


Detection and reconstruction of short-lived particles produced by neutrino interactions in emulsion



IN LIBRARIES, GENEVA

LIBRARIES, GENEVA



Johan
Uiterwijk

**Detection and reconstruction of short-lived
particles produced by neutrino interactions
in emulsion**

Proefschrift

ter verkrijging van
de graad van Doctor aan de Universiteit Leiden,
op gezag van de Rector Magnificus Prof. Mr. P.F. van der Heijden,
volgens besluit van het College voor Promoties
te verdedigen op dinsdag 12 juni 2007
klokke 15:00 uur

door

Johannes Wilhelmus Edmond Uiterwijk

geboren te Maastricht

in 1969

promotor: Prof. Dr. M. de Jong
co-promotor: Dr. J. Panman
referent: Prof. Dr. E. Koffeman
promotiecommissie: Prof. Dr. A. Achúcarro
Prof. Dr. P.J. van Baal
Prof. Dr. J.M. van Ruitenbeek
Dr. J.L. Visschers

The work described in this dissertation is part of the research program of 'het Nationaal Instituut voor Kernfysica en Hoge-Energie Fysica (NIKHEF)' in Amsterdam, the Netherlands. The Author was financially supported by 'de Stichting voor Fundamenteel Onderzoek der Materie (FOM)' and the European Center for Nuclear Research (CERN) in Geneva, Switzerland. FOM is funded by 'de Nederlandse Organisatie voor Fundamenteel Onderzoek der Materie (NWO)'

ISBN: 978-90-8666-034-6

Contents

Introduction	vii
1 Neutrino masses and oscillations	1
1.1 Neutrino history	2
1.1.1 The electron neutrino	2
1.1.2 The muon neutrino	2
1.1.3 The tau neutrino	3
1.2 The standard model	4
1.2.1 Particles and forces	4
1.2.2 Parameters and constants	6
1.2.3 Neutrinos	6
1.2.4 Neutrino–matter interactions and cross-sections	7
1.2.5 Beyond the standard model	10
1.3 Neutrino masses and oscillation	10
1.3.1 Direct mass measurements	10
1.3.2 Neutrino oscillation	11
1.3.3 Oscillation detection methods	14
1.4 Neutrino oscillation hints	15
1.4.1 Cosmology	15
1.4.2 Solar neutrinos	16
1.4.3 Atmospheric neutrinos	20
1.4.4 The 1998 Super-Kamiokande result	22
1.5 The CHORUS oscillation search	23
1.5.1 Motivation	23
1.5.2 Excluded parameter space	24
1.6 Current neutrino experiments	25
1.6.1 Super-Kamiokande	26
1.6.2 SNO	29
1.6.3 KamLAND	30
1.7 Status of oscillation research	31
1.7.1 Atmospheric neutrinos	31
1.7.2 Solar and reactor neutrinos	35
1.7.3 Accelerator experiments	40
1.7.4 Cosmic neutrinos	42
1.7.5 Three-flavour oscillation	43
1.8 Discussion and outlook	43

2	The CHORUS experiment	47
2.1	Detection principle	48
2.1.1	Tau identification in emulsion	48
2.1.2	Background processes	50
2.2	Neutrino beam	51
2.3	Experimental setup overview	54
2.4	Emulsion target and electronic tracking detectors	55
2.4.1	Emulsion target considerations	55
2.4.2	Interfacing emulsion and electronic tracking detectors	58
2.4.3	Tracking detector	59
2.4.4	Target region experimental setup	60
2.5	Trigger	62
2.6	Downstream detectors	64
2.6.1	Hadron spectrometer	64
2.6.2	Streamer-tubes and honeycomb detector	65
2.6.3	Calorimeter	66
2.6.4	Muon spectrometer	66
2.7	Online monitoring	68
2.8	Track reconstruction and scanning predictions	70
2.9	Emulsion and scanning techniques	71
2.9.1	Hybrid experiments and automatic scanning	71
2.9.2	Emulsion characteristics	72
2.9.3	Scanning microscopes	76
2.9.4	Automatic track recognition	76
2.10	Reconstructing tracks and vertices in emulsion	77
2.10.1	Alignment	77
2.10.2	Interaction location by scan-back	79
2.10.3	Vertex reconstruction with net-scan	82
3	Honeycomb tracker	85
3.1	Motivation and requirements	86
3.2	Detection principle	87
3.3	Design and mechanical construction	88
3.3.1	Monolayer	89
3.3.2	Honeycomb module	91
3.3.3	Honeycomb tracker	92
3.3.4	Prototype measurements	93
3.4	Data-acquisition and read-out electronics	95
3.4.1	The bit-stream principle	95
3.4.2	Chambercards	96
3.4.3	Clockcard	100
3.5	Read-out system	101
3.5.1	Read-out protocol	103
3.5.2	Read-out program	104
3.5.3	Data compression	106

3.6	Tracking	106
3.6.1	Drift-time calibration	107
3.6.2	Track finding per module	108
3.6.3	Reconstructing 3-D tracks	109
3.7	Performance and discussion	111
3.7.1	Resolution	111
3.7.2	Read-out electronics	111
3.7.3	Conclusion	112
4	Track finding in emulsion	113
4.1	Introduction	114
4.1.1	Microscope optics and stages	115
4.1.2	Tracking input characteristics	118
4.1.3	Algorithm restrictions and requirements	118
4.1.4	Toolkit abstraction	119
4.2	Multi-dimensional ordering containers	119
4.2.1	Find-in-range algorithm	119
4.2.2	Search trees	120
4.2.3	Hash table	123
4.2.4	Implementation	124
4.2.5	Timing performance	127
4.3	Track-finding algorithm	129
4.3.1	Concept	130
4.3.2	Implementation	133
4.3.3	Tracking time	137
4.4	Tracking efficiency for simulated data	140
4.4.1	Hit generators	141
4.4.2	Acceptance criteria	143
4.4.3	Efficiency criteria	145
4.4.4	Results and discussion	145
4.5	The track trigger	149
4.5.1	Concept	149
4.5.2	Implementation	149
4.6	Application in CHORUS emulsion scanning	150
4.6.1	Tracking configuration	151
4.6.2	Prediction matching on the interface sheets	152
4.6.3	Local-alignment procedure	153
4.6.4	Tracking results	154
4.7	Conclusion and discussion	156
5	D^{*+} production	157
5.1	Introduction	158
5.2	Charm-quark production and fragmentation	158
5.3	Vertex reconstruction and charm sample selection	161
5.4	Event selection of $D^{*+} \rightarrow D^0 \pi^+$	163

5.4.1	D^0 secondary vertex selection	163
5.4.2	Primary π^+ selection	163
5.4.3	Monte-Carlo simulation	164
5.4.4	Signal extraction	165
5.5	Efficiency and background evaluation	167
5.5.1	Background evaluation	167
5.5.2	Detection efficiency	168
5.6	Results and conclusion	168
Bibliography		171
Summary		184
Samenvatting		186
Acknowledgments		188
Curriculum Vitae		190

Introduction

Several experiments have detected discrepancies in the fluxes of solar and atmospheric neutrinos. Neutrino oscillations provide a possible explanation for the measured differences. This dissertation presents one of the first generation, specific, neutrino-oscillation experiments using a man-made neutrino source, the CHORUS experiment at the European laboratory for particle physics (CERN)¹.

CHORUS was designed to look for $\nu_\mu \rightarrow \nu_\tau$ oscillation. Such oscillations can only happen if the neutrino eigenstates of the weak interaction do not coincide with the mass eigenstates. In the standard model of elementary particles and interactions, the neutrinos are assumed to be massless. However, there is no fundamental reason why neutrinos would be the only elementary fermion with no mass.

In 1998, the Super-Kamiokande experiment published clear evidence for neutrino oscillation [86]. That result unfortunately implied that CHORUS would not observe any oscillation signal. Since then, the emphasis of the CHORUS data analysis has shifted to exploiting the unique potential of the large number (about 100,000) recorded ν_μ interactions in the emulsion to analyse in detail the production and decay of charmed particles.

In Chapter 1, the key aspects of neutrino masses and oscillations are discussed. A description of the role of the CHORUS experiment is given, followed by an overview of the current status of neutrino-oscillation experiments. Chapters 2 and 3 describe the layout of the CHORUS experiment and the design of one particular sub-detector, the honeycomb tracker. The algorithms for reconstructing particle tracks contained in the emulsion images and their implementation are the subjects of Chapter 4. The last chapter, Chapter 5, reproduces a published charmed-meson production study. An introduction to the phenomenology of charm-quark production and fragmentation is given. This chapter highlights the possibilities of emulsion and current-day scanning techniques to reconstruct neutrino-interaction vertices and subsequent decay of short-lived particles.

As a final note, in this dissertation all particle masses, branching ratios, and other particle characteristics have been taken from the Particle Data Group review of 2004 [1].

¹The abbreviation stands for ‘Conseil Européen pour la Recherche Nucléaire’, the original committee that proposed to build the CERN laboratory

Chapter 1

Neutrino masses and oscillations

The neutrino was introduced as a hypothetical particle in 1930 by Pauli to solve the energy crisis in nuclear β -decay. Because of its extremely small cross-section to interact with matter, experimental neutrino physics only started much later. The basic interaction was first described by Fermi's theory and later unified in the electro-weak theory of the standard model of elementary particles.

When anomalously low fluxes of solar and atmospheric neutrinos were measured, neutrino oscillations were presented as a possible explanation for the missing neutrinos. In the last two decades, the experiments searching for neutrino oscillations have become more sensitive, culminating in clear evidence for oscillation of atmospheric neutrinos in 1998. This was followed in 2001 to 2003 by confirmation of the solar-model predicted flux for ^8B neutrinos and determination of the oscillation parameters causing the solar-neutrino disappearance.

This chapter is mostly dedicated to the physics of neutrino oscillations. It will introduce several aspects of the physics involving neutrinos, namely: the historic discoveries of the neutrinos; neutrino masses and mixing, including a simplified theory of oscillation; the parameter space explored by the CHORUS experiment; and an overview of the current results of oscillation experiments. This chapter concludes with a brief discussion of the current knowledge about neutrino mixing and an outlook to the remaining questions to be answered in the future.

