measures the positions of the three fiducial marks which are closest to the predicted position. These fiducial marks define a local system of reference which is used to store found track segments.

If a track, which has been found in at least three consecutive TPs, is lost for two subsequent upstream ones, the first TP where this occurs is defined as the *vertex plate*. If the loss of the track is not due to scanning inefficiency, then in that TP there is a vertex, which is either the primary one of the neutrino interaction, a secondary re-interaction, or a decay.

4.3.2 Searching for decay topologies

There are two major classes of searches for decay topologies (kinks) once the vertex plate has been identified:

- The *short decay path search*, where the primary and the decay vertices are assumed to appear both in the vertex plate.
- The *long decay path search*, where in the "vertex plate" there is only the decay kink, while the primary vertex appears one or more plates upstream.

The treatment of the different decay topologies in the case of the muonic channel is shown in Fig. 4.1

For the short decay path search, the method which is followed depends on the number of tracks reconstructed by the TS in the sheet downstream the vertex plate. If there is at least one additional track segment that can be matched to the TT predictions (Fig. 4.1, topology a), the impact parameter of the potential tau daughter is evaluated as the minimum distance between this scan-back track and any other track of the event. Events with impact parameter greater than $2\,\mu$m are selected for visual inspection. If no matched track segment is found, then digital images of the microscope fields over the whole depth of the vertex plate are recorded. An off-line (*video image*) analysis is performed afterwards to check for the existence of a decay topology.

The long decay path search is performed according to the kink angle. If it is larger than the scan-back angular tolerance, the scan-back procedure stops, and the kink plate is assumed to be the vertex plate (Fig. 4.1, topology b). If no other emulsion tracks

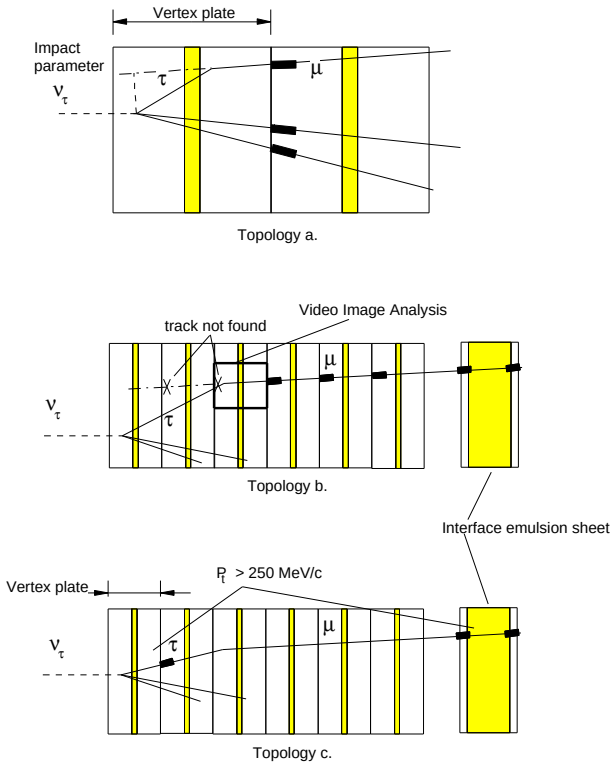

Fig. 4.1: Search for kink topologies, as described in the text.

can be matched to the TT predictions to apply the short decay path search, the video image analysis is undertaken. In case the kink angle is smaller than the scan-back angular tolerance (Fig. 4.1, topology c), the vertex plate is indeed the interaction plate and decays can be detected by measuring a track direction in the plate immediately downstream of the vertex plate which is not compatible with measurements in the other detectors (SS, CS, TT).

The treatment of the data taken from 1995 and onwards has been dominated by the *parent track search* method: Whenever a large angle decay is found, the kink plate is assumed to be the vertex plate (Fig. 4.1, topology b). For the upstream part of the vertex plate *general scanning* (search for any track segment with a polar angle of less than 400 mrad) is performed. This is done in order to check the possibility of associating a segment with the scan-back track. If the minimal distance between a track segment and the scan-back track is less than 15 μm, the former is considered to be the *parent* of the latter.

4.3.3 General search for the production of short-lived particles

For the *Phase II analysis* of the CHORUS data, a more general method in searching for decay topologies is adopted. This method, quoted as "NET-scan", has been developed

also in view of studying the production of charmed particles, including neutral ones, in neutrino interactions.

For every event the largest possible number of electronically reconstructed tracks will be followed until the primary interaction vertex is reached. As soon as the TP with the primary vertex has been identified, general scanning is performed for the seven subsequent downstream plates. The scanning area covered per plate is such that a cone is formed, which starts from the primary vertex, and has an opening angle of 400 mrad.

The results of the general scanning is the set of all track segments found in real-time by the TS. These are analysed off-line to check for the existence of decay topologies. From the track segments reconstructed by the TS in the various depths of the emulsion target, passing-through tracks are discarded. For the rest, it is checked whether they can be originating from common vertices. Therefore, not only the primary interaction vertex is fully reconstructed, but one can also check for the existence of a secondary re-interaction or decay topology.

The driving force behind the NET-scan method is the performance of the TS technology, which has been increasing continuously over the last years. This was also the case in the past with the adoption of the parent track search method for the data taken in 1995 and onwards.

5. THE SCANNING FACILITY AT CERN

The scanning facility at CERN [112] was set up in mid 1996. The first microscope was very similar to the existing microscopes in Nagoya, which served as a training device to the scanning techniques. The early status and performance of this system is described in Ref. [113]. Since then the CERN and NIKHEF groups of CHORUS built another three scanning stations, made of completely new hardware elements. Moreover, new ideas in the scanning strategies have been applied. Ref. [114] is an early status report of the new microscopes. However, neither the hardware configuration, nor the data acquisition scheme described there represent the current situation. This chapter describes the hardware and software developments which made the scanning facility at CERN operational.

5.1 Hardware of the microscopes

5.1.1 The laboratory layout

The scanning facility at CERN occupies a total area of $45\,\mathrm{m}^2$. A windowed wall separates the clean room where the microscopes are installed from the rest of the laboratory. In the clean room the environmental conditions are controlled and monitored. The temperature there is kept at $(21 \pm 1)\,^{\circ}\mathrm{C}$, the relative humidity at (70±5)%. The room is kept at a small over-pressure to avoid dust floating in the air. For safety reasons, smoke detectors have been installed as well. Controlled environmental conditions are essential for the emulsion sheets, which are stored in aluminum boxes in the clean room.

In the main working area of the laboratory there are the computers driving the microscopes and a few other PCs which are used for software development and to run monitoring programs. The old NIKON microscope, which now serves as a development system, is in a small separate room with a standalone air-conditioning. All the computers which need not be inside the scanning laboratory, such as the database server, the dispatcher host, and the Linux farm, are installed in a room next to the scanning laboratory.

5.1.2 Mechanics of the microscope stage

Microscopes based on a NIKON stage can only scan less than half of the surface of a CHORUS emulsion sheet at a time. A sheet has dimensions of up to $73 \times 37\,\mathrm{cm}^2$, while the acceptance of the NIKON stage is $34 \times 34\,\mathrm{cm}^2$. This implies that the mounting of the emulsion sheets and the calibration procedure have to be performed twice. In addition some of the surface cannot be covered with two separate mountings. This increases the

total scanning time and introduces inefficiencies. To overcome this problem, the CERN group of CHORUS decided to develop in collaboration with MICOS[1] a new stage, which covers a transverse area of approximately $80 \times 40\,\mathrm{cm}^2$, therefore allowing for the scanning of an entire CHORUS emulsion sheet. The vertical axis of these systems, carrying the objective and the camera, can move with a stroke of $\sim 20\,\mathrm{cm}$.

The desirable precision in the local motion of the stage is better than 1 μm. To preserve such a precision throughout the scanning procedure, the environmental conditions have to be kept well under control. The microscope stages and their granite supports have high temperature stability by design.

Controls

For every microscope stage, linear movement of each of the three axes is achieved with precision screws driven by micro-stepper motors. Linear encoders ensure the precise motion of the stage by continuously monitoring its absolute position with a resolution of $0.2\,\mu$m.

The whole stage is managed by a controller which is linked to the PC through an RS-232 serial interface. The commands issued to the controller by the PC and the corresponding replies of the controller are in ASCII format. The baud rate of the RS-232 communication is set by dip switches on the controller's panel. It can be set up to 19200 bps, but this value has lead to frequent errors and crashes. A rate of 9600 bps, which is recommended by the manufacturer, is used instead.

The controller is used to start, complete or interrupt an absolute or relative movement, read the current position of the stage and enable or disable joystick control. It is also used to set the velocity and acceleration values of the stage, the boundaries of allowed motion and the units expressing position, velocity and acceleration. The controller can also redefine the stage coordinate system by assigning specific coordinate values to a certain position.

For every axis, the stepper motor completes a full turn of the screw in 40000 steps. The pitch size of the screw is known to the stage controller and for all the available screws the pitch size is either 2 or 4 mm, corresponding to a step size of 0.05 or 0.1 μm respectively. This is in both cases smaller than the resolution of the linear encoder. The pitch size is used by the controller to transform the position and velocity from the base units of steps and steps/s to the more relevant units of μm and μm/s. It is also useful for the communication software running on the host PC to transform the value of constant rotational velocity into linear velocity.

The stage can move with a maximum speed of 80 mm/s. However, to ensure smooth movements without stressing the screws, we operate the stages at a speed of 30 mm/s. The acceleration used to reach the top speed can be within the range $6.4 - 640\,\mathrm{mm/s}^2$. High acceleration values result in oscillations of the long horizontal screw at the end of the movement. On the other hand, making use of the high stability of the screw driving the vertical axis, when issuing a command for a movement with a constant velocity along this axis (a typical action during image data taking), the acceleration is set temporarily to

[1] http::/www.micos-online.com

high values to ensure that the top speed is reached as soon as possible.

Every point on the horizontal plane is always approached from the same direction to avoid hysterisis effects of the stage linear encoder. For this reason, every movement on the horizontal plane is performed in two steps. The stage moves first towards a point which is translated by a *backlash* vector from the actual destination. The backlash movement ensures that the arrival of the stage to the actual target position.

When the controller has been turned off, it is necessary to calibrate the stage and reset the values of allowed range, velocity and acceleration. After slowly moving the stage towards the lower hardware limit switches of every axis, the controller marks the position where the stage has stopped as the origin of the coordinate system. This position becomes the initial lower motion limit. Such a limit is a software limit switch in the controller, as robust and safe as a hardware limit switch on the screw. These limits are automatically updated when the stage reaches the hardware lower or upper limit switches, but they are usually set to define a smaller range of motion, to avoid dependencies on the accuracy and reproducibility of the hardware switches. The communication software ensures that the software limits are properly recalculated every time the stage coordinate system is transformed.

The lower limit of every axis is set 1 mm away from the stage position after the calibration procedure. The upper limits can in principle be set in a similar way, by driving the stage towards the complementary hardware limit switches. Instead, nominal values for the upper limits are explicitly set in the controller's memory. This prevents the objective from accidentally moving too low, thus damaging the emulsion or the glass holder. On the other hand, when the microscope is at the lower limit of the vertical axis, it should still focus on a plane below the bottom emulsion layer. For the horizontal axes, the upper limits were set to at least 1 mm away from the hardware switch.

Mounting of the emulsion sheets

After development and cleaning, all emulsion sheets are framed in plastic holders. When they have to be scanned, they are mounted on the microscope stage on top of an aluminum framed glass plate. They are kept flat and pressed onto the glass plate using a vacuum system connected to the aluminum frame. The layout of the plastic holders and the glass plate is shown in Fig. 5.1.

The holders have dimensions $88 \times 45\,\mathrm{cm}^2$ wide. They consist of two plates of 1 mm thickness, glued together. The upper plate has a window corresponding to the size of the emulsion sheet. The bottom plate has a set of channels, 3 mm wide, which carry the vacuum. The emulsion sheets are mounted on the plastic holders using thin transparent tape applied on their edges.

The glass plate is supported by an aluminum frame which is mounted with high precision on the stage. The frame is 96 cm long and 56 cm wide, while the 4 mm glass plate covers an area of $79.5 \times 45\,\mathrm{cm}^2$. The space between the glass and the frame provides for the vacuum path. The vacuum itself is fed from one of the short sides of the frame. On the frame an O-ring prevents vacuum leakage between the surfaces of the aluminum frame

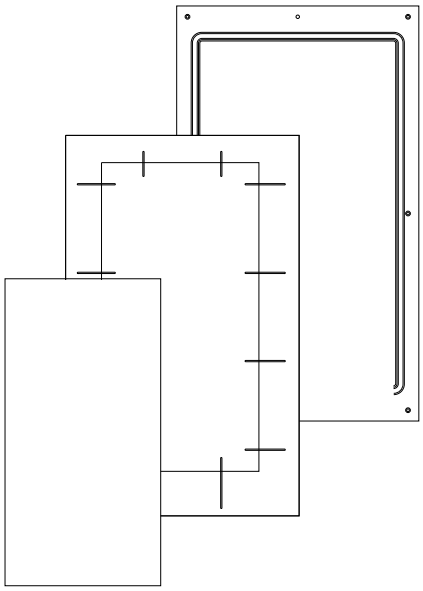

Fig. 5.1: The glass plate of the stage and the plastic holder of the emulsion sheet. The rectangle in front is a CHORUS emulsion sheet. Behind there is the plastic holder, shown with its vacuum channels. The back rectangle represents the glass plate which is drawn with the positions of the screws used to mount it on the stage, the vacuum protecting O-ring and the vacuum channel.

and the plastic holder.

All microscopes are connected to a central vacuum tank. In the vacuum path between the tank and the microscopes there are oil traps to prevent the optical oil leaking into the vacuum system from entering into the tank and ending eventually at the vacuum pump. The main requirement on the pump was to be able to cope with large leakage. The diaphragm vacuum pump[2] which is used can exhaust up to 3.3 m^3 of air keeping the pressure in the tank at 2 mbar.

5.1.3 The optical system

The optics for the microscopes were developed in collaboration with JENOPTIK[3]. These systems replaced the standard NIKON optics using the Tiyoda 50× objective lens, which were used initially in all microscopes.

[2] MD 4T diaphragm vacuum pump by Pfeiffer Vacuum GmbH, Emmeliusstrasse 33, D-35614 Asslar, Germany.

[3] JENOPTIK Laser, Optik, Systeme GmbH, Jena, Germany

Requirements for the optical system

Generally, the requirements for an optical system to be used for nuclear emulsion readout do not only depend on the physics which is aimed at with the use of nuclear emulsion in an experiment, but also on the design principles of the rest of the readout system. A very detailed description of the optical properties of such a system exists in Ref. [115]. The most relevant points are presented here.

As it has already been mentioned, a nuclear emulsion detector has an intrinsic three-dimensional spatial resolution of $\sim 1\,\mu$m, which is determined by the size of the developed grains and their dispersion around the actual particle trajectory. To exploit this feature a microscope which is used to read out the emulsion has to perform accordingly: a grain in the emulsion should be distinguishable not only from nearby grains in the same plane, but also from grains in different depths along the optical axis. The first requirement is easily met because the grain spatial separation is larger than the microscope's *two-line separation*

$$\delta = \frac{\lambda}{2NA} \tag{5.1}$$

where λ and NA are the monochromatic light wavelength and the numerical aperture of the system. Even for moderate values of the numerical aperture, and using visible light, a two-line separation of better that $1\,\mu$m is achievable.

The optical parameter which is more relevant to emulsion scanning is the total depth of focus (DOF) which depends not only on the optics itself but also on the detector which reads out the image:

$$d = \frac{n\lambda}{NA^2} + \frac{M}{NA}e \tag{5.2}$$

where n is the refractive index of the medium, and e is the smallest distance resolved by an image detector placed in the image plane of an objective with magnification M.

The DOF gives a measure not only of the effective size of the grain along the optical axis, but also of the tolerance for defining two grains as coplanar. It can be therefore thought of as the third component of the actual spatial resolution which is achieved with emulsion.

Automatic emulsion scanning generally requires small DOF, for increased tracking efficiencies. The DOF determines the number of independent planes that can be taken for a fixed thickness of emulsion layer. When searching for tracks it is important that the ratio of track to fog grains is maximized, and a small DOF leads to this direction.

From (5.1) and (5.2) it can be inferred that the most relevant optical parameter to resolution issues, which should be maximized is the numerical aperture

$$NA = n \sin\alpha \tag{5.3}$$

where α is the half angle of the most oblique rays reaching the front surface of the objective. From this definition it can be noticed that the NA is maximized, by minimizing the distance between the sample and the objective (working distance). A limitation in the achievable

NA for an optical system used for emulsion readout comes from the fact the working distance should be at least 1 mm which is the thickness of an emulsion sheet.

The numerical aperture of the condenser is defined similarly. For the definitions of the microscope resolution and DOF the effective numerical aperture of the system

$$NA_{\text{effective}} = 0.5 \cdot (NA_{\text{objective}} + NA_{\text{condenser}}) \tag{5.4}$$

should be considered.

Among the requirements on the optical system which is dependent on the rest of the emulsion readout system is the size of the field of view (FOV). We believe that there are several advantages using large FOV.

When scanning the sheets for track segments to be used as alignment data, having many tracks within the same view increases the probability of a proper matching between two consecutive sheets. For TP scanning (eg. scan-back), a background track in the same view as the predicted track is essential to correct for distortion effects. Larger fields have an increased probability of finding more than one tracks within a single view. During CS and SS scanning, a significant amount of overhead is reduced since the same area can be covered using less views, while the fraction of the overlapped area is decreased. Moreover, the acceptance for large angle tracks is increased.

A large FOV is also useful when general tracking methods are applied for the scanning of target sheets (see 5.3.6), since larger volumes are obtained. This increases the efficiency of identifying an electron from the nearby conversion of a bremsstrahlung photon, as well as the probability of picking up photons and neutrals coming from the primary vertex, through their conversions and decays, respectively.

Achieving large FOV has been our main reason for introducing MPixel CCD cameras (see 5.1.4). However, their capability can only be fully exploited with optics delivering good quality over a large (Ø 500 μm) FOV. An example emulsion image taken with the MPixel CCD, using the Tiyoda objective, is shown in Fig. 5.2. Even though in the central region the grains in focus are round with a high contrast, it may be observed that towards the edges the grains are deformed into thin ellipses. The focal plane itself is also curved. The difference in depth between an grain in focus in the centre of the view and a grain in focus at the edges is about 5 μm. The flat region of the FOV delivered with the Tiyoda objective has a diameter of 200 μm and a DOF of $\sim 2.6\,\mu$m.

In the same image one may notice the difference in the brightness along the height. This is an effect of light accumulation in the lines of the CCD when they are transferred from the image zone[4] to the memory zone[5]. Given the constant transfer rate on a CCD chip, this effect becomes more pronounced with larger CCD chips. Without the presence of a shutter, which would prevent light accumulation during the image transfer to the memory zone, this effect could be moderated by using larger exposure times. However, The latter is undesirable, whenever the microscope moves, and when one aims for high camera rates. In our case, a shutter in the light path of the optics is therefore required.

[4] the zone of the CCD chip which is exposed to light

[5] the zone of the CCD chip which is next to the image zone. It is protected from the light to let the digital image be read out without loss of its original pattern due to continuous light accumulation

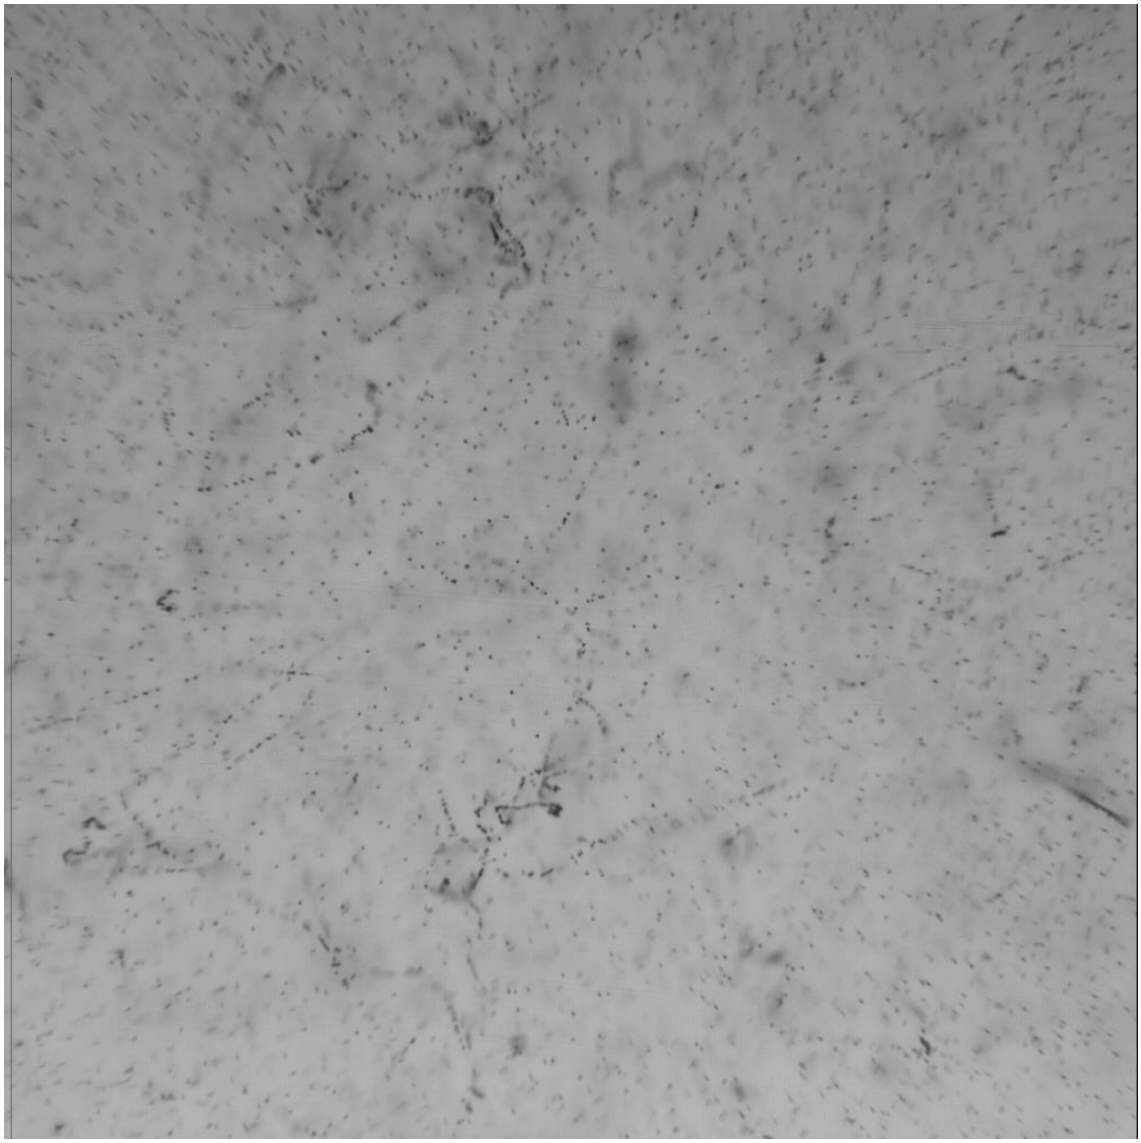

Fig. 5.2: An emulsion image ($280 \times 280 \mu m^2$) taken with the MPixel CCD and the 50× Tiyoda objective.

The implementation at CERN

The optical systems at CERN were constructed by Jenoptik with the basic requirements described above. To achieve high spatial resolution in three dimensions (high resolution in plane and small DOF), the objective should have high numerical aperture, working with a light of relatively small wavelength and bandwidth.

The thickness of an emulsion sheet poses the additional requirement of a large (more than 1 mm) free working distance (FWD).

At CERN, already before the introduction of the new optical systems, we use immersion

oil with different refractive index for each sheet type. The refractive index of the oil is matched to the one of the material that dominates the optical path along the emulsion sheet: plastic for CS/SS-like sheets, emulsion for TP-like sheets. This optimization does not only aim at the best possible image quality. It also minimizes systematical errors during the measurement of the position of the emulsion surfaces, which originate from the mismatching of the refractive indices. The new optical system was designed to cope with the variable refractive index and thickness of each component in the package comprising the immersion oil, the two emulsion layers, and the plastic base of the emulsion sheet.

The main requirements for the illumination system, apart from the wavelength of the light source, include a variable attenuator for matching brightness, a shutter that can be synchronized with the CCD, and a high condenser NA, similar to the objective NA.

The imaging system

The imaging system consists of a micro-objective lens and three exchangeable tube lens systems. It was designed to have a numerical aperture of 1.05, a 0.5 mm diameter of the FOV, and a FWD of 1.3 mm.

There are eleven lens elements in seven groups in the micro-objective lens. The particular requirement of the variable refractive index results in spherical aberration, which is compensated by shifting a lens group inside the micro-objective within a range of 1.5 mm. To prevent any other aberration due to decentering (like coma) during the shifting, the accuracy of the mechanical components had to meet stringent criteria: the tolerance of the horizontal displacement of the moveable lens group is 2 μm over the entire length of shifting.

The three tube lenses are separate mechanical units, exchangeable to achieve different magnifications of the microscope without changing the overall length from the object to the camera. The tubes have focal lengths of 200 mm, 300 mm, and 400 mm for magnifications 40×, 60×, and 80× respectively, and an identical image field size of 20 mm. Under construction is currently a set of tubes for 28× magnification as well.

The illumination system and its controls

The illumination system is a special microscope illuminator of Köhler type with integrated shutter and attenuator (Fig. 5.3).

The light source is based on a 200 W mercury arc lamp of type HBO. A mirror behind the lamp is used to increase its useful luminosity. Using a special filter package the g-line ($\lambda = 436\,\mathrm{nm}$) of the spectrum is isolated with a bandwidth of about 10 nm. It also removes the infrared component protecting the rest of the system and the emulsion from overheating. This component is also undesirable for the CCD camera which is far more sensitive to this radiation than to the g-line. A diffuser in front of the lamp ensures brightness uniformity in the illuminated field.

The optical system was designed with fixed focus of the illumination system. Good illumination homogeneity is therefore necessary not simply in the horizontal plane of the

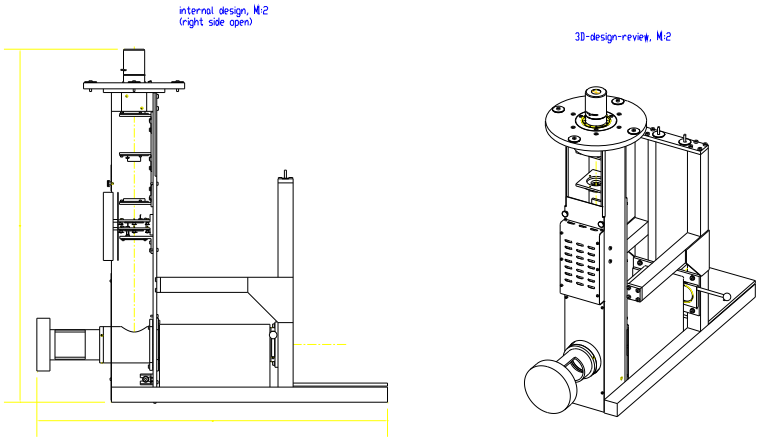

Fig. 5.3: The illumination unit of the optics. The housing of the mercury arc lamp is installed in the free space shown in the right-bottom corner of the left picture. The light is partly deflected upwards by a semi-transparent mirror. Before reaching the condenser lens, shown at the top of both pictures, it passes through the attenuator-shutter unit, the field stop iris and the aperture stop iris. The light which is not reflected upwards, is projected on a screen shown at the left-bottom corner of both pictures. It is used to align the lamp with the rest of the optical system.

focal point, but throughout a cylindrical volume with diameter of 0.5 mm, corresponding to the size of the field of view, and height of 1 mm, which is the typical thickness of an emulsion sheet. To achieve this, the illumination system, especially the high-aperture condenser, had to fulfill stringent requirements on imaging quality, telecentricity and illuminated field. To prevent vignetting during scanning along the vertical axis, the illuminated field is larger (Ø 1.5 mm) than what is needed for the actual imaging. A field stop consisting of an iris is used to adapt the illuminated field to the imaged field after a change of magnification.

The condenser lens was designed to have an $NA = 0.9$, to match as much as possible the NA of the objective. The aperture stop, which is an iris, can be used to reduce the numerical aperture of the condenser if needed.

The LC[6]-Unit consists of two combined LC-cells: the LC variable attenuator (ALM) and the LC light shutter (ALM-Sh). The ALM serves for light control by reducing the light intensity accordingly. It consists of a twisted nematic liquid crystal display (TN-LCD). The LC glass plates carry a transparent indium tin oxide (ITO) electrode which is covered by the rubbed polyimide. The nematic LC molecules align parallel to the buffing direction and are inclined to the glass plate. The rubbing direction of the upper and lower glass plates are perpendicular and hence the molecules perform a 90° twist over the thickness of the LC layer. Polarized light follows the 90° twist on its way through the LC material. The

[6] Liquid Crystal

polarization state of a light beam can be adjusted by re-orientation of the liquid crystal molecules with the application of an electric field. By mounting the modulator between crossed polarizers (polarizer and analyzer), amplitude modulation of the light throughput can be achieved. The ALM is driven by a DC signal of the range 0-5 V. At maximum voltage an attenuation by a factor of 30 is achieved.

The function of the ALM-Sh, in design similar to the ALM, is shuttering the light beam. The nematic LC is 90° twisted between the glass plates, and the light polarized by the ALM follows the twist of the ALM-Sh on its way through the LC layer. At the end of its path the light passes a polarizer, which is crossed to the analyzer of the ALM. In this orientation state, which is the default one, the light passes the second LC cell unchanged, without being absorbed. If the applied electrical voltage is sufficiently high, a full deformation of the LC layer is reached, and the light can no longer pass the last polarizer, being attenuated by a factor similar to one achieved by the ALM. The activation of the shutter is done with a TTL[7] signal. The switching time from minimum to maximum of the transmission is about 7 ms. The complementary procedure has to be as fast as possible, so that the switching time is much smaller than the exposure time of the CCD camera. A rise time of about 0.3 ms has been achieved.

For the ALM operation a Digital-to-Analog converter is needed to deliver the required DC signal. Every microscope has been equipped with a CerealBox[8], which is a small device that communicates with the PC via an RS-232 serial interface line, and provides 8 analog output channels, 8 analog input channels and 24 digital input/output channels. The update rate is 30 Hz and an analog signal is generated or read with 12 bit precision. An analog output channel may provide a DC signal in the range 0-4095 mV, while an analog input channel may sample a DC signal in the range 0-5000 mV. The 24 digital signals are grouped in 3 sets of 8. Every set can be configured entirely as input or output. In order to have easy access to all analog channels, they are connected to an interface box which provides a Lemo connector for every channel.

The TTL signal for the shutter activation is provided by the interface card, which receives the corresponding control signal from the C80 card (see 5.1.4). With this scheme, the shutter and camera operations are synchronized.

The immersion oil system

For the immersion oil, a containment system was developed at CERN to keep it always confined in a small region around the objective lens. The system is based on a ring embracing the objective with rubber edges, which ensure that the emulsion sheet below slides smoothly when the horizontal stage axes are moving. It also protects the oil from dust and humidity. The oil circulates in a closed circuit including the ring itself, a pair of plastic tubes, an oil tank and a peristaltic pump[9]. The tubes have an inner diameter of

[7] Transistor-Transistor Logic

[8] CerealBox LV824-H, manufactured by BG SYSTEMS INC., Palo Alto, CA 94306, USA.

[9] MasterFlex C/L model 77120-62, manufactured by Cole-Parmer Instrument Company, Illinois 60061, USA.

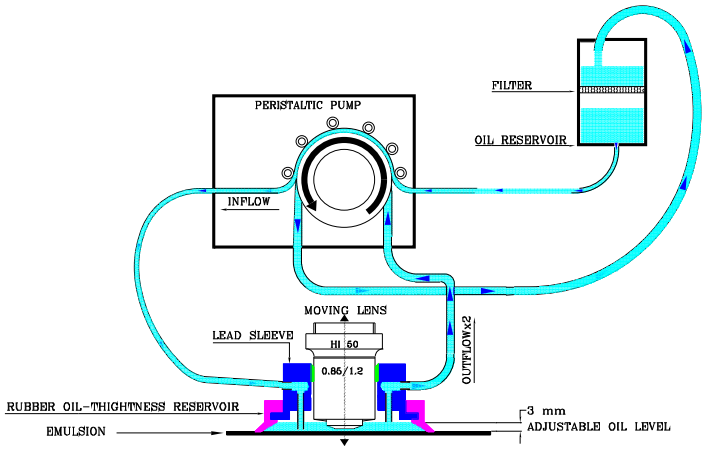

Fig. 5.4: The oil system. The version shown in the figure is the one used with the Tiyoda objective. For the Jenoptik objective the transverse dimension of the ring was larger.

1.3 mm and 1.8 mm for the inlet and the outlet respectively. With this system a continuous oil flux of about 1 cm^3/hour is insured, keeping the losses to a minimum. The mean oil consumption per microscope is reduced from about 2 l/day to less than 100 ml/day. The larger diameter of the outlet tube ensures that the ring never overflows. To avoid emptying and refilling the tubes with a different type of oil when a different type of emulsion sheet is scanned, there are two oil systems per microscope, with dedicated rings. A schematic layout of the oil system is shown in Fig. 5.4.

Optics tuning

To achieve the best optical quality in the emulsion images the imaging (upper) and illumination (lower) parts of the optics have to be well aligned. Moreover, to achieve a true high numerical aperture for the condenser, its focal plane should be positioned in the middle of the base of an emulsion sheet mounted on the microscope, and the amount of focused light should be maximized.

A small scale tuning involving at least the bottom part of the illumination system is needed also whenever the arc lamp is replaced. Tuning is also needed whenever the type of emulsion to be scanned changes. It is not only the different refractive index for which the moveable lens group in the objective has to be shifted accordingly. Due to the slightly different thicknesses of the sheets the height of the condenser needs to be adjusted as well.

A large scale tuning of the optics is necessary, whenever the magnification tube of the imaging system is replaced. Since such an operation involves the unmounting of the whole imaging system from the microscope stage, any previous alignment is destroyed.

The tuning procedure is done with the assistance of a computer program which continu-

ously displays various image quality factors on the screen (see 5.3.3). The overall procedure is accomplished with the following steps:

Lamp alignment This is a necessary procedure whenever the lamp is replaced. There are five screws on the housing of the lamp which can be used to move lamp and the mirror behind it. The lamp and its mirrored image are projected on the screen behind the semi-transparent mirror of the optics. Good alignment is achieved when the two images can be seen with the same brightness, placed symmetrically around the centre of the screen. The images should be close to each other, but to protect the lamp from overheating by its mirrored illumination, they should not overlap.

Checking the LC-Unit Even though there is no actual tuning to be done with the LC elements, it is good practice to check whether they are in good condition. Units which have been damaged and need replacement, affect the final image quality apart for attenuating its brightness.