1.1 Neutrino history

In the beginning of the 20th century, there was a fundamental problem in physics: energy seemed not to be conserved in certain radioactive decays. The energy of electrons emitted in β -decays did not have a definite value but instead a continuous spectrum. In 1930, Wolfgang Pauli, in his famous letter to the ‘radioactives’, postulated that an unknown neutral particle was present inside the nucleus which escaped unobserved in β -decays, carrying away the missing energy and momentum [2]. In 1934, Enrico Fermi developed a comprehensive theory of radioactive decays [3] involving Pauli’s hypothetical particle which he named in Italian the ‘little neutron’. Henceforth, this has stayed the particle’s name: the ‘neutrino’. The symbol used for this particle is the greek letter ν .

1.1.1 The electron neutrino

Fermi’s calculations showed that a neutrino interacts only very weakly (or equivalently very rarely) with matter. At the time, the neutrino–matter interaction cross-section was considered too small for the neutrino to be detected. It was only 23 years after its postulation by Pauli, that the first experiment to detect electron neutrinos was set up by Reines and Cowan [4]. At a nuclear reactor, a prolific source of $\bar{\nu}_e$ from β -decays of the neutron-rich fission products, anti-neutrinos were detected via the inverse β -decay process:¹

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


The detector consisted of liquid scintillator and dissolved cadmium-chloride. The light produced by fast electrons in the liquid scintillator was detected by photo-multiplier tubes. A neutrino interaction was identified by the delayed coincidence of two pulses: the first one originating from the two gammas from the e^+ annihilation; the second, required to arrive several microseconds later, being the signal from the neutron captured by a cadmium nucleus. Within a year the first signals of neutrinos were seen [5], followed by statistically stronger evidence in 1956 [6].

1.1.2 The muon neutrino

The muon, a heavier version of the electron, was first detected with cloud-chambers in cosmic-ray experiments around 1935 [7–11]. In its decay, $\mu \rightarrow e$, the emerging electron shows a continuous energy spectrum with similar polarization as in radioactive β -decay. Assuming that Fermi’s theory could also describe muon decay, there had to be three particles in the final state. It was natural to assume that the two undetected particles were both neutrinos. One of these neutrinos was associated with the electron (now known as the ν_e). The absence of electro-magnetic decay of the muon, $\mu \rightarrow e + \gamma$ via the process shown in Figure 1.1, was an indication for the existence of a second neutrino species associated with the muon, nowadays indicated as ν_μ .

¹The symbols ‘ p ’ and ‘ n ’ stand for the proton and neutron. The symbols ‘ d ’ and ‘ α ’ will be used for the deuteron and the helium-4 nucleus. Occasionally, when the focus is on isotope composition, the proton, deuteron and helium-4 nuclei will be given by the corresponding nuclear symbols: ^1H , ^2H and ^4He .

Figure 1.1: Unobserved possible Feynman diagram of $\mu \rightarrow e + \gamma$ as considered in the 1960s [12]. The interchanged heavy boson indicated by X is nowadays known as the W^\pm .



In 1962 at the Brookhaven national laboratory, using the first accelerator-generated neutrino beam ever built, muon neutrinos were identified as the neutral product in pion decays. It was known that the pion decayed to a muon and a neutrino, but the exact nature of this neutrino was unknown at that time. A 10-ton spark-chamber detector was placed behind a 17 m thick shield, made from recycled armour of scrapped battle ships. All particles produced by interactions of a 15 GeV proton beam in a beryllium target were stopped inside the shielding, only the neutrinos from pion decays could pass through it. If these neutrinos were the same as the neutrinos from nuclear β -decay, the experiment should have observed as much electrons as muons. Instead, the experiment observed several neutrino interactions in the detector, all of which were accompanied by a muon [13]. This result proved that the muon and electron were part of two separate families of leptons, each with their own neutrino. Instead of families, the terms generations or flavours² are also used.

1.1.3 The tau neutrino

After the discovery of the tau lepton in 1975 [14], it was assumed that also the τ had its associated neutrino, the ν_τ . Further indications came from the following arguments: the τ decay kinematics [15]; the number of light neutrino species measured to be 2.994 ± 0.012 [1] from the invisible Z^0 width [16]; the non-observation of ν_e or ν_μ induced τ production (see also section 2.1); the deduction from $\tau \rightarrow \rho$ decays that the ν_τ has spin 1/2 [17]; the measurement of the weak isospin of the τ^- with as result $T_3 = -1/2$ from neutral-current lepton-pair production ($Z^0 \rightarrow \ell^+ \ell^-$) at LEP [18] and the parity violation in the same process for polarized $e^+ e^-$ at SLC [19]. This last results implies that the τ is part of a weak-isospin doublet with a weak-isospin partner, the ν_τ .

It took until 2001 before the detection of a charged-current ν_τ interaction was made by the DONUT experiment at Fermilab [20]. This delay is due to the lack of abundant sources of ν_τ neutrinos. Only by placing a detector capable of identifying τ tracks behind a beam dump of a 800 GeV proton beam could four charged-current ν_τ interactions be identified. Inside the beam dump, only short-lived particles can decay before they re-interact. All other particles are absorbed or are swept out by the applied magnetic field. Behind the beam dump, the relative ν_τ flux is therefore enhanced. The ν_τ originate mostly from the decay $D_s^+ \rightarrow \tau^+ \nu_\tau$ or $D_s^- \rightarrow \tau^- \bar{\nu}_\tau$ (branching ratio of $[6.4 \pm 1.5] \%$) and the subsequent τ decay.

²Strictly speaking, flavour is used to indicate one of the six different quarks, one of the three different charged leptons or one of the three neutrinos.

The DONUT experiment had a similar setup as the CHORUS experiment (described in Chapter 2) with emulsion and electronic detectors. To analyze the DONUT emulsion data, the FKEN laboratory in Nagoya had to develop new, much faster, scanning hardware and new event reconstruction techniques than was foreseen for the CHORUS emulsion analysis. The new event reconstruction, called net-scan, can reconstruct the τ track as well as the primary and decay vertices. These new developments, explained in section 2.10.3, have later been adopted to do a much more extensive analysis of CHORUS emulsion data and made the charm-quark study in Chapter 5 possible.

1.2 The standard model

A complete discussion of the standard model of elementary particles and interactions (SM) can be found in many of the standard text-books on particle physics, for example Ref. 21. This section presents a brief overview of the particles and parameters of the SM.

1.2.1 Particles and forces

In Figure 1.2, the twelve fermions in the SM are shown: six quarks and six leptons organized in three families. There are thirteen bosons transmitting the three elementary forces included in the SM: the Higgs gives mass to particles; the electro-magnetic interaction is transmitted by the photon; the strong interaction is transmitted by eight coloured gluons; and the weak interaction is transmitted by the W^\pm and Z^0 bosons in charged-current and neutral-current weak interactions, respectively. The fourth elementary force, gravity, is not (yet) described by the SM.

All the forces are symmetric under the parity (coordinate inversion) P , the charge-conjugation C (particle \leftrightarrow anti-particle interchange) and the time-reversal T operators, except for the weak interaction. The W^\pm couples only to the left-handed chirality states of the fermions (right-handed for anti-particles), while the Z^0 couples differently to left and right-handed fermions. The weak interaction therefore violates P symmetry. Furthermore, the charged-current weak interaction is particular as it is the only interaction that transforms the fermion. The maximum violation of P symmetry by the weak interaction, proposed by Lee and Yang [22], was first demonstrated in the β -decay of spin-polarized ^{60}Co [23] and later in the kinematics of the pion and muon decays [24, 25]. More recently, it was discovered that also the combined CP symmetry, long thought to be conserved in weak decays, was violated in K^0 decays [26]. The W^\pm has a different coupling with each quark. The charged-current weak interaction changes one quark into another, violating both quark-type and family-number conservation laws. These laws are obeyed separately by all other interactions, including the neutral-current weak interaction. Strangely enough, the charged-current does obey family conservation laws in the lepton sector.

The quarks in Figure 1.2 are named: up ‘ u ’, down ‘ d ’, strange ‘ s ’, charm ‘ c ’, bottom ‘ b ’ and top ‘ t ’. The charged-current preferably couples the quarks within a family ($u \leftrightarrow d$, $s \leftrightarrow c$, $b \leftrightarrow t$), known as Cabibbo-enhanced transformations. The transformation of quarks between the first two families ($u \leftrightarrow s$, $d \leftrightarrow c$) are known as Cabibbo-suppressed transformations. Cabibbo discovered (before the charm quark was found) that the suppression in the decay rate of strange-quark containing particles, could be described by a single parameter, now known as the Cabibbo angle θ_C [27]. It turned out that the