Alignment of the system For the alignment of the optical axes of the illumination imaging parts of the optics it is assumed a priori that they are parallel. This assumption is based on the high precision mechanics and the initial installation of the system by experts. After closing the field stop and checking that the aperture stop is open, we focus on the iris. Alignment is achieved when its image is in the centre of the field. Transverse movement of the whole illumination part is allowed after loosing the screws which are used to mount it on the stage.

Condenser height adjustment The precise mechanics of the optics allow for adjustment of the height of the condenser focal plane, by turning the unit containing the condenser lens, with only a small perturbation to the already achieved alignment. Only in rare cases we had to repeat the previous step. After mounting an emulsion sheet on the stage, we check whether the light intensity is maximized in the middle of the base and whether the light conditions are similar on both emulsion sides.

Correction of the refractive index This is done by shifting the moveable lens group in the objective. To shift the lens group one turns accordingly a ring on the mechanical structure of the objective. At the optimal position, the relevant quality factor displayed by the assisting computer program is maximized.

Performance of the optical system

A typical emulsion image obtained with the new optics using the 40× magnification tube, is shown in Fig. 5.5. Thanks to the diffuser and the shutter, the brightness and the quality of the image are uniform. Contrary to the image shown in Fig. 5.2, in this one there is no evident deformation or curvature. The field is known to be flat with a tolerance of less than $1\,\mu$m over the area covered by the camera $350 \times 350\,\mu\text{m}^2$.

The DOF is about $1.2\,\mu$m. However, this is the value which is obtained by (5.2) if the nominal value for the numerical aperture is assumed. If the optics is not very well tuned,

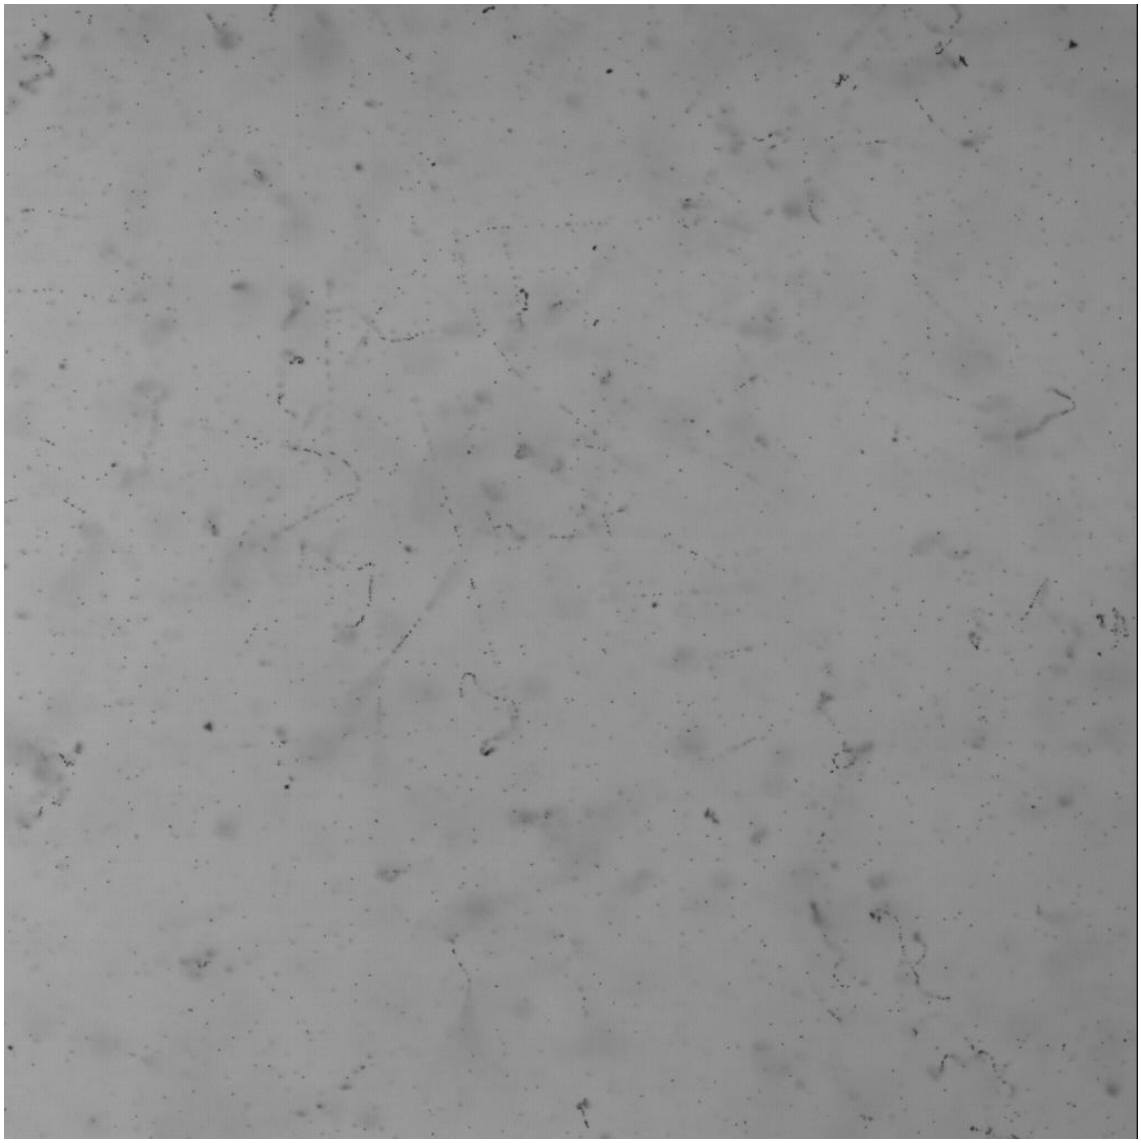

Fig. 5.5: An emulsion image ($350 \times 350 \mu m^2$) taken with the MPixel CCD and the 40× objective by Jenoptik.

then the effective NA is smaller and the DOF becomes larger. The effective NA is also decreased with the introduction of not perfectly transparent material along the light path, before reaching the objective. This happens because the various scattering centres (the silver grains in emulsion) are deflecting the light from its focused direction towards the front objective lens. As a consequence, the optical quality of the top emulsion surface is always better than the bottom one. We have also noticed that sheets with thick emulsion layers produce in general a poorer image compared to those with thin emulsion layers.

What is mostly relevant to the actual scanning is the minimum distance in the emulsion

which gives independent information. This is given by a convolution of the effective DOF, the sensitivity of the digital filter applied on the emulsion images to identify the grains, and the length of the camera exposure time. This minimum distance is eventually less than 2 μm.

The above arguments hold equally for the old NIKON optics with the Tiyoda objective. Compared to these optical systems the new ones by Jenoptik allow us to analyse more than six times larger fields with a two times higher sampling rate along the optical axis.

5.1.4 Camera and DSP hardware

The emulsion readout system developed at CERN was conceived with a modular design, based on off–the–shelf industrial products wherever possible. Such a design allows for upgrades with short development cycles. An extensive use of software solutions has been made to allow for maximum flexibility in automatic scanning. This is a rather different but complementary approach to the one adopted by Nagoya for their systems, where hardware solutions and application specific implementations are preferred in order to achieve a high scanning speed.

The essence of the microscope DAQ system was realized with a MPixel digital CCD camera, and two PCI[10] boards hosting the C80 and C620 DSP[11] processors. For each of these devices, the specific hardware features which are used for the DAQ cycle are described subsequently.

The digital camera

Every microscope is equipped with a THX 7888A camera by Thomson[12]. The sensitive area of this digital 12-bit CCD camera is covered by 1024×1024 useful pixels. The pixels are square and their size is $14\,\mu\text{m} \times 14\mu\text{m}$.

The data rate of the camera depends on the readout mode. When every CCD line is read out using one amplifier, the maximum data rate is 15 Hz. The readout of the camera is performed transferring 8-bit words. If two amplifiers are used, each of them reading half a line, then the data rate increases to 30 Hz. In this case 16-bit word transfers are used during the readout, with each amplifier output occupying 8 bits. Byte reordering is needed to retrieve the original image before it can be used for processing or display.

Out of the 12 data bits only 8 are read out. This number of bits is adequate for a digital filter to distinguish the pixels belonging to grains from the background. Eight bits also fit in a single byte, which makes the digital processing simpler and faster. It also complies to the standard definition of a digital monochromatic image being a matrix of 8-bit unsigned values (0-255). When running the microscopes with the NIKON optical systems, the most significant bits (4-11) where used. With the introduction of the systems by Jenoptik we had to use bits 1-8 in order to have the same brightness in the digital images. This was due

[10] Peripheral Component Interconnect.
[11] Digital Signal Processor
[12] http::/www.thomson-csf.com

Name (Thomson Designation)	Function
RESET	Resets camera kit when low
CCD_RQ	Exposure control. Length determines exposure, and falling edge starts image to memory zone transfer
BINV	Activates vertical binning
BINH	Activates horizontal binning
INTER	Activates interlaced mode
SFT	Selects image to memory zone transfer
LINK_TEST	Selects high or low impedance of digital video outputs
MODE	Selects one or two amplifier readout

Tab. 5.1: The camera control signals. The BINH, BINV, and INTER signals are presently not used.

to the lower quantum efficiency of the camera in the blue light delivered by the mercury arc lamp, and the light absorption in the various elements of the illumination system of the optics.

Apart from the 2×12 bit data signals, the camera also provides frame, line and pixel synchronization signals. A set of digital input channels can be used to set the camera modes and to trigger the exposure. The exposure time itself is defined by the length of a digital pulse. The control signals from the interface card are RS-422. A list of the controls in use is shown in Table 5.1.

The data signals are transferred from an I/O driver card to an interface card as LVDS[13] to keep the integrity of the signals over the 15 m from the camera to the interface, being faster with less slew and skew than RS-422. The interface card delivers the data to both the C80 board and the frame grabber, as soon as the frame-enable and line-enable synchronization signals are set.

The C80 DSP board

The image data acquisition is managed by a PCI/C81 board from LSI[14], which incorporates the Texas Instruments MVP[15] TMS320C80 fixed point processor, referred to as the C80 DSP. Image processing was also being performed using this DSP before the introduction of the C620 board. A Digital Video Module (DVM), piggy backed onto the DSP board, is responsible for the control and data taking from the camera. One of its available control signals is used by the interface card to generate the TTL signal for the shutter of the optics.

The C80 is a single chip multi-processor, operating at 50 MHz. It has a floating-point

[13] Low Voltage Differential Signal
[14] Loughborough Sound Images, and presently Blue Wave Systems
[15] Multimedia Video Processor

32-bit RISC Master Processor (MP) and four parallel Advanced DSPs. The MP has on-chip data and instruction caches, of 4 kb each. It has a 32-bit instruction bus and a 64-bit bus for other on-chip memory accesses. It performs the overall management of tasks run on the C80.

The Parallel Processors (PPs) are 32-bit integer digital signal processors. They each have an on-chip instruction cache and can access the 4 kb on-chip data and parameter RAMs. They have a 64-bit instruction word and can perform operations simultaneously or independently in a Multiple Instruction Multiple Data configuration. Each PP can perform in a single clock cycle up to two address operations, one multiplication, and one general purpose ALU[16] operation. A PP has eight 40-bit wide registers which can be used for ALU and the Multiplier units. For each of the two address units there are four associated registers. Moreover, there are registers allowing for three hardware controlled loops in a program.

The memory provided on the PCI/C81 board consists of 32 Mb of SDRAM[17] and 512 kbytes of Flash Memory. In addition, the board is fitted with PCI and I/O FIFOs. The PCI FIFO is used to handle DMA[18] data transfers between the C80 and other PCI devices such as the C620 DSP board and the host PC itself. The transfer speed is about 100 Mb/s.

The I/O FIFO is used to transfer the images captured with the CCD camera from the DVM module to the DSP memory. The typical transfer speed in this case is 400 Mb/s. Communication and synchronization with the PC is achieved with signals and messages. A signal is generated and received within 50 μs.

The C620 DSP board

The task of identifying the emulsion grains in a digital image is handled by two Texas Instruments TMS320C6201 fixed point processors, referred to as the C620 DSPs, operating at 200 MHz. They are integrated in the PCI/C6600 board from BWS[19].

The core of the C620 processor is based around the Texas Instruments VelociTI VLIW[20] architecture. There are two sets of functional units, each of them equipped with a multiplier, three ALUs, and 16 general purpose registers. Up to eight 32-bit instructions can be performed during a clock cycle. This is achieved using 256-bit wide packet fetches to/from the 64 kb internal program/cache memory. The C620 has a 64 kb internal data memory, and its on-chip peripherals include a Host Port Interface (HPI), an External Memory Interface (EMIF), and a four channel DMA controller.

In terms of the effective clock speed the C620 and C80 processors are equivalent (200 MHz vs 4×50 MHz). However, the C620 can perform up to eight instructions per cycle, compared to four in the case of a C80 PP. Moreover, the C620 has larger internal

[16] Arithmetic and Logical Unit
[17] Synchronous Dynamic RAM
[18] Direct Memory Access
[19] Blue Wave Systems, http::://www.bluews.com
[20] Very Long Instruction Word

data memory (64 kb vs 4 kb) and a more flexible register file (32 general purpose registers vs 16 registers of restricted usage),which allow executing algorithms, such as digital filters, in fewer cycles. The flexibility provided by the large number of general purpose registers in the C620 processor, largely overwhelms the fact that there are no hardware controlled loops possible, as in the case of the C80 PPs.

For our purposes, in terms of performance in sheer image processing per processor, an overall improvement in speed by a factor of two has been observed with the C620 DSP. The fact that two C620 processors can operate simultaneously and independently on the PCI/C6600 board, resulted to an additional factor of two in terms of processing speed.

A bank of 16 Mb SDRAM is available to each processor on the board. There are also two 1 Mb banks of shared SRAM[21] accessible by the processors via two local buses on the PCI/C6600 board. A PCI interface chip is connected to the first bus, so that the two processors and one of the shared memory banks are visible in the PCI address space. This allows other PCI devices to perform master/slave accesses on the internal memory of the C620 processors and the shared memory bank. The processors can interact by interrupting each other, and by sending single word messages. Interrupts are very useful for synchronizing the tasks running in them, and for controlling the data flow on the board. Finally, the board has an I/O port and a PMC[22] site, which are connected to the internal buses. In principle they could be used for digital data input.

The PC configuration

A PC driving a microscope (DAQ PC) is currently based on a dual Pentium III processor at 500 MHz, running the Windows NT operating system. It is equipped with 256 Mb of RAM, a hard disk of 5 Gb, and a (Matrox Millennium) display card. Its two serial ports are used for the communication with the motor controller and the CerealBox for the light control. A PCI frame grabber board (Matrox Pulsar) is used to display on a secondary screen the output of the digital camera, which is read in via a digital video module ISA card. The Pulsar also allows for graphics overlaying. It therefore serves as a monitoring tool during data taking. Finally, a PCI based 100 Mbit/s Fast-Ethernet card is used to connect the PC to the network. In total five PCI slots and one ISA slot are required for the PC motherboard.

5.2 The microscope DAQ implementation

5.2.1 Overview

The flow of data and control signals between the microscope, the DAQ PC, and the file server managing the database where the output of the emulsion readout is stored, is shown schematically in Fig. 5.6.

[21] Static RAM
[22] PCI Mezzanine Card

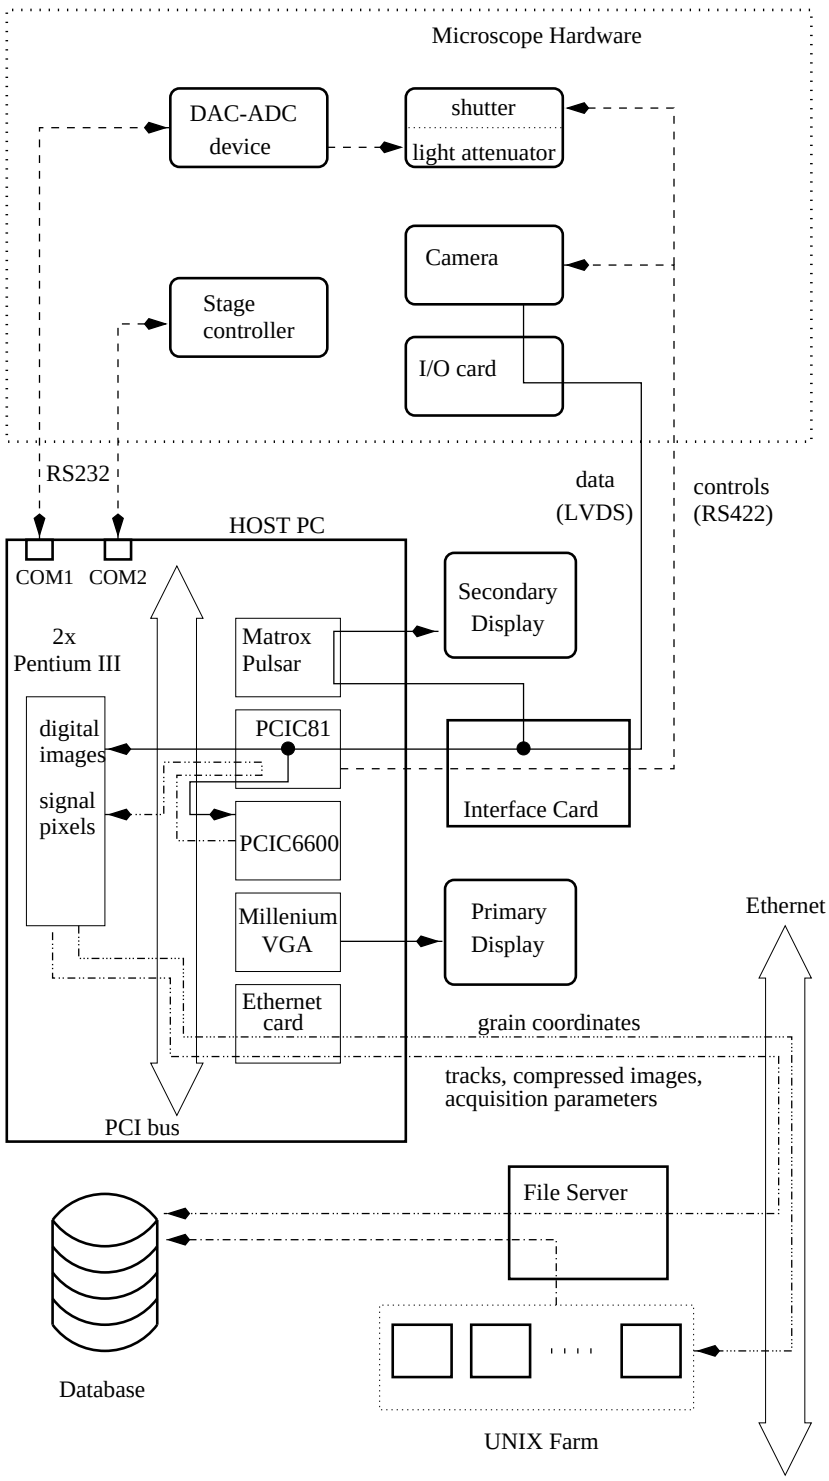

Fig. 5.6: Flow of data and control signals between the microscope, the PC the database and the off-line processing farm. For simplicity purposes we have omitted the flow of dispatcher and run-control messages, between the PC and the dispatcher server.

One of the PC's serial ports is used to send commands and receive data from the stage controller, while the other is connected to the CerealBox, which is responsible for controlling the light attenuator of the optical system.

The controls to the camera and the shutter unit of the optics are sent by the DVM of the C80 board through the interface card. The camera digital signals are sent by the driver card to the interface card as LVDS. From there, the emulsion image data are distributed to both the DVM of the C80 board for image analysis and the digital video input of the frame grabber for display.

The C620 board acts as a slave device of the processor on the C80 board. Its sole purpose is the application of a digital filter on the captured emulsion image to identify the pixels belonging to grains. Once an image is analysed the signal pixels are packed and sent back to the C80 board.

The acquisition cycle keeps running continuously on the DSP cards, independently of the host PC. For on-line track finding the Pentium processor in the PC requests from the C80 the data of a number of successive images. This number may vary from 10 to 200 planes, depending on the type of scanning being performed. For normal prediction scanning a set of about 20 planes is used. The pixels are clustered into grains in the shadow of the data taking, and while the stage moves to the next view an on-line tracking algorithm is performed to identify candidate tracks. Those candidates are stored in an Objectivity/DB[23] database together with all data taking parameters.

The flexibility offered by the use of DSP boards allows for data taking along the whole depth of the emulsion sheet, using the maximum number of planes. This is the case for example when there is an indication for the existence of an interaction vertex, or of a decay topology. Instead of performing on-line tracking on the data acquisition PC, the positions of the identified grains are sent via the network to a Linux farm, performing track and vertex finding on the events they receive from the microscopes. The output of this general tracking procedure is stored in the same database as the output from the on-line tracking in prediction scanning. In special cases the raw data is stored together with the filtered output. The digital images are JPEG compressed and written to the database, to be used later in interactive checks off-line.

5.2.2 The realization of the DAQ cycle using the DSPs

Tasks running in the C80 and response to the PC commands

As has already been mentioned, the C80 processor on the PCI/C81 board is the manager of the DAQ cycle concerning the emulsion readout. It has the responsibility to trigger the camera for image capturing, to collect the image data from the DVM, to process the image for grain recognition, and if asked, to send the analysis output and/or the raw image data to the processor of the host PC. This DAQ cycle is repeated continuously, at a rate of 15 or 25 Hz, according to the readout mode of the camera. There are two basic design requirements on the C80 operation.

[23] Reference for Objectivity/DB

1. The acquisition cycle should be stable and occur in a constant rate regardless of the input, because during data taking the successive images have to be equidistant (see 5.3.6). This poses the requirement on the image processing algorithm that it should be executed always in a fixed number of clock cycles for any image of fixed size. The output should also be of fixed size to avoid fluctuations in the transfer times. When the image processing rate is higher than the camera frame rate, then the period of the acquisition cycle is determined by the camera readout time. In general, it is good practice to remove *intelligence* from the programs running on the DSPs, which very often requires branching with different time responses. These processors should be mainly used as *number crunchers*, since their architecture is optimized for this.

2. The continuous repetition of the DAQ cycle should be stable regardless of the interaction with the host PC. To this end no supervision of the host PC processor should be involved. In case of interaction between the DSP boards and the PC, which is the case when the processor in the PC asks for the processing output of a series of camera frames, no assumptions should be made concerning the relative speed of execution of the processes running on either of the processors. Moreover, the DSPs should use the absolute minimum of the system resources in the PC, to avoid possible clashes with the operating system. In fact, as it will be explained later, it is only the PCI bus which is needed for the continuous running of the DAQ.

The above requirements are fully met by the program running on the C80 processor. The Multitasking Executive operating system which runs on the RISC processor (MP) of the C80, allows the running of parallel processes (tasks) with different priorities. The tasks can be synchronized using semaphores, while information exchange can be achieved with the use of shared global variables.

There are three tasks running continuously on the MP. Their functions and interactions during a single DAQ cycle are shown in Fig. 5.7. These are the Main task (MT), the Image Readout task (IR), and the Signal Handling task (SH). The hierarchy of the priorities has been set assigning the highest priority to the task which occupies the processor least. To ensure stability, the program running on the MP was written such that the more time critical a task is, the least it occupies the processor.

The SH task, which has the highest priority, simply waits for a signal from the PC to arrive. With such a signal the PC draws the attention of the C80 to issue subsequently a specific request. Once such a signal is detected, the task sets a global flag to notify the MT task of the arrival of this signal. It waits then on a semaphore which is signaled by the MT once the request from the PC is serviced. Given the fast response of the C80 to signals from the PC, this tasks occupies the MP of the C80 for less than 1 ms per DAQ cycle.

The IR task, which has an intermediate priority, has the responsibility of transferring using DMA the image which is captured by the camera, from the FIFO connected to the DVM module to a block of memory in the SDRAM of the C80 processor. Before starting the transfer the IR task waits on a semaphore which is signaled by the MT as soon as the camera is ready to send the data. At the end of the transfer it signals a different

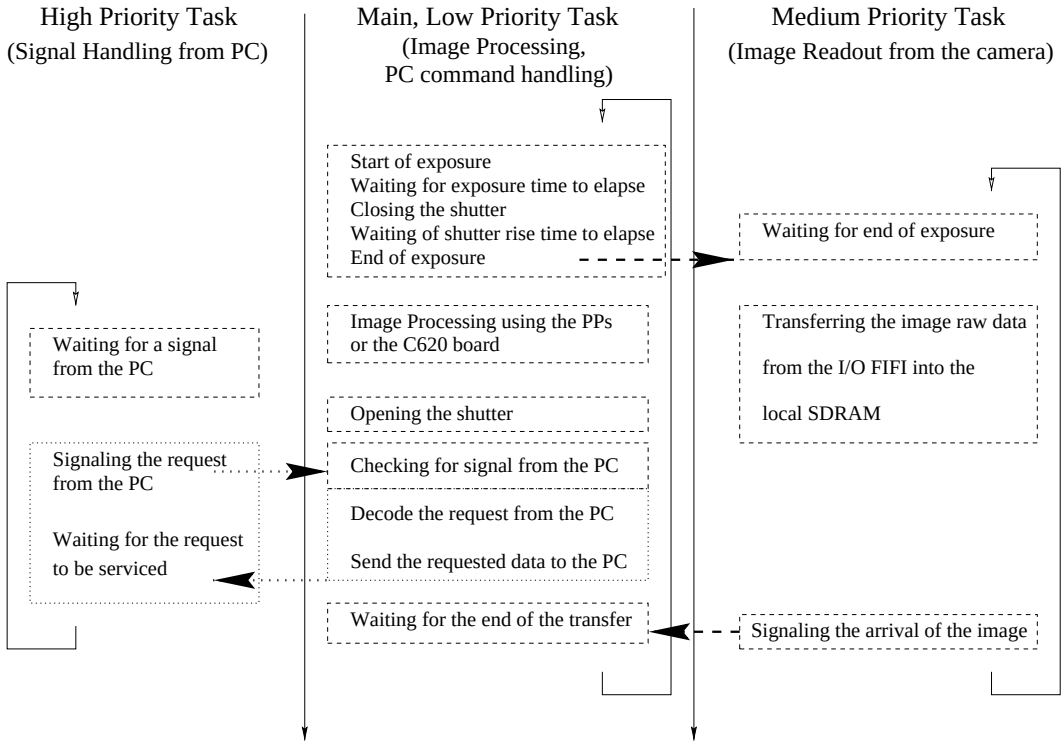

Fig. 5.7: Tasks on the MP of the C80 processor, and their synchronization.

semaphore to notify the MT that it can process the newly captured image and trigger the camera for a new one. While the transferring takes 63 or 33 ms depending on the camera readout mode, the MP of the C80 is only used for about 2.5 ms, which is the time needed to transfer 1 Mb of raw data with a rate of 400 Mb/s.

The MT task has the responsibility to trigger the camera, process the previously captured image, and if required, to serve requests from the host PC such as sending the results of the image processing. The DAQ cycle starts with the MT task sending a control signal to the camera to start the exposure. After waiting for a number of clock cycles corresponding to the exposure time, it triggers the closing of the shutter. The task waits for a number of clock cycles corresponding to the shutter rise time and then it signals the end of the exposure. At this moment the camera starts transferring the captured image to the DVM of the C80 board. The MT task signals then the corresponding semaphore for the IR task to start transferring the image to the local SDRAM.

While image N is being transferred, the processing of image $N-1$ may start. Prior to the introduction of the C620 board the image processing was performed in the PPs of the C80. One quarter of the image was processed by each of the four PPs. A PP applied to each line a high pass digital filter (see 5.2.3) coded in native assembly language. Applying a threshold cut on the filter output, the pixels belonging to grains were identified. At the end of the cycle the result of the treatment of image $N-1$ was available to the host PC.

In that case, two alternating 1 Mb buffers in the SDRAM were used to dispose the raw data of image N from the camera, and to analyse image $N-1$.

With the introduction of the C620 board the processing is no longer performed using the PPs of the C80. Instead of that, image $N-1$ is transferred to the C620 board after receiving the filter output corresponding to image $N-3$. In this case the pipelining is longer by two camera frames, and there are four 1 Mb buffers to dispose the image raw data from the camera instead of two.

Before processing an image either in the PPs or in the C620 board, the MT masks a set of noisy or dirty pixels. The list of these pixels is set by the PC during initialization and can be updated at any time. The masking is performed by replacing the gray value of the pixel with that of the previous one in the same line.

The MT task services any potential request from the PC, and repeats the DAQ cycle after it has been signaled by the IR that the transferring of image N from the camera into the SDRAM is complete.

Servicing requests from the host PC

In general, the following way of handshaking is used for every interaction between the PC and the C80. For every signal, message, or block of data sent by the PC to the C80, a signal or a message is sent back to the processor of the PC upon successful reception and service request from the C80. This is also the way the two processors synchronize during the initialization phase. This method does not only ensure stability irrespective of the relative speed of the two processors, but it also makes potential problems easier to spot.

Once the DAQ loop has started, the PC can request data from the C80 at any time. It first draws the attention of the C80 by issuing a signal which is caught by the SH task of the C80. The MT task in case of such a signal, undertakes the rest of the handshaking by returning a verification signal to the PC, and subsequently receives a series of messages containing the information needed for the execution of a specific command.

The simplest commands are those which set running variables on the C80, such as the length of the exposure signal, the digital filter threshold and coefficients, the image region of interest, the list of noisy pixels which have to be masked, etc. Such requests are always serviced within a single DAQ cycle. Other commands executed within a single DAQ cycle require transferring data from the C80 to the host PC. Such are the transferring of the raw data of the most recently disposed image, and the transferring of the total number of pixels above the filter threshold corresponding to the most recently analysed image. In all single cycle commands the semaphore on which the SH task is waiting, is signaled at the end of the cycle, and the PC can issue another command immediately.

The command which is most directly related to the automatic track searching in the emulsion is executed in more than one cycle. This is the request of the filter output for a series of camera frames. For a number of cycles equal to the number of the images in the series, the C80 sends to the PC via DMA transfers the filter output of the most recently filtered image. The transfers are handled exclusively by the C80 without requiring that the PC keeps up with the continuous reception of the images using synchronization signals.

Since the semaphore on which the SH task is waiting, is signaled only at the end of the transfer of the whole series, the PC cannot issue another command during the data taking.

The PC can be informed of the progress of the transfers by peering an array of status words which are transferred via DMA together with the filter output. The status words contain the indices of the already transferred frames. Therefore, the PC can execute in parallel any other task related to the data taking at its own pace. The status words also include the time of exposure of every frame relative to the time of the first handshaking signal when the PC has asked for the data taking of a series. As it will be explained in 5.3.4, this time is used during the data taking to determine the absolute position in space where an emulsion image is captured. Moreover, for every frame there is also the corresponding number of pixels above the filter threshold. In 5.3.4 it is shown how this information is used to automatically update the position of the edges of the emulsion layers.

The C80 performs the DMA transfers using guide tables containing the addresses of the destination buffers in the memory of the host PC. These guide tables are filled by the PC processor and sent to the C80 during initialization after allocating the necessary buffers. It is possible to change in the default number of frames of the series data taking, during the DAQ loop. In this case the buffers in the PC memory are reallocated accordingly, new guide tables are created and they are sent to the C80 to be used in subsequent DMA transfers.

Depending on the data taking mode, it is possible to send the raw data together with the filter output for a series of camera frames. Such a feature is interesting when regions around a potential primary or decay vertex are analysed. Taking the raw data and storing the emulsion images, allows for off-line visual inspection of the emulsion without the need of an actual microscope. The maximum number of consecutive frames that can be taken as a single set is limited by the amount of SDRAM in the C80, which is filled with the guide tables for the DMA transfers. The C80 boards which are used in the microscopes allow for a maximum number of 250 images that can be taken as a single set.

The DAQ cycle with the C620 board

During automatic scanning, the tracks that lie within the scanning acceptance of a microscope view should be reconstructed before the microscope can take a new view. To this end, the track finding procedure should be parallelized as much as possible to the image data taking from the camera and the motion control of the stage. These tasks require in general only little processing power. The various steps of the track finding procedure fill the idle time intervals of the available processing power.

While track finding can be completed only upon the availability of the last image of the microscope view, the grain recognition in an image can start already after its disposal from the camera. For this reason, in most implementations of the automatic scanning principle, the aim is to complete the grain recognition within a camera cycle.

The C620 board was introduced to exploit the 16-bit readout mode of the camera. Every image has to be analysed for grain recognition within 33 ms, instead of 63 ms which is the case for the 8-bit readout mode. For the digital filter we use (see 5.2.3) the PPs in

the C80 processor can analyse a 1024×1024 digital emulsion image within 48 ms. Their processing power was therefore adequate only for the 8-bit mode operation, and their functionality was taken over by the two C620 processors on the C620 board.

The use of the C620 board and its integration into the DAQ system was dictated by its hardware architecture. The following boundary conditions had to be satisfied:

- No appropriate interface was available to control the camera and to send the image data directly into the C620 board via its I/O port or its PMC site. The C80 would continue controlling the camera and the image data being disposed into the C80 memory. The C620 board would accept the data from the C80 using fast transfers over the PCI bus.

- Contrary to the C80, the C620 processors cannot perform themselves DMA transfers over the PCI, because the board was designed as a slave PCI device. Data transfers via DMA between the C620 board and the PC are possible if they are handled by the PC processor, which is synchronized with C620 processors with interrupts. However, this is contrary to one of the main requirements of the DAQ system, that no supervision from the PC processor is involved in the DAQ cycle. All transfers have to be managed by the C80 board which was designed as a master PCI device.

- The memory resources of each processor on the C620 board can be accessed by other PCI devices through an 8-Mb window, whose start address can be set by the PCI interface chip of the board. However, DMA transfers of data from other PCI devices are executed at the PCI top speed only when the target position is in the shared SRAM or the internal memory of each processor. The slower accesses of the external 16-Mb SDRAM bank of a C620 processor from other PCI devices is due to the fact that these memory banks are not directly connected to the shared local bus to which the PCI interface chip is connected. It should be noted that unlike the C620, for the C80 board there is no slave interface and therefore there is no accessibility to its own memory resources.

- A device can use the shared bus on the C620 board only exclusively. Therefore, only one of the C620 and C80 processors can access the shared SRAM on the C620 board at a given moment.

To meet these conditions, the C620 board is used in the DAQ system as a slave device to the C80 board. The C80 processor sends the image raw data and retrieves the filter output from the two C620 processors through the shared SRAM.