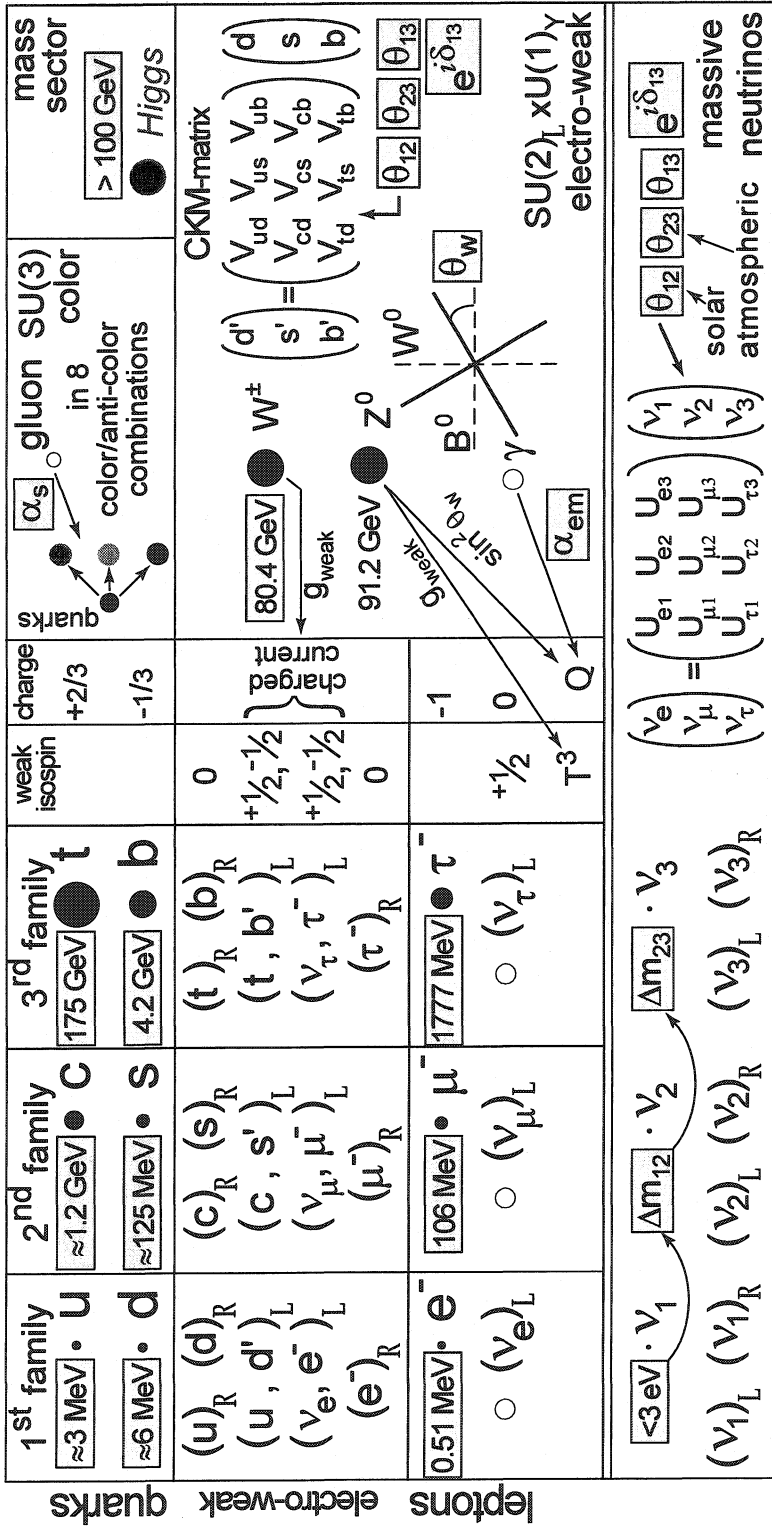


Figure 1.2: Schematic representation of the standard model of elementary particles and forces. The fermions from the three families are shown on the left. The four bosons responsible for the three basic interactions are shown on the right and their coupling with the fermions is indicated by the arrows. Open circles indicate massless particles, all other masses are parameters of the standard model. The parameters which can only be determined from measurements are shown inside the gray boxes. The Higgs particle is drawn in gray, because it has not yet been detected and its mass is only known indirectly. The extension of the standard model required by massive neutrinos is indicated below the line at the bottom.

coupling between the quarks in the first two families could be described by a rotation matrix with θ_C as the rotation angle. The W^\pm couples (in the approximation with only two-families) to the linear combination of $|d\rangle \cos \theta_C + |s\rangle \sin \theta_C$, with $\sin^2(\theta_C) \approx 0.05$. In general, the quark couplings by the charged-current interaction can be described by a 3×3 unitary matrix, known as the Cabibbo, Kobayashi, Maskawa (CKM) matrix [28]. The CKM matrix is parametrized by a complex phase for CP -violation and three mixing angles, one of which is θ_C . The CKM matrix describes how the quark eigenstates of the strong and electro-magnetic interaction connect to the charged-current weak-interaction eigenstates. The CKM matrix, shown on the right-hand side of Figure 1.2, is by convention applied to the lower half of the quark doublets.

Interestingly, the charm quark was proposed on symmetry and theoretical grounds, before its discovery in the J/Ψ resonance [29, 30]. A fourth quark was necessary to explain the anomalously slow decay of the K^\pm [31, 32]. In hind-sight, charmed mesons had already been seen earlier in emulsion events [33]. In a similar fashion, the third family was postulated before any quark or lepton of the third generation was discovered. Instead, the motivation was a desire to explain CP violation within the Cabibbo scheme. This requires a complex number in the rotation matrix, but with only two families such a term can always be eliminated by a suitable redefinition of the quark phases.

1.2.2 Parameters and constants

The SM has been very successful in predicting the outcome of many experiments and allows the calculation (in theory) of many high-energy particle interactions. However, the SM contains quite a few parameters (constants in the Lagrangian) which have to be measured by experiment and are an input to the SM. These parameters are indicated by the gray boxes in Figure 1.2 and include the masses of the particles and the coupling constants of the forces. The goal of the grand-unified theory (GUT) is to derive these parameters from basic principles and include the gravitational force. The unification of the electro-magnetic and weak interaction demonstrated the relation between several parameters which were considered independent before. The Weinberg weak-mixing angle relates the weak and electro-magnetic coupling constants and the W and Z mass. One could argue that the particular structure of the families (also its number) and interactions, which follows as a consequence of the $U(1) \otimes SU(2)_L \otimes SU(3)$ group structure of the SM, introduces another concept to be explained by theory. For example, the electron charge is by definition -1 because the electro-magnetic coupling strength α contains the basic electric charge e . However, there might be a deeper correlation between the fractional charge of the quarks and the family structure (3 quark colours with the electron and neutrino in a single $SU(5)$ multiplet with total charge zero).

1.2.3 Neutrinos

In the SM, neutrinos appear as the weak-isospin partners of the left-handed charged leptons ℓ (e , μ or τ). The coupling to the W^\pm bosons is via the weak charged-current:

$$J_\mu^{CC} = \bar{\psi}_{\nu_\ell} \gamma_\mu \frac{1}{2} (1 - \gamma^5) \psi_\ell \quad , \quad (1.1)$$

where γ_μ are the Pauli-Dirac spin matrices and ψ are the 4-component spinor wave-functions of the particles. The coupling of ψ_{ν_ℓ} to ψ_ℓ in equation (1.1) indicates that the

neutrino flavour is uniquely identified by the production of the corresponding charged lepton ℓ in charged-current interactions. The maximum parity violation is expressed by the term $(1 - \gamma^5)$ which projects out the left-handed chirality component of the ψ_ℓ wave-function.

The weak neutral-current coupling of the Z^0 boson to the fermions has the form

$$J_\mu^{NC} = \bar{\psi}_f \gamma_\mu \frac{1}{2} (c_V^f - c_A^f \gamma^5) \psi_f, \quad (1.2)$$

where c_V and c_A are the vector and axial-vector couplings. For fermions, they are given by $c_V^f = T_f^3 - 2Q_f \sin^2(\theta_W)$ and $c_A^f = T_f^3$. Here Q_f is the electric charge of the fermion f and T_f^3 is the third component of the fermion's weak isospin. The neutral-current reflects the mixing between the weak and electro-magnetic forces via the Weinberg or electro-weak mixing angle θ_W . For neutrinos $Q = 0$ and $T^3 = 1/2$ and equation (1.2) becomes

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

which has the same structure as equation (1.1), describing a coupling with the left-handed helicity states only. The right-handed leptons are $SU(2)_L$ weak-isospin singlets with $T^3 = 0$ and couple only via the $2Q_f$ term in the vector coupling c_V^f in equation (1.2). A right-handed neutrino has both $Q = 0$ and $T^3 = 0$ and therefore does not couple via either the electro-magnetic or the charged-current and neutral-current weak interactions. As all leptons are colour-less, a right-handed neutrino also does not interact via the strong force. Right-handed neutrinos are therefore completely sterile in the SM and the question of their existence is meaningless if they are massless as well.

1.2.4 Neutrino-matter interactions and cross-sections

To understand how experiments detect neutrinos, it is important to consider the interactions of neutrinos and anti-neutrinos with matter and how the cross-sections and kinematics of these reactions depend on energy. In this section, the basic principles will be outlined. As described above, neutrinos or anti-neutrinos can interact with matter either via the charged-current or the neutral-current weak interaction.




Figure 1.3: First order Feynman diagrams for scattering of ν_e , ν_μ and ν_τ and the anti-neutrinos on electrons. There are two diagrams for ν_e and $\bar{\nu}_e$ scattering because both the charged-current and neutral-current contribute. For ν_μ and ν_τ and their anti-particles only the neutral-current is possible, which is the same for all neutrino and anti-neutrino types.

Neutrino–electron scattering

The theoretically best understood reaction is between (anti-)neutrinos and electrons, because the electron can be considered as an elementary and free particle. The first order Feynman diagrams are shown in Figure 1.3. For ν_e and $\bar{\nu}_e$, the neutral-current and charged-current amplitudes interfere, while for ν_μ , $\bar{\nu}_\mu$, ν_τ , and $\bar{\nu}_\tau$ only the neutral-current contributes. The cross-section is therefore about a factor 6 larger for electron (anti-)neutrinos than for muon or tau (anti) neutrinos.

Furthermore there is a difference between neutrinos and anti-neutrinos due to the maximum parity violation of the weak interaction. For neutrinos the interaction has zero total angular momentum and the scattering angle is isotropic. For anti-neutrinos, the total angular momentum is one and, as can be easily seen from helicity arguments, backward scattering is suppressed. For anti-neutrinos the angular distribution is given by $1 - \cos\theta$, with θ the angle between the direction of the neutrino before and after the scattering. Consequently, the cross-section for anti-neutrino scattering is about a factor three smaller than for neutrino scattering. Figure 1.4 shows the cross-sections for neutrino–electron and anti-neutrino–electron scattering as function of the Weinberg weak-mixing angle θ_W [34]. A determination of the cross-sections for the different neutrinos was actually used to determine θ_W .




Figure 1.4: Standard model calculated cross-sections for neutrino and anti-neutrino scattering on electrons as function of the weak-mixing angle θ_W . The vertical line at $\sin^2 \theta_W = 0.23$ indicates the actual value of θ_W .