The flow of data and signals between the processors on the DSP boards and the shared memory of the C620 board is shown in Fig. 5.8. Once the C80 has triggered the camera for the capturing of image N, it retrieves from the shared SRAM the filter output of image $N-3$ and sends the raw data of image $N-1$. It then notifies the C620 processor which disposed the filter output into the shared SRAM, of the arrival of a new image. The C620 processor transfers the image into its external SDRAM, notifies the other processor on the

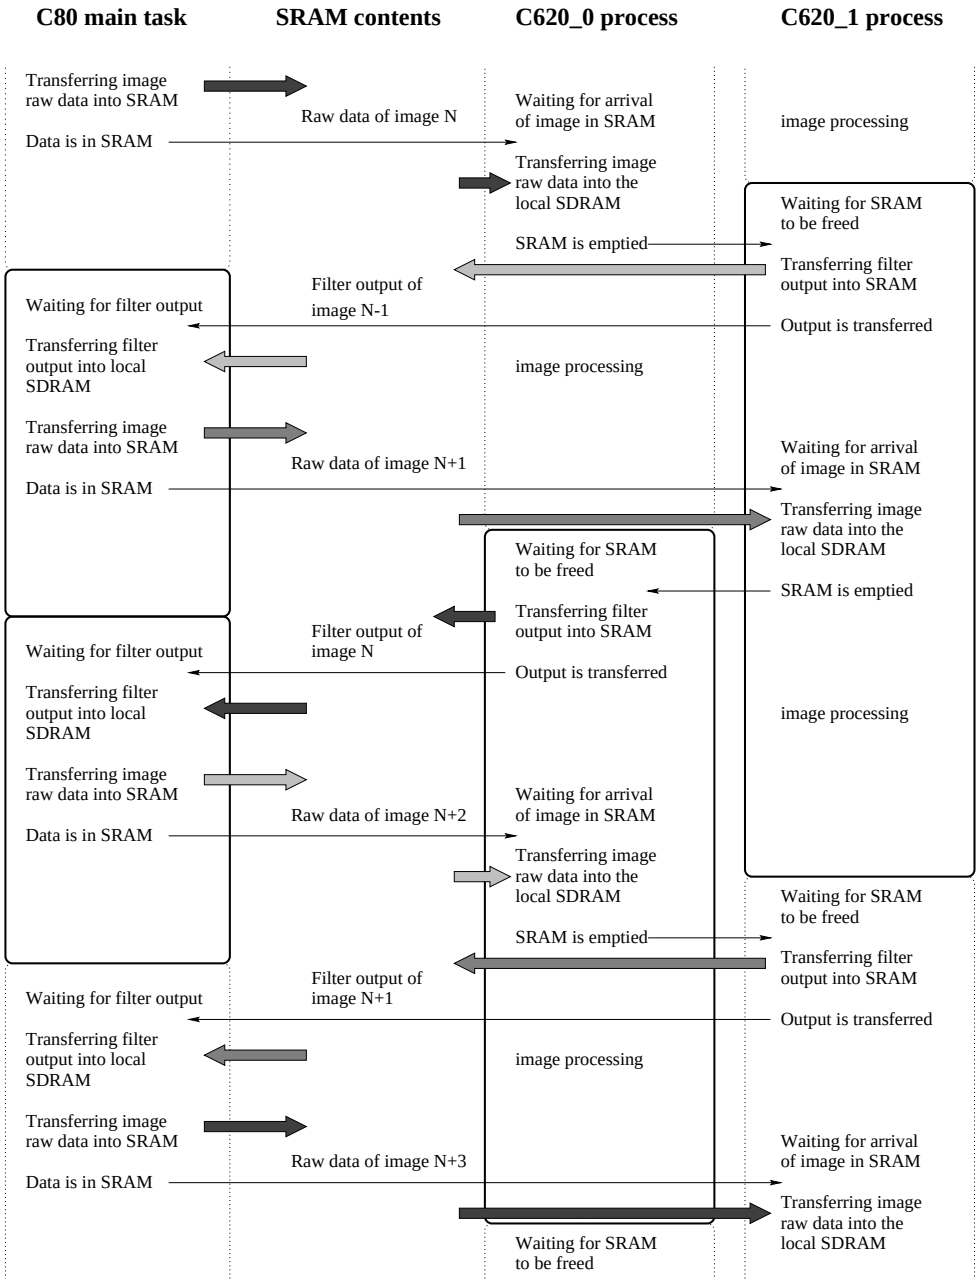

Fig. 5.8: Data flow between the C620 and C80 boards, and synchronization of the processes running on these processors. The block arrows correspond to flow of actual data. For the right and the left pointing arrows the data transferred is the actual image raw data and the filter output respectively. Arrows with the same gray level correspond to the same image. Line arrows correspond to signals and interrupt sending from one processor to another. The boxes surrounding a sequence of actions in the processors define their internal DAQ cycles. For the C620 processors the cycle is two times longer that the one for the C80 processor.

board of the shared SRAM availability, and starts the image processing. In the subsequent DAQ cycle the C80 exchanges data with the other C620 processor.

It is already mentioned that a C620 processor cannot directly access resources on the C80 processor, and therefore they cannot send themselves the necessary synchronization signals to the C80. Instead, they set accordingly variables stored in their internal memory, and the C80 peers then via PCI word read accesses, whenever needed. Using PCI word write accesses the C80 is also able to set the values of variables stored in the C620 internal memory. Such variables are the filter threshold and coefficients and the image region of interest, which are set by the host PC by issuing the corresponding commands to the C80. For the latter to be able to access the internal memories of both C620 processors and the shared SRAM, it needs to know their base address in the PCI address space. During initialization, the PC retrieves these PCI addresses, and the addresses of the variables in the C620 internal memory, and sends them to the C80.

Once a C620 processor is informed of the arrival of new image raw data in the shared SRAM, it has to transfer it to its external 16-Mb SDRAM bank as fast as possible. This is achieved by starting two DMA transfers in parallel, transferring the image raw data from the shared SRAM into the SDRAM through a line buffer in the internal memory. The reverse procedure is performed when the C620 processor disposes its filter output.

The image processing on a C620 processor has to be completed within two DAQ cycles. The image is treated line by line. Every line is transferred first from the SDRAM into a line buffer in the internal memory. If the 16-bit mode of the camera is used, byte reordering is performed before filtering. After the application of the digital filter, the output is packed so that every pixel is represented by a single bit whose value is set according to whether the filter output is above or below the threshold. This way the size of the image after processing shrinks by a factor of eight. The decoding of the packed output is done on the PC processor during data taking.

Performance of the DAQ system

The implementation of the low-level microscope DAQ system using the DSPs as described above, has proven to meet the basic requirement of stability. Automatic scanning sessions lasting for more than a week have never been interrupted because of problems occurring at the DSP level of the DAQ.

The duration of a DAQ cycle is the sum of the camera readout time (63(33) ms for the 8(16)-bit mode of the camera), the shutter rise time (0.5 ms), and the exposure time. In general, we aim for the shortest possible exposure times. Fast exposures do not contribute significantly to the overall data taking time. If the fraction of the exposure time in the DAQ cycle is large, the grains in the image loose their sharpness when the stage is moving, and the digital filter is less efficient in recognizing them. As discussed in 5.1.3, large exposures increase the minimum distance of independent information in emulsion. For these reasons, the exposure time is as long as it is needed for the camera to collect enough light to produce an image with high contrast. Depending on the emulsion transparency, typical exposure times are within the range 5-10 ms. Summing all the contributions we end up with a

typical length of the DAQ cycle of 70(40) ms for the 8(16)-bit mode of the camera, which corresponds to 14(25) Hz frame rate during the data taking.

From Fig. 5.8 one can deduce what is the maximum frame rate that the DAQ implementation can handle. Given the PCI speed of about 100 Mb/s, each DMA transfer from/to the shared SRAM lasts about 10-12 ms and 2-3 ms for the image raw data and the packed filter output respectively. During data taking the filter output should be also transferred to the PC. Therefore, even for negligible exposure times a DAQ cycle can not be shorter that 30 ms. The necessity of packing the output data becomes now obvious. If no packing would be performed, then exploiting the 16-bit mode of the camera would not be possible, even with infinite processing power of the C620 processors.

In a C620 processor about 30 ms are spent in data transfers, or waiting for transfers to be completed. Here we have also included the data transfers from/to the SDRAM to/from the processors internal memory. Given the fact that a C620 cycle is two times longer than the C80 cycle, for the 8(16)-bit mode of the camera there are about 50(110) ms left for pure image processing and output packing. To keep up with these numbers, all the processes occurring in the C620 DAQ loop have been coded in native assembly language.

While the DAQ system has been developed to fully exploit the 16-bit mode operation of the camera, it eventually turned out that it could not be used. The reason for this is that the camera developed noise, when running in that mode. The digital filter produced fake grains whenever the noise appeared. As a consequence, the higher level scanning procedures were failing so often, that automatic scanning was practically impossible using this camera mode.

5.2.3 Construction and implementation of digital filters

The scanning laboratories in CHORUS have developed various methods of grain identification in emulsion images. The approach which was adopted by our group largely inherits from the principles underlying the Nagoya TS operation. Studies on the properties of emulsion images [116] have shown that the grain information is enhanced by selecting the relevant regions in the frequency domain (filtering).

The output of the CCD camera after digitization is a discrete (digital) signal x_n, where n loops over the pixels along a line. The convolution of such a signal with a coefficient series c_k

$$y_n = \sum_{k=-\infty}^{\infty} c_k \cdot x_{n-k} \tag{5.5}$$

is defined as a *Finite Impulse Response* (FIR) filter. This class of filters can be extended such that previous results are re-used in the linear sum:

$$y_n = \sum_{k=-\infty}^{\infty} c_k \cdot x_{n-k} + \sum_{k=-1}^{\infty} d_k \cdot y_{n-k}\,. \tag{5.6}$$

Such filters are defined as *Infinite Impulse Response* (IIR), or recursive.

In both cases the output y_n is the transformed digital sequence after enhancing part of the frequency spectrum (which can be obtained by applying a discrete Fourier transform), and suppressing the rest. The behaviour of the filter in the frequency domain is determined by the filter coefficients c_k and d_k . In order to determine a set of coefficients which are optimal for a specific application, one makes use of the *z-transform*.

The z-transform of a digital sequence u_n is the complex polynomial

$$Z(z) = \sum_{k=-\infty}^{\infty} u_k \cdot z^k \,. \tag{5.7}$$

The z-transform evaluated on the unit circle $z = e^{-2\pi if}$ corresponds to the discrete Fourier transform of u_n, and can be used therefore to represent a digital sequence in the frequency domain. If $Y(z)$, $C(z)$, and $X(z)$ are the z-transforms of the sequences y_n, c_n, and x_n respectively, then it can be shown that convolution in the space domain corresponds to multiplication in the frequency domain

$$Y(z) = C(z) \cdot X(z) \tag{5.8}$$

and $C(z)$ becomes the transfer function of the filter in the frequency domain. By constructing a transfer function which meets the specifications for the behaviour of the filter in the frequency domain, the filter coefficients can be determined by inverting $C(z)$.

The most important parameters that have to be specified for the transfer function in the frequency domain are the cutoff frequencies, the strength, and the slope of the filter. The last two quantities are increased with the number of filter poles (coefficients). In an application this number is limited by the available processing power.

Many of the methods used for the construction of the transfer function, and hence the determination of the filter coefficients [117], inherit from analog filter design. The basic difference is that the z-transform is used instead of the Laplace transform. For our purposes we resorted to the Butterworth method using the interactive **mkfilter** program [118].

Our requirements are dictated by the appearance of the grains in a typical emulsion image. A grain is a rather small object compared to the length of a CCD line. It does not occupy more than five pixels along a line of 1024 pixels. Moreover, grains in focus have sharper edges than those which are out of focus. For this reason, the first requirement for the filter is that it should be high-pass with a cutoff frequency of about $0.1\,\mathrm{pixel}^{-1}$.

We also decided to use an IIR filter. While FIR filters are by definition stable (i.e. for finite input values, there always finite output values), IIR filters are more advantageous if stability is ensured, because they need less coefficients to achieve the same strength and slope.

We did not pose strict requirements on the strength of the filter. Given the number of poles one can increase it only at the cost of the achieved slope. Since all pixel gray values are in the range 0-255, an attenuation of 3 dB serves well our purposes. Such a strength suppresses low frequencies sufficiently that grains in focus can be selected after filtering with a threshold cut.

The following 3-pole high-pass IIR filter was constructed by the mkfilter program for the above specifications.

$$y_n = (x_{n+2} - x_{n-1}) + 3 \cdot (x_n - x_{n+1}) + c_{-3} \cdot y_{n-3} + c_{-2} \cdot y_{n-2} + c_{-1} \cdot y_{n-1} \tag{5.9}$$

The coefficients depend on the actual cutoff frequency. In Table 5.2 the values of the coefficients are shown for various cutoff frequencies. Using the 40× objective, it turned

Cutoff frequency (pixel^{-1})	c_{-3}	c_{-2}	c_{-1}
0.100	0.2781	-1.1829	1.7600
0.125	0.1978	-0.9104	1.4590
0.150	0.1378	-0.6959	1.1619
0.175	0.0921	-0.5344	0.8682
0.200	0.0563	-0.4218	0.5772
0.225	0.0267	-0.3553	0.2881

Tab. 5.2: Coefficients of the high-pass filter described by (5.9), for various cutoff frequencies

out that using a cutoff frequency of $0.15\,\text{pixel}^{-1}$ yields reasonable results.

In Fig. 5.9 a CCD line is shown before and after the filtering. One can observe the effect of the inefficiencies in charge transfers along a CCD line, which makes a threshold cut on the gray values meaningless. On the other hand, a cut on the filter output around 25 suffices to select exclusively and efficiently the pixels belonging to grains in focus. It should be noted that given the circular shape of the grains in an emulsion image it suffices to process the image only along the lines in order to identify the pixels belonging to grains.

To implement this filter on the C80 and the C620 DSPs we had to face the additional difficulty that these processors operate on fixed point values. To this end, the coefficients are left-shifted by seven bits (multiplied by 128) before they are used in the multiplications of (5.9). The input values are left-shifted by 14 bits, and once the final output value is produced, it is right-shifted back by seven bits before it can be used for the subsequent calculation.

While filtering we also perform the threshold cut and count the total number of pixels above the threshold. These operations are done without any additional requirement for processing time, by filling empty slots in the instruction cycles needed for the processing of a single pixel. On the C80 processor nine cycles per pixel were needed to perform all calculations, while on the C620 this was reduced to six cycles. Given the fact that the effective clock speed of the two processors is the same ($200\,\text{Hz} = 4 \times 50\,\text{Hz}$), an increase in speed by 50% was achieved by using the C620 processor. The time which was gained was used to perform the packing of the output as well. By fully exploiting the C620 processing capabilities we managed to have the packing done at a speed of 1 pixel/cycle. The byte reordering needed for the 16-bit mode of the camera is performed at the same speed. Eventually, using the C620 processor more is done compared to the C80, and at a higher

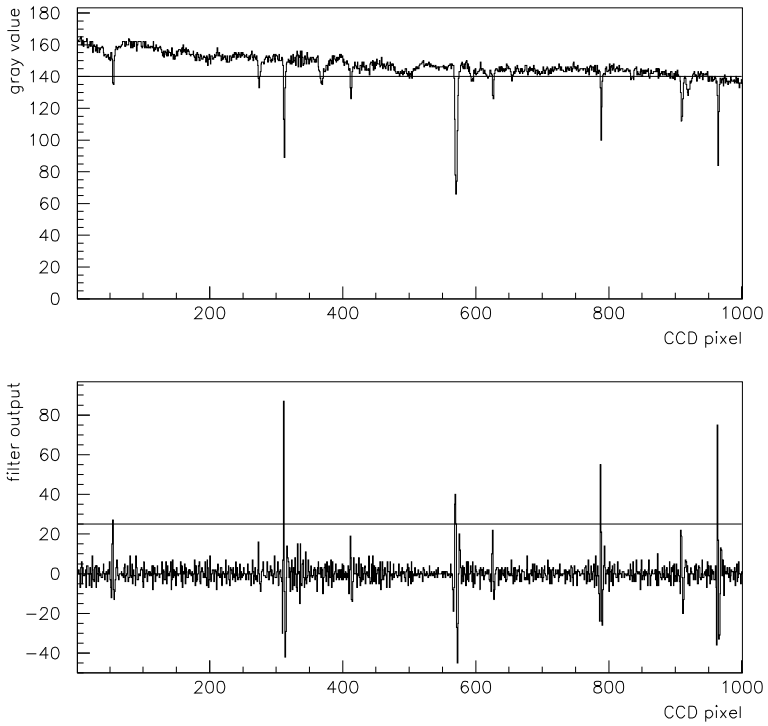

Fig. 5.9: A CCD line before (up) and after (down) filtering. The horizontal line in the upper plot indicates a potential threshold cut on the gray values. This is to be compared with the corresponding line in the lower plot which cuts on the output of the high-pass digital filter. The peaks that do not reach the threshold value correspond to grains out of focus, which have a lower contrast and hence a smaller high-frequency component.

speed. Given the fact that there are two processors in the C620 board, the improvement which was achieved was more than enough to cope with the 16-bit mode of the camera.

5.2.4 Real time tracking

The final goal in automatic scanning is to find particle tracks in emulsion. Tracks have to be found both efficiently and fast. A track finding algorithm can be performed either on-line at the end of every microscope view, or off-line on one or more computers of the Linux farm. On-line tracking is performed when the microscope searches for a track with predicted angle and position. It is a stripped version of the general tracking program which is run off-line for cases where an emulsion volume is scanned for any track with arbitrary direction and position. In this case, the position of the grains obtained in a microscope view are sent over the network to the Linux farm (see Fig. 5.6).

The track finding algorithm inherits its basic concepts from the CHARM II tracking [119]. It is described in detail in Ref. [120]. Its main features are outlined here.

The algorithm makes extensive use of a generalization of the binary tree concept to D dimensions to store grain positions. For any tracking method there is always the need to search for the existence of *hits* (grains in our case) within some volume. Such an operation is performed many times. If there are in total N hits, the time for a linear search scales with N. For binary tree searches, this time scales with $\log_{2^D} N$. If many searches have to be performed using the same binary tree, this improvement in speed largely compensates the time needed to construct it.

The first step of the on-line prediction track finding emulates, at least in its result, the Nagoya TS operation. During data taking whenever the data of a new plane has arrived from the C80 board, the pixels above the filter threshold are clustered into grains. The positions of the grains on the plane are stored in a two-dimensional binary tree. For every grain, a cone is opened towards the subsequent plane. The cone has the direction of the predicted slope and an opening angle given by the scanning angular acceptance. If within the projection of the cone onto the subsequent plane a grain is found, then a sum counter associated to the original grain is incremented, and the latter is replaced by the newly found one. If no grain is found, then the original grain is promoted to the subsequent plane along the predicted direction, without incrementing its sum counter. Once all grains on a plane have been processed, a next iteration starts in the subsequent plane. Once the last plane has been reached there may exist grains above a summation threshold. For those grains a thin cylinder around them is defined, which has an orientation of the predicted direction. All grains within the volume defined by this cylinder, the *track region*, are used as input for a general tracking session. The session decides whether the sum above the threshold was a random coincidence, or whether the positions of the grains are consistent with the scenario of a straight track segment.

Thanks to the technology of the binary tree, the summation for a data taking of 25 planes is performed in less than 100 ms. The actual tracking may take 0.5-5 s, depending on the required angular acceptance, the emulsion grain density, and the tightness of the various tracking parameters and thresholds.

For a general tracking session the grains are treated as points in space. A session starts by inserting the positions of the grains from all planes in a three-dimensional binary tree. From every grain a link is created to any other grain which lies within a small cube of $\sim 15\,\mu\mathrm{m}^2$, with the initial grain in the centre of the bottom surface (acceptance in forward direction only). Looping over all the grains the links are followed to develop segments. A segment is grown by a new link as long as the latter is within some acceptance defined by the previous links of the segment. Segment growing stops when no other link can be followed. If its length is below some threshold, it is discarded. Among multiple segments starting from the same grain, only the best is kept. The quality factor of a segment is determined by the number of grains it includes and the dispersion of the grains around a fitted straight line.

If there are N grains in total to be used as input for general tracking, the tracking time scales as $N \log_8 N$. If instead of binary trees linear searches would be used, the

tracking time would scale as N^2. To achieve high speeds all calculations involve fixed point operations. To this end, the relative positions of the grains are expressed in unit of $\frac{1}{u}\mu$m, where typical values of u are 20-50. Calls to trigonometric functions are avoided. Vector algebra (dot and cross products) is used instead. Moreover, thanks to the use of object-oriented technology in the coding, the track finding algorithm is adaptable under any input condition and flexible to use various different criteria for the segment growing step.

5.2.5 Database implementation

The track segments found in an emulsion sheet comprise the ultimate output data set during the scanning. Logically this data set belongs to a specific sheet of a target module which was exposed to the neutrino beam in some specific time period. Therefore, the scanning output could be stored under a logical tree of runs, target modules, and sheets. However, for any physics analysis such a hierarchy is useless. Segments found in an emulsion sheet, belong to tracks of particles created in neutrino interactions, and a tree structure starting from an event basis is more relevant. Furthermore, segments from various sheets can be combined into tracks only when inter-sheet calibration and alignment has been achieved. Usually this is done using a subset of track segments which do not necessarily belong to a specific event or an electronically reconstructed track.

For all these reasons an object oriented database is more appealing to use for storing and analysing the scanning output, compared to a relational database or a hierarchical storage. The use of Objectivity/DB has allowed us to store polymorphic data with many logical links between the data sets.

The definition of the data objects to be stored in Objectivity/DB is specific to a *federated database*. Apart from the *schema*, the federation keeps a record of the *databases* it contains, which are actual files of the operating system. A database can contain one or more *containers*. During a transaction with the federated database, read/write locks are set at the container level. Persistent *objects* are stored in containers.

We use a single federation for both the scanning data and the analysis results. This is done mainly because a program can have transactions with only one federation at a time. However, using a single federation is convenient for the software development, because one has to maintain and evolve only one schema. Therefore there are databases related to the actual scanning, and other analysis specific.

For every reconstructed event in CHORUS, scanning predictions are generated for only the target module, where the primary vertex is expected to be located. For every target module which is to be scanned, there is a database specific to it. The database contains emulsion sheet and event containers. In this way all necessary associations between the track segments found in the emulsion sheets and the event structures where the tracks belong to, are kept within a single database. Every sheet in the database contains its physical description, the nominal position of the fiducial marks on it, and a set of predictions to be scanned. For every prediction, all the candidate track segments found during the scanning are stored together with their track regions (see p. 97). Associated to every sheet there

are calibration and alignment objects, which are used to transform the position and the slope of every segment in the sheet to a system of reference common for all sheets in the module.

Online acquisition databases keep track of the microscope hardware and on-line tracking parameters during a scanning session. Every segment is stored with an association to the acquisition object of the scanning session during which it was found. Currently, the hardware parameters that are stored concern the microscope stage, the optics, and the DAQ system.

Other databases are used to store the output of the off-line tracking, and the results of an analysis. Thanks to the object oriented technology of the database, one can have direct links of such results to the actual raw data without having to build a tree structure.

5.3 Scanning procedures

5.3.1 Software implementation

It has already been mentioned that it was a design choice to construct the microscopes with modular hardware. To really benefit from such a design the corresponding software had to be based on object oriented programming. This would ensure that only short development cycles are introduced when upgrading a hardware component, ideally without affecting high level programs, such as the one for the automatic emulsion scanning. Therefore, the adoption of the C++ programming language was a design requirement rather than a matter of taste for the software development. Moreover, the intrinsic features of the language, such as the exception handling mechanism, suit very well the requirement of the DAQ system for stability. The use of the C++ programming language was the main reason for choosing Objectivity/DB for the database implementation.

The design of the software components which are used for the emulsion scanning is shown in Fig. 5.10 following the UML[24] notation.

For all the hardware components which either could be controlled or accessed directly by the microscope PC, or the settings of which were used in on-line or off-line analysis algorithms, the corresponding software object was constructed. A software instance of the microscope (a CMicroscope class) is a collection of virtual interface components, such as a stage, a light control unit, an image taker, and a videocard. For every hardware subsystem physically present in the microscope its corresponding class is used to implement the interface. For example, a CStageNikon, a CStageMicos, or a CStageDummy object is used whenever the microscope has a NIKON, a MICOS, or no stage component respectively. Such a scheme is followed by all hardware components. It ensures short software development cycles whenever a new technology is introduced, without affecting higher level classes and programs, and without the need of the physical presence of all hardware components in a system.

For most applications the microscope is used together with an emulsion sheet. For

[24] Unified Modeling Language

this reason the CMicroscope class is extended to a CScanner, which implements additional functionalities related to the scanning procedure, such as surface detection, light control, and data taking methods.

Interface classes, a feature of the object oriented software technology, are also used in the scanning program to handle all possible types of scanning. As shown in Fig. 5.10 the program is based on generic algorithms which only assume that an emulsion sheet consists of a base and two emulsion layers, but of no predefined thicknesses.

For the data stored in the database the persistent classes inherit from the same data classes as the corresponding transient classes do. Moreover, a thin interface layer is built on top of the low level access functions to the database. With this design the scanning and the analysis programs use only transient objects, without the need of extensive knowledge of the database structure.

All programs are written such as to accept as input a settings file. A CParameterReader class holds its contents as pairs of parameter names and values grouped in various *sections*. The name of each section is usually indicative to a specific class, which uses the values under the section to initialize its parameters. Thanks to the use of such a class it is possible to have all the code free of numerical constants.

5.3.2 Handling of messages and controls

In the CHORUS DAQ system a dispatcher program was used for message transmission between processes (see p. 54). The same program runs on a server of the scanning facility. It handles messages sent by/to the programs running on the DAQ PCs.

For stability purposes all programs interacting with the microscope hardware, such as the scanning program, are not implemented with a user interface. Instead, they listen to dispatcher messages. This allows running of run control processes and graphical user interfaces in different computers than the DAQ PCs. The link to the dispatcher is created during the startup of a program, which defines the tags of the messages that it has to receive from other processes. Most often these tags are of the format PANEL0x, where x is the microscope identifier (1,2, or 3).

During the initial phase of the software development the dispatcher was also used to distribute grain data from the DAQ PCs to the Linux farm. Having the dispatcher as a mediator decouples the grain data taking from the processing in the farm. This limits the overall dead-time in the scanning program to the time needed for the transferring of the grain data to the dispatcher host over the network. At a later stage of the software development a more efficient use of Objectivity/DB was achieved. The grain data were stored directly into the database and consequently this use of the dispatcher was abandoned.

5.3.3 Hardware calibration and quality controls

Whenever a maintenance operation is made to a microscope, or a new emulsion sheet is mounted, a series of hardware calibration programs are run, before starting an automatic scanning session.

Initialization of the microscope

For every program using the microscope, a CMicroscope object is constructed. For a microscope with all the hardware components present the pointers to the interface classes are initialized accordingly:

1. The VStage pointer is initialized with the address of a CStageMicos object. The running parameters from the stage controller (such as the range limits and the pitch size) are read and mirrored in the software object.

2. The VImageTaker pointer is initialized with the address of a CImageTakerDSP object. The DSP programs indicated by the settings file are downloaded to the C80 and C620 processors. The object launches the microscope DAQ, by starting and synchronizing the processes on the processors. The class reads from the settings file the list of pixels that have to be masked, and the filter coefficients and threshold.

3. The VVideocard pointer is initialized with the address of a CVideocardPulsar object. It initializes and configures the Matrox Pulsar graphics card to start displaying the camera output on the secondary screen.

4. The VLightControl pointer is initialized with the address of a CLightControlDACLCD object. It initializes the CerealBox and activates the analog output channel which is used to control ALM element of the LC-Unit of the optics. If it is specified in the settings file, it loads a mapping between the nominal light intensity, which ranges from 0 to 1, and the input DC voltage to the ALM element.

5. The CMicroscope class inherits from the CObjective class which keeps the transformation between the camera coordinate system (pixels and lines) and the stage coordinate system (μm in X and Y). This transformation is read from the settings file.

6. If a CScanner class is used, the settings file contains all the necessary parameters related to light control, surface detection and data taking procedures.

Image quality control

The ObjectiveRing program was designed to assist the optics tuning. After the construction of the CMicroscope object a second thread in the program is set up to continuously display quality factors on the secondary screen as overlay text on top of the camera output. At the same time the program allows the user to use the stage joysticks and to modify the light intensity via dispatcher messages. The second thread continuously transfers an emulsion image from the C80 board, analyses it, and updates the quality factors on the screen.

The transferred image is analysed with the same digital filter which runs on the DSPs, implemented on the Pentium processor. Whenever a grain segment is identified, the program calculates the integral of the filter output above the threshold over the segment pixels. The mean value of these segment integrals is the first quality factor. High values of this

number indicate sharp grains and high contrast, hence good optical quality. When we shift the lens group in the objective which corrects for the emulsion refractive index, the optimal position is found where this number is maximized.

The *brightness* of the image is calculated not as the mean gray value of its pixels, but as the gray value for which only a small fraction (typically 10%) of the pixels are brighter. This is the second quality factor and it is used to assist the tuning of the condenser height. For CS/SS type of sheets at the optimum position the brightness is the same in both emulsion layers, and reaches a maximum at the centre of the base. For TP type of sheets, the brightness is maximized close the the edge of the emulsion side that is scanned.

The third displayed quantity is the number of pixels above the filter threshold, as it arrives from the DSPs. For a given position of the focal plane in the emulsion the number of pixels above threshold is increased when the optical quality improves, because a grain occupies more pixels in the digital image.

The program kills the second thread and exits once it receives the appropriate message from the dispatcher.

Identification of noisy and dirty pixels

The FindNoisyPixels program was written for robust identification of the pixels of the CCD which are noisy, dead, or covered by dirt. The list of such pixels is passed to the C80 board for masking (see p. 88).

The program needs and emulsion surface to be already mounted on the stage. It starts by identifying all emulsion and base surfaces (see 5.3.4). Afterwards it focuses in the middle of the base. Since one would not expect to find there any emulsion grain or dust particle, any pixel with filter output above the filter threshold is noisy, dead, or covered by dirt.

In principle one could retrieve all the pixels from the DSPs, as the list of pixels above the filter threshold. However, single noisy or dead pixels have relatively big contrast to the image background and the high-pass digital filter resonates strongly. As a consequence, more than one pixel may pass the filter threshold. To avoid masking more pixels than what is absolutely necessary, a differential filter running on the Pentium processor is used instead. This filter searches for segments enclosed between two sharp and opposite direction slopes. The threshold on the slope and the maximum length of a segment are set from the settings file.

The program continuously transfers emulsion images, searches for noisy pixels, and updates the list of pixels to be masked on the C80. It terminates when no more pixels above the filter threshold are found by both the DSPs and the Pentium processor.

Camera calibration

The positions of the grains identified in an emulsion image by the DSPs are transformed from the camera coordinate system to the stage coordinate system before they can be used as input for the tracking. The CameraCalibration program calculates the corresponding

affine transformation[25] that is loaded from the settings file during the construction of a CMicroscope object.

After the user has positioned the microscope such that it focuses inside an emulsion layer, the microscope stage moves horizontally with a given translation. If the stage moves by a distance smaller than one view, some of the grains that were visible before the movement are still present in the camera frame but in a different position. Their translation in the camera frame of reference is calculated by maximizing the overlap between the filter output before and after the stage movement. This is repeated eight times in total and from the pairs of translations in each reference system, the optimum affine transformation is calculated. This procedure is iterated until the calculated difference of the distance travelled by a point brought from the top left to the bottom right corner of the camera, is within a predefined tolerance (typically 1 μm).

With this procedure, a possible rotation of the camera with respect to the stage is also corrected. However, we always try to have the axes of the camera and the stage as much as possible parallel with respect to each other.

5.3.4 Scanning overhead

In all programs using the CScanner a significant fraction of the time is spent in surface detection and light control. The size of this overhead is proportional to the size of the DAQ cycle.

Surface detection

The criterion whether the microscope is focusing inside the emulsion or not, is determined by the number of pixels above the filter threshold, which is proportional to the number of visible grains. Therefore, at any moment the DSPs can inform the higher level procedures whether the emulsion is in focus by comparing this number of pixels against some threshold. Typical values of this threshold are 200-800 pixels depending on the emulsion grain density.

All surface detection procedures are based on an algorithm which detects the transition while the stage is moving along the vertical axis with a constant velocity. The position of the transition point is determined by the stage speed and the DAQ cycle length. A surface is always found in two steps. First it is detected with a coarse resolution of about 5 μm. Then the stage starts moving slower in the opposite direction to determine the position of the surface with a resolution of about 1 μm. In both cases the length of the DAQ pipeline is taken into account when calculating the position of a surface.

There are four surfaces in total: the two outer emulsion edges and the two base edges. Their position is mirrored in the CScanner class so that it is fast to re-measure.

[25] Affine is called the mathematical transformation $X \mapsto R \cdot X + V$, where X and V are two-dimensional vectors, and R is a 2×2 matrix.

Light control

The light control is essential to have the optical conditions in both emulsion layers similar. This procedure aims at setting the light intensity (input voltage of the ALM element of the LC-unit) such that the brightness of the resulting image is within the accepted range. The optimal light intensity for the top and bottom emulsion layer is always monitored in the CScanner class to make subsequent calls to the procedure execute faster.

Care is taken to avoid over-illuminating the emulsion, because it results in camera overflows. In that case a high increase in the number of pixels above the filter threshold is observed. Hence, camera overflow can be detected indirectly by comparing this number to a threshold value. In the light control procedure it is checked by this method whether camera overflow is occurring before calculating the brightness of an emulsion image transferred from the C80 board. In case of overflow the brightness is assumed to have the maximal possible value (255), and the algorithm will therefore try smaller light intensities.

LC-Unit calibration

The response of the light intensity to the voltage applied to the ALM is non-linear, therefore a lookup table is used to achieve a better mapping. This lookup table is automatically constructed during the initialization of the scanning program.

Once all emulsion surfaces have been measured, the microscope focuses inside the base. The CScanner takes image raw data for increasing light intensities and calculates the resulting image brightness, until camera overflow is detected or the highest light intensity (equal to 1.0) is reached. From the obtained mapping between intensity and brightness the lookup table is constructed such that the mapping becomes linear.

5.3.5 Definition of the sheet coordinate system

In 3.3 it is mentioned that for the alignment of the CS and SS sheets special X-ray guns are used which create on the sheets round marks with a diameter of 1 mm. Moreover, a grid of fiducial marks is printed on the TP and SS sheets using ultra-violet light before they are developed. The X-ray and fiducial marks are used during the scanning to determine a coordinate system on the sheet, referred to as *the nominal frame*. Measuring the marks precisely and efficiently is essential, because all predictions and track segments are stored in the database in the nominal frame.

X-ray mark finding

Given the fact that an X-ray mark is larger than the size of the optical field, one can define a simple trigger for the detection of a mark. Once the emulsion surfaces are detected, the microscope brings in focus the side where the mark is printed. Spiraling around the starting position the CScanner object searches for the mark by checking whether the mean gray value of the emulsion image is below some threshold.

Once the X-ray mark has been identified the CScanner object measures its position in a way analogous to the surface detection. Moving slowly along each of the two horizontal directions, the mean gray value of the captured emulsion images is compared to the threshold value. This way we measure the position of four edges of the mark, and assuming that it is round, the position of its centre is determined with a resolution better than 10 μm.

Fiducial mark finding

There are two types of fiducial marks. They are either simple round blobs with a radius of 30 μm, or they are complemented with concentric ring with inner and outer radii of 180 μm and 240 μm respectively. Using 40× magnification, the central blob can be contained within the field, but not the entire ring. However, if part of it is visible, the fiducial mark finding algorithm is able to estimate the position of the central blob, and guide the microscope accordingly.