Neutral–current interactions

The neutral-current scattering of (anti-)neutrinos with nucleons or nuclei is similar to electro-magnetic electron–nucleon interactions, but with much smaller coupling strength. Instead of a photon, a (virtual) Z^0 boson is exchanged. Like the electro-magnetic interaction, neutral-current reactions have no intrinsic energy threshold. For high-energy neutrino experiments, one is mostly concerned with deep-inelastic scattering, where the momentum and energy transfer is so high that the Z^0 couples to a quark inside a nucleon and produces a shower of secondary particles. However, the SNO experiment (discussed in section 1.6.2) uses low-energy neutral-current interactions, where the Z^0 dissociates a deuterium nucleus with a threshold energy of 2.224 MeV.

Charged-current interactions

For charged-current (anti-)neutrino interactions with nuclei, the interactions can be categorized more or less by the energy and momentum transferred by the (virtual) W^\pm boson. First of all, charged-current interactions have a threshold energy because of the produced lepton in the $\nu_\ell \rightarrow W^\pm + \ell$ reaction. Typically, ν_μ (ν_τ) neutrinos below 106 MeV (1.8 GeV) do not interact via charged-current reactions.

At energies less than several MeV, the W^\pm interacts with the nucleus or nucleon as a whole and causes proton \leftrightarrow neutron transitions. This process — the inverse of radioactive β -decay — is used in low-energy solar-neutrino experiments. At medium energy ($E \approx 1$ GeV), the W^\pm knocks out the converted proton or neutron from the nucleus. These processes are called quasi-elastic interactions (QE). In an experiment, QE events are characterized by a single lepton track with possibly a short, low-energy, proton track. The cross-section for QE reactions as function of the neutrino energy E_ν is shown in Figure 1.5a [34]. It rises up to $0.9 \cdot 10^{-5}$ nb at 20 GeV above which it stays approximately constant due to the nucleon form-factor.

At energies above 1 GeV, deep-inelastic scattering (DIS) becomes dominant. In deep-inelastic scattering the W^\pm couples to a quark inside a nucleon. The struck quark then changes flavour according to the CKM matrix couplings and is knocked out of the nucleon. The struck quark and the remains of the nucleon lead to a particle shower inside a detector. The flavour-changing interaction leads to a significant production of strange and charmed quarks. The cross-section for DIS scattering of neutrinos and anti-neutrinos is shown in Figure 1.5b [34].

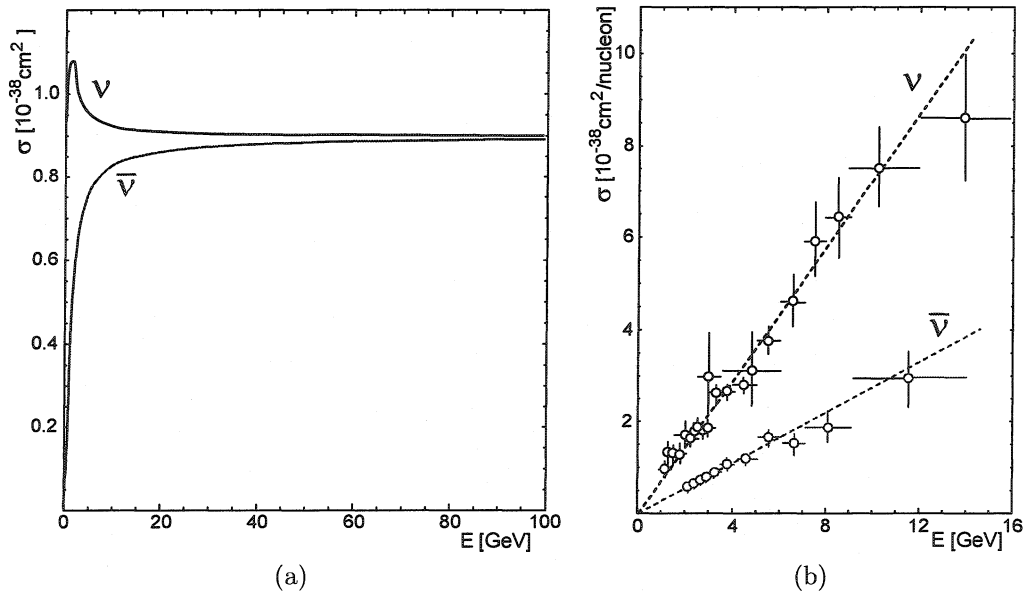


Figure 1.5: Cross-sections for (anti-)neutrino \leftrightarrow nucleon scattering for (a) quasi-elastic and (b) deep-inelastic scattering as a function of the neutrino energy.

1.2.5 Beyond the standard model

In 1998, the Super-Kamiokande experiment (see section 1.4.4) published convincing evidence for neutrino oscillation. As oscillations require neutrinos to have mass, the neutrino masses need to be added to the SM. Masses and oscillations of neutrinos can be accommodated in the SM Lagrangian, but they add at least another 7 parameters, as indicated at the bottom of Figure 1.2. Furthermore, right-handed neutrinos exist as well because helicity is not a conserved quantity for particles with a non-zero mass. The question is then what this right-handed neutrino state corresponds to. Depending on the exact form of the mass term, the right-handed neutrino state can correspond to the anti-neutrino or to an independent neutrino state. In the first case, the neutrino is a Majorana particle, in the second case it is a Dirac particle. The difference between these two cases will be briefly discussed in section 1.8.

1.3 Neutrino masses and oscillation

Neutrino oscillation was first proposed by B. Pontecorvo in 1957 [35,36]. In this section, neutrino-oscillation phenomenology is discussed within the context of the simplest case of two-flavour mixing. First, an overview of the current limits from direct neutrino-mass measurements is given.

1.3.1 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 one weak-interaction eigenstates ν_e , ν_μ and ν_τ respectively. Below, mass limits are quoted for these weak eigenstates since they correspond to the neutrino species investigated in the experiments.

Limit on m_{ν_e}

The square of the electron neutrino mass $m_{\nu_e}^2$ is measured in tritium β -decay experiments by fitting the shape of the electron spectrum near the endpoint. The current best limits come from the Mainz and Troitsk experiments [37,38]. Because the flux of electrons near the end-point of the tritium decay spectrum diminishes very rapidly, a high-intensity tritium source is needed. At the same time, though, no material can be introduced between the source and the spectrometer. To get enough statistics for a meaningful determination of the mass, these experiments require an extremely accurate spectrometer with very large acceptance. The PDG review gives a combined upper limit on the electron neutrino mass of $m_{\nu_e} < 3 \text{ eV}$ at 95 % confidence level [1].

Limit on m_{ν_μ}

The ν_μ 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, this yields the present best limit of $m_{\nu_\mu} < 190 \text{ keV}$ at 90 % confidence level [1,39].

Limit on m_{ν_τ}

Upper limits on m_{ν_τ} are obtained by studying the kinematics of hadronic τ -decays. The analysis involves the calculation of the hadronic invariant mass and fitting the energy spectrum of the hadronic system [40]. The best result comes from the ALEPH collaboration quoting an upper limit for the tau neutrino mass $m_{\nu_\tau} < 18.2 \text{ MeV}$ at 95 % confidence level [41].

Mass differences

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 limits. This happens because the neutrino mass limits follow the pattern of the corresponding leptons: $m_e \ll m_\mu \ll m_\tau$. As is shown in the following sections, neutrino oscillations provide the means to probe mass differences which may be smaller than the lowest neutrino mass limit. Mass differences determined in neutrino oscillation experiments therefore establish upper limits on the absolute mass of the heavier neutrinos very close to the upper limit on the mass of the lightest neutrino.

1.3.2 Neutrino oscillation

If neutrinos have mass, the weak eigenstates ν_e , ν_μ , and ν_τ do not necessarily coincide with the mass eigenstates ν_1 , ν_2 , and ν_3 and lepton flavour is no longer a conserved property of the neutrino. Hence flavour transitions may take place in vacuum. The original proposition of neutrino oscillation considers oscillation between $\nu \leftrightarrow \bar{\nu}$ (the muon and tau neutrino had not yet been discovered). The same idea applies, though, to flavour oscillations [42].

To illustrate the effect of mixing in the propagation of a massive neutrino, only the simplest case assuming two flavours and two mass eigenstates is considered here. Despite its simplicity, this is a basic assumption underlying the data analysis carried out by most oscillation experiments. The results of such an analysis can be presented in a single plot of a two-dimensional oscillation parameter space. In the more realistic situation with three flavours and three masses, the two-flavour analysis result still provides a good description in the limiting case where the third flavour consists only of the third mass eigenstate or in the case that the mass differences are very different such that the phase in equations (1.8) and (1.9) for one mass difference is small with respect to the other.