The algorithm starts by binning 8×8 pixels of the original emulsion image into a single pixel. The rest of the pattern recognition is performed on the binned image which has a size of 128×128 pixels. The binning is done for the algorithm to execute fast. It steps over the following procedures:

1. The image is binarized using a threshold cut on the gray value.

2. Neighbouring pixels are clustered to form blobs. Very small and very big blobs are discarded; the limiting numbers of pixels are read from the settings file. If no blob satisfies the size criterion, the algorithm returns a negative result.

3. For the blobs satisfying the size criterion the perimetric pixels are identified. A blob may have inner perimeters as well. This is the case when there are holes in the blob. If an inner parameter is very big, this blob is discarded. At the end of this step there is a set of external blob perimeters. If this set is empty, the algorithm returns a negative result.

4. The pixels belonging to the blob perimeters are used to fill an empty image. In this last step the algorithm tries to maximize the overlap between those pixels and a circle which has the radius of a fiducial mark. Given the fact that three points define uniquely a circle, three pixels of the perimeters are sampled many times to generate candidate positions. For each of them an overlap calculation is performed. If the candidate with the highest overlap exceeds a threshold, the algorithm decides that it has identified the fiducial mark and returns the position of the best circle as the reconstructed mark position.

5. In case the central blob is not found, the algorithm tries to detect part of the outer ring. In this case, the algorithm requires a large overlap with two circles which have radii equal to the inner and outer radii of the fiducial ring.

In case the reconstructed position is far from the camera centre, which is the case when a part of the outer ring is identified, the stage moves accordingly and the procedure is repeated. To achieve a measurement resolution of better than $2\,\mu$m, the algorithm is repeated on an image which has per axis two times smaller binning; this image has a size of 256×256 pixels.

5.3.6 Operations during normal oscillation search

CS, SS scanning

For a physical emulsion sheet there are one or two yearly exposures which define its position in the experiment. For every exposure there are several CS periods. A CS period is defined as the time period in the actual experiment between two consecutive replacements of the CS sheets. By definition, for a CS sheet there is a single CS period for a single exposure.

After the measurement of the X-ray marks, the program loops over the yearly exposures, the CS periods, and for every period it loops over the corresponding predictions. Before scanning a prediction, its position and its slope are transformed from the nominal frame to the stage coordinate system.

The area which has to be covered around the predicted position is determined by the extrapolation errors from the electronic trackers. However, this is done only after a good calibration of the plate has been achieved. To this end, the first several hundred predictions are scanned covering a larger area. Matching the patterns of predicted positions and slopes to the ones of the candidate tracks, the position of the sheet with respect to the trackers can be defined with a resolution of better than $20\,\mu$m.

The emulsion sheet is mounted such that the downstream side is on top. Track segments are searched for in the downstream side first. The area to be scanned is not covered by spiraling around the predicted position, but using diagonal movements. In 5.1.2 the need of a backlash movement is described. Setting the backlash vector equal to the view step along a diagonal only one movement is effectively performed, thus minimizing the time spent in stage movements.

When scanning a nuclear emulsion sheet, a number of images is taken at various depths across the emulsion layer, keeping the distance between two consecutive images constant. This procedure has to be as fast as possible, ideally limited only by the camera frame rate and the mechanical stability. Therefore, rather than taking a set of images moving in steps, the stage initializes a motion of constant velocity along the vertical axis and it stops only after the image data taking procedure has finished.

For CS/SS type of scanning the thin emulsion layer has to be covered fully and sampled sufficiently. The data taking method which is implemented such that for the given number of planes that are taken (typically 25) there are always a few of them in the beginning and at the end (typically five in total) which are outside the emulsion (empty). The number of planes inside the emulsion is allowed to differ by one from the nominal value. If for a view there are no empty planes before or after the emulsion, or if they are too many, the data taking for this view is repeated after modifying the starting position and the speed

of the stage. These values are updated after every view to correct for small variations of the thickness and the position of the emulsion layer. With this feedback procedure, the two surfaces defined by the emulsion layer edges are continuously updated, eliminating the need for repeating the surface detection procedure during the scanning of a prediction.

Candidate segments found in the downstream emulsion layer are put in a stack in the memory. As soon as the stack contains more than 100 segments, or when the downstream area is fully covered, the microscope starts searching for the corresponding segments in the upstream side. Every segment found downstream is extrapolated to the upstream side of the sheet with the predicted slope. If it belongs to a track which is within the angular acceptance, a segment should be found close to the middle of the view. For this reason and to speed up the tracking, when scanning the upstream side the filtered area of the image is reduced to the central region of 300×300 pixels.

If the candidate segment is validated, it is stored with a position on the horizontal plane given by the base intercept of the fitted track in the downstream side. The two base intercepts and the measured base surfaces are used to calculate the slope of the track segment. Before storing a segment in the database it is transformed from the stage coordinate system to the nominal frame. The upstream and downstream track regions, and the positions of the four surfaces are stored as well. The stored segment keeps a reference to the most recent measurement of the marks and the object with the acquisition parameters of the scanning session.

Scan-back in the bulk

As is explained in 4.3.1, tracks found in the CS and SS sheets are *scanned-back* in the TP sheets until a decay or the primary interaction vertex is reached. The predictions and the track segments are stored in the database in the nominal frame of the three nearest fiducials. The use of a local rather than a global coordinate system on the emulsion sheet removes more efficiently the effects of large-scale distortion, which is larger in TP-like sheets than in CS/SS-like emulsion sheets. To speed up the overall scanning of an emulsion sheet, instead of looping over predictions, the corresponding procedure in the scanning program loops over the fiducial marks. For every fiducial mark, it loops over the yearly exposures, the CS periods, and the predictions associated to each period.

After the three closest fiducial marks have been measured, the prediction for a track is scanned using a single microscope view at the upstream edge of the emulsion. Data taking is performed using always the same stage speed, which corresponds to a fixed plane spacing determined in the settings file. Judging from the optical quality of the emulsion sheet the spacing is set accordingly between $1.5 - 3\,\mu$m. For this type of scanning the fraction of time spent for image data taking is rather small compared to the case of CS/SS type of scanning, where it accounts for about 50% of the scanning time. The largest fraction is spent in fiducial mark measurements, stage movements, and the actual tracking, which has to cope with larger numbers of grains for the same emulsion volume.

The alignment procedure with respect to the downstream sheet, precedes the scan-back procedure. It is done by scanning large square areas in eight places on the sheet for tracks

of beam related muons.

Vertex trigger and off-line tracking

Whenever a scan-back track is no longer found in a TP, there are three possibilities concerning the fate of the track.

1. The segment exists but is not reconstructed due to scanning inefficiency.
2. The primary neutrino interaction vertex from which the track originates exists in the sheet, downstream the region where the data taking takes place.
3. A large angle kink along the particle trajectory occurring downstream the region where the data taking takes place, brings the scan-back track out of angular or spatial acceptance.

In case a scan-back track is not reconstructed automatically in the upstream layer of the sheet, a *vertex trigger* is set. Grain data taking is performed covering the whole emulsion thickness. This is achieved with two views (upstream and downstream layer) of about 80 planes each. The grain maps and some relevant information concerning the position of the view are compressed and sent to the dispatcher.

At least one PC in the Linux farm should be running at that moment the general tracking program. The program waits for the dispatcher to send the grain data sets and starts the processing. In the meantime if another data set has arrived at the dispatcher's shared memory, it either stays there until the processing of the previous set is finished, or it is distributed to another PC. At least three PCs are foreseen to be able to handle the data rate from a single microscope.

The general tracking algorithm should eventually decide between the above three scenarios. In case of recovered scanning inefficiency it generates the prediction on the subsequent upstream TP sheet. If a decay kink is identified, it stores the relevant information to the electronically reconstructed event in the database. It also updates the slope and the position of the scan-back track accordingly generating the prediction on the subsequent upstream TP sheet. Finally, in case an interaction vertex is reconstructed, the trajectories of all produced MIPs are stored. Afterwards they can be followed-down towards the electronic trackers to obtain a complete reconstruction of the neutrino event.

5.3.7 Procedures for other types of scanning

Volume scan

In 4.3.3 it is mentioned that for the Phase II analysis of the CHORUS data, the NET-scan technique will be applied to search for all decay vertices in a large volume downstream the primary interaction vertex. For this type of scanning a normal scan-back procedure of all well reconstructed tracks attached to the vertex is done first. With the Nagoya TS system for the subsequent downstream sheets only a thin emulsion layer will be analysed. This is

due to the hardware limitation of the TS that it can take only 16 planes per microscope view.

Our model for the Phase II analysis is based on the flexibility from the use of the DSPs and the relatively easy scalability of processing power in PCs. As soon as the primary interaction is identified, for the subsequent downstream sheets a volume scan will be performed, similar to the case of the vertex trigger. Every microscope view is covering a sub-volume for which the general tracking algorithm has to find all secondary vertices and passing through tracks. The number of PCs in the farm can scale easily to cope with the incoming data rate. The 100 Mbit/s network is also able to transfer sufficiently fast large amounts of grain data. Therefore, the overall speed for this type of scanning is determined mainly by the image data taking rate.

Surface scan

The notion of "predictions" is also used to define the centre of a relatively large area to be scanned for all track segments which have a slope within some angular window. This is the case of the alignment data taking for the TP sheets in CHORUS. During a surface scan tracks are searched for in a relatively thin emulsion layer at the top surface of an emulsion sheet which is mounted on the stage. For CS/SS-like sheets, a surface scan is the same as the normal scanning procedure with the exception that the area covered per prediction is much larger.

The surface scan method is best exploited by using the microscope only for grain data taking and performing tracking in PCs of the Linux farm. This strategy has various advantages. First of all it maximizes the scanning speed because there is no processing load due to on-line tracking. If the tracking results are not very satisfactory, there is no need to re-scan the same sheet with different on-line tracking parameters. Since the grain positions are stored in the database, one has only to re-run in off-line mode the tracking program. Such a type of scanning was used for the white kink measurement described in chapter 6.

5.3.8 Scanning Performance

Measurement resolution

The positions and the slopes of the track segments found by the automatic scanning program on a sheet, are parameters which are measured with a finite resolution. Three basic factors determine the resolution of the measurements.

1. the precision of the stage movements and the readout resolution of its linear encoder.
2. the precision of the camera calibration with respect to the stage.
3. the tracking algorithm.

A calculation of the measurement resolution of slopes an positions in emulsion is described in 6.3.3, p. 127.

Scanning efficiency

The automatic scanning program does not necessarily find all tracks even if they lie within the angular acceptance. The overall scanning efficiency is determined mainly by the following factors:

1. The *grain finding efficiency.* After the development of an emulsion sheet, the grains along a track segment are spaced in the direction normal to the sheet by at most $1.5\,\mu$m. This value is similar to the effective depth of focus. For the image data taking the distance between successive planes is of the same order. Therefore, for any track in a microscope view, the number of reconstructed grains along the track are expected to be equal to the number of planes taken inside the emulsion. However, fluctuations in the grain distance distribution along the track and the emulsion sensitivity may result in losing a small number of grains.

   The distributions of the number of grains along the found tracks in CS/SS sheets are shown in Fig. 5.11. There are plots for the cases when 19, 20, and 21 microscope images were taken inside the emulsion. A lower bound on the "grain finding efficiency" of the digital filter running on the DSP can be retrieved by performing binomial fits to these distributions. They yield an average grain finding probability of (86±1)% which is a product of the grain finding efficiency and the probability that a grain exists on a given tomographic plane. The binomial fits are not excellent because the planes are not completely independent. As a consequence the resulting distributions are broader.

2. The *tracking efficiency.* The efficiency of the tracking depends mainly on the signal (grains belonging to the track) to noise (background grains) ratio of the input. The signal can be increased by finer sampling of the emulsion. This is the reason that we use 20-25 planes to sample an emulsion layer of $40 - 50\,\mu$m, compared to the 16 planes taken by the Nagoya TS for the same thickness. The number of background grains depends on the quality of the emulsion and the development procedure, as well as the accumulated particle radiation. High emulsion and development quality ensures that there are mainly grains belonging to track segments and relatively few randomly distributed fog grains. For the CHORUS experiment, CS sheets have fewer grains than SS sheets, which in turn have fewer grains than TP sheets.

   The results of extensive studies concerning the tracking efficiencies can be found in [120].

3. Specifically for the CHORUS TP sheets a track may not be found due to inefficiency of the fiducial mark measurement. Since tracks are searched for using a single microscope view, an inaccurate measurement of a nearby fiducial mark may guide the microscope far from the actual position of the track segment. The fiducial mark finding efficiency is estimated to about 95% for a triplet of marks.

Speed considerations

The scanning speed depends on the type of the scanning performed and the emulsion sheet.

A successful microscope view, i.e. a view where the correct number of planes have been taken inside, before, and after the emulsion layer, takes about 2.1 s when 25 frames are taken. The data taking of the 25 frames takes about 1.6 s, assuming 8-bit mode operation of the camera and 5 ms exposure time. The rest of the time is spent in waiting for the stage to reach constant speed (about 200 ms), and in reading out the position of the stage from the controller. The readout of the position is performed four to six times during a view and each of them take about 50 ms. The large time spent in reading the position from the stage controller is due to the RS-232 interface and the fact that numbers are sent as ASCII characters without any compression.

Once the view has been taken the stage moves to the position of the next view. Meanwhile, the on-line tracking starts its calculations or grain data are sent to the dispatcher. The movement of the stage from one view to the next takes about 1 s. The lower limit on this time interval is posed by the maximum acceleration that the stage can use without losing its stability. If the grains are simply sent to the dispatcher, or the scanning is performed on a CS/SS-like sheet with low grain density, the tracking finishes in time and the microscope can start the data taking of the next view. In this case the whole procedure for a view is completed in about 3 s. If the grain density is large, which is the case for TP-like sheets, then the tracking may take in total up to 5 s. Therefore depending on the emulsion sheet and the available processing power in the DAQ PC, the complete measurement of a view is performed in 3 to 6 s.

For TP-like sheets, the fiducial marks are measured before a prediction is scanned. Moreover, a prediction is scanned using a single view. For this reason, most of the time is spent in movements, and in fiducial mark finding. The latter involves a lot of surface detection, which as already mentioned scales with the length of the image DAQ cycle. Given the fact that the on-line tracking has to deal with a large number of grains, it takes about 5 s on the current DAQ PCs. Overall, in a scan-back session scanning a prediction takes on average about 10 s.

For the reasons described above, it is foreseen in the future to apply the following improvements.

- A new camera type based on CMOS[26] technology will be used to achieve higher frame rates. Initially, the new cameras will operate at 40 Hz. Updating the DSP hardware will allow us to operate them at 120 Hz or even more, depending on the limitations from the microscope mechanics.

- A new stage controller can be used. This will allow faster readout of the position and faster movements without becoming unstable.

- The on-line tracking will be eliminated from the scanning program. As explained in p. 109 the microscopes will be used only for stage control and grain data taking.

[26] Complementary Metal Oxide Semiconductor.

Replacing the processors of PCs with more powerful ones, it is expected that the scanning speed will be determined only by the camera frame rate and the speed of movement of the stage.

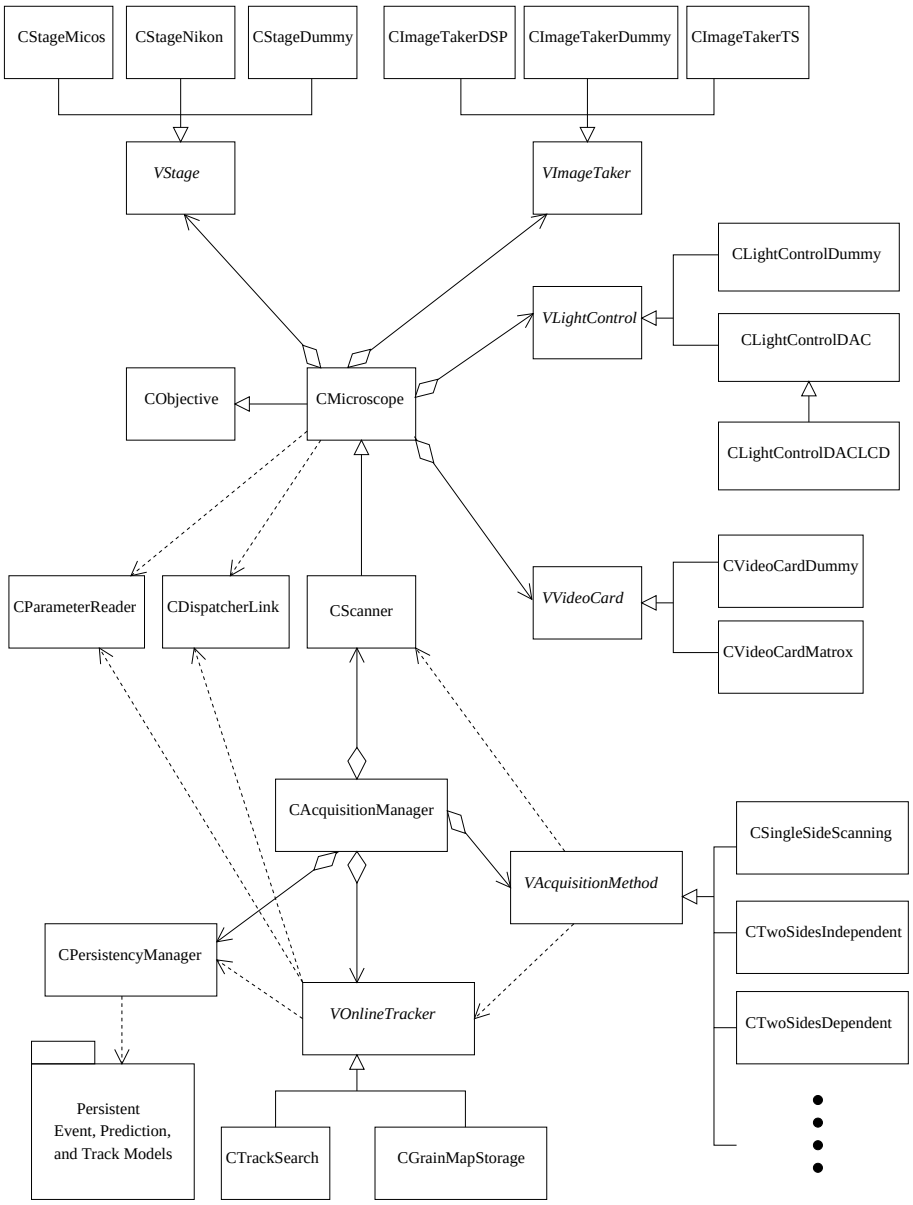

Fig. 5.10: Class diagram for the software responsible for the scanning procedures, following the UML notation.
The boxes represent classes or class categories. Wherever the names of the classes appear with italic fonts, they are abstract interfaces.
Dashed arrows represent use relationships,
arrows with diamond-shaped starting point represent containment,
lines with large open arrowheads represent realizations of interfaces or generalizations.

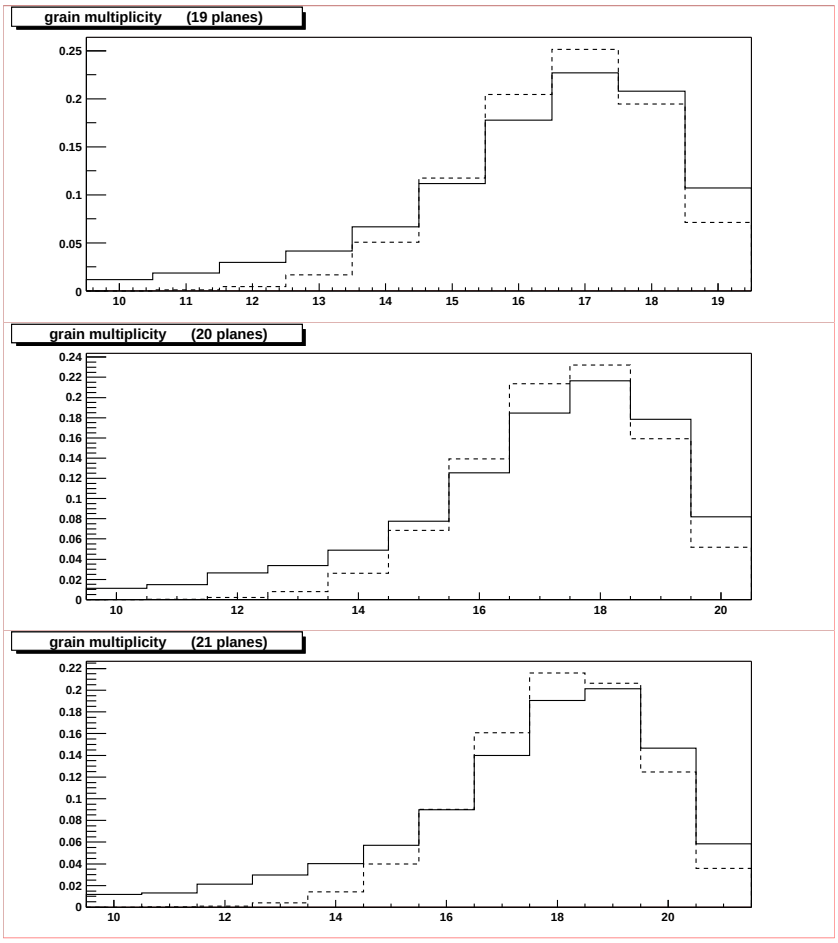

Fig. 5.11: Distribution of the number of reconstructed grains along a track. The solid line represents the fitted binomial distribution.

6. THE TEST EXPERIMENT AT THE PS AND THE WHITE KINK MEASUREMENT

In September 1997 a test experiment was performed at the CERN-PS in view of testing new detector configurations based on emulsion and high precision electronic detectors. A pion beam with momentum up to 15 GeV/c was directed towards an emulsion target followed by a spectrometer where silicon strip and emulsion trackers were used. The experimental setup, and in particular the emulsion target and tracker, were used also for a measurement of the hadronic *white kink* interaction rate.

As mentioned in 3.5, the white kink is the most significant background process in the $\nu_\mu \rightarrow \nu_\tau$ oscillation search in CHORUS. This chapter describes the test experiment at the CERN-PS, and the analysis of the data taken for the white kink measurement. Emphasis is given to the momentum measurement of the white kink daughter particles using the emulsion tracker.

6.1 Aims of the test experiment

The test experiment at the PS was part of the detector R&D for the TOSCA experiment [121]. TOSCA was proposed to search for $\nu_\mu \rightarrow \nu_\tau$ oscillations in a parameter space wider than what is expored by CHORUS and NOMAD, towards smaller mixing angles (by at least an order of magnitude) and Δm^2 values.

To achieve this goal, the design of the TOSCA detector was driven by the requirement for high background rejection power. This would be achieved by combining the ν_τ CC detection techniques developed in CHORUS and NOMAD: Neutrino interactions would take place in emulsion, where topological criteria can be applied to identify ν_τ CC candidates. Background processes would be suppressed with the kinematical analysis of the particle tracks reconstructed in an advanced spectrometer.

According to its conceptual design TOSCA is a magnetic detector consisting of five independent modules. Each module has a 9-10 cm thick emulsion target followed by silicon strip detectors, an emulsion and an electronic tracker. The longitudinal size of each module is about 100 cm, so that all modules fit inside the UA1 magnet, which was used also by NOMAD, leaving also some space for electromagnetic calorimetry downstream the last module. This toroidal magnet can deliver a magnetic field with strength of 0.7 T.

The three main aims of the test experiment at the PS were:

1. To check the compatibility of the basic elements of a TOSCA module (thicker emulsion targets, with respect to the CHORUS one, silicon detectors, and emulsion track-

ers) in a magnetic field of 0.7 T.

2. To test the performance of a compact tracker made of emulsion sheets.

3. To measure the white kink interaction rate. This would determine whether the hadronic channel would be included in the $\nu_\mu \rightarrow \nu_\tau$ oscillation analysis of TOSCA, and if this is the case (like in CHORUS), which kinematical cuts should be applied in order to suppress it efficiently.

6.2 Beam and experimental setup

6.2.1 The pion beam

The test experiment took place at the T9 beam line of the PS. In this beam, particles are available with momentum up to 15 GeV/c.

The experiment was designed to trigger on every pion interaction during a PS proton spill. The DAQ system was able to accept about 20 triggers per spill. As a consequence, the beam needed to have very low intensity, so that enough particle trajectories would be recorded by the electronic detectors without the emulsion sheets being over-exposed. To meet this requirement on the beam a special tuning of the elements of the beam line had been performed before the actual test beam. It involved mainly the optimization of the opening of the various collimators and the choice of the production target.

Three sets of runs were taken during the time allocated to the test experiment, corresponding to its three main aims.

1. An exposure with 15 GeV/c pions to test the whole detector configuration. This choice of momentum was made because pions of such energy give interactions in emulsion which resemble NC interactions of 20-30 GeV/c neutrinos in terms of track multiplicity and total transverse momentum.

2. Two exposures with pions of momenta 5 and 10 GeV/c to measure the momentum resolution of the emulsion tracker. This is the range of momenta relevant to the analysis of neutrino interactions.

3. Two exposures with pions of momenta 3 and 5 GeV/c to measure the white kink interaction length. A pion daughter of a tau produced in a ν_τ CC interaction would most likely have a momentum within this range.

The momentum of the pions was accurate to better than $\pm 1\%$ and the beam spot was wide enough to obtain a uniform particle density in the fiducial area. Beam profiles and intensities were monitored using delay wire chambers and a scintillation counter telescope.

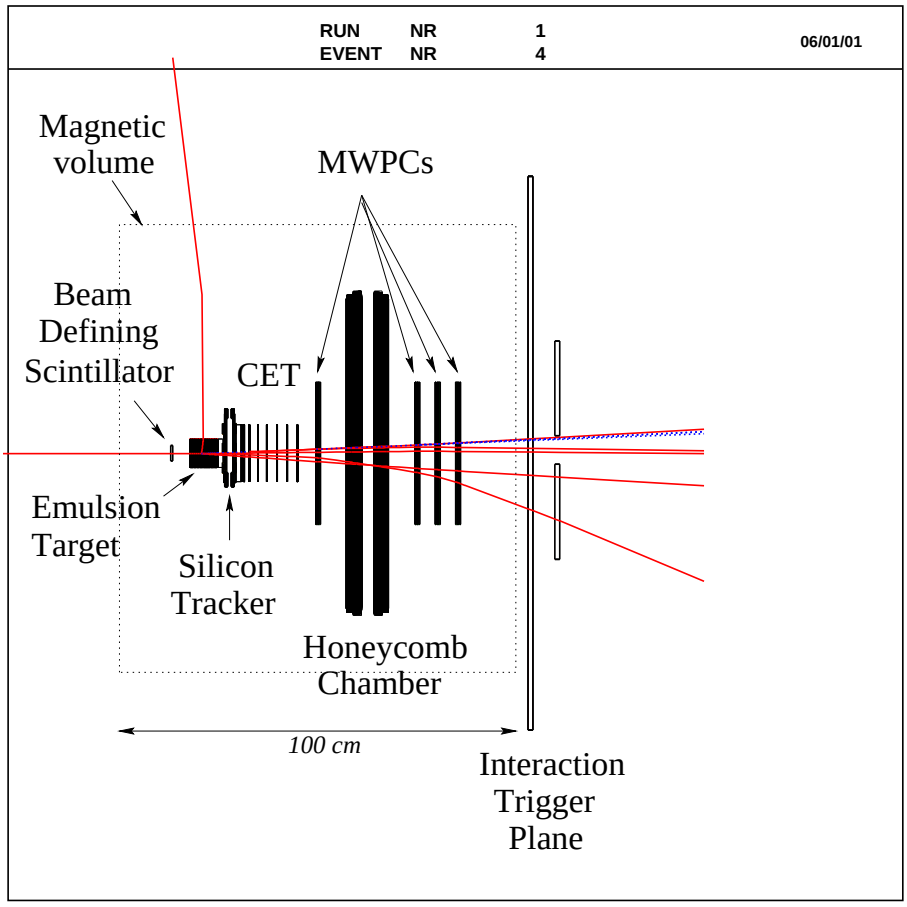

Fig. 6.1: The detector configuration of the T9 test beam (top view, bending plane) and a simulated pion interaction.

6.2.2 Detector setup

A schematic layout of the detector configuration of the T9 test beam is shown in Fig. 6.1, together with a simulated interaction of a 15 GeV/c pion in the emulsion target.

Most detectors were installed in the 50 cm gap of the poles of a C-magnet which had a size of 1×1 m^2. Its magnetic field had a strength of 0.7 T and the uniformity close to the centre, where the emulsion tracker was installed, was better than $8 \cdot 10^{-3}$.

The detectors inside the magnetic volume were the emulsion target, the silicon and emulsion trackers, the Multi-Wire Proportional Chambers (MWPC), and the Honeycomb chamber. The interaction plane of the trigger system was installed downstream the magnet.

The emulsion target consisted of 110 sheets. They were identical to the CHORUS target plates in their longitudinal structure (see 3.3) but with smaller transverse dimensions of 9×9 cm^2.

The silicon tracker consisted of 4 planes of silicon strip detectors. They were arranged in Y-X-X-Y readout configuration, where X is the axis along the magnetic field and Y the

bending coordinate. Each plane had an active surface of 18×18 cm^2. The silicon tracker is described in detail in Ref. [122].

There were four planes of MWPCs per projection. Each plane covered a surface of 41×41 cm^2 using for the readout 336 strips with a 1.2 mm spacing.

The honeycomb chamber was constructed similarly to the ones of CHORUS (3.4.2). In this case though, the wire spacing was 19 mm. There were three subsets of four honeycomb planes. The three subsets were mounted in the chamber at -5°, 0° and 5°. This configuration was driven by the space constraints of the dipole magnet used in the experiment.

The trigger system consisted of two scintillators along the beam line, a 5×5 cm^2 tile scintillator 100 mm upstream the target, and the interaction trigger plane downstream the magnet. The latter consisted of two pairs of scintillator strips arranged in such a way that a square hole was formed in the middle of the plane. The hole had adjustable dimensions (5.5×5.5 cm^2-6.5×6.5 cm^2) to let non-interacting pions pass through the plane without hitting any of the strips. The hole was centered along the X axis, with an adjustable offset of about 1 cm along the Y to take into account the beam curvature in the magnetic field of a 15 GeV/c pion beam. A pion interaction trigger was defined by the coincidence of a hit in the target defining scintillator and a hit in any of the strips comprising the downstream scintillator plane.

For the 5 and 10 GeV runs for the second phase of the experiment (measurement of the momentum resolution of the emulsion tracker) the emulsion tracker was replaced with an identical one and the emulsion target was removed.

Finally, for the white kink measurement, a new emulsion target was installed, the silicon tracker was removed and the emulsion tracker was moved to be as close as possible to the target.

6.2.3 The Compact Emulsion Tracker

The emulsion tracker which was tested in this experiment is referred to as the *Compact Emulsion Tracker* (CET). It is an alternating arrangement of emulsion sheets and lightweight spacers. Given the sub-micron precision of nuclear emulsion, the momentum of a particle can be accurately determined even with a short lever arm or a moderate magnetic field.

In a neutrino experiment like TOSCA the CET would be essential in order to measure momenta of low energy particles with large angle tracks. These are likely to fall outside the acceptance of the downstream electronic spectrometer and the CET provides the only means to save the kinematical information of such particles. Moreover, once a track is reconstructed in the CET it can be used to improve the prediction accuracy of its slope and intercept on the first emulsion sheet of the target. As it is explained in 3.3 when the prediction accuracy becomes better, the possibility of picking up background tracks when searching for the primary neutrino interaction vertex is reduced.

The design of the CET was driven by the following basic requirements:

- The working magnetic field strength should be 0.7 T.

- The device should be able to determine the sign of the charge and the momentum of particles in the range 1-20 GeV/c, with a resolution better than 15%. This is required by the kinematical analysis needed by an oscillation experiment such as TOSCA.

- Starting from the most upstream sheet, where track predictions could be available by a silicon detector, one should be able to follow down the track and find the corresponding track segments in all sheets within a single microscope view of about $200 \times 200\,\mu\text{m}^2$. This is required to avoid large background rate and extended microscope time.

- Given the fact that the track segment finding efficiency per sheet is less than 100%, the number of sheets used should be the smallest possible. A small number of sheets is desirable also to minimize the effects of multiple scattering in the momentum measurement.

A large set of possible configurations were simulated and it was checked if they satisfy the above requirements. For every candidate configuration it was assumed that the sheets had been already aligned with respect to each other, and the track segments in the most upstream plane was found. The position and slope accuracy of a measured segment was assumed to be 1 μm and 1 mrad respectively. Once a segment was found in a sheet, it was extrapolated to the immediately downstream one using the information from the upstream found segments. In the bending plane the track was assumed to follow a parabolic line of unknown curvature. In that respect the momentum measurement was continuously being updated and improved from sheet to sheet. The CET which was eventually used in the test experiment is shown in Fig. 6.2. It is composed of eight sheets and seven light plastic foam spacers placed in a honeycomb box. The sheets consist of two 100 μm emulsion layers coated on either side of an 800 μm plastic base, like the CHORUS SS/CS emulsion sheets. They were individually vacuum packed in light-tight paper. The separator and light box materials were chosen to minimize the mass density in order to reduce multiple scattering. The spacers are coated with a 127 μm fibre glass to ensure sufficient flatness, thickness uniformity and mechanical rigidity. The required flatness was better than ± 20 μm over an area of about 200×200 mm^2. The distance between successive emulsion sheets increases in the beam direction from 8 mm to 32 mm. For the purpose of the test the transverse size of the sheets is limited to 180×180 mm^2.

For the purpose of the alignment of the sheets with respect to each other, in each spacer four holes were drilled to host bi-directional X-ray guns. A concentrated liquid solution, of about 1 MBq/cm^3 activity, of ^{55}Fe was used. The high intensity of the source compensated for the short duration of the test. The guns produced circular marks with a diameter of about 1 mm. These were used during the scanning procedure to assign a coordinate system on the sheet, and an initial alignment between the sheets at the level of a few hundred microns.

The analysis of the CET data taken during the exposure without the upstream target has been presented in Ref. [123]. The momentum resolution for 5 and 10 GeV/c pions was measured $(13 \pm 1)\%$. This value is independent of the momentum, as expected, since

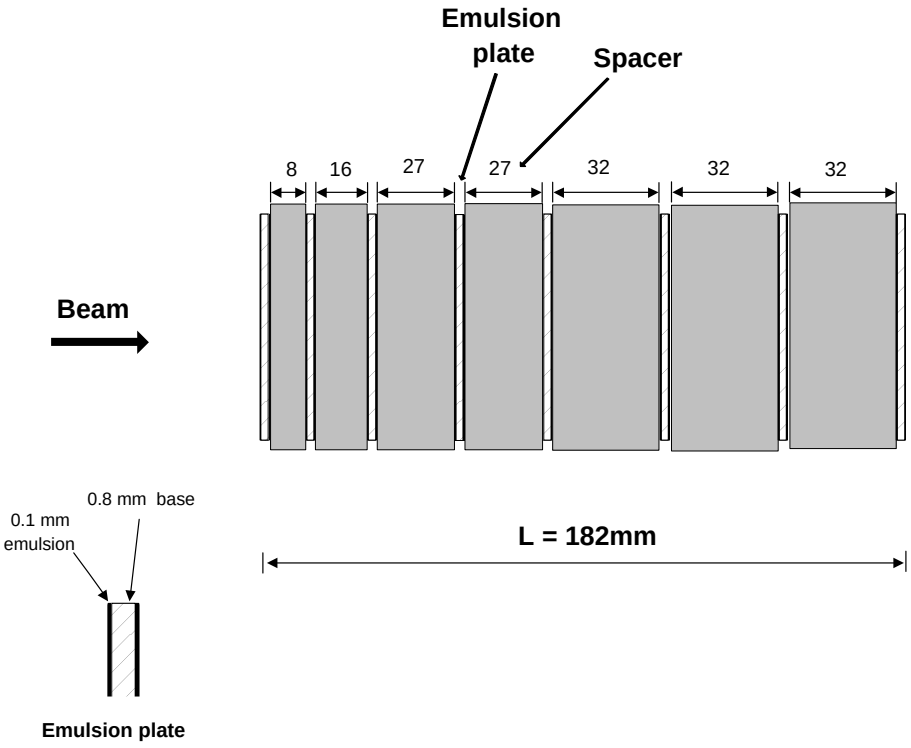

Fig. 6.2: Side view of the emulsion spectrometer used in the test. The thickness of the emulsion is not drawn to scale. Units are in mm.

the multiple scattering errors dominate. It should be noted that Monte Carlo simulations predicted $\frac{\Delta P}{P} < 15\%$ for charged particles up to 20 GeV/c.

6.3 The white kink measurement

6.3.1 Strategy

The white kink measurement was performed using exclusively emulsion detectors. A single emulsion target was used for both the 3 and 5 GeV/c exposures. On the contrary, the CET was replaced after the 3 GeV exposure by a second identical one. Moreover, during the second exposure the magnet polarity was reversed. As a consequence the non-interacting pions of the two energy sets had slopes of the opposite sign on the bending plane.

The measurement would was performed the following steps below:

1. An initial sample of tracks in the upstream SS was followed in all target sheets until the downstream SS. The number of tracks in the sample and the size of the emulsion target determines the total track length in emulsion which was probed.

2. For every particle which interacted in the emulsion target the interaction type was determined by manually examining the vertex.

3. For every interaction that was classified as white kink candidate, the daughter track was followed down to the downstream SS of the target. A white kink candidate should have a kink angle less than 400 mrad and more than 50 and 83 mrad for the 5 and 3 GeV sets respectively. The upper limit comes from the scanning acceptance covered by the scanning stations within CHORUS, while the lower limit corresponds to the 250 MeV cut in the transverse momentum for the oscillation search. The CET was used then to measure the momentum of the daughter particle. For the cases that the momentum magnitude was very close to the one of the pion beam, then the interaction would be verified as a white kink.

Detailed simulation studies [124] have shown that the overwhelming majority of the white kink interactions that could have appeared in the CHORUS target are in nature pion elastic scattering on heavy nuclei of the atoms composing the emulsion medium with low momentum transfer. In this case, the speed which is acquired by the struck nucleus is so small that it cannot deposit its energy within a range around the interaction vertex in order to render one ore more silver grains developable.

White kinks can also be inelastic interactions. Either a neutron receives most of the momentum transfer and is kicked out of the nucleus, or the initial pion is absorbed by the nucleus and a proton is kicked out. In both cases the daughter particle has a kinetic energy significantly smaller than that of the incoming pion. Moreover, the struck nucleus becomes radioactive and has a probability to decay within the duration of the exposure. The emulsion target sheets of CHORUS were left in the beam for two years and most of such nuclei would decay leaving visible signs of activity in the emulsion. On the contrary, the emulsions of the test experiment were developed only a few hours after the exposure. A possible nuclear decay would therefore not be registered with the formation of a latent image.

For the above reasons an interaction would be verified as a white kink in case the momentum of the daughter particle, as it was measured in the CET, was similar within the errors to the incoming pion beam momentum.

6.3.2 Summary of the target analysis

The analysis of the emulsion target started with the scanning of 1×1 cm^2 area in the middle of all the 112 sheets (110 TP and 2 SS) of the emulsion target. During the scanning the on-line tracking procedure was set up to search for track segments with slopes of absolute value less than 50 mrad per projection. The opposite magnet polarity in the two exposures separated completely the two sets of track segments in the angular space in the upstream SS. Therefore despite the gradually mixing of the sets in the angular space due to multiple scattering, it has been possible to assign an unambiguous momentum value to every track that was followed down to the downstream SS.

The sheets were mounted on the microscope such that the three axes of the stage and experimental coordinate systems were parallel. The z-direction was along the beam, the x-direction along the dynamic lines of the dipole magnetic field but always oriented towards the top, and the y-direction was defined to form a right-handed coordinate system.

Following tracks from the upstream SS down to the downstream SS was possible only once the 112 sheets were aligned with respect to each other. The alignment consisted in defining an optimal affine transformation between every pair of subsequent sheets. The optimum transformation was found by maximizing the number of matched track segments and minimizing the distances between the positions of the segments on the reference sheet and the projected positions of the segments on the other one. The RMS value of these distances for any sheet pair was $5 \pm 1\,\mu$m. Given this value and the mean track density of $\sim$3.5 tracks/$(100 \times 100\,\mu\text{m}^2)$, the probability of matching segments belonging to different tracks was very low.

To cross-check the validity of the measurement of the white kink, we performed a measurement of the inelastic scattering mean free path λ_{I}. Taking into account the number of tracks which have been followed, the fiducial volume in the emulsion sheets and the track finding efficiency, the effective total track length was calculated 106.6±7.5 m. We observed 299 inelastic interactions which yields $\lambda_{\text{I}} = 35 \pm 3\,$cm. Previous measurements[125] reported $\lambda_{\text{I}} = 32.4 \pm 4.0\,$cm for the same energy region.

The effective total track length for the white kink measurement was calculated 98±18 m. It is different from the corresponding value used for the inelastic scattering mean free path measurement because the detection efficiency of a white kink interaction is slightly lower. It was required that the scattering angle is larger than 30 mrad. Below this cut it was ambiguous whether an interaction took place, or whether it was a local effect from emulsion distortion.

In Table 6.3.2 there is a summary of the six white kink candidates which have been found in our track sample. For every candidate there is an estimation of the transverse

Track identifier	Parent momentum (GeV/c)	kink angle (mrad)	P_T (MeV/c)
7638	5	266±5	1331±26
7527	5	65±10	325±50
3269	3	101±5	303±13
3247	3	94±10	282±30
7811	5	199±10	995±50
9096	5	327±10	1585±57

Tab. 6.1: The white kink candidates. The transverse momentum assumes elastic scattering.

momentum under the assumption of elastic scattering.

6.3.3 Momentum measurement of the white kink daughters

Emulsion scanning and track segment reconstruction

For the downstream SS and all of the CET sheets of both sets, a square area of 4.5×4.5 cm^2 around the sheet centre was analysed with the microscopes. We used the "surface scan" method (see 5.3.7, p.109), i.e. no on-line tracking was performed but only the grain data

were retrieved and stored directly into the database. The rest of the scanning procedure was identical to a normal CS/SS scanning (see 5.3.6, p.106), with microscope views of 25 frames and the requirement that 20±1 fully cover the emulsion layer while the other are taken outside the layer marking both its surfaces.

The analysed area was divided into 49 regions, 5×5 mm^2 each. The division of the total area was done to avoid loss of the data of the whole sheet in case the on-line program abandoned the scanning of a region. A typical case where such a thing happens is when a big stain on the sheet was preventing the system of detecting the emulsion and base surfaces.

The effective view size was 315×315 μm^2 and data taking for a view took about 3 s. Thanks to this performance the total area of 20.25 cm^2 per sheet was scanned in about 21 hours.

Contrary to normal CS/SS scanning, the segments were stored not in the nominal frame defined by the X-ray marks but in the stage coordinate system. The X-ray marks were measured only to define the central position on the sheet and to reset the horizontal coordinates of the microscope stage to (0,0) at this point before the starting of the actual scanning. Due to the fact that the exposure was rather short and high intensity X-ray sources had to be used, the marks did not have a well defined round shape, and their centre could not be measured with a precision of better that 1 mm. If a microscope-to-nominal affine transformation would be used to define the nominal reference frame, fake rotations and shrinkage factors could be introduced because of the large errors in the mark measurement. Such effect would make the alignment of the sheets more complicated and the momentum measurement of the tracks in the CET less precise.

After the sheets had been scanned, the stored grain maps were processed for track segment finding. Assuming a nominal emulsion shrinkage factor of 2.0 the off-line tracking program processed every microscope view for the reconstruction of all track segments with slopes up to 300 mrad per projection. Processing a view was taking on average 0.5 s on a 600 MHz Pentium III processor.

Matching of the emulsion segments on a CET sheet

For every sheet the two emulsion sides were scanned independently and therefore the off-line tracking program produced two sets of track segments, one per emulsion side. The segment parameters that were stored in the database contained the following information:

- the number of digital images inside the emulsion layer taken in the microscope view where the segment was found.
- the number of images with an identified grain for this segment.
- the positions of the first (closest to the plastic base) and last (closest to the external surface) grain of the segment in the real 3-dimensional cartesian coordinate system of the microscope stage.
- the positions of the two boundaries of the emulsion layer.

The positions of the emulsion boundaries were used to reconstruct the sag of the microscope's glass plate onto which the emulsion sheet was mounted. The scanned rectangular area was always very small with respect to the overall area of the glass plate. The shape of the sheet mounted on the microscope could therefore be approximated very well with a parabolic surface, where every cartesian point (x, y, z) satisfies an equation

$$z = Ax^2 + Bx + Cy^2 + Dy + E \tag{6.1}$$

where A, B, C, D and E are the surface parameters.

The intercept points of the track segments onto the surface of the closest plastic base were used to fit the parameters of the two base surfaces. From the latter an averaged surface was calculated corresponding to the middle plane of the sheet when this was mounted on the microscope. The parameters of the surfaces have been used to transform the slope of any segment defined in the sheet from the reference frame of the microscope stage to that of a flat sheet. One could observe differences in the slope of the track segments before and after the correction of up to 2.5 mrad.

In order to increase the resolution and remove possible biases and systematic effects in the angular measurement, a *base measurement* (see 4.3.1 and 5.3.6) had to be performed for every track from the two corresponding emulsion segments. This has been achieved with an iterative procedure which aimed at the alignment of the two emulsion layers and the calibration of the slopes of the emulsion tracks.

Each iteration included the following steps:

1. Every segment in both emulsion sides has its slope undergo a linear transformation:

$$\overrightarrow{\theta}' = f_s \overrightarrow{\theta} + \overrightarrow{\theta_c} \tag{6.2}$$

   where $\overrightarrow{\theta} = (\frac{\partial x}{\partial z}, \frac{\partial y}{\partial z})$, f_s is a shrinkage correction factor and $\overrightarrow{\theta_c}$ is a constant offset. For each emulsion side a different transformation is used. Initially the values for f_s and $\overrightarrow{\theta_c}$ are such that they define a unity transformation.

2. Every segment in the upstream side was extrapolated to the downstream side. It was checked whether a segment existed in the downstream side within a position and angular acceptance. The acceptance cut was set at three standard deviations for each of the two space and two angular projections. For each projection the value of the standard deviation was entirely defined by the error of the slope measurement in the emulsion layers. This was given by the quadratic sum of a constant term and another which is proportional to the magnitude of the slope:

$$\sigma_\theta = \sigma_c \oplus \sigma_p |\overrightarrow{\theta}| \tag{6.3}$$

   The initial values for σ_c and σ_p are set to 5 mrad and 0 respectively.

   A pair of segments matched this way are used to perform a base measurement of the track slope and to define a *segment on the sheet* of which the position is defined as

the intercept of the track with the middle plane of the plastic base. It may occur that several candidate segments in the downstream emulsion side survive the acceptance cut. The segment which is kept is the one with the minimum slope difference from the corresponding base measurement.

The same procedure is repeated extrapolating the segments in the downstream emulsion side to the upstream one. Only the matched pairs which have been found with the previous procedure as well are eventually kept for further processing. Requiring this coincidence in the pair forming suppresses the background from random combinations.

3. From the collection of segments on the sheet as they are defined above, the emulsion slopes are calibrated to match the base measurement, fitting the constants of the transformation (6.2).

4. The parameters of (6.3) are re-calculated: The angular spectrum of the track segments as it is obtained by the base measurement is divided into seven regions of equal integrals. For each region a gaussian fit is performed to all three distributions of slope differences between the two emulsion slopes and the slope from the base measurement. From the fit results a mean resolution can be calculated for the base and each of the emulsion slopes valid for the current region. The parameters of (6.3) are therefore fitted using seven points for each of the three slopes.

The final parameters of (6.2) and (6.3) are stored in the database for later usage. Together with them the final set of *sheet segments* was stored with all the calculated segment parameters from the base measurement. It is this set of segments that was used from each sheet throughout the rest of the analysis. On average 60000 segments where found per CET sheet corresponding to about 38 tracks/mm^2 and 4 tracks per microscope view.

The above procedure has improved the angular definition of the segments as it can be seen from the graphs in Fig. 6.3, where the differences in the slope values between the emulsion and the base measurement are plotted. Plotting these differences as function of the absolute slope, reveals the dependency of the resolution in the slope measurement from the magnitude of the slope.

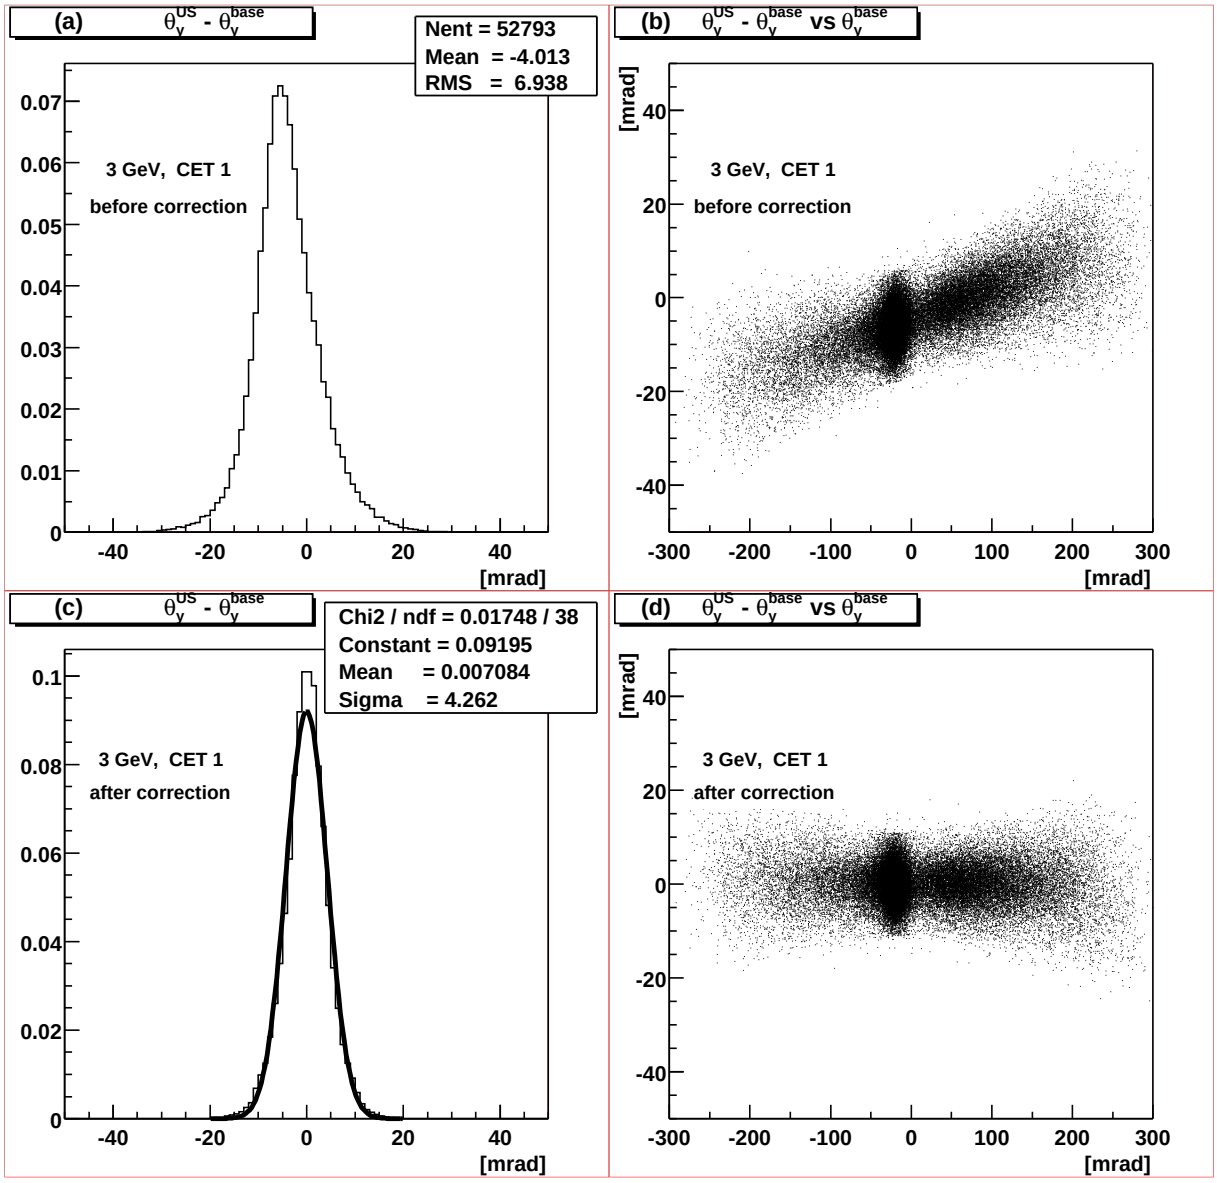

Fig. 6.3: Slope differences (y-projection) for matched segments in the two emulsion sides. Data are shown for the most upstream CET sheet from the 3 GeV exposure.
(a) The difference in the slope measured in the upstream emulsion layer and the slope from base measurement.
(b) The same as (a), as a function of the slope from the base measurement.
(c), (d) The corresponding plots after the slope calibration.
The increased population around 20 mrad corresponds to the beam pions which did not interact in the target.

	3 GeV exposure		5 GeV exposure	
Sheet No	σ_c (rad)	σ_p	σ_c (rad)	σ_p
1	0.0011	0.012	0.0010	0.014
2	0.0009	0.008	0.0011	0.013
3	0.0010	0.010	0.0010	0.012
4	0.0013	0.012	0.0012	0.013
5	0.0009	0.010	0.0008	0.013
6	0.0010	0.013	0.0011	0.011
7	0.0011	0.009	0.0012	0.008
8	0.0012	0.012	0.0009	0.015

Tab. 6.2: The slope resolution constants for the CET sheets.

The calculated constants of (6.3) that determine the slope resolution from the base measurement are presented in Table 6.2. Under the assumption that the thickness of the plastic base is 800 μm it can be computed from the calculated values for σ_c that the position of an emulsion grain is measured with a precision of $\sigma_h = (0.42 \pm 0.03)\mu$m per projection on the horizontal plane, The existence of σ_p is because of the total error in the measurement of the position of the base surfaces along the microscope optical axis. The contributions to this error are the effective depth of focus and the precision of the position measurement using the encoders of the microscope stage while it is moving. The calculated values for σ_p indicate an error of $\sigma_z = (6.5 \pm 1.3)\mu$m.

Corrections applied to the segments of each CET sheet

As mentioned the parameters of the parabolic surface defined in (6.1) are used to transform the position and the slope of each track segment from the coordinate system of the microscope stage to the system where the sheet is flat:

$$\begin{pmatrix} x \\ y \end{pmatrix} \mapsto \begin{pmatrix} \int_0^x \frac{1}{cos(atan(2At+C))}\, dt \\ \int_0^y \frac{1}{cos(atan(2Bt+D))}\, dt \end{pmatrix} \tag{6.4}$$

and

$$\begin{pmatrix} \theta_x \\ \theta_y \end{pmatrix} \mapsto \begin{pmatrix} \theta_x + 2Ax + C \\ \theta_y + 2By + D \end{pmatrix} \tag{6.5}$$

where the integrals in (6.4) were calculated numerically on a grid over a range in (x, y) which covers the scanned area. The nodes of the grid were spaced by 50 μm. For the rest of the (x, y) points the corresponding values were calculated using interpolating linearly the values from the neighbouring nodes.

The segments from more than one CET sheet can be used for any combined analysis only if they are transformed to the reference frame of the experimental setup. During the

exposure, there has been an effort to have the emulsion sheets as flat as possible with the use of light spacers (see 6.2.3). However, it appears from the data that the sheets were not perfectly flat. If the sheets were perfectly flat no correlation would be observed between the measured slopes and positions for the track segments. The scatter plot in Fig. 6.4a shows the measured slope versus the measured position along x-axis for the track segments of the SS sheet. The largely populated central band corresponds to the pions which did not interact in the emulsion target. While their angular spread remains constant, the mean value of their slope changes by about 80 μm from the one extreme of the scanned region to the other. Given the fact that the beam was uniformly distributed over the scanned area, the dependence between the slopes and positions indicates a deviation of the sheet's shape from flatness.

The deformation of the sheets was corrected by assuming that during the exposure they had a shape approximated by surface with equation:

$$z = Ax^4 + Bx^3 + Cx^2 + Dx + Ey^4 + Fy^3 + Gy^2 + Hy \tag{6.6}$$

where the constants describing the surface were calculated by fitting the distributions $\theta_x = f(x)$ (as the one shown in Fig. 6.4a) and $\theta_y = f(y)$ to fourth-order polynomials.

The surface (6.6) is used to define a transformation which transforms segments from the reference system of the deformed sheet to the system where the sheet is flat:

$$\begin{pmatrix} x \\ y \end{pmatrix} \mapsto \begin{pmatrix} \int_0^x \frac{1}{cos(atan(4At^3+3Bt^2+Ct+D))}\,dt \\ \int_0^y \frac{1}{cos(atan(4Et^3+3Ft^2+2Gt+H))}\,dt \end{pmatrix} \tag{6.7}$$

and

$$\begin{pmatrix} \theta_x \\ \theta_y \end{pmatrix} \mapsto \begin{pmatrix} \theta_x + 4Ax^3 + 3Bx^2 + 2Cx + D \\ \theta_y + 4Ey^3 + 3Fy^2 + 2Gy + H \end{pmatrix} \tag{6.8}$$

where the integrals in (6.7) were calculated following the same procedure for the calculation of the ones of (6.4).

The segments of every sheet undergo first the transformations (6.4) and (6.5) and then the inverse of (6.7) and (6.8). With this procedure one obtains all track segments in the reference frame of the sheet with the shape that it had in the experimental setup.

The plots presented in Fig 6.4 show the slope distributions before and after the corrections for the non-bending plane of the SS sheet, which had the most severe deformation. The central peak in the angular distribution of Fig 6.4d corresponds to the pions which did not interact in the emulsion target and has a width of about 8 mrad, which can be fully attributed to the multiple scattering in the target.

Relative alignment of the CET sheets

Reconstruction and track fitting in the CET is possible only after transforming the segments of all sheets into a common reference system. This is done by applying the proper

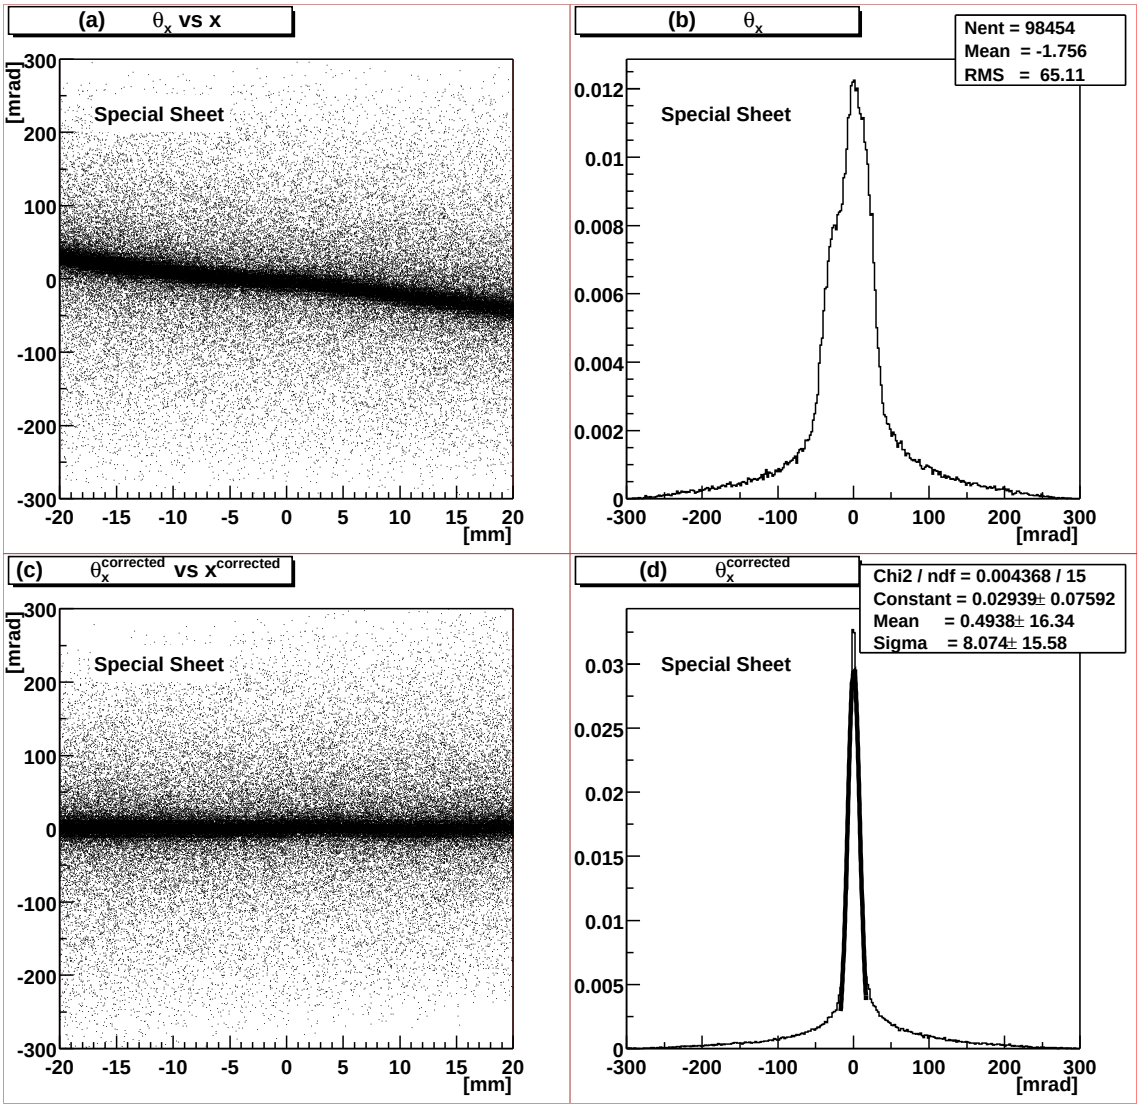

Fig. 6.4: Slope distribution in the non bending coordinate of the SS.
(a) Slope versus position before correcting for the sheet deformation.
(b) Slope distribution before correcting for the sheet deformation.
(c), (d) The corresponding plots after correcting for the sheet deformation.

transformation to every sheet and its segments, after they have been corrected with the transformations discussed above.

The first step in the alignment/calibration procedure is to define the form of the transformation. After trying out several forms, it turned out that the transformation with the smallest number of free parameters giving satisfactory results includes the following operations on a sheet and its segments.

1. The sheet is displaced longitudinally (along the z-axis) by dz with respect to its

nominal position.

2. A tilt correction is applied to the orientation of the sheet to accommodate the difference of its plane from normality to the z-axis. This is expressed as a slope offset in both projections applied to every segment on the sheet:

$$(\theta_x, \theta_y) \mapsto (\theta_x, \theta_y) + (\theta_x^t, \theta_y^t) \tag{6.9}$$

where (θ_x^t, θ_y^t) is the tilt of the sheet.

3. A general affine transform is applied to the positions and the slopes of the segments on the sheet:

$$\begin{pmatrix} x \\ y \end{pmatrix} \mapsto \begin{pmatrix} R_{xx} & R_{xy} \\ R_{yx} & R_{yy} \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} + \begin{pmatrix} V_x \\ V_y \end{pmatrix} \tag{6.10}$$

and

$$\begin{pmatrix} \theta_x \\ \theta_y \end{pmatrix} \mapsto \begin{pmatrix} R_{xx} & R_{xy} \\ R_{yx} & R_{yy} \end{pmatrix} \begin{pmatrix} \theta_x \\ \theta_y \end{pmatrix} \tag{6.11}$$

where the matrix (R_{ij}) is not necessarily unitary. This operation accommodates rotations around the z-axis, shrinkages and offsets in the XY plane.

It can be shown that transformations of this kind form a mathematical group. The parameters of the unity transformation are all set to zero except $R_{xx} = R_{yy} = 1$. One can start the alignment procedure by calculating the relative transformations between neighbouring sheets and derive eventually for each sheet a global transformation with respect to some reference. For each of the two exposures this reference was chosen to be the most upstream CET sheet. This choice was made because this sheet was pressed against the honeycomb side of the CET box, which was mechanically aligned with the target box. It has therefore the highest probability of having been normal to the beam direction (z-axis) during the exposure. Moreover, it is the one closest to SS. Choosing the most upstream CET sheet to be the reference plane minimizes the errors in the position and the slope of an extrapolated track reconstructed in the CET onto the SS plane.

For every pair of neighbouring sheets the relative alignment procedure would calculate the parameters of a transformation applied to the downstream sheet with the two following properties: the maximum number of pairs of segments appearing as parts of the same track; the minimum total position and angular difference between matched track segments, after the downstream one is extrapolated to the XY plane of the upstream sheet.

As mentioned before, the majority of the tracks belong to pions which did not interact in the emulsion target. Since these were beam particles, they were concentrated in a relatively small region of the angular space as it can be shown from the plots in Fig. 6.5. Given also the fact that in the CET sheets the segments of only a single exposure were present, there were relatively few constraints for the nine parameters of the transformation, especially the longitudinal offset dz. In order to compensate for this, the knowledge of the

momentum of the tracks was used when extrapolating the segments from a downstream sheet to the upstream one. These tracks were selected by identifying the peaks in the angular distributions of both projections, and by cutting at three standard deviations from the centre of each peak. This cut was efficient in separating the segments from the two exposures in the SS, as it can be seen from Fig. 6.5a.

The relative alignment was an iterative procedure of the following steps:

1. The segments corresponding to non-interacting pions were selected from both sheets. The current relative transformation, initially the unity transformation, was applied to the downstream sheet and its segments.

2. The tilt of the downstream sheet was calculated. For the first iteration this is done by comparing the centre of the peak in the angular distribution of the downstream segments with the ones which are obtained by extrapolating the upstream segments to the downstream sheet assuming the nominal beam momentum. The tilt correction is formed by the differences in the peak centres. For the subsequent iterations one obtains the matched segment pairs and calculates the tilt correction by the centres of the distributions

$$\theta_x^{US} - \theta_x^{DS,extr} \tag{6.12}$$

and

$$\theta_y^{US} - \theta_y^{DS,extr} \tag{6.13}$$

where $\theta_{x/y}^{DS,extr}$ is the slope of the extrapolated downstream segment to the upstream sheet. No tilt correction is applied if these distributions are centered to zero, which is the case in the plots shown in Fig. 6.6.

The pairs of matched segments are obtained by extrapolating each downstream segment to the upstream sheet together with its angular errors. The latter are calculated from the sheet's resolution constants which appear in Table 6.2. Segments from the upstream sheet are searched in the neighbourhood of the four-dimensional extrapolated position. The boundaries of the acceptance neighbourhood are defined by cutting at three standard deviations from the central extrapolated value. In case more than one candidate is found, the one with the smallest difference in position is selected.

The procedure that looks for matching segments calculates at the same time a quality value of the current alignment. Every extrapolated downstream segment contributes to a sum by the squared difference of its extrapolated position from the position of the matched segment on the upstream sheet. In case no match was found it is checked whether any upstream segment exists in an extended square area of $500 \times 500 \mu\text{m}^2$ around the extrapolated position. If no segment is found this means that the search was done in a non-overlapping or dead (from the scanning point of view) area. In

this case no contribution to the sum is made. Otherwise the sum is incremented by the square of the diagonal of the extended square area. The quality value is formed by the square of the number of matching tracks divided by the sum.

3. The optimum affine transformation is found. First the parameters V_x and V_y which define the transverse offset are estimated. This is done by trying out randomly various (x, y) points over a given area, selecting at the end the point which gives the highest quality value. At the first iteration a coarse search is performed over an area of $1500 \times 1500 \mu\text{m}^2$ which is followed by a fine search over an area of $100 \times 100 \mu\text{m}^2$ around the current best point. The number of random trials for the coarse search is such that the mean distance between tried points is less than 100μm while for the fine search the mean distance is less than 10μm. In the subsequent iterations, only the fine search is performed. From the matched segment pairs which correspond to the best values for V_x and V_y found with this procedure, a subset of them is selected by cutting at three standard deviations from the centre of the peak in the distributions (6.12), (6.13),

$$x^{US} - x^{DS,extr} \tag{6.14}$$

and

$$y^{US} - y^{DS,extr} \tag{6.15}$$

to be used for the calculation of the optimum affine correction. This is done by fitting the affine transformation which maps the positions of the downstream segments to the positions of the extrapolated upstream segments.

4. The longitudinal offset dz is calculated by applying a linear fit in the scatter plot $x^{US} - x^{DS,extr}$ vs θ_x^{US} for the matched segment pairs identified at the previous step. The offset dz has the value of the fitted slope.

5. The current transformation is updated with the calculated corrections.

The relative alignment is evaluated by inspecting the distributions (6.12), (6.13), (6.14) and (6.15). Each of them should have the shape of a normal distribution with a spread consistent with the assumed slope and position measurement errors.

The angular differences shown in Fig. 6.6 are consistent with the error assumptions as they are derived from the resolutions in Table 6.2. Comparing the spread of the angular and position difference distributions, and taking into account the distance between the matched sheets in each case, it can be deduced that the position differences are dominated by the error in the slope measurement.