Vacuum oscillation

The flavour eigenstates ν_ℓ and $\nu_{\ell'}$ can be described as a superposition of the mass eigenstates ν_1 and ν_2 according to a 2×2 unitary mixing matrix:

$$\begin{pmatrix} \nu_\ell \\ \nu_{\ell'} \end{pmatrix} = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \end{pmatrix},$$

where θ is known as the mixing angle. At production (for example in a weak decay) neutrinos are created in a definite flavour eigenstate ν_ℓ . Assuming that a neutrino is

produced in a weak decay as ν_ℓ with momentum p , its wave-function at $t = 0$ is:

$$|\nu_\ell\rangle = \cos(\theta) |\nu_1\rangle + \sin(\theta) |\nu_2\rangle \quad . \quad (1.3)$$

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

$$|\nu_\ell(t)\rangle = \cos(\theta) e^{-iE_1 t} |\nu_1\rangle + \sin(\theta) e^{-iE_2 t} |\nu_2\rangle \quad . \quad (1.4)$$

Here

$$E_i = \sqrt{p^2 + m_i^2} \approx p + \frac{m_i^2}{2p} \quad , \quad (1.5)$$

assuming that the neutrino is highly relativistic such that $p^2 \gg m_i^2$. From equation (1.5), it follows that

$$E_1 - E_2 \approx \frac{\Delta m^2}{2p} \approx \frac{\Delta m^2}{2E} \quad , \quad (1.6)$$

with $\Delta m^2 = m_1^2 - m_2^2$. The probability that the neutrino will still be found as ν_ℓ after traveling some distance $L = ct$ is then given by:

$$P_{\nu_\ell \rightarrow \nu_\ell}(L) = |\langle \nu_\ell(0) | \nu_\ell(t) \rangle|^2 \quad . \quad (1.7)$$

Substituting the expressions from equations (1.3), (1.4), and (1.5) in equation (1.7) and using (1.6), one obtains

$$P_{\nu_\ell \rightarrow \nu_\ell}(L) \approx 1 - \sin^2(2\theta) \sin^2\left(\frac{\Delta m^2 L}{4E}\right) \quad . \quad (1.8)$$

Similarly, the transition probability $P_{\nu_\ell \rightarrow \nu_{\ell'}}$, which is the probability for the neutrino to interact as $\nu_{\ell'}$ rather than ν_ℓ , is given by:

$$P_{\nu_\ell \rightarrow \nu_{\ell'}}(L) = \sin^2(2\theta) \sin^2\left(\frac{\Delta m^2 L}{4E}\right) \quad . \quad (1.9)$$

The quantity $P_{\nu_\ell \rightarrow \nu_{\ell'}}$ given in equation (1.9) is known as the oscillation probability. From the above equations, one can define the oscillation length

$$L_{osc} \equiv 4\pi \frac{E}{\Delta m^2} \quad , \quad (1.10)$$

which is the travel distance needed to go from ν_ℓ to a maximally mixed state 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 then given respectively by

$$P_{\nu_\ell \rightarrow \nu_{\ell'}}(L) = \sin^2(2\theta) \sin^2\left(1.267 \frac{\Delta m^2 L}{E}\right) \quad (1.11)$$

and

$$L_{osc} = \frac{\pi E}{1.267 \Delta m^2} \quad . \quad (1.12)$$

In the case that $L_{osc} < \Delta L$, with ΔL the range of detected flight lengths, or if there is a wide spectrum of neutrino energies in the source, the L/E ratio can take on many values at the detector. In this case the oscillation pattern is lost, hence the probabilities in equations (1.8) and (1.9) become equal to their time averaged values

$$P_{\nu_\ell \rightarrow \nu_\ell} = 1 - \frac{1}{2} \sin^2(2\theta) \quad (1.13)$$

and

$$P_{\nu_\ell \rightarrow \nu_{\ell'}} = \frac{1}{2} \sin^2(2\theta) \quad , \quad (1.14)$$

respectively. This happens in most accelerator generated neutrino beams (section 2.2) where the distance between the source and the experiment is of the same order as the length of the decay tunnel. It is also the case for atmospheric and (most) solar neutrinos (see section 1.4) which have both wide energy spectra and a large spread in flight lengths.

Matter enhanced oscillation

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

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

takes place only via a neutral-current interaction if $\ell' \neq \ell$, while if $\ell' = \ell$ the transition amplitudes of the neutral and charged-current interfere. When neutrinos propagate in matter, they undergo elastic interactions with electrons. For large regions, this can lead to coherent effects in the oscillation. The elastic forward scattering can be described in terms of a potential energy

$$V = V_Z + V_W \quad ,$$

where V_Z and V_W are the potential energies due to Z^0 and W^\pm 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 at rest with constant electron density N_e , the potential V_W is given by:

$$V_W = \sqrt{2}G_F N_e \quad ,$$

where G_F is the Fermi weak coupling constant.

Wolfenstein, Mikheyev, and Smirnov [43, 44] pointed out that the effect of coherent forward scattering can change the oscillation pattern of neutrinos traveling through matter. For such a mixing, with given Δm^2 and $\sin^2(2\theta)$ in vacuum, the observed oscillation pattern in matter could be described by the effective values

$$\Delta m_{\text{eff}}^2 = \Delta m^2 \sqrt{(a - \cos(2\theta))^2 + \sin^2(2\theta)} \quad (1.15)$$

and

$$\sin^2(2\theta_{\text{eff}}) = \frac{\sin^2(2\theta)}{(\cos(2\theta) - a)^2 + \sin^2(2\theta)} \quad , \quad (1.16)$$

where

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

In the limit that $a \rightarrow \cos(2\theta)$, one gets as effective oscillation parameters from equation (1.15) and (1.16):

$$\lim_{a \rightarrow \cos(2\theta)} \Delta m_{\text{eff}}^2 = \Delta m^2 |\sin(2\theta)| \quad (1.17)$$

and

$$\lim_{a \rightarrow \cos(2\theta)} \sin^2(2\theta_{\text{eff}}) = 1 \quad , \quad (1.18)$$

describing an effective maximal mixing in matter. This resonance behaviour, the so-called MSW effect, may take place if N_e has the appropriate value. In this case the observed Δm_{eff}^2 is smaller than the vacuum Δm^2 .

1.3.3 Oscillation detection methods

Neutrinos created in a flavour eigenstate will oscillate to other eigenstates during propagation. Such oscillation can in principle be detected in two ways, appearance and disappearance.

Appearance experiments

An experiment which detects a neutrino flavour absent in the source is known as an appearance experiment. The sensitivity of these experiments depends on the number of positive events with respect to the expected background. The most convincing oscillation signal is to unequivocally detect a neutrino flavour from its charged-current reaction in a channel where there is no background.

Because of the energy-threshold in charged-current reactions, these experiments are usually only sensitive to oscillation with high Δm^2 unless they use very long baselines. On the other hand, an appearance experiment can detect small mixing angles if it can acquire a large event sample. The last requirement usually implies a short baseline to create a high neutrino flux and therefore reduces further the sensitivity to low Δm^2 .

Disappearance experiments

In a disappearance experiment, one tries to detect a deficit in the neutrino flux with respect to the theoretical flux for a given source. Most experiments with natural neutrino sources are disappearance experiments, because the neutrino energy is below the threshold for muon or tau production. One exception is the SNO experiment's neutral-current measurement, discussed in section 1.6.2.

These experiments are only sensitive to large mixing angles, because a small deficit will be drowned in the statistical and theoretical uncertainties of the neutrino flux. On the other hand, they can detect very small Δm^2 by tuning the neutrino energy and baseline such that $L/E \approx 1$. The most convincing evidence for neutrino oscillation in a disappearance experiment is to measure the characteristic sinusoidal disappearance of the neutrino flux as function of L/E .

Parameter extraction

In general, for any oscillation experiment, the only observable is the number of neutrino interactions. Depending on the experiment, these can be categorized by the neutrino flavour (for charged-current interactions) and, in some cases, the energy. In some circumstances, the incoming neutrino direction can be reconstructed. In such a case, an analysis as function of L/E becomes possible, for example for high-energy atmospheric neutrinos.

From the number of interactions for ν_ℓ and $\nu_{\ell'}$ (the last only for appearance experiments), constraints can be placed on the range of allowed values [45] of the oscillation probabilities $P_{\nu_\ell \rightarrow \nu_\ell}$ (disappearance) and $P_{\nu_\ell \rightarrow \nu_{\ell'}}$ (appearance), usually as function of energy. The interpretation of this experimental value in terms of theoretical parameters, such as neutrino masses and mixing angles, depends on the underlying theoretical model and the experimental constants, in particular the neutrino-energy spectrum and the flight-length distribution. Most experiments restrict the analysis to the simplest case of two-flavour mixing, which has two mixing parameters as given in equation (1.11), $\sin^2(2\theta)$ and Δm^2 . The experimental results can then be presented as excluded or favoured regions in the $\sin^2(2\theta)$, Δm^2 plane in a simple 2-D contour plot, see for example Figure 1.11.

1.4 Neutrino oscillation hints

In this section, some of the hints that something was going on with neutrinos will be discussed. The recent experimental results will be discussed in section 1.7 in which a summary of the current neutrino-oscillation results is presented.

1.4.1 Cosmology

The contributions from different particles and forces to the total energy density of the universe determines its evolution in time. These contributions are usually expressed as fractions $\Omega_i \equiv \rho_i/\rho_c$ of the critical energy density $\rho_c = 3H_0^2/(8\pi G_N)$ required for a flat universe. At the critical density, the expansion rate, given by the Hubble constant H_0 , is exactly balanced by the gravitational attraction of which the strength is given by Newton's constant G_N . Figure 1.6 gives a schematic view of the components of the total energy density.

The requirement for dark matter in Figure 1.6 is based, among others, on the discrepancy between the amount of luminous matter and its measured velocities in spiral galaxies. The rotational velocities of stars around a galaxy's center are too high to keep the galaxy together based on the gravitational influence of the visible matter. Assuming that the structure of galaxies is long-lived, more matter must be present inside the orbits of the stars to keep the galaxy together. Similar measurements on the relative motion of galaxies inside galaxy clusters or super-clusters indicate the presence of vast amounts of non-visible matter. However, dark matter has been an ever shifting target over the last 30 years. Different quantities of dark-matter have been proposed with yet again different fractions of hot (relativistic and therefore exerting a positive pressure) and cold (non-relativistic) matter. The cosmological constant causes the expansion to speed up and was reintroduced because of supernova type-I measurements.

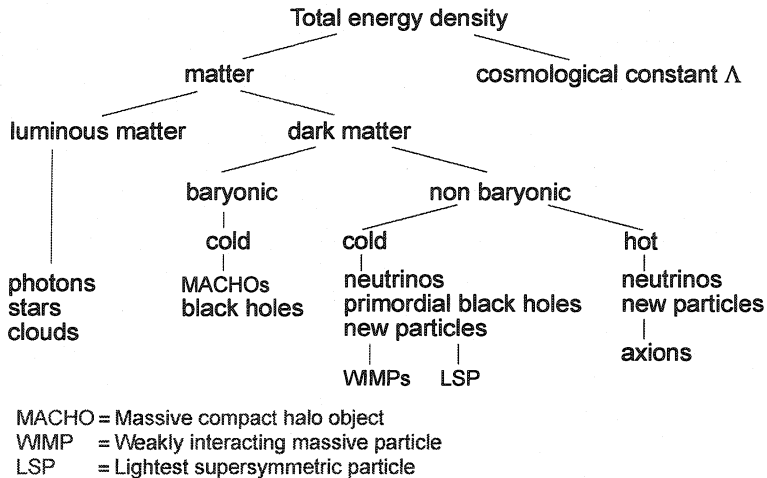


Figure 1.6: Schematic view of the total energy density of the universe, showing the different contributions to the energy density.