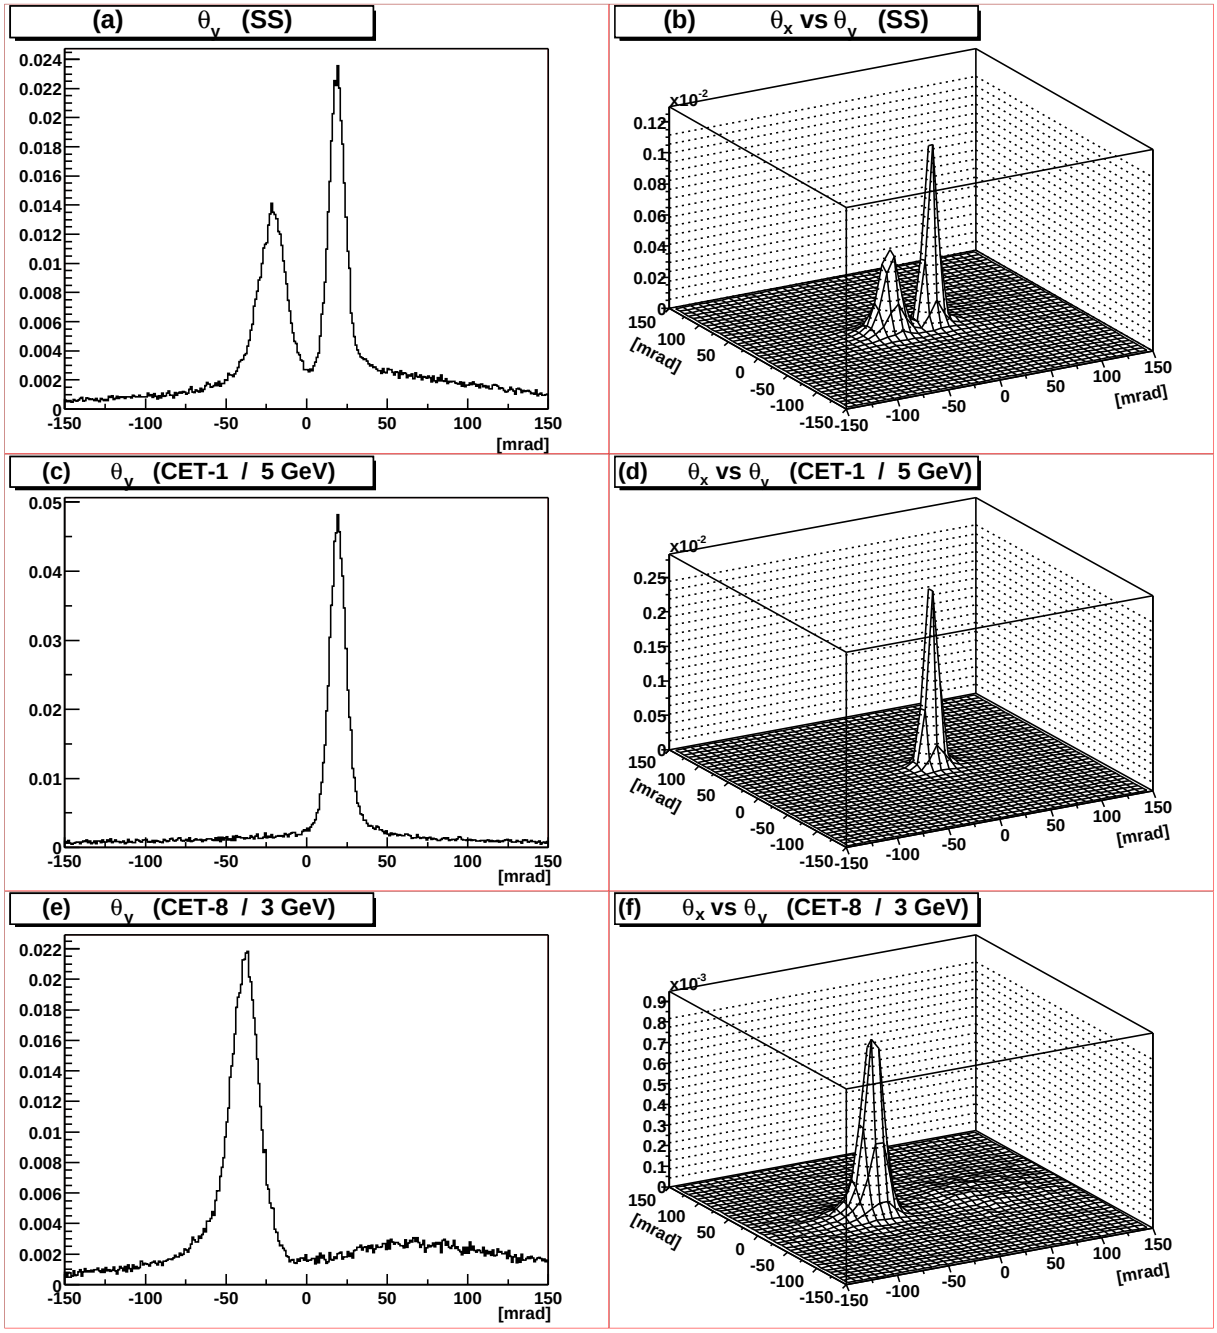

Fig. 6.5: Angular distribution of segments on the CET sheets.
(a) The distribution of θ_y on the SS. The peak centered at the negative values correspond to the 3 GeV exposure, and the other to the 5 GeV exposure.
(b) The same as (a) in both projections (θ_x, θ_y).
(c), (d) As (a) and (b) for the most upstream CET sheet of the 5 GeV exposure.
(e), (f) As (a) and (b) for the most upstream CET sheet of the 3 GeV exposure.

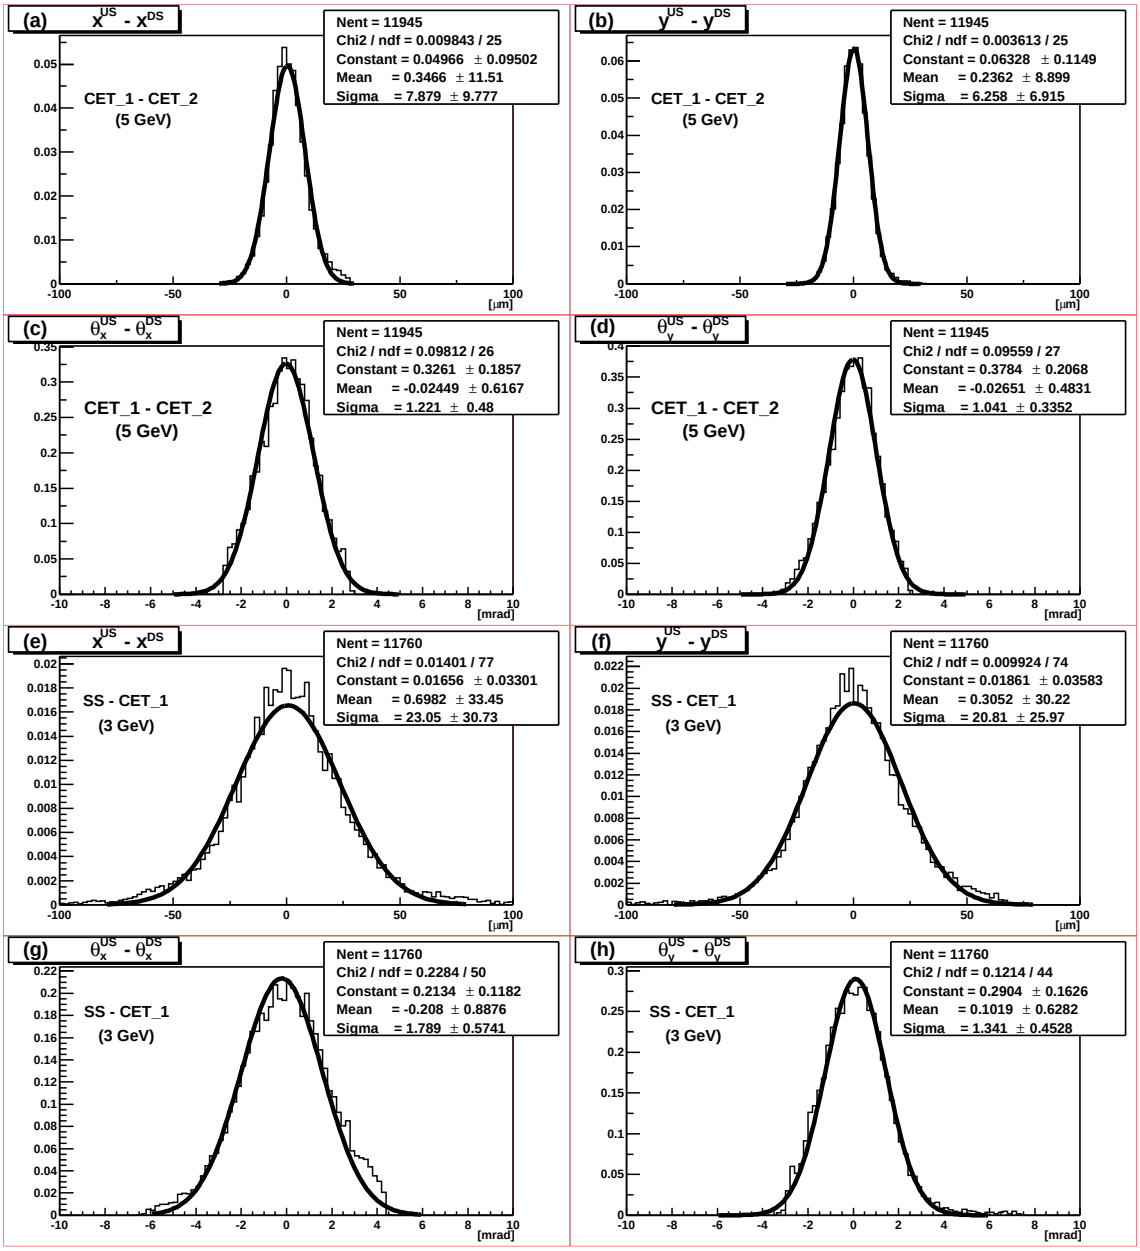

Fig. 6.6: Sheet to sheet matching residuals.
(a) The position difference in the x-coordinate between the upstream segment and the downstream segment extrapolated to the upstream sheet after applying on the downstream sheet the relative transformation. The distribution is the result of the matching of the two most upstream CET sheets of the 5 GeV exposure.
(b) The same as (a) for the y-coordinate.
(c), (d) show the corresponding slope differences.
(e), (f), (g) and (h) are the corresponding plots for the matching of the SS and the most upstream CET sheet for the 3 GeV exposure.

Track fitting in the CET

Given the group properties of the transformations which are applied to the sheets one can deduce the global transformations of all sheets with respect to the most upstream one of each CET. With this procedure the segments from all sheets are transformed to a common reference frame whose origin is on the reference sheet. Tracks can be formed by combining segments from different CET sheets. Such tracks can be fitted to a helix in order to deduce their parameters.

For the purposes of global track fitting in the CET with fixed or unknown momentum for a variable number of identified segments in the eight CET sheets, a machinery was developed based on the Least Squares Method (LSM). The method is extensively described in Ref. [126]. Here the main points are summarized and the problem specific implementation is discussed.

In the most general case the track of a charged particle in a magnetic field can be represented as a point $\mathbf{p}$ in a five-dimensional hyperplane. The track can also be represented in the measurement space through a transformation $\mathbf{f}(\mathbf{p})$ by solving the equations of motion. The LSM works on a linear expansion of the equations $\mathbf{f}(\mathbf{p})$ around a point $\overset{0}{\mathbf{p}}$:

$$\mathbf{f}(\mathbf{p}) = \mathbf{f}(\overset{0}{\mathbf{p}}) + \mathbf{A} \cdot (\mathbf{p} - \overset{0}{\mathbf{p}}) + O\left((\mathbf{p} - \overset{0}{\mathbf{p}})^2\right) \tag{6.16}$$

where $A_{ij} = \partial f_i / \partial p_j$ at $\mathbf{p} = \overset{0}{\mathbf{p}}$. Assuming that the measurement errors are gaussian and constant in the neighbourhood of $\overset{0}{\mathbf{p}}$, if the covariance matrix $\mathbf{V}$ of the measurement vector is known, then one can obtain an estimation of $\mathbf{p}$ by minimizing the function

$$\chi^2 = \left[\mathbf{f}(\overset{0}{\mathbf{p}}) + \mathbf{A} \cdot (\mathbf{p} - \overset{0}{\mathbf{p}}) - \mathbf{m}\right]^T \mathbf{W} \left[\mathbf{f}(\overset{0}{\mathbf{p}}) + \mathbf{A} \cdot (\mathbf{p} - \overset{0}{\mathbf{p}}) - \mathbf{m}\right] \tag{6.17}$$

where $\mathbf{W} = \mathbf{V}^{-1}$ and $\mathbf{m}$ is the actual measurement vector. The estimated parameter vector of the track can be calculated analytically as

$$\tilde{\mathbf{p}} = \overset{0}{\mathbf{p}} + \left(\mathbf{A^T W A}\right)^{-1} \mathbf{A^T W} \left[\mathbf{m} - \mathbf{f}(\overset{0}{\mathbf{p}})\right] \tag{6.18}$$

with a covariance matrix

$$\mathbf{V}_{\tilde{\mathbf{p}}} = \mathbf{A^T V A}. \tag{6.19}$$

In summary, for the purposes of the track fitting one needs the solutions of the equations of motion $\mathbf{f}$ as a function of the track parameters, their first partial derivatives, which would allow the calculation of the matrix $\mathbf{A}$, and the covariance matrix $\mathbf{V}$.

To obtain the equations $\mathbf{f}(\mathbf{p})$ one starts from the general equation of motion for a charged particle in a magnetic field, neglecting any energy loss from material effects or bremsstrahlung:

$$\frac{d^2\vec{r}}{ds^2} = \frac{\kappa\, q}{P} \frac{d\vec{r}}{ds} \times \vec{B}(\vec{r}(s)) \tag{6.20}$$

where $\vec{r}$ is the position of the particle, s is the path length along its track, P its momentum and $\vec{B}$ the magnetic field strength. If the position is given in metres, the momentum in GeV/c and the magnetic field in Tesla then $\kappa = 0.2999792458\,\mathrm{T}^{-1}\cdot\mathrm{m}^{-1}\cdot\mathrm{GeV/c}$.

Solving the equation in the cartesian coordinate system where the magnetic field is uniform and parallel to the x-axis one obtains:

$$x = x_0 + \sin\lambda\,\frac{P}{\kappa\,q\,B}\,\left[\arccos\left(\frac{z}{\cos\lambda}\frac{\kappa\,q\,B}{P} + \cos\phi_0\right) - \phi_0\right] \tag{6.21}$$

and

$$y = y_0 + \frac{P\,\cos\lambda}{\kappa\,q\,B}\,\left[\sqrt{1 - \left(\frac{z}{\cos\lambda}\frac{\kappa\,q\,B}{P} + \cos\phi_0\right)^2} - \sin\phi_0\right] \tag{6.22}$$

where $(x_0,\, y_0)$ is the position of the track on the reference plane at $z = 0$, λ is the slope angle, and ϕ_0 is the azimuth angle of the point $(x_0,\, y_0,\, 0)$ with respect to the helix axis.

Expanding into a Taylor series for z, approximating the particle trajectory on the $X-Z$ and $Y-Z$ plane with a straight and parabolic line respectively, one obtains:

$$x \approx x_0 - \frac{\tan\lambda}{\sin\phi_0}\cdot z \tag{6.23}$$

and

$$y \approx y_0 - \frac{1}{\tan\phi_0}\cdot z - \frac{1}{P_b}\cdot\frac{\mathrm{sign}(\phi_0)}{2\,\cos\lambda\,\sin^3\phi_0}\cdot z^2 \tag{6.24}$$

where $P_b = P/(\kappa q B)$. These are actually the equations $\mathbf{f}(\mathbf{p})$ with

$$\mathbf{p} = \begin{pmatrix} x_0 \\ y_0 \\ \lambda \\ \phi_0 \\ P_b \end{pmatrix}. \tag{6.25}$$

From (6.23) and (6.23) one can calculate analytically the partial derivatives $\partial\mathbf{f}(\mathbf{p})/\partial\mathbf{p}$ and therefore the matrix $\mathbf{A}$ can be calculated for every point $\overset{0}{\mathbf{p}}$.

In the case of the CET, every sheet with a reconstructed segment contributes to the measurement by a (x, y) position and a (θ_x, θ_y) slope at a z-coordinate corresponding to the middle plane of the sheet. From these quantities one can construct two (x, y) positions at $z \pm 400\,\mu m$. Therefore the number of measurements for a track in the CET is $4 \cdot n$, where n is the number of sheets with reconstructed segment, which makes the dimension of the matrix $\mathbf{A}$ equal to $4 \cdot n \times 5$.

This track fitting machinery can be modified so that the momentum is a constant. In this case the track parameter vector is four-dimensional:

$$\mathbf{p} = \begin{pmatrix} x_0 \\ y_0 \\ \lambda \\ \phi_0 \end{pmatrix} \tag{6.26}$$

and the dimension of $\mathbf{A}$ becomes $4 \cdot n \times 4$.

The actual fit is performed in two steps. In the first step a pre-fit is performed by fitting independently the $X - Z$ and $Y - Z$ projections, without taking strictly into account the errors, in order to determine the initial parameter vector $\overset{0}{\mathbf{p}}$. In the second step the matrix $\mathbf{A}$ is constructed for the calculated $\overset{0}{\mathbf{p}}$, and the result is obtained performing the matrix multiplications (6.18) and (6.19).

The covariance matrix is constructed as a sum of two terms:

$$\mathbf{V} = \mathbf{V}_e + \mathbf{V}_{ms} \tag{6.27}$$

where $\mathbf{V}_e$ is the contribution from the errors of the position and slope measurement, and $\mathbf{V}_{ms}$ is the contribution from the multiple scattering using the pre-fitted momentum for the calculation of the mean scattering angle.

The covariance matrix $\mathbf{V}_e$ has the form

$$\mathbf{V}_e = \begin{pmatrix} \mathbf{V}_{s,1} & \mathbf{0} & . & . & . & \mathbf{0} \\ \mathbf{0} & \mathbf{V}_{s,2} & . & . & . & \mathbf{0} \\ . & . & . & . & . & . \\ \mathbf{0} & \mathbf{0} & . & . & . & \mathbf{V}_{s,n} \end{pmatrix} \tag{6.28}$$

where $\mathbf{V}_{s,i}$ is the contribution from the sheet i, and $n \leq 8$. $\mathbf{V}_{s,i}$ is the covariance matrix of the four measurements x_u, y_u, x_d, y_d, where u and d denote the upstream and downstream emulsion side respectively. It has the form

$$\mathbf{V}_{s,i} = \begin{pmatrix} \sigma_s^2 + \sigma_{x,i}^2 & 0 & \sigma_s^2 & 0 \\ 0 & \sigma_s^2 + \sigma_{y,i}^2 & 0 & \sigma_s^2 \\ \sigma_s^2 & 0 & \sigma_s^2 + \sigma_{x,i}^2 & 0 \\ 0 & \sigma_s^2 & 0 & \sigma_s^2 + \sigma_{y,i}^2 \end{pmatrix} \tag{6.29}$$

where

- σ_s the error in the position of the sheet. It is set globally for all sheets and its exact value was determined from the alignment procedure;

- $\sigma_{x,i}$ is the total error in the position measurement along the x-coordinate of a segment in an emulsion layer. It is calculated from the errors σ_h and σ_z, which were discussed in p. 127, and the measured slope $\theta_{x,i}$ along the x-coordinate on sheet i:

$$\sigma_{x,i}^2 = \sigma_h^2 + \sigma_z^2 \cdot \theta_{x,i}^2 \cdot \frac{2}{L^2} \tag{6.30}$$

  where L=800 μm is the assumed thickness of the base;

- $\sigma_{y,i}$ is the corresponding total error along the y-coordinate.

The multiple scattering covariance matrix is constructed assuming a scattering centre at the middle plane of each sheet. Each scattering centre i located at z_i^{sc} contributes to the total error in the position measurement on a sheet j located at z_j by

$$\sigma_{i,j} = (z_j - z_i^{sc}) \cdot \theta_0 \cdot \Theta(z_j - z_i^{sc}) \tag{6.31}$$

where $\Theta(x)$ is the step function and θ_0 is the mean scattering angle on a plane which is equal to

$$\theta_0 = \frac{0.0136}{\beta\, P} \sqrt{l_{rl}}\ (1 + 0.038 \log l_{rl}) \tag{6.32}$$

where P is momentum of the particle in GeV/c and l_{rl} is the thickness of the scattering centre in radiation lengths. For normally incident particles it has been calculated that $l_{rl} \approx 0.015$. This value has been calculated using the data in Table 6.3. For the calculation of θ_0 the mean slope of the track is taken into account, as it is estimated with the pre-fit procedure, by scaling the quantity l_{rl} accordingly.

material	thickness (μm)	radiation length (cm)	density (g/cm^3)
emulsion	$8 \times 2 \times 100$	2.97	3.7
plastic base	8×800	42	1.032
fibre glass	$8 \times 2 \times 127$	10.56	2.56
spacer foam	$7 \times$ 8-32 $\cdot 10^3$	676	0.06

Tab. 6.3: Material thicknesses and properties in the CET.

Since a scattering centre affects all measurements on the sheets downstream, the covariance matrix $\mathbf{V}_{ms}$ is constructed such that it takes into account the corresponding correlations. This is done by constructing the matrix in the following form:

$$\mathbf{V}_{ms} = \begin{pmatrix} \Sigma_{11} & 0 & \Sigma_{12} & 0 & . & . & . \\ 0 & \Sigma_{11} & 0 & \Sigma_{12} & . & . & . \\ \Sigma_{21} & 0 & \Sigma_{22} & 0 & . & . & . \\ 0 & \Sigma_{21} & 0 & \Sigma_{22} & . & . & . \\ . & . & . & . & . & . & . \\ . & . & . & . & . & . & \Sigma_{2n\,2n} \end{pmatrix} \tag{6.33}$$

where n is the number of sheets with identified segment and

$$\Sigma_{ij} = \sum_k (z_i - z_k^{sc})(z_j - z_k^{sc})\, \theta_0^2\, \Theta(z_i - z_k^{sc}) \Theta(z_j - z_k^{sc}) \tag{6.34}$$

with the sum looping over the scattering centres.

The effect of multiple scattering dominates the momentum measurement since the values for the error in (6.31) can take values up to 150 μm for 3 GeV particles.

Global alignment using reference tracks

From the matched segment pairs found with the relative alignment procedure one can identify reference tracks with identified segments on all of the CET sheets. These tracks were fitted to a helix with fixed curvature (momentum) and they were used for a fine tuning of the alignment.

The momentum which was assumed was that of the beam subtracted by 50 MeV, which is the mean energy loss due to ionization in the emulsion target. For the alignment purposes the multiple scattering process is not treated, because the latter assumes a reference plane where the errors are minimized. This is not valid for the alignment problem. An alignment process with the multiple scattering treatment enabled would not handle correctly displacements of the most downstream sheets.

The alignment procedure is an iterative process consisting of the following steps:

1. The current global transformations were applied to the sheets and their segments. Initially, these are the ones constructed from the sheet-to-sheet relative transformations.

2. The reference tracks are identified. For each pair of neighbouring sheets the relative transformation is calculated and the corresponding matched segment pairs are identified. From the combination of the pairs one can identify the reference tracks as those which have segments on all sheets.

3. The reference tracks are fitted with the fixed momentum.

4. In the sample of the reference tracks various cuts are applied to identify a set of tracks with reduced background coming from random segment combinations. A three-sigma cut is used to isolate the peaks in the momentum distribution and the residual distributions $x^{measured} - x^{fitted}$, $y^{measured} - y^{fitted}$, $\theta_x^{measured} - \theta_x^{fitted}$ and $\theta_y^{measured} - \theta_y^{fitted}$ for each CET sheet.

5. For the reduced track sample, the tilt was corrected comparing the centres of the peaks in the distributions $\theta_x^{measured} - \theta_x^{fitted}$ and $\theta_y^{measured} - \theta_y^{fitted}$ for each sheet with respect to corresponding values for the reference sheet.

6. The affine transformation was corrected by comparing the fitted positions of the tracks on each sheet with respect to the measured ones. Similarly to the tilt correction, a group subtraction operation is performed to the resulting relative transformation correction by the corresponding affine transformation correction for the reference sheet.

7. The longitudinal offset correction is calculated identically to the case of the relative alignment.

8. The transformations are updated for the next iteration.

Once the CET sheets have been aligned, the multiple scattering covariance matrix is taken into account and the reference tracks are fitted without the momentum constraint. The distribution of the reconstructed momentum of the reference tracks is shown in Fig. 6.7. The calculated momentum resolution is $\frac{\Delta P}{P} = (12.9 \pm 0.4)\%$, in agreement with what is quoted in Ref. [123] (see 6.2.3). In Fig. 6.7 there is also the distribution of the upper tail χ^2 probability, which at the end of the alignment procedure should ideally be uniform. In our case there are also small peaks towards the minimum and and maximum values. This is due to the fact that the error parameters are used globally. On the other hand the achieved result is already satisfactory and a further complication of the alignment and track fitting problems would not improve considerably the momentum measurement.

The error on the sheet position σ_s, which was discussed in the previous subsection, was determined from the variance of the difference $x_1^{measured} - x_1^{fitted}$ where $x_1^{measured}$ is the measured position along the x-coordinate on the most upstream CET sheet, and x_1^{fitted} the intercept of the track with this sheet when it is fitted taking into account only the seven downstream sheets. Several iterations of the alignment procedure were needed in order to converge to the value of $\sigma_s = 3.5\,\mu$m which was used also in the subsequent steps.

The SS sheet was aligned separately from the rest of the CET sheets because its mechanical connection to the CET was worse than the CET sheets among themselves, it had the most severe systematics to cope with, and being exposed to both 3 and 5 GeV pions, it contained two times more segments. The relative alignment with respect to the most upstream CET was used as a starting point for determining its global transformation and its fine tuning was performed by extrapolating the fitted reference tracks to the SS. The alignment procedure was the same as for the rest of the CET sheets with the difference that the fitted tracks did not take into account the SS measurement.

The residuals between the extrapolated positions and slopes from the CET tracks and the measured positions and slopes on the SS are shown in Fig. 6.8. The fact that for the 5 GeV set there are two times more matched segments is due to the fact that one of the sheets in the CET of the 3 GeV exposure was by accident half-scanned, combined to the requirement for a reference track to have reconstructed segments in all eight CET sheets.

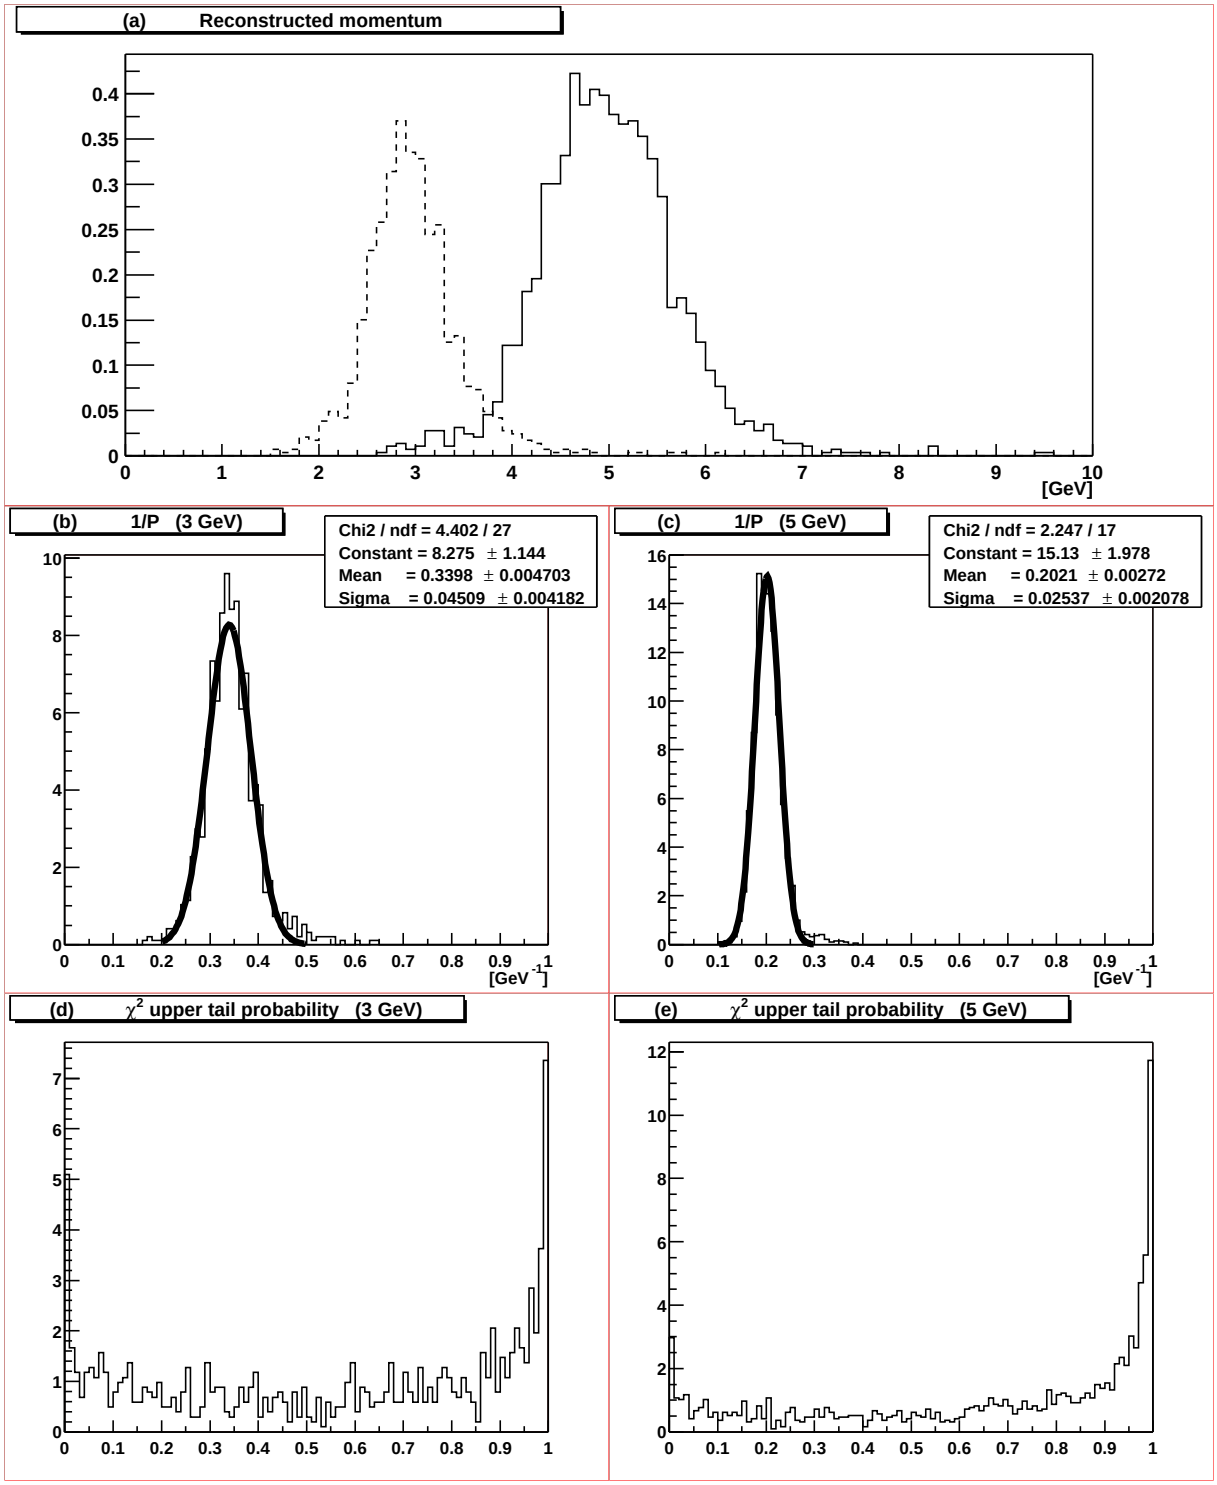

Fig. 6.7: Momentum of the reference tracks in the CET.
(a) The momentum distribution from both exposures.
(b) The fitted $1/P$ distribution for the 3 GeV exposure.
(c) The fitted $1/P$ distribution for the 5 GeV exposure.
(d) The upper tail χ^2 probability distribution for the 3 GeV exposure.
(e) The upper tail χ^2 probability distribution for the 5 GeV exposure.

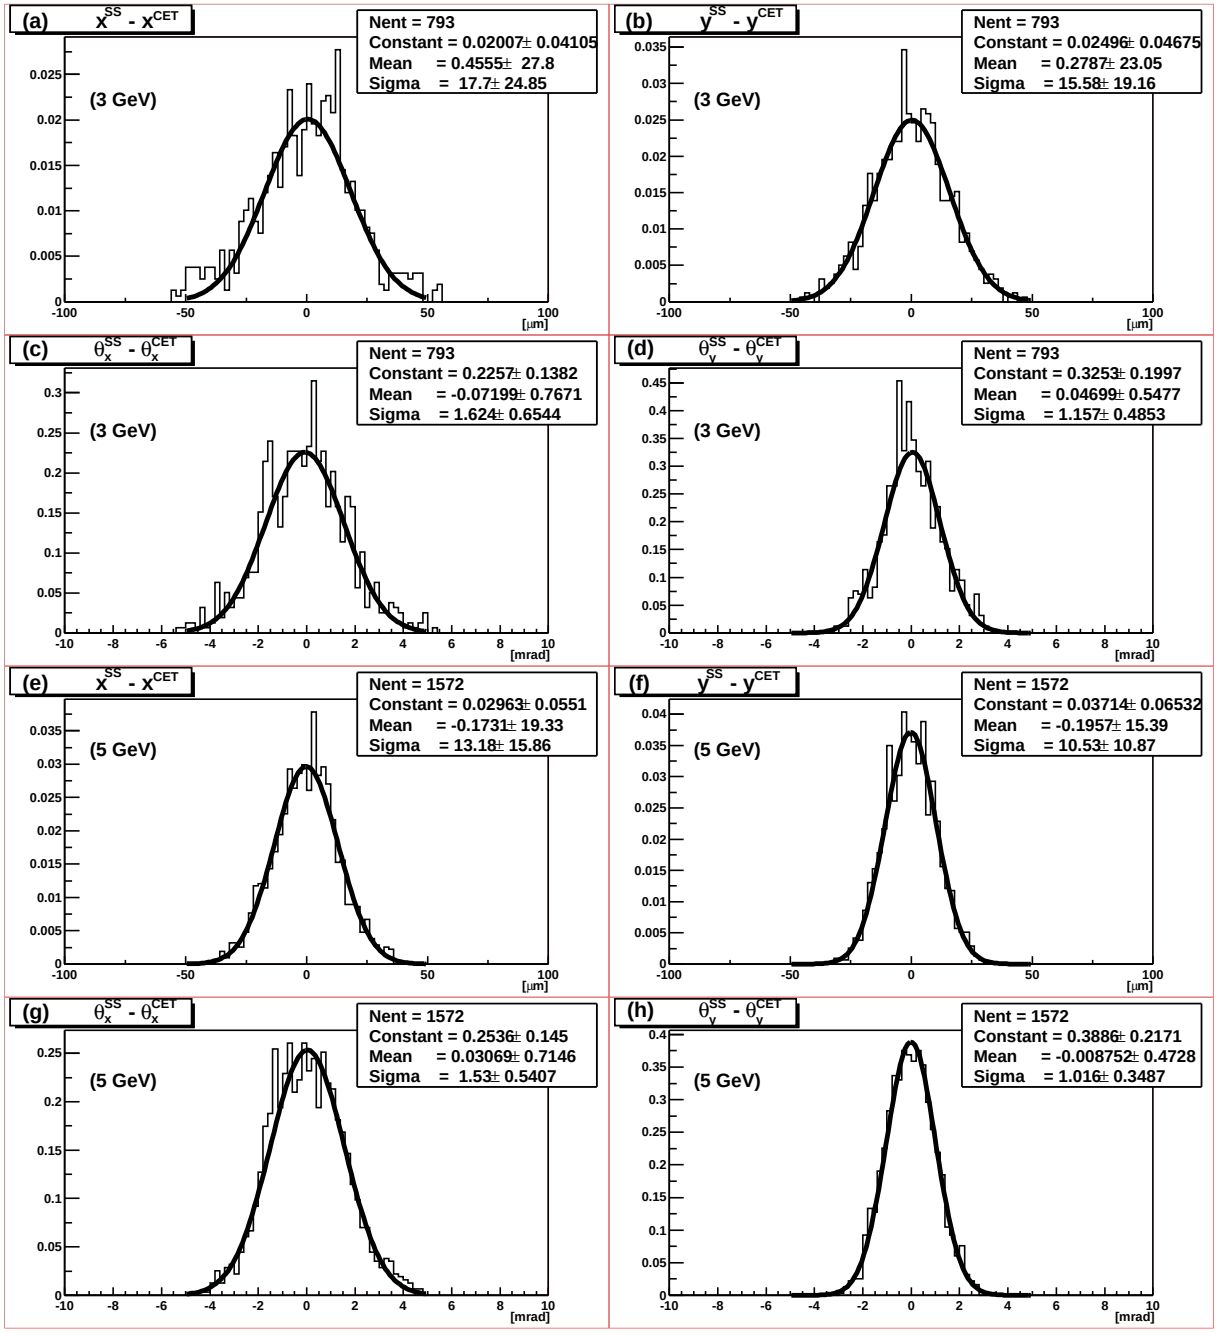

Fig. 6.8: Matching residuals between SS segments and CET reference tracks.
(a), (b) The position difference between the SS segment and the extrapolated CET track in x and y respectively for the 3 GeV exposure.
(c), (d) The slope difference between the SS segment and the extrapolated CET track in θ_x and θ_y respectively for the 3 GeV exposure.
(e), (f), (g), (h) The corresponding distributions for the 5 GeV exposure.

General track finding in the CET and final alignment

Having performed the global alignment in the CET is has become possible to apply a pattern recognition algorithm on the segments of all sheets in order to reconstruct tracks of charged particles with unknown momentum. The basic steps of the track finding algorithm are outlined here:

1. Links are established between segments which potentially belong to the same track. For every segment the algorithm searches for links to other segments on each of the three downstream sheets. Two segments are linked if they meet within errors in space after being extrapolated linearly to the middle parallel plane of the corresponding sheets. Extrapolating to the middle plane cancels out the effects of a uniform magnetic field and removes the problem of the a priori unknown momentum.

   The error in the extrapolated position comes from both the position and slope error measurement. A cut of $3 \times \sqrt{2}$ standard deviations per projection is used to define the acceptance area around the extrapolated position of the upstream segment where extrapolated positions of downstream segments are searched for.

   To reduce background from random links there is an additional requirement that the difference in the non-bending slope θ_x should not be more than $3 \times \sqrt{2}$ standard deviations.

2. The links are followed to build up tracks. Each segment not already belonging to an identified track is a seed for a new track. A track is grown by following the links of the most downstream segment until no or more than one links exist. When a link is followed the end point is checked whether it is compatible with the track which has already been formed from the previous points. The new point is accepted if it lies within three standard deviations from the extrapolated position and slope of the track at the z-coordinate of its plane. Whenever there is a case where two or more links can be followed, multiple track segments are generated. The one which gives the best χ^2 in the track fit is followed and the rest are put into a stack and are followed later.

   Out of all the track candidates which have been generated from a given seed segment, the best one is kept on the basis of the number of sheets with segments. In case more than one have the maximum value their χ^2 from the fit is used as the ultimate quality criterion.

   A track is eventually accepted if it has segments on at least four sheets.

3. The final set of reconstructed tracks are fitted and those which have a χ^2 upper tail probability greater than 0.001 are stored in the database for subsequent analysis.

After the tracks in the CET were reconstructed, a global alignment iteration was performed which was similar to the one used for the alignment with the reference tracks (described in p. 139). In this case the selection in the momentum is much looser. All

tracks with absolute reconstructed momentum of less than 10 GeV are selected. With this procedure we ensure that the final calibration/alignment of the sheets is not biased with the nominal momentum of the beam.

The track finding and alignment procedures are repeated a few times and the final alignment constants converge rapidly to their final values.

The plots in Fig. 6.9 present the results of the CET reconstruction. Negative momentum values correspond to positively charged particles. The majority of the reconstructed tracks have the momentum of the beam, which means that they correspond to pions which did not interact in the emulsion target. By isolating the peak in the inverse momentum distribution corresponding to these particles, one can calculate that the momentum resolution in the general case is $\Delta P/P = (14 \pm 1)\%$.

The flat χ^2 upper tail probability distribution shown for both sets proves that the construction of the covariance matrix from the assumed errors and the multiple scattering factors was rather successful not only in the case of the reference tracks where the momentum was known. It was found out that the distribution does not depend neither from the momentum of the tracks nor from the multiplicity N_s of the sheets with reconstructed segment. The difference in the distribution of N_s for the two exposures can be attributed to the fact that one sheet of the 3 GeV exposure CET was half-scanned.

The plots in Fig. 6.10 show a few residual distributions after the final alignment of the CET sheets. The distributions are shown for the cases that the smaller and larger widths are expected, i.e. the most upstream sheet of the 5 GeV exposure and the most downstream of the 3 GeV exposure respectively.

A final alignment iteration of the SS with respect to the CETs was performed using the tracks found with the general tracking procedure. Every CET track was extrapolated to the SS and it was matched with a segment if its position and angular difference was not more than three standard deviations. The latter was determined by the measurement error on the SS and the errors of the extrapolated track parameters. The residual distributions for this final alignment iteration are shown in Fig. 6.11.

Efficiency and background

As it can be inferred from the statics shown in the plots of Fig. 6.11 and Fig. 6.9 any track reconstructed in the CET has about 75% probability of having and identified counterpart segment in the SS. Equally important to know is how many of the CET tracks are fake and what is the efficiency of any real track with a segment on the SS to be found in the CET.

The tracks of both CETs were extrapolated together with their errors on the SS. For every segment on the SS it was searched whether a matching CET track exists for both exposures. It turned out that less than 50% of the SS segments were identified with a track in any of the two CETs. On the other hand the probability that a SS segment is matched to at least one track from each CET was smaller that 10^{-4}. The probability that a SS segment with an identified track had at least another track candidate from the same CET was of the order of 1%.

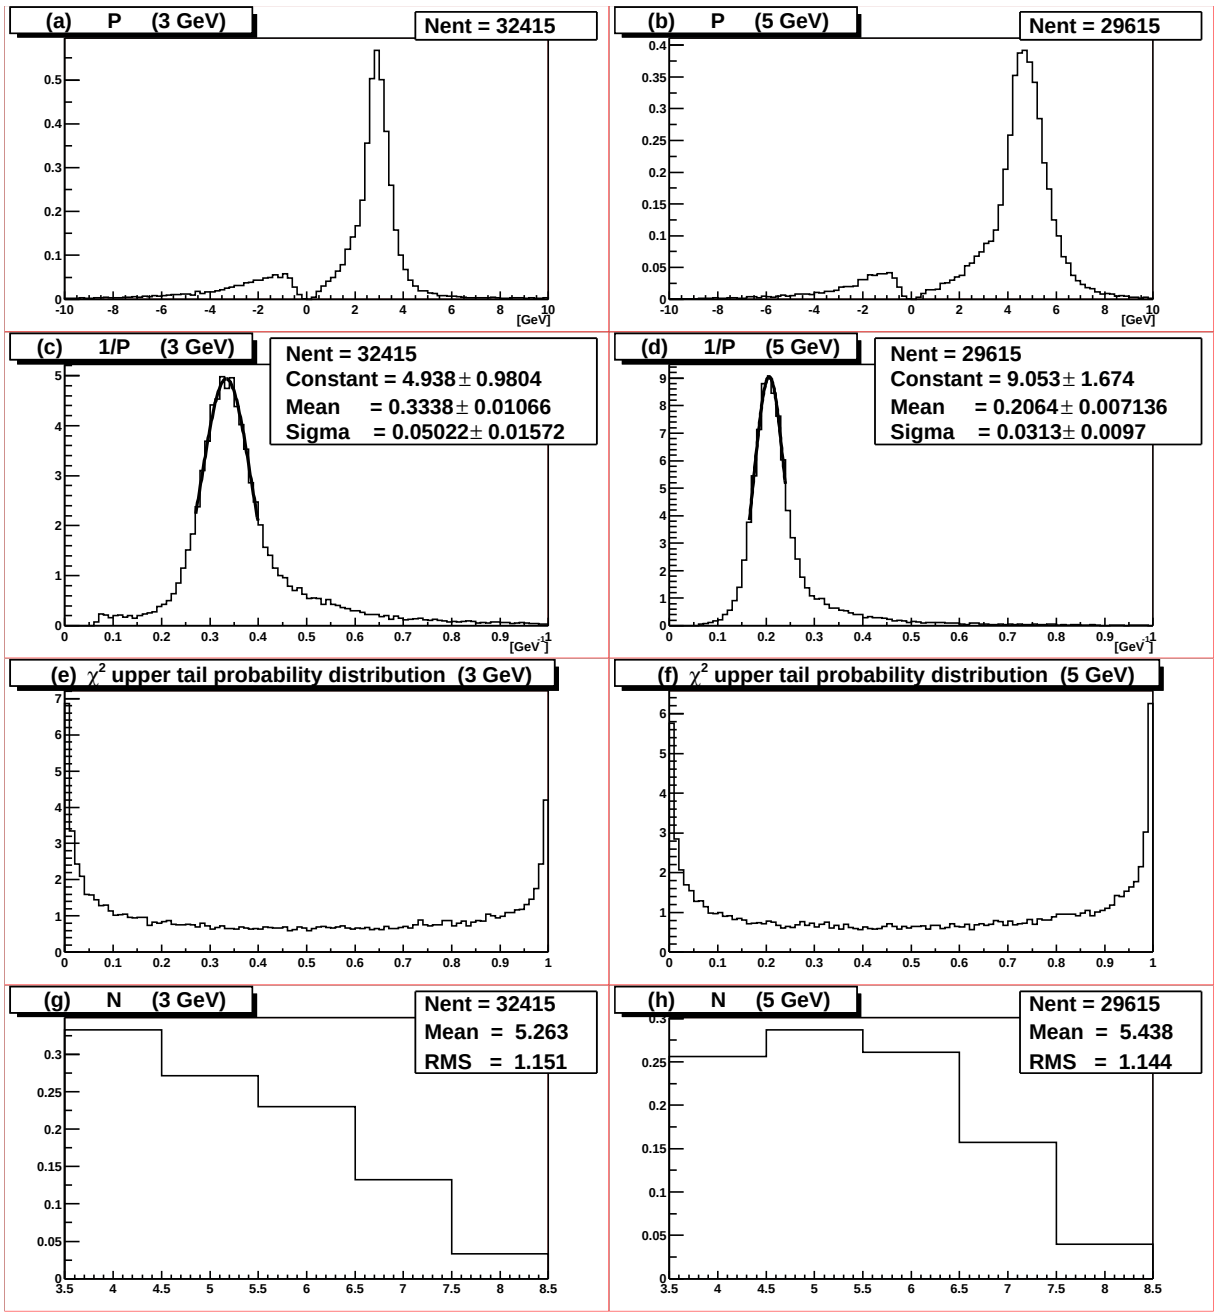

Fig. 6.9: Results of general tracking in the CET.
(a), (b) The reconstructed momentum for the 3 GeV and 5 GeV exposures.
(c), (d) The inverse reconstructed momentum ($1/P$) for the 3 GeV and 5 GeV exposures.
(e), (f) The χ^2 upper tail probability for the 3 GeV and 5 GeV exposures.
(g), (h) The number of sheets with reconstructed segments (N_s) for the 3 GeV and 5 GeV exposures.

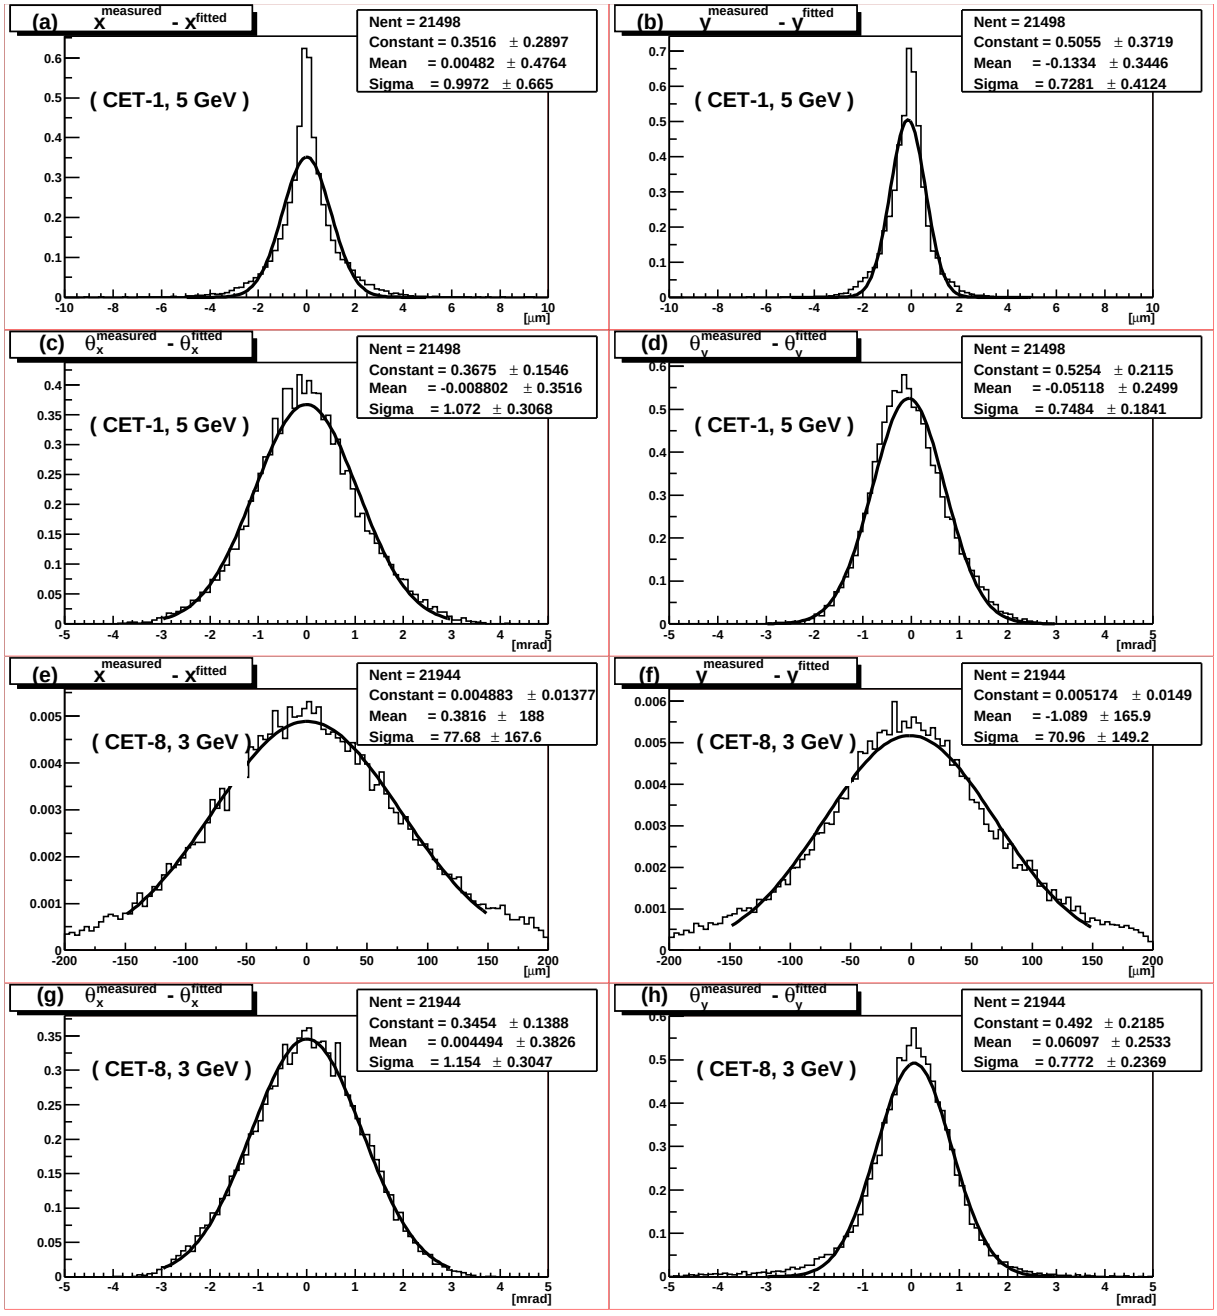

Fig. 6.10: Track fit residuals in the CET, after the final alignment and track finding iteration.
(a), (b) The distributions $x^{measured} - x^{fitted}$ and $y^{measured} - y^{fitted}$ for the most upstream CET sheet of the 5 GeV exposure.
(c), (d) The distributions $\theta_x^{measured} - \theta_x^{fitted}$ and $\theta_y^{measured} - \theta_y^{fitted}$ for the most upstream CET sheet of the 5 GeV exposure.
(e), (f), (g), (h) The corresponding plots for the most downstream CET sheet of the 3 GeV exposure.

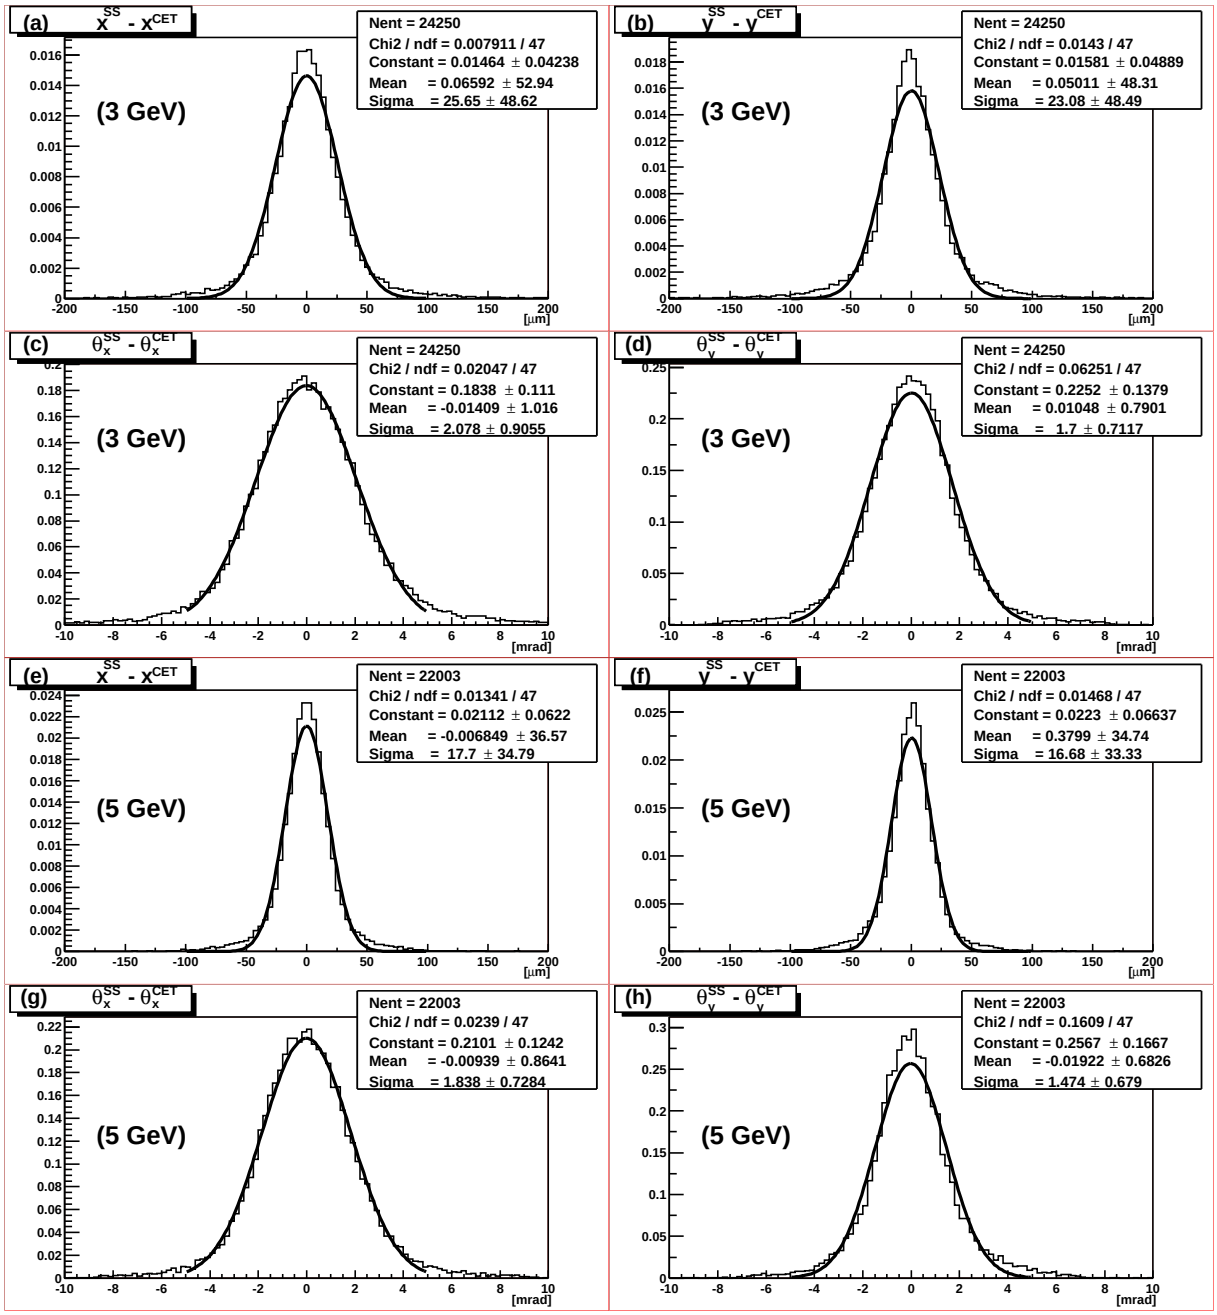

Fig. 6.11: Matching residuals between SS segments and CET tracks.
(a), (b) The position difference between the SS segment and the extrapolated CET track in x and y respectively for the 3 GeV exposure.
(c), (d) The slope difference between the SS segment and the extrapolated CET track in θ_x and θ_y respectively for the 3 GeV exposure.
(e), (f), (g), (h) The corresponding distributions for the 5 GeV exposure.

Identification of the white kink daughters

The tracks of the six white kink daughter candidates listed in Table 6.3.2 had to be followed manually downstream the kink point until the SS. This was not achieved for one (id. 7638) out of the six tracks, because it eventually escaped out of the angular acceptance. This is a strong evidence that the momentum of the corresponding particle was much too low, and as a consequence it could not have been a white kink.

For the rest of the five tracks a manual measurement of the position and the slope of the corresponding segment on the SS was performed. The results of this manual measurement were used to identify these segments with a set of five from those which were found with the automatic procedure of the grain map data taking and the off-line general tracking in the emulsion. All of the segments were found with a single candidate at a distance less than 75 μm from their nominal position and a slope difference of less than 5 mrad.

To make sure that no background segment was picked, a set of eight tracks belonging to *black kink* daughters were followed similarly until the SS and the same identification procedure was followed. *Black kinks* are the scattering topologies where there is evidence of small activity (for example a second grain, or a blob) around the kink point. Applying the same cuts as for the white kink daughter candidates, a single identified segment on the SS was found for each one of the eight tracks as well.

For the measurement of the momentum of the remaining five white kink daughter candidates we could try to match the corresponding segments on the SS with one of the reconstructed tracks of the corresponding CET. However, due to the strict quality criteria in the track finding, no CET track was found. We resorted then to the following method:

1. Each segment on the SS belonging to a white kink daughter candidate was linked to segments from all of the eight CET sheets similarly to the first step of the general track finding procedure in the CET (see p. 143). The only difference was that the extrapolated positions could meet in the middle plane within twice larger errors.

2. From the segments which were accepted it was tried manually to form a set which would be consistent with a track pointing to the corresponding SS segment. A successful set would contain segments whose θ_x slopes would not differ by 6 mrad from the slope of the SS segment, and the θ_y slopes would follow a monotonic behaviour.

3. The selected segments were used to form a track that was fitted and a corresponding momentum was calculated.

Out of the five candidates, for three of them (ids. 7811, 9096 and 7527) not a single segment could be found with the links. This is a strong indication that the corresponding tracks were eventually out of the angular acceptance of the CET. This means that their momentum was far below the beam momentum, and could therefore not be classified as white kinks.

From the remaining tracks the number of found segments and the measured momenta are summarized in Table 6.4.

track identifier	beam momentum (GeV)	number of segments in the CET	measured momentum (GeV)
3269	3	3	2.6
3247	3	4	3.1

Tab. 6.4: Final list of white kinks

Recalling that the momentum resolution in the CET is about 14% we can ascertain that the measured momenta are consistent with the corresponding beam momenta. Therefore the tracks can been validated as white kink daughters.

6.3.4 Result of the measurement and comparison with other estimations

We recall from (6.3.2) that the total effective path length which was tried for pions of 3 and 5 GeV in emulsion is $L_{tried} = 98 \pm 18$ m. We observed two events with topology and kinematics consistent with the definition of a white kink background in the CHORUS experiment. The number of observed events follow a Poisson distribution with a mean value

$$\mu_n = \frac{L_{tried}}{\lambda_{wk}} \tag{6.35}$$

where λ_{wk} is the mean free path length for the white kink interaction.

Given an experimental outcome one can construct the corresponding $1-\sigma$ confidence integral for μ_n. For the two events that were observed and using the ordering principle described in [35] one obtains that

$$\mu_n \in [0.74, 4.25] \quad \text{at } 68.27\% \text{ C.L.} \tag{6.36}$$

We can therefore calculate that

$$\lambda_{wk} = 49^{+83\,(stat)}_{-26} \pm 9^{(syst.)}\,\text{m}.$$

This result is consistent with the calculation using the NC data from CHORUS. The corresponding value is 24.0±8.5 m and was calculated from the number of white kink interactions which were accumulated in the emulsion target of the experiment and which lied topologically and kinematically outside the *tau* signal region (see following chapter).

Another dedicated experiment [127] performed at the PS with pion beams with momenta of 2, 3 and 5 GeV, predicts a value of about 48±12 m.

The combination of the above values for λ_{wk} is significantly lower than the value predicted by an experiment at KEK with a pion beam of 4 GeV (the experiment is described in [128]). After introducing updated values for the inelastic cross-section in the analysis of the experimental data [124], the value which is derived is 103±35 m. It should be noted though, that the number of the observed events with high values (above 200 MeV) of transverse momenta was very small.

The simulation studies described in [124], which were based on the FLUKA hadronic generator [129], predict a mean free path length of 22 m. However, it should be mentioned that the simulation does not take into account precisely the photographic process of the emulsion development. Moreover, it is based on energy cuts of the recoil particles to determine whether a halide crystal would render developable, instead of pure range cuts. As a consequence, one should be rather careful with the use of absolute event rates predicted by a simulation program. On the other hand, the simulation can be used to produce qualitative results such as a prediction of the shape of the distribution of the transverse momentum of the white kink daughters.

7. RESULTS OF THE OSCILLATION SEARCH IN CHORUS

Results on the search for $\nu_\mu \rightarrow \nu_\tau$ oscillation performed with the CHORUS experiment are presented in this final chapter. The analysis is fully described in [102]. Here the main points are outlined.

7.1 The accumulated statistics from electronic and emulsion data

As mentioned in 3.5, the $\nu_\mu \rightarrow \nu_\tau$ oscillation search in CHORUS has been performed using the following decay modes of the τ lepton:

$$\tau^- \rightarrow \mu^- \overline{\nu}_\mu \nu_\tau$$

and

$$\tau^- \rightarrow h^- n\pi^0 \nu_\tau.$$

where the production and decay of the τ lepton occur in the emulsion target. The information of the electronic detectors is used to define two data sets, the 1μ and 0μ samples, distinguished by the presence or absence of one reconstructed muon of negative charge.

A kinematical selection is applied to each sample in order to reduce the number of events that need to be scanned while keeping a high sensitivity to the two decay modes of the τ. An event is selected for scanning if it contained at least one negative track as a possible decay product of a τ. Moreover, the interaction vertex should have a reconstructed position inside the emulsion target.

An event belongs to the 1μ sample if it contains one reconstructed muon track of negative charge. The muon identification and reconstruction is largely based on the response of the muon spectrometer. A fraction of those muons that did not reach the spectrometer have been identified in the calorimeter. Simulation studies [130] have shown that the selection efficiency for ν_μ CC events is about 80%. The resulting 1μ sample consists of 713,000 events. A typical reconstructed 1μ event is shown in Fig. 7.1.

The 0μ sample contains events with no identified muon and a reconstructed vertex inside the emulsion target. The sample consists of 335,000 events with a calculated contamination of about 140,000 misidentified CC interactions and about 20,000 interactions of neutrinos other than ν_μ. A typical reconstructed 0μ event is shown in Fig. 7.2.

The background processes and the kinematical cuts which are applied to the two event samples are discussed in detail in Section 3.5. In addition to these, an improved analysis of the 0μ sample has been used, which resulted in a reduction of the number of expected

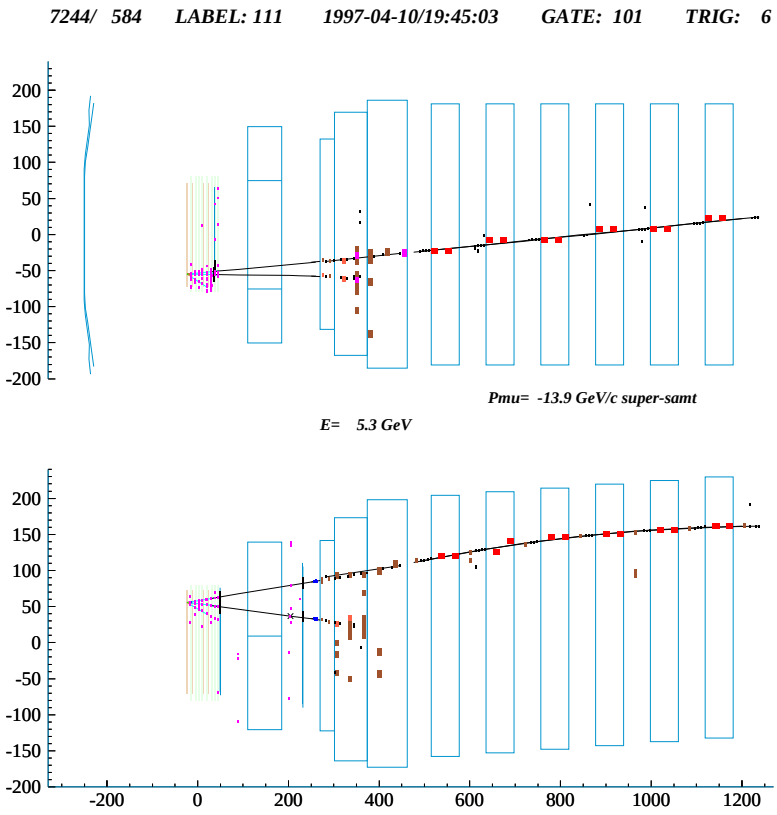

Fig. 7.1: A reconstructed 1μ event in CHORUS.

background events by a factor of 3.25, keeping at the same time more than 70% of the τ signal. The determination of the cuts was a result of an optimization [130] which maximized the oscillation sensitivity of the experiment. In particular there were two requirements on the topology of the decay vertex:

1. The decay kink should occur within a certain number of emulsion sheets downstream from one where the neutrino interaction vertex is located. This number is no longer fixed to three, but it is a function of the momentum of the scan-back hadron p_h. For each p_h a fixed fraction of 80% of the τ efficiency is required.

2. The angle $\Phi_{(\tau-h)}$ in the plane transverse to the beam axis, between the direction of the parent of the scan-back particle and the hadronic shower axis, should be less than 90°. For true ν_τ interactions this angle is close to 180°, owing to the momentum conservation at the neutrino interaction vertex. On the contrary, when the particle with the kink is part of the hadronic shower as in the white kink or charm cases, the angular distribution is flatter, sloping up towards 0°.

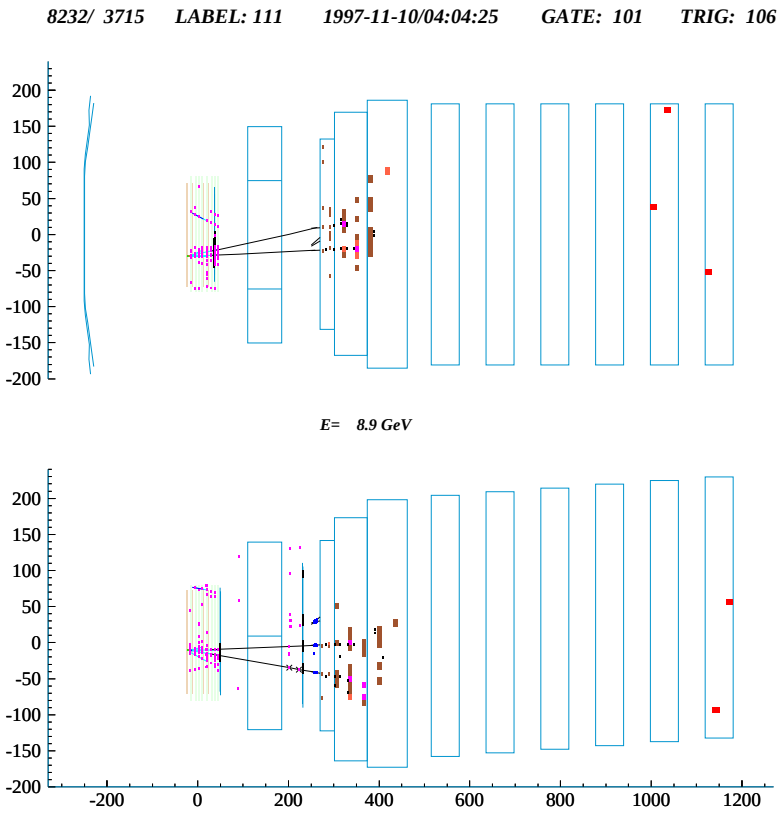

Fig. 7.2: A reconstructed 0μ event in CHORUS.