The amount of dark-matter required to explain the velocities of stars far exceeds the visible baryonic matter in the universe. From all proposed contributions to the dark-matter in Figure 1.6, only the neutrino is known to exist.³ Neutrinos are highly abundant in the universe with about 100 neutrinos of each flavour per cubic-centimeter. For a long time, massive neutrinos were considered a very good candidate for dark matter. Nowadays, the current results on neutrino masses from oscillation experiments and upper limits on the total neutrino mass from the microwave background measurements exclude that neutrinos are a substantial part of the dark matter.

Only in recent years, with the accurate measurements of the cosmic microwave background anisotropy by the WMAP satellite, have the relative contributions of matter and other-fields to the total density of the universe settled down a bit. The current WMAP results give a cold dark-matter contribution to the energy content of the universe of 23% plus 4% normal (baryonic) matter of which again only about 1/10th is in visible (luminous) matter [51]. The total neutrino mass is now limited to less than 0.68 eV. For more details, the reader is referred to reviews in the literature, like Refs. 1, 52, 53.

1.4.2 Solar neutrinos

When it became clear in the beginning of the 20th century that the age of the solar system should be measured in billions of years, it was realized that gravitational contraction could not have sustained the sun's energy output for this time. When it also became known that the sun is mainly composed of hydrogen and helium (discovered in solar spectral lines), it was sir Arthur Eddington who proposed that nuclear fusion could be the sun's energy source. It took until 1938 before a theory of the nuclear reactions within the sun was developed by Bethe and Chritchfield [54, 55]. The sun's main energy source is provided by the fusion of four protons plus two electrons into one tightly-bound helium-4 nucleus and two electron neutrinos, liberating 26.73 MeV in energy.

³Some MACHOs [46, 47] have been observed from micro-lensing effects, but far too few to contribute significantly to the baryonic dark matter [48–50].

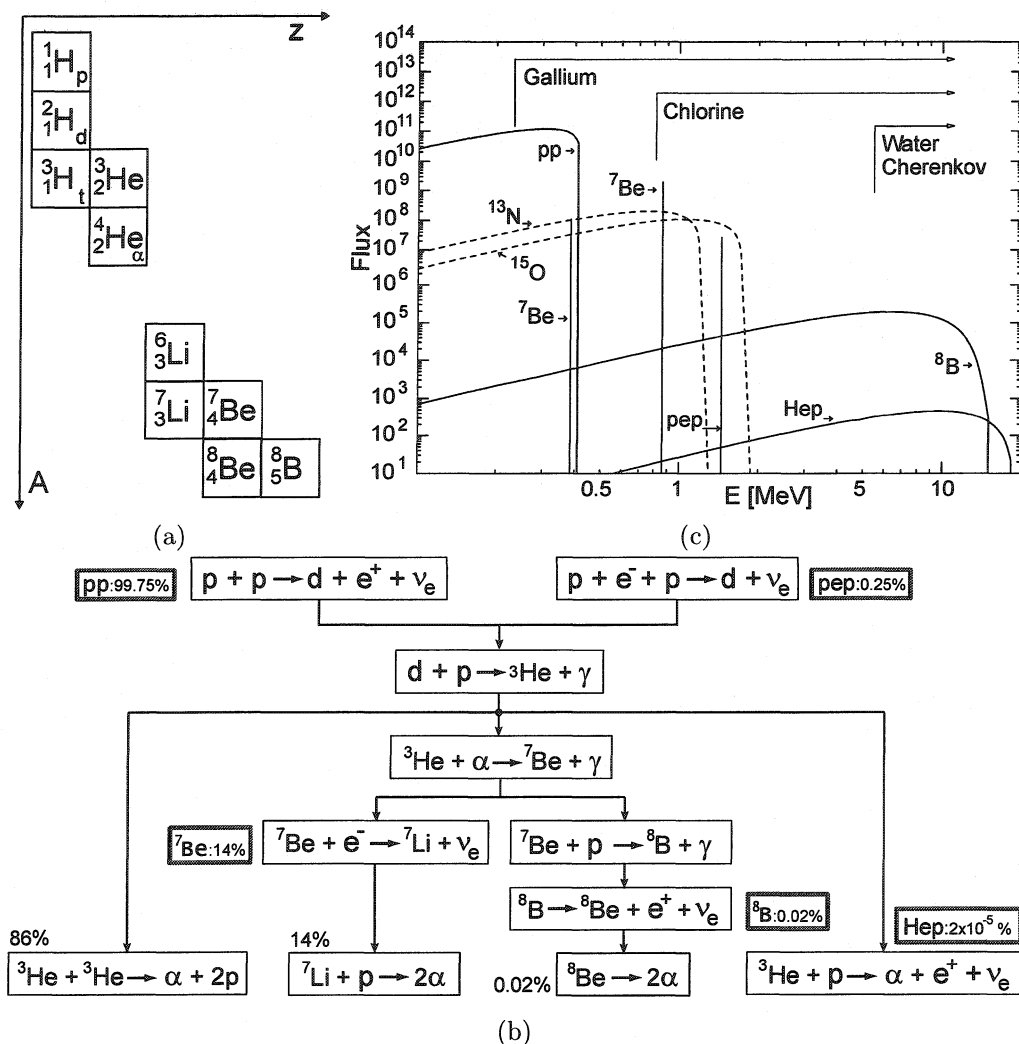


Figure 1.7: (a) The different isotopes involved in the solar fusion reactions. (b) Schematic diagram of the different pathways for the same overall fusion reaction: $4p \rightarrow \frac{4}{2}\text{He} + 2e^+ + 2\nu_e$. The relative abundance for the ν_e producing reactions are also given. (c) The ν_e spectrum.

The fusion actually proceeds via four different paths with several intermediate steps, as illustrated in the diagram of Figure 1.7b. The relative strength of the paths depends strongly on the exact energy levels of the nuclei and on the temperature and density in the interior of the sun. These dependencies lead to uncertainties in the overall ν_e spectrum. This spectrum, shown in Figure 1.7c, consists of line-spectra from electron-capture reactions and continuous spectra from $p \rightarrow n$ reactions. A less important contribution is due to fusion reactions catalyzed by ^{12}C , known as the **CNO cycle**. In the standard solar model (SSM), the neutrino flux is calculated from the equation of state. The SSM has been verified by checking its sound-velocity predictions with measurements from helio-seismology which agree to within 0.1% RMS [56, 57].

As 99.75% of the fusion paths start with the pp reaction, its ν_e flux is strongly correlated with the total energy production in the sun, which in turn is directly connected with the observed total luminosity. Therefore, the pp ν_e flux should be proportional to the sun's luminosity. However, the diffusion time for photons and energy from the sun's core to its surface is about $1.7 \cdot 10^5$ and $3 \cdot 10^7$ years, respectively [58, 59], while the neutrinos escape immediately. Since the sun is currently in a very stable phase of its stellar evolution, the equilibrium between fusion energy and luminosity is not influenced by the time-lag between photons and neutrinos.

As can be seen in Figure 1.7c, most of the solar neutrinos (99.9%) have an energy below 1 MeV. As a consequence, these solar neutrinos can only be detected with a reaction with a threshold lower than 1 MeV. One process sensitive to such low energy ν_e is the inverse β -decay reaction on some nucleus X: ${}^A_Z\text{X} + \nu_e \rightarrow {}^A_{Z+1}\text{Y} + e^-$. The energy threshold for this reaction is lowered by the mass difference between neutron and proton and the difference in nuclear binding energies. Experiments for solar-neutrino detection based on the ${}_Z\text{X} \rightarrow {}_{Z+1}\text{Y}$ reaction require a large amount of the target nucleus X and therefore its (relative) isotope abundance on earth must be high. The two nuclei involved, X and Y, must also fulfill other criteria. First, the mass-difference $M(\text{Y}) + 511 \text{ keV}/c^2 - M(\text{X})$ must be small or negative to have an energy threshold below the typical energy of the solar neutrinos. Second, in order to count the number of neutrino interactions, the created atoms Y must be counted and therefore the elements must be separable by some means. Since the number of target atoms must be very large, direct separation by, for example, a mass spectrometer is not possible. Another physical or chemical difference is required for the separation. If counting relies on the radioactive decay of nucleus Y, its half-life must be long enough so that its concentration can grow in the target before decaying and short enough so that its decays can be measured after extraction. Experiments based on this technique are known as radio-chemical experiments.

The first experiment to measure solar neutrinos, Homestake, was started by Davis Jr. in the 1960s using the reaction ${}^{37}\text{Cl} + \nu_e \rightarrow {}^{37}\text{Ar} + e^-$ [60, 61]. This reaction has a threshold energy of 814 keV and is therefore only sensitive to the ${}^7\text{Be}$ and ${}^8\text{B}$ neutrinos, as indicated in Figure 1.7c. The chlorine was contained in cleaning fluid (C_2Cl_4) from which the noble gas argon can be extracted by flushing with helium. The isotope ${}^{37}\text{Cl}$ has a natural relative abundance of about 24% and ${}^{37}\text{Ar}$ decays by electron-capture with a half-life of 34.8 days.