The expected number of white kink background events has been calculated assuming a mean free path length of $\lambda_{wk} = 24.0 \pm 8.5$ m, i.e. the calculation from the NC events in CHORUS. The expected number of background events is proportional to the total length of tracks which were scanned in the emulsion, inversely proportional to λ_{wk}, and proportional to the white kink efficiency after all the topological and kinematical cuts have been applied. This efficiency has been calculated from simulation studies.

After all selections the 1μ and the 0μ data sets eligible for emulsion scanning consist of 477,600 and 122,400 events respectively. The events sent to the automatic scanning procedure were fewer (355,935 and 85,211 respectively), mainly because of fiducial volume cuts imposed by the scanning technique and the bad quality of a few emulsion sheets.

With the scanning procedure described in 4.3.1, the primary interaction vertices were located of 143,742 and 20,081 events from the 1μ and the 0μ sample respectively. The events with the located vertices were searched for *kinks* (τ decays) close to the interaction vertex, according to the procedures described in 4.3.2. The kink finding efficiency was estimated using simulations of the automatic scanning procedure. Their results were in good

agreement with an independent estimation studying hadrons coming from re-interactions in the emulsion of other hadrons produced by neutrino interactions.

1μ: events in initial sample	713,000
1μ: events after kinematical cuts	477,600
1μ: events scanned	355,395
1μ: events with located vertex	143,742
1μ: N_τ^{max}	5,014
1μ: charm induced background events	0.1
1μ: WK induced background events	0
1μ: total background events	0.1
1μ: observed ν_τ events	0
0μ: events in initial sample	335,000
0μ: events after kinematical cuts	122,400
0μ: events scanned	85,211
0μ: events with located vertex	20,081
0μ: N_τ^{max}	2,004
0μ: charm induced background events	0.3
0μ: WK induced background events	0.8
0μ: total background events	1.1
0μ: observed ν_τ events	0

Tab. 7.1: Summary of the results of the oscillation search in CHORUS.

A summary of the event selection, the scanning results, and the background expectation for the current statistics are presented in Table 7.1. For each sample also the number N_τ^{max} is reported. This is the number of ν_τ events which would be observed in case all incident ν_μ had converted into ν_τ. For the 1μ sample this number is

$$(N_\tau^{max})_{1\mu} = N_{1\mu}^{loc} \cdot r_\sigma \cdot r_A \cdot \epsilon_{kink} \cdot Br_\mu \tag{7.1}$$

where:

- $N_{1\mu}^{loc}$ is the number of located 1μ events ($N_{1\mu}^{loc} = 143,742$);
- $r_\sigma = \langle \sigma_\tau^{CC} \rangle / \langle \sigma_\mu^{CC} \rangle$ is the neutrino energy weighted CC cross-section ratio. A value $r_\sigma = 0.53$ has been used; it takes into account quasi-elastic interactions, resonance production and deep inelastic reactions;
- $r_A = \langle A_\tau \rangle / \langle A_\mu \rangle$ is the cross-section weighted acceptance ratio for ν_{tau} and ν_μ interactions. A_τ and A_μ take into account the effect of geometrical and kinematical selections applied before scanning and the reconstruction and location efficiencies. It has been calculated that $r_A = 0.97$;

decay mode	Br	ϵ_{kink}	r_A	$Br \times \epsilon_{kink} \times r_A$
$\tau^- \to \nu_\tau h^- n h^0$	0.495	0.11	2.88	0.157
$\tau^- \to \nu_\tau \overline{\nu}_e e^-$	0.178	0.05	2.21	0.020
$\tau^- \to \nu_\tau \overline{\nu}_\mu \mu^-$ (0μ)	0.174	0.10	0.69	0.012
$\tau^- \to \nu_\tau \overline{\nu}_\mu \mu^-$ (1μ)	0.174	0.39	0.97	0.066

Tab. 7.2: Detection efficiency of the ν_τ for the τ decay modes contributing to the oscillation analysis samples.

- ϵ_{kink} includes the efficiency of the decay search procedure and that of the geometrical and kinematical selections applied after the kink is found. Its average value is 0.39;
- $Br_\mu = 17.4\%$ is the branching ratio of the decay $\tau^- \to \nu_\tau \overline{\nu}_\mu \mu^-$.

The formula giving N_τ^{max} for the 0μ sample contains the contributions of different decay modes of the τ lepton. The expression is

$$(N_\tau^{max})_{0\mu} = N_{0\mu}^{loc} \cdot r_\sigma \cdot \sum_{i=1,3} r_{A_i} \cdot \epsilon_{kink_i} \cdot Br_i \tag{7.2}$$

where $N_{0\mu}^{loc} = 20,081$ is the number of located 0μ events and r_σ is the same cross-section ratio appearing in (7.1). The three decay modes which contribute to the 0μ category are the τ decay in a negative hadron plus neutrals: $\tau^- \to \nu_\tau h^- n h^0$; the decay $\tau^- \to \nu_\tau \overline{\nu}_e e^- n\gamma$ $(n \geq 0)$, when the electron behaves similarly to a hadron (no early showering); the decay $\tau^- \to \nu_\tau \overline{\nu}_\mu \mu^- n\gamma$ $(n \geq 0)$, for the fraction of events where the muon is not identified but is still selected by the 0μ criteria.

The efficiency parameters for the various decay modes of the τ lepton are summarized in Table 7.2.

7.2 Determination of the oscillation sensitivity

Based on the fact that no candidate was observed in both the 1μ and 0μ samples, a limit on $\nu_\mu \to \nu_\tau$ oscillation has been determined.

As explained in 2.3.2, for large Δm^2 values the energy dependent term of the oscillation probability averages to 1/2 and, the upper limit of the probability is obtained from the equation:

$$P_{\mu\tau} \leq \frac{N_\tau}{(N_\tau^{max})_{1\mu} + (N_\tau^{max})_{0\mu}} \tag{7.3}$$

where N_τ is the upper limit on the number of τ decay candidates at 90% C.L.

When Δm^2 is comparable or smaller than the ratio E/L the spectrum of the ν_τ resulting from oscillation is modified by the energy dependent term. Then, to compute limits on

the oscillation parameters, appropriate integrations are performed to take into account the effect of energy dependent cross-sections, acceptances and efficiencies. Fig. 7.3 shows the efficiency product $A_\tau \cdot \epsilon_{kink}$ of ν_τ interactions as a function of the neutrino energy.

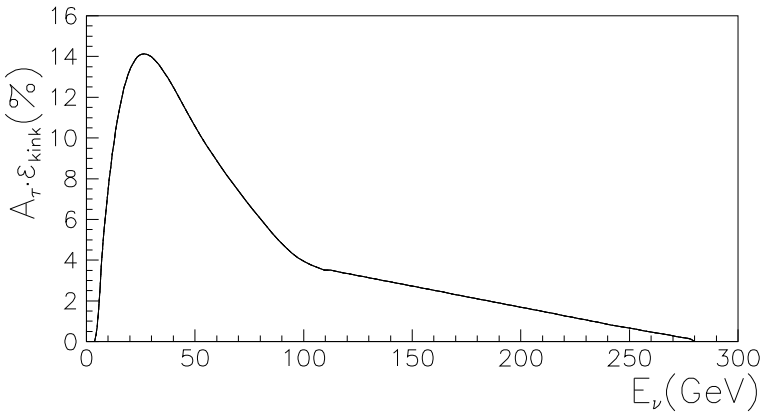

Fig. 7.3: Total efficiency for the detection of ν_τ interactions as a function of the neutrino energy.

The value of N_τ which has been used is derived using the method proposed by Junk [131] which allows the combination of different channels, taking into account the errors on the background and the signal. The difference between this and the F.C. method [35] lies in the choice of the ordering principle. A value of $N_\tau = 2.4$ has been obtained and the corresponding 90% C.L. upper limit curve on $\nu_\mu \rightarrow \nu_\tau$ oscillation is shown in Fig. 7.4. In the same plot the results from previous experiments are overlayed. It should be noted that the result from NOMAD was obtained using the F.C. statistics. Applied to the CHORUS data this statistical method would yield $N_\tau = 1.4$ and the 90% C.L. upper limit curve shown with the dashed-dotted line.

It should be noted that if one would use the result on λ_{wk} derived in the previous chapter, the number of total background events would change from 1.1 to 0.7 However, since no event was observed the final oscillation limit would practically not be affected.

In Fig. 7.5 the current limit is compared with the limit which would be obtained with the treatment of the 0μ sample only. On the same plot the proposed sensitivity to $\nu_\mu \rightarrow \nu_\tau$ oscillation is drawn as well. The proposed sensitivity is hoped to be reached with the *Phase II* analysis, described in Section 4.3.3, which involves upgraded scanning methods and improved reconstruction algorithms, leading to an increased overall ν_τ signal efficiency.

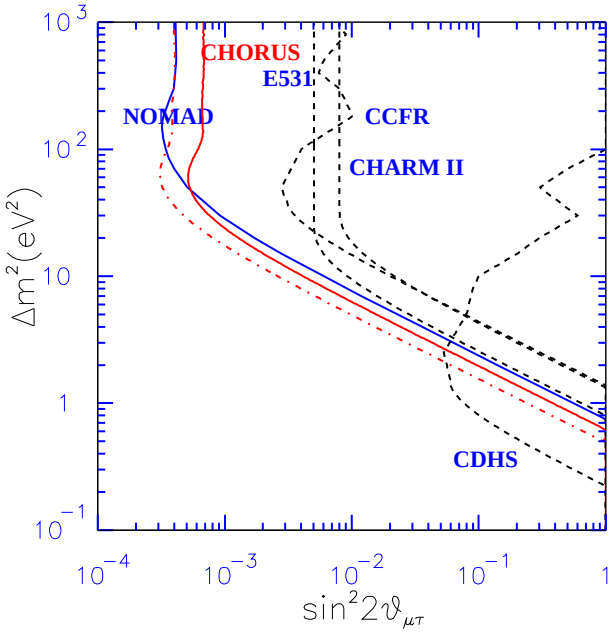

Fig. 7.4: Present limit on $\nu_\mu \rightarrow \nu_\tau$ oscillation from CHORUS, compared to the results of NOMAD [66] and previous experiments [132, 133, 134, 135].

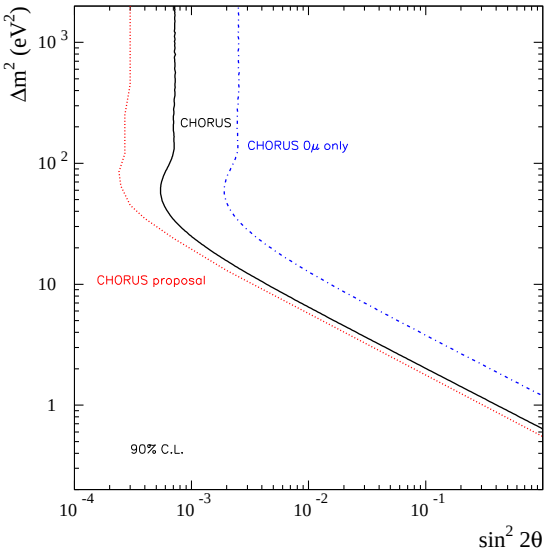

Fig. 7.5: Present limit on $\nu_\mu \rightarrow \nu_\tau$ oscillation from CHORUS (solid line), compared to the proposed sensitivity (dotted line) and the limit which would be obtained treating the 0μ sample only.

8. SUMMARY AND CONCLUSIONS

The subject of neutrino mixing and oscillations has gained interest in light of recent results from solar, atmospheric and accelerator neutrino experiments. If the existing positive results on neutrino oscillations are confirmed, then neutrinos can no longer be considered massless. This will have a large impact on High Energy Physics since a non-zero neutrino mass can be related to an energy scale where unknown phenomena occur. Massive neutrinos are also interesting for Cosmology because they could be good candidates for part or all of the dark matter in the universe.

In 2002, the future of neutrino oscillation physics at accelerators looks for the moment promising and exciting. Most major High Energy Physics laboratories around the world have included a neutrino oscillation research programme in their activity planning exploiting conventional neutrino beams to confirm or contradict the existing claims of positive results. At the same time new ideas in the production of intense neutrino beams like a neutrino factory based on a muon storage ring receive a lot of attention and relevant R&D activities are already in progress.

The CHORUS experiment is performing a search on $\nu_\mu \rightarrow \nu_\tau$ oscillation probing mixing angles in the range of $10^{-4} - 10^{-3}$ and at mass differences in the range of 1-10 eV, a region in the parameter phase space which is relevant to cosmology. CHORUS has adopted the hybrid technique combining electronic detectors with nuclear emulsions. The emulsion technique has been proven to be the ideal detector to perform neutrino experiments at accelerators where the ν_τ is involved. Therefore, the technology of automatic emulsion scanning has been evolving continuously over the last years with the participating laboratories of CHORUS performing a lot of R&D in the field.

At CERN a scanning facility has been set up trying out new ideas on the emulsion scanning. New advanced optics have been designed in order to make emulsion scanning faster and the results of the track finding and fitting from the silver grains in the emulsion sheets more efficient and precise. A flexible data acquisition system based on Digital Signal Processors has allowed the full exploitation of the capabilities of the optical system enhancing the physics scopes of the usage of the microscopes. The software needed for the operation of the microscopes and the analysis of the scanning results was designed and implemented using object-oriented technology. This ensured the smooth transition towards the new hardware components. It also helped us create a stable environment for the development of new scanning methods and analysis algorithms.

The microscopes built at CERN were used to scan not only the emulsion sheets of CHORUS, but also the ones of the target and a new tracking detector used in a test experiment at the PS. Within the context of this experiment a dedicated measurement of

the hadronic white kink interaction length has been performed. This has been achieved after having analysed the data in the test emulsion target and after having measured the momenta of the outgoing particles in the new emulsion tracker. The resulting estimate on the expected background event rate is compatible with the internal calculations using the data of CHORUS.

After having completed the so-called *Phase I* analysis in CHORUS no ν_τ event has been observed in the search for $\nu_\mu \rightarrow \nu_\tau$ oscillation. The null result of CHORUS has been used to exclude a region in the oscillation parameter phase space as it is shown in Fig. 7.4. The on-going *Phase II* of the oscillation analysis in CHORUS is expected to increase the final sensitivity to $\nu_\mu \rightarrow \nu_\tau$ oscillation.

BIBLIOGRAPHY

[1] F. Reines and C. L. Cowan, Jr., *Phys. Rev.* **92**, 830 (1953).

[2] C. L. Cowan, Jr., et al., *Science* **124**, 103 (1956).

[3] G. Danby et al., *Phys. Lett.* **9**, 36 (1962).

[4] K. Kodama et al., *Phys. Lett.* **B504**, 218 (2001).

[5] T. D. Lee and C. N. Yang, *Phys. Rev.* **104**, 254 (1956).

[6] C. S. Wu et al., *Phys. Rev.* **105**, 1413 (1957).

[7] R. L. Garwin et al., *Phys. Rev.* **105**, 1415 (1957).

[8] J. I. Friedman and V. L. Telegdi, *Phys. Rev.* **105**, 1681 (1957).

[9] CHARM II Collaboration, P. Vilain et al., *Phys. Lett.* **B335**, 246 (1994).

[10] The LEP Collaborations, the LEP Electroweak Working Group, and the SLD Heavy Flavor Group, CERN/PPE/97-154.

[11] D. A. Krakauer et al., *Phys. Lett.* **B252**, 177 (1990).

[12] A. M. Cooper-Sarkar et al., *Phys. Lett.* **B280**, 153 (1992).

[13] E. Majorana, *Nuovo Cimento* **14**, 171 (1937).

[14] R. Plaga, 'A demonstration that the observed neutrinos are not Majorana particles', hep-ph/9610545.

[15] M. Gell-Mann, P. Ramond and R. Slansky, *'Supergravity'*, edited by D. Freedman et al., North Holland (1979).

[16] T. Yanagita, *Prog. Theor. Phys.* **64**, 1103 (1980).

[17] A complete listing of the limits on neutrinoless double beta decays using various nuclei are given by the Particle Data Group: *Eur. Phys. Jour.* **C3**, 323 (1998).

[18] J. Ellis, Proc. of the *17th International Conference on Neutrino Physics and Astrophysics (Neutrino 96)*, Helsinki, Finland, 13-20 June 1996.

[19] The Particle Data Group, *Eur. Phys. Jour.* **C3**, 313 (1998).

[20] Ch. Weinheimer et al., *Phys. Lett.* **B460**, 219 (1999).

[21] K. Assamagan et al., *Phys. Rev.* **D53**, 6067 (1996).

[22] OPAL Collaboration, R. Akers et al., *Z. Phys.* **C65**, 183 (1995).

[23] ALEPH Collaboration, R. Barate et al., *Eur. Phys. Jour.* **C2** 395 (1998).

[24] G. Gelmini and E. Roulet, *Rep. Progr. Phys.* **58**, 1207 (1995).

[25] B. Pontecorvo, *J.E.T.P* **33**, 429 (1957);
B. Pontecorvo, *J.E.T.P* **34**, 172 (1958).

[26] Z. Maki, M. Nakagawa and S. Sakata, *Prog. Theor. Phys.* **28**, 870 (1962).

[27] C. Giunti and C. W. Kim, *Phys. Rev.* **D58**, 17301 (1998).

[28] L. Wolfenstein, *Phys. Rev.* **D17**, 2369 (1978).

[29] S. P. Mikheyev and A. Y. Smirnov, *Sov. J. Nucl. Phys.* **42**, 913 (1985).

[30] J. N. Bahcall, S. Basu and M. H. Pinsonneault, *Phys. Lett.* **B433**, 1 (1998).

[31] Y. Declais et al., *Phys. Lett.* **B338**, 383 (1994).

[32] M. Honda et al., *Phys. Lett.* **D52**, 4985 (1995).

[33] R. D. Cousins and V. L. Highland, *Nucl. Instr. Meth.* **A320**, 331 (1992).

[34] The Particle Data Group, *Eur. Phys. Jour.* **C3**, 174 (1998).

[35] G. J. Feldman and R. D. Cousins, *Phys Rev.* **D57**, 3873 (1998).

[36] J. Neyman, Phil. Trans. Royal Soc. London Series A, **236**, 333-80 (1937).

[37] Proceedings of the Workshop on Confidence Limits, edited by F. James, L. Lyons and Y. Perrin, CERN 2000-005, Geneva 17-18 January 2000.

[38] L. L. Chau and W. Y. Keung, *Phys. Rev.* **D16**, 1425 (1984).

[39] S. M. Bilenky and B. Pontecorvo *Phys. Rep.* **41**, 225 (1978);
L. Wolfenstein *Nucl. Phys.* **B 186**, 147 (1981).

[40] G. L. Fogli et al., *Phys. Rev.* **D52**, 5334 (1995).

[41] R. Davis Jr., D. S. Harmer and K. C. Hoffman, *Phys. Rev. Lett.* **20**, 1205 (1968).

[42] J. N. Bahcall, N. A. Bahcall and G. Shaviv, *Phys. Rev. Lett.* **20**, 1209 (1968).

[43] W. Hampel et al., *Phys. Lett.* **B388**, 364 (1996).

[44] V. Gavrin et al., in *Neutrino 98*, Proceedings of the XVII International Conference on Neutrino Physics and Astrophysics, Takayama, Japan, 4-9 June 1998, edited by Y. Suzuki and Y. Totsuka. To be published in *Nucl. Phys. B* (Proc. Suppl.)

[45] Y. Fukuda et al., *Phys. Rev Lett.* **77**, 1683 (1996).

[46] Y. Suzuki for the Super-Kamiokande Collaboration, in *Neutrino 98*, Proceedings of the XVII International Conference on Neutrino Physics and Astrophysics, Takayama, Japan, 4-9 June 1998, edited by Y. Suzuki and Y. Totsuka. To be published in *Nucl. Phys. B* (Proc. Suppl.)

[47] B. T. Cleveland et al., *Astrophys.* **J496**, 505 (1998).

[48] J. N. Bahcall, P. I. Krastev and A. Yu. Smirnov *Phys. Rev.* **D58** 96016 (1998).

[49] G. Ewan et al., Sudbury Neutrino Observatory Proposal *SNO-87-12* (1987).

[50] Q. R.Ahmad et al., *Phys. Rev. Lett.* **87** 071301 (2001).

[51] C. Aprasella et al., BOREXINO proposal, Vols. 1 and 2, eds. G. Bellini, R. Raghavan et al., (Univ. of Milano, Milano, 1992).

[52] M. Aglietta et al., *Europhys. Lett.* **8**, 611 (1989).

[53] K. Daum et al., *Z. Phys.* **C66**, 417 (1995).

[54] R. Becker-Szendy et al., *Phys. Rev.* **B46** 3720 (1992).

[55] W. W. M. Allison et al., *Phys. Lett.* **B449**, 137 (1999).

[56] Y. Fukuda et al., *Phys. Lett.* **B335**, 237 (1994).

[57] T. Toshito for the Super-Kamiokande collaboration, "Super-Kamiokande atmospheric neutrino results", talk presented on XXXVIth Rencontres de Moriond Electroweak Interactions and Unified Theories, 10-17 March 2001.

[58] M. Apollonio et al., *Phys. Lett.* **B420**, 397 (1998);
M. Apollonio et al., "Limits on Neutrino Oscillations from the CHOOZ Experiment", *hep-ex/9907037*, (1999).

[59] A. Suzuki for the KamLAND Collaboration, in the 8th International Workshop on Neutrino Telescopes, Feb. 23-26, Venice, Italy, 1999.
http://www.awa.tohoku.ac.jp/html/KamLAND/index.html.

[60] C. Athanassopoulos et al., *Phys. Rev Lett.* **77**, 3082 (1996).

[61] C. Athanassopoulos et al., *Phys. Rev Lett.* **81**, 1774 (1998).

[62] K. Eitel for the KARMEN Collaboration, in *Neutrino 98*, Proceedings of the XVII International Conference on Neutrino Physics and Astrophysics, Takayama, Japan, 4-9 June 1998, edited by Y. Suzuki and Y. Totsuka. To be published in *Nucl. Phys. B* (Proc. Suppl.)

[63] K. Eitel, "Compatibility Analysis of the LSND Evidence and the KARMEN Exclusion for $\overline{\nu}_\mu \rightarrow \overline{\nu}_e$", *hep-ex/9909036*, 1999.

[64] M. de Jong et al., *CERN-PPE/93-131* (1993).

[65] P. Astier et al., *CERN-SPSLC/91-21; CERN-SPSLC/91-48; CERN-SPSLC/91-53; CERN-SPSLC/92-51* (1992).

[66] J. Astier et al. *Phys. Lett.* **B483**, 387 (2000).

[67] R. P. Thun and S. McKee, *Phys. Lett.* **B439**, 123 (1998).

[68] E. Church et al., *LA-UR-97-2120*, (1997).

[69] N. Armenise et al., *CERN-SPSC/97-21; SPSC/I-216* (1997).

[70] K. Nishikawa et al., KEK-PS Proposal E-362 (1994).

[71] E. Ables et al., MINOS Collaboration, Fermilab Proposal P-875 (1995).

[72] J. Hylen et al., Fermilab-TM-2018 (1997).

[73] G. Acquistapace et al. *CERN Report* **98-02**, CERN, Geneva, 19 May 1998.

[74] K. Kodama et al., *CERN-SPSC-98-25; SPSC-M-612; LNGS-LOI-8-97 Add.I* (1998).

[75] J-P. Revol et al., ICARUS proposal, LNGS-94/99-I (1994).

[76] B. Autin et al. *CERN Report* **99-02**, CERN, Geneva, 30 April 1999.

[77] C. M. Ankenbrandt et al. (Muon Collider Collaboration), *Phys. Rev. ST Accel. Beams* **2** 081001 (1999).

[78] The HARP experiment, http://www.cern.ch/harp.

[79] The MICE experiment, Letter of Intent to Paul Schrerrer Institute, 28 November 2001.

[80] P. Zucchelli, "A novel concept for a $\overline{\nu}_e$ neutrino factory", CERN-EP/2001-056, July 2001, (Submitted to *Phys. Lett. B*).

[81] J. Ellis, J. L. Lopez, D. V. Nanopoulos, *Phys. Lett. B* **292** 189 (1992).

[82] E.,Eskut et al., *Nuc. Instr. Meth.* **A401**, 7 (1997).

[83] E. Heijne, CERN Yellow Report, CERN 83-06, July 1983.

[84] G. Acquistapace *et. al.*, CERN Internal Note, CERN-ECP 95-14, July 1995.

[85] L. Casagrande *et. al.*, CERN Yellow Report, CERN 96-06, August 1996.

[86] K. Bernier *et. al.*, CERN Yellow Report, CERN 97-07, 3 July 1997.

[87] S. Sorrentino "GBEAM, the neutrino beam simulation", CHORUS internal note 98005, 12/5/1998.

[88] B. Van de Vyver, and P. Zucchelli *Nuc. Instr. Meth.* **A385**, 91 (1997).

[89] M.C. Gonzalez-Garcia and J.J. Gomez-Cadenas, *Phys. Rev.* **D55** 1297 (1997).

[90] P. Annis *et al.*, *Nuc. Instr. Meth.* **A409**, 361 (2000).

[91] S. Aoki *et al.*, *Nuc. Instr. Meth.* **A447**, 629 (1998).

[92] F. Bergsma *et al.*, *Nuc. Instr. Meth.* **A357**, 243-248 (1995).

[93] J. W. E. Uiterwijk *et al.*, *Nuc. Instr. Meth.* **A409**, 682-686 (1998).

[94] E. Di Capua *et al.*, *Nuc. Instr. Meth.* **A378**, 221 (1996).

[95] H. Abramovitz *et al.*, *Nuc. Instr. Meth.* **180**, 429 (1981).

[96] J. P. Dewulf *et al.*, *Nuc. Instr. Meth.* **A252**, 443 (1986).

[97] M. G. van Beuzekom *et al.*, *Nuc. Instr. Meth.* **A427**, 587-606 (1999).

[98] A. Artamonov *et al.*, "The Data Acquisition System of the CHORUS Experiment", to be published as CERN-EP preprint, and submitted to *Nuc. Instr. Meth.*, (2000).

[99] R. Gurin and A. Maslennikov, Control Host project, CHORUS internal notes (1995); R. Gurin and A. Maslennikov, "ControlHost: field-tested package for Distributed Data Handling", HEPiX95 Meeting, Institute of Physics of the Academy of Sciences (FZU), Prague, Czech Republic, 31 May to 2 June 1995; http://www-hep.fzu.cz/computing/HEPiX/HEPiX95/talks/conhost.ps.

[100] E. Eskut et al., *Phys. Lett.* **B424**, 202-212 (1998).

[101] E. Eskut et al., *Phys. Lett.* **B434**, 205-213 (1998).

[102] E. Eskut et al., *Phys. Lett.* **B497**, 8-22 (2001).

[103] E. Pesen, PhD thesis, Middle East Technical University, Ankara, Turkey, 1997.

[104] P. Righini, PhD thesis, University of Cagliari, Italy, 1998.

[105] The Particle Data Group, *Eur. Phys. Jour.* **C3**, 287 (1998).

[106] D. Allasia et al., *Nucl. Phys.* **B176**, 13 (1980).

[107] S. Aoki et al., *Nucl. Instr. Meth.* **A274**, 64 (1989).

[108] N. Ushida et al., *Nucl. Instr. Meth.* **224**, 50 (1984).

[109] K. Kodama et al., *Nucl. Instr. Meth.* **B93**, 340 (1994).

[110] http::/flab.phys.nagoya-u.ac.jp/flab.html

[111] T. Nakano *Ph.D. thesis*, Nagoya Univ., Japan, 1997.

[112] http::/choruswww.cern.ch/CERN-Microscope

[113] B. Van de Vyver, 'Scanning at CERN', CHORUS internal note, March 1, 1997.

[114] A. di Bartolomeo et al., 'Status of the automatic microscopes at CERN', CHORUS internal note 98-006, June 19, 1998.

[115] B. Van de Vyver, PhD thesis, Inter-University Institute for High Energies (ULB-VUB) Brussels, Belgium, 2002.

[116] P. Zucchelli, 'On microscope images, filtering and compression', CHORUS internal note, 1997.

[117] R. W. Hamming, 'Digital Filters', Prentice Hall International Inv, 1989.

[118] http://www-users.cs.york.ac.uk/ fisher/mkfilter, 'Interactive Digital Filter Design', A. J. Fisher.

[119] J. Panman, private communication.

[120] J. W. E. Uiterwijk, *Ph.D. thesis in preparation*, NIKHEF, Amsterdam, The Netherlands.

[121] A. S. Ayan et al., 'A High Sensitivity Short Baseline Experiment to Search for $\nu_\mu \rightarrow \nu_\tau$ Oscillation', CERN-SPSC/97-5, SPSC/I213, March 1997.

[122] M. G. Catanesi et al., 'A Silicon Tracker for Track Extrapolation into Nuclear Emulsions', *Nucl. Instr. Meth.* **A434**, 218 (1999).

[123] S. Buontempo et al., 'The Compact Emulsion Spectrometer', CERN-EP-2000-060, May 9, 2000. To be published in *Nuc. Instr. Meth.*

[124] F. Cassol 'Monte Carlo simulation of the White Star Kinks', CHORUS Internal Note 98002, April 15, 1998.

[125] P. Abrahamson et al. *Nuovo Cimento* **Vol. XII. N. 1**, 939 (1959).

[126] R. K. Bock et al., 'Data analysis techniques for high-energy physics experiments', edited by M. Regler, Cambridge University Press 1990.

[127] A. Bülte and the CHARON Collaboration (http://www.cern.ch/charon).

[128] K. Kodama et al., FERMILAB P803 Proposal (1993), Appendix D.

[129] A. Ferrari, T. Rancati and P.R. Sala, SARE-3 Workshop, KEK-Tsukuba, KEK Report Proceedings 97-5, p. 165 (1997).

[130] A. Satta, PhD thesis, Università Degli Studi di Roma "La Sapienza", 2001.

[131] T. Junk, *Nucl. Instr. Meth.* **A434**, 435 (1999).

[132] N. Ushida et al., *Phys. Rev. Lett.* **57**, 2897 (1986).

[133] F. Dydak et al., *Phys. Lett.* **B134**, 103 (1984).

[134] M. Gruvé et al., *Phys. Lett.* **B309**, 463 (1993).

[135] K. S. McFarland et al., *Phys. Rev. Lett.* **75**, 3993 (1995).

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

It has been a great pleasure having to write this small part of my thesis. I guess that writing the "acknowledgements" has always been a pleasure for every Ph.D. candidate in any field for two reasons. The first one is obvious but rarely openly admitted: Nobody can ask for corrections, modifications or improvements, not even for matters concerning the language of the text. Nowadays the automatic spelling and grammar checker can do it for you. Forbidden words such as *"I"*, *"my"* and *"mine"* are at last eligible to be used. It is therefore the sense of liberty which makes the writing of this piece of text so attractive.

The second reason is that this is the very last piece of text that needs to written after having spent God-knows-how-many hours, days and nights over the keyboard of a computer typing and correcting the chapters of the thesis. It happens right after the moment when you feel like shouting "I'm done, at last!!!". It also coincides with the moments when you remember how often you were ready to give up any hope of succeeding; how strongly you wanted to break into million pieces the screen of your computer because this unbelievably stupid and dump slave does not understand what his master commands; when the experimental data seemed to have arranged the world's largest conspiracy against your mental integrity; when all kinds of problems appear together at the very same time, which is usually when you think that you have almost accomplished your goal.

The fact that I am writing this piece of text is a proof that all of these adventures are now over. I eventually did not regret for doing a PhD in experimental Particle Physics. Thanks to few people not only did I succeed in finishing my thesis and in surviving my small odyssey, but also I remained spiritually alive and mentally healthy (or at least I think so...). And to be really honest, I enjoyed it a lot!

I am fully indebted to my two supervisors, Dr. Christine Kourkoumelis from the University of Athens, and Dr. Jaap Panman from CERN for the time that they had to spend dealing with me. They provided me with moral support and have been remarkably patient and tolerant despite the fact that I kept breaking my promises for the date of delivery of my thesis.

Jaap tried very hard over a period of five years to cool down my greek hot temperament and transform me into a reasonable professional. If I have advanced so far in my career, this is mainly thanks to his mentorship and its positive impact on my personality.

Christine made sure that the contents of my thesis would meet the standards required by the University of Athens in order to obtain my Doctorate. She has also been my best ally when I had to fight against the formal administrative procedures governing the University of Athens.

This thesis would never have been finished without the help of my dearest friends Bart Van de Vyver, Johan Uiterwijk and Erhan Pesen. We spent countless hours together during days and nights in the scanning laboratory at CERN, sweating and sometimes even bleeding in order to build the automatic microscopes, develop the data acquisition and analysis algorithms and scan the emulsion sheets from CHORUS and the test-experiment.

I managed to master the operation of the digital cameras and the DSP boards thanks to Bev Friend and Misha Chizhov. I have also profited from the presence of Oliver Melzer and Mathieu Doucet who also participated in the microscope project at CERN. Helge Meinhard was always available to fix any problem related to computing infrastructure.

I enormously enjoyed the stimulated discussions on neutrino physics with Piero Zucchelli, Maarten De Jong and Pasqualle Migliozzi.

I feel very lucky that I had the chance to meet and discuss with the world masters of automatic emulsion scanning Niwa-sama and Nakano-san.

I am indebted to the whole CHORUS Collaboration, the people involved in the setting up and analysis of the test-experiment at the PS, and everybody who helped me in any way with the analysis of the WK measurement. In particular I would like to thank Stefania Ricciardi, Etsuko Niu, Maarten Litmath, Aimo Bülte, Rui Ferreira, Jürgen Brunner, Franca Cassol, Alessia Satta, Roberta Santacessaria, Gino Saitta, Paolo Strolin, Antonio Ereditato, Ubaldo Dore, Petr Gorbunov and Klaus Winter.

Emmanuel Tsesmelis, Lucie Linssen and Luigi Di Lella played a crucial role during the early years of my career at CERN. They have been the first people to introduce me to experimental Neutrino Physics when I was a CERN Summer Student in 1994 and when I worked on my master thesis in 1996-1997.

During the last two years that I worked for the HARP experiment, I had the great pleasure to work with Friedrich Dydak, Simone Giani, Dirk Düllmann and Emilio Radiccioni who helped me acquire both knowledge on neutrino physics and technical skills in software development and data handling that have been useful for the work needed by my thesis as well.

I would also like to thank the professors of my jury: Pavlos Ioannou, Thanasis Lahanas, George Tzanakos, Manolis Dris, Paris Sphicas and Dimitris Fassouliotis, for carefully reading the final draft of my thesis, for their comments on it, and for making me enjoy the defence of my thesis.

Last and certainly not least I would like to thank the persons I love most in the world, my wife Katerina and my daughter Danae. I have given them a promise that I will keep at any cost: after obtaining my Doctorate I will become a normal human being again. No more working during the weekends or public holidays. No more returning late from work. No more talking on work related matters during a dinner with friends. I will increase the range of my interests and free-time activities, and I will become a more interesting person for somebody who wants to chat with me.

There are certainly plenty of many nicer things to do in one's life besides gazing at numbers and plots on a computer screen. Yes, even in a physicist's life.

Ioannis M. Papadopoulos