A reaction with a lower threshold (233 keV) is between gallium and germanium: ${}^{71}\text{Ga} + \nu_e \rightarrow {}^{71}\text{Ge} + e^-$ [62]. As germanium is normally a solid, they are actually used as liquid metallic gallium or dissolved GaCl_3 and gaseous GeCl_4 . The isotope ${}^{71}\text{Ga}$ has a natural relative abundance of about 40% and ${}^{71}\text{Ge}$ decays by electron-capture with a half-life of 11.4 days. Even though the pp flux is dominant, the above reaction is only caused for about 55% by pp. The ${}^7\text{Be}$ and ${}^8\text{B}$ neutrinos contribute for approximately 27% and 10%, respectively, because of their higher energy. Two experiments, Gallex and SAGE, started measuring this reaction in 1990. Gallex used 30 tons of gallium in solution of GaCl_3 , while SAGE used from 30 to 60 tons of liquid metallic gallium.

The results of radio-chemical experiments are usually expressed in solar-neutrino units (SNU) which corresponds to one neutrino capture per 10^{36} target atoms per second. The predictions for the chlorine and gallium experiment from one solar model are given in Table 1.1. For 30 ton of material, the quoted interaction rates correspond at most to some tens of created Y atoms after exposure for two to three half-lives.

reaction	Standard Solar Model		
	flux [$\text{cm}^{-2}\text{s}^{-1}$]	^{37}Cl [SNU]	^{71}Ga [SNU]
pp	$5.94 \cdot 10^{10}$	-	69.6
^7Be	$4.80 \cdot 10^9$	1.15	34.4
pep	$1.39 \cdot 10^8$	0.2	2.8
^8B	$5.15 \cdot 10^6$	5.9	12.4
hep	$2.10 \cdot 10^3$	0.0	0.0
cno	$1.14 \cdot 10^9$	0.5	9.8
Total		$7.7^{+1.2}_{-1.0}$	129^{+8}_{-6}

Table 1.1: Fluxes and interaction rates for the chlorine and gallium experiments from the solar model of Ref. 63.

The Homestake experiment has been running almost continuously from 1970 until 1994. From the beginning its results were lower than expected from solar models. This low solar-neutrino flux has since been known as the “solar-neutrino problem”. Homestake measured a flux of below 3 SNU [64, 65], which is about $1/3^{\text{th}}$ of the expected flux. The difficulty of this type of experiment is illustrated by the fact that in the 24 years that Homestake has run, only an estimated 2200 argon atoms were produced in the tank which contained $2 \cdot 10^{30}$ chlorine atoms!

Both Gallex and SAGE needed more than a year to reduce an unexpected background from long-lived ^{68}Ge (half-life 271 days) created from the gallium by cosmic-rays when it was still at the surface. In the initial five runs, SAGE detected almost no signal [66] and set an upper limit of 79 SNU (90 % confidence level). Later results of SAGE were in better accordance with the Gallex measurement of $83 \pm 19(\text{stat}) \pm 8(\text{sys})$ SNU [67], which was the first detection of pp neutrinos from the sun.

The upgraded Kamiokande nucleon decay experiment, Kamiokande-II (KamII), started measuring solar neutrinos in 1986. In contrast to the radio-chemical experiments, KamII could measure the interactions in real-time and measure the direction of the incoming neutrinos. However, because of the energy threshold of 7.5 MeV, it was only sensitive to the high-energy component of the solar spectrum, namely the ^8B neutrinos. Detection was done using the neutrino–electron scattering reaction in which the electron direction is correlated with the incoming neutrino direction. The KamII detection technique was later used by the, much larger, follow-up experiment, Super-Kamiokande. A detailed description can be found in section 1.6.1. In the KamII measurement [68], the correlation of the signal’s direction with the position of the sun proved the existence of solar neutrinos (see also Figure 1.18). KamII measured a ratio for the ^8B flux of $[46 \pm 5(\text{stat}) \pm 6(\text{sys})]$ % with respect to the solar model of Ref. 69.

The results from all the solar-neutrino experiments could be reconciled if neutrino oscillation is responsible for the disappearance of solar neutrinos. Because of the high density of the sun, the MSW effect could effectively lead to maximum mixing which explained the low rates for Homestake and KamII. Time variation measurements (day versus night for MSW effects in the earth or seasonal for baseline variation) by Homestake and KamII and energy-dependence measurements by KamII further excluded certain regions in the oscillation parameter space [70]. Four regions in the $\Delta m^2, \sin^2(2\theta)$ plane could explain all the data. In three of these the MSW effect is important and one is due to pure vacuum oscillation [71, 72].

1.4.3 Atmospheric neutrinos

Atmospheric neutrinos are generated by the interaction of cosmic rays with the earth's atmosphere. Cosmic rays were discovered in the beginning of the last century. Electrically charged objects lose their charge even though the air surrounding them is an insulator, so something must be ionizing the air. First thought to be natural radioactivity, it was Hess who showed in 1912 that the charge leakage increases with altitude and is therefore due to something arriving from space [73]. The cosmic rays detected at the earth's surface are actually the remainders (mostly muons) of showers generated when high-energy primary particles, mostly protons and helium nuclei, strike the top of the atmosphere. Most of the secondary particles in the shower are pions which either reinteract or decay before reaching the earth's surface. Typical cosmic-ray showers start at an altitude between 25 km and 20 km. Because of the low atmospheric density, most of the low-energy ($E < 10$ GeV) π^\pm ($c\tau = 7.8$ m) and K^\pm ($c\tau = 3.7$ m) mesons produced in the showers decay before interacting:

$$\begin{aligned} \pi^+ &\rightarrow \mu^+ + \nu_\mu & , & & \pi^- &\rightarrow \mu^- + \bar{\nu}_\mu & \text{ (branching ratio 99.99\%)} & , \\ K^+ &\rightarrow \mu^+ + \nu_\mu & , & & K^- &\rightarrow \mu^- + \bar{\nu}_\mu & \text{ (branching ratio 63.51\%)} & . \end{aligned} \quad (1.19)$$

Most of the low-energy muons ($c\tau = 659$ m) also decay before reaching the ground:

$$\begin{aligned} \mu^+ &\rightarrow e^+ + \nu_e + \bar{\nu}_\mu & \text{ (branching ratio 100\%)} & . \\ \mu^- &\rightarrow e^- + \bar{\nu}_e + \nu_\mu \end{aligned} \quad (1.20)$$

As the neutrinos hardly interact with the earth's matter, neutrinos will be incident on a detector from all directions. Atmospheric neutrinos therefore have flight lengths from about 15 km from directly overhead to 13000 km from straight below.

The observed energy spectrum of incoming primary particles reaches up to 10^{20} eV. Above $5 \cdot 10^{19}$ eV, the GZK cut-off due to interactions with the cosmic microwave background comes into play [74, 75]. The incoming flux of primary particles can be considered isotropic, but at energies below several GeV the earth's magnetic field influences the primary flux. Consequently, atmospheric neutrinos cover a wide range of energies and the flux is isotropic above a few GeV, but not below. Furthermore, the sun's solar wind also influences the primary flux and leads to fluctuations coupled to the eleven year sunspot cycle.

Early Monte-Carlo calculations, predicted the flux of each neutrino species with an error of about 20%. However, many of the uncertainties cancel in the flux ratio


$$r = \frac{\phi(\nu_\mu) + \phi(\bar{\nu}_\mu)}{\phi(\nu_e) + \phi(\bar{\nu}_e)} , \quad (1.21)$$

which at low energies should equal 2, according to equations (1.19) and (1.20). The ratio r is predicted with an error of about 5%. The value of r and its energy dependence from recent calculations is plotted in Figure 1.8 [101]. At higher energies ($E > 3$ GeV) the probability of muon decay is reduced and r rises.

Most oscillation searches using atmospheric neutrinos are performed by looking for possible deviations of the measured flux ratio compared to the Monte-Carlo. This is usually expressed as the double ratio

$$R_{\text{atm}} = \frac{r_{\text{data}}}{r_{\text{Monte-Carlo}}} , \quad (1.22)$$

Figure 1.8: The flux ratio of $\nu_\mu + \bar{\nu}_\mu$ to $\nu_e + \bar{\nu}_e$ versus neutrino energy averaged over all zenith and azimuth angles. Solid, dashed and dotted lines show the predictions by three different flux models.



In this, the detection efficiency for each neutrino species in a particular experiment is taken into account. The low flux of the atmospheric neutrinos makes these measurements difficult as they require a large target mass and very good rejection of the backgrounds, like the much higher atmospheric-muon flux due to the decays given in equation (1.19). The measured values of R_{atm} in the early 1990s did not give a clear picture, the water-Cherenkov experiments yielded values below 1, while the iron calorimeters showed no difference from the expectations.

Early atmospheric-flux results

The first measurement came from the Kamioka Nucleon Decay Experiment (Kamiokande) in Japan, which measured a value of R_{atm} of 0.6 [76]. Later measurements by the water-Cherenkov detectors, all at energies below 1.5 GeV, also yielded values of R_{atm} less than 1 [77–80]. The detection and analysis of atmospheric-neutrino interactions in the water-Cherenkov detectors is explained in sections 1.6.1 and 1.7.1. The 1992 result of the Kamiokande-II experiment is based on 310 fully-contained, single-ring, quasi-elastic, neutrino interactions, classified as 159 electron-like (e -like) and 151 muon-like (μ -like) events [77]. Taking into account the detection efficiencies for ν_e and ν_μ interactions, the measured ratio was

$$R_{\text{atm}} = 0.60_{-0.06}^{+0.07}(\text{stat}) \pm 0.05(\text{sys}) ,$$

which is more than 4σ different from one. Similarly, the 1992 IMB result [80], is based on 610 events and gave $R_{\text{atm}} = 0.70$ and 2.6σ below one. Looking at the rate of ν_e and ν_μ interactions individually, the experiments indicated that the measured deficit was mostly due to a low $\nu_\mu + \bar{\nu}_\mu$ rate. Both experiments performed a cross-check on the muon identification by tagging stopped muons from the $\mu \rightarrow e$ decay, with essentially identical results. Both experiments did not see any dependence on either momentum, $100 \text{ MeV}/c < p < 1500 \text{ MeV}/c$, or on neutrino flight length (zenith angle). At the same time, the iron-calorimeters, NUSEX and FREJUS, measured values of R_{atm} compatible with 1 [81–84]. The first measurement for higher-energy muons events ($E > 1 \text{ GeV}$) by IMB, based on upward-going muons, showed also no difference from the Monte-Carlo expectation [85].

1.4.4 The 1998 Super-Kamiokande result

In 1998, the Super-Kamiokande experiment, described in detail in section 1.6, published its high-statistics measurements of atmospheric neutrinos [86]. Not only did these results confirm the earlier low value of R_{atm} and extended it to higher energies, but they also showed a clear dependence of the ν_μ flux on zenith angle and energy. These dependencies, shown in Figure 1.9, are fitted very well with the expectation of $\nu_\mu \rightarrow \nu_\tau$ oscillation using $\sin^2(2\theta) = 1$ and $\Delta m^2 = 2.1 \cdot 10^{-3} \text{ eV}^2$ [86].

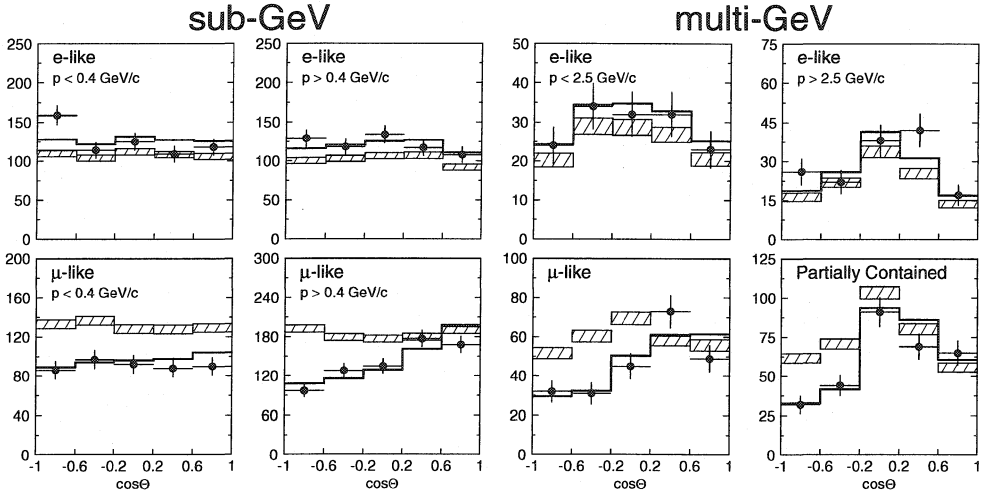


Figure 1.9: Zenith angle distribution for e -like and μ -like atmospheric-neutrino events in Super-Kamiokande. The partially contained events are mostly high-energy ν_μ interactions, where the muon leaves the inner detector. The points show the data, the dashed boxes the Monte-Carlo expectation without oscillation and the solid lines the best-fit expectation for $\nu_\mu \rightarrow \nu_\tau$ oscillation with $\sin^2(2\theta) = 1$ and $\Delta m^2 = 2.1 \cdot 10^{-3} \text{ eV}^2$.

At lepton momenta below $0.4 \text{ GeV}/c$, there is basically no correlation between the lepton and neutrino direction, therefore there is no zenith angle dependence and only an overall ν_μ disappearance can be seen. The zenith angle is connected to the neutrino flight length L , with $L \approx 13 \text{ km}$ for $\cos \theta = 1$ and $L \approx 13000 \text{ km}$ for $\cos \theta = -1$. With the fitted Δm^2 , the oscillation length is $L_{\text{osc}} \approx 1200 E/\text{GeV km}$. For $p < 0.4 \text{ GeV}/c$ and maximum mixing, one then expects to see 75% of the expected ν_μ flux, the average of 50% fully-oscillated and 100% non-oscillated fluxes, as is indeed the case for the μ -like, $p < 0.4 \text{ GeV}/c$ events. At higher energies, where there is good correlation between lepton and neutrino direction, one can clearly see the ν_μ disappearance at long flight lengths, reaching 50% of the expected flux for $\cos \theta = -1$. There is no disappearance of ν_e , nor is there any clear sign of ν_e appearance where ν_μ disappears. Therefore, no oscillation is expected from ν_μ to ν_e (also excluded by reactor and solar-neutrino results). Thus the oscillation is expected to be from ν_μ to ν_τ .

The oscillation hypothesis can be better analysed looking at the L/E dependence. The measurement of e -like and μ -like events as function of L/E is depicted in Figure 1.10 [86]. Here it is clear that e -like events are unaffected, while μ -like events show the characteristic fall-off with L/E , going to an average of 1/2 when L/E is much larger than the oscillation length.




Figure 1.10: Super-Kamiokande events plotted as function of L/E show the characteristic disappearance of μ -like events with larger L/E , while the e -like events are unaffected. The dashed lines show the best $\nu_\mu \rightarrow \nu_\tau$ oscillation fit to the data.

Since 1998, the Super-Kamiokande results have been confirmed by continuing and new measurements. An overview of the current neutrino-oscillation experiments and results is discussed in the sections 1.6 and 1.7.

1.5 The CHORUS oscillation search

The lower than expected neutrino fluxes from natural sources, led to several searches for neutrino oscillations in artificial sources, like nuclear reactors and accelerator beams. These sources have the advantage that their neutrino fluxes are more controllable and better understood, thus removing the uncertainties from modeling either the solar nuclear processes or the atmospheric cosmic-ray showers.

1.5.1 Motivation

Many of the earlier accelerator experiments were not primarily designed to study neutrino oscillation. A dedicated experiment could substantially expand the excluded parameter region or actually detect neutrino oscillation. Accelerator neutrino beams, described in Chapter 2, consist primarily of muon neutrinos with a few percent electron neutrinos and are quasi-free of tau neutrinos. It is therefore natural to try to detect the oscillation of a muon neutrino into a tau neutrino as this gives a basically background free result.

Around 1990, a proposal to search for $\nu_\mu \rightarrow \nu_\tau$ oscillation with a new experiment, called CHORUS, was proposed. The detection technique used is discussed in Chapter 2. At the time, the expectation was that the mixing angles for neutrino oscillation were small, in accordance with quark mixing (section 1.2). Low mixing angles imply that the weak eigenstates are almost identical to the mass eigenstates. It was also expected that the mass eigenstates would have the same hierarchy as the leptons and quarks, such that $m_{\nu_e} \ll m_{\nu_\mu} \ll m_{\nu_\tau}$. With a limit for $m_{\nu_e} \lesssim 10$ eV at that time, cosmic-relic neutrinos were a good candidate to explain dark-matter if the heaviest mass eigenstate would be around 15 eV [87].

CHORUS was approved to take data in 1994 and 1995, later extended for two more years. A second complementary experiment, NOMAD, using a different detection technique based on kinematical cuts, was approved as well [88].

1.5.2 Excluded parameter space

In the simplest analysis of the CHORUS data, the main experimental observables are the total number of neutrino interactions N_{total} and how many of these are ν_τ interactions N_{ν_τ} . The CHORUS experiment did not detect any ν_τ in a sample of $N_{\text{total}} = 5 \cdot 10^5$. The expected background in this sample is about one event. The region CHORUS can then exclude for $\nu_\mu \rightarrow \nu_\tau$ oscillation is shown in Figure 1.11 together with an earlier result of NOMAD. The final exclusion limits for oscillation from the NOMAD experiment can be found in Ref. 89.




Figure 1.11: Exclusion plot for the result of the CHORUS experiment [90] and results from NOMAD and earlier experiments. The two curves for CHORUS use a different statistical analysis. The left-most curve is based on the same statistical approach as used by NOMAD. Using the same approach makes it possible to compare the individual results.

The shape of the excluded region can be understood by examining how the oscillation parameters influence the detected number of events. The oscillation probability, given by equation (1.11) for the two-flavour case, is proportional to:

$$P_{\nu_\mu \rightarrow \nu_\tau} \propto \frac{N_{\nu_\tau}}{N_{\text{total}}} \cdot \varepsilon(\nu_\mu, \nu_\tau) \quad ,$$

where $\varepsilon(\nu_\mu, \nu_\tau)$ accounts for all differences in cross-sections and detection efficiencies between ν_μ and ν_τ interactions. The excluded region for N_{ν_τ} compatible with the expected background is limited by two asymptotes in the log-log contour plot. The vertical asymptote at high Δm^2 determines the sensitivity to the mixing angle. Its position is inversely proportional to N_{total} . The other asymptote determines the minimum Δm^2 value at maximum mixing. This asymptote moves upward linearly towards smaller mixing angles, because the minimum Δm^2 as function of $\sin^2(2\theta)$ is proportional to

$$\Delta m_{\text{min}}^2 \propto \left\langle \frac{E}{L} \right\rangle \frac{1}{\sin^2(2\theta) \cdot \sqrt{N_{\text{total}}}} \quad . \quad (1.23)$$

The behaviour of the exclusion limit at the bending point depends on the distribution of neutrino energies and flight lengths in the beam. From equation (1.23), it is clear that the minimum detectable Δm^2 for CHORUS only decreases inversely proportional to $\sqrt{N_{\text{total}}}$. To detect smaller Δm^2 , one has to move to lower energies or longer baseline without letting N_{total} decrease as well.

1.6 Current neutrino experiments


The current evidence for neutrino oscillations come from basically three experiments which have similar experimental setups: in Japan, Super-Kamiokande (SuperK) and the Kamioka Liquid-scintillator Anti-Neutrino Detector (KamLAND); in Canada, the Sudbury Neutrino Observatory (SNO). Below, the general characteristics of these experiments are described first, before discussing each detector's particularities. The combined oscillation results are given in section 1.7.

The detectors consist of an inner and outer detector consisting of a large volume of liquid monitored by photo-multiplier tubes (PMTs). The inner detector serves as neutrino target, while the outer detector acts as an active veto. Accurate timing on the PMTs allows to reconstruct the interaction vertex and the direction of a produced muon or electron shower. A set of PMTs facing the outer detector detects charged particles entering the detector. Schematic drawings of the detectors are given in Figure 1.12. The experimental setups are described in detail in Ref. 94 for SuperK, Ref. 93 for SNO and Ref. 91 for KamLAND.


As the product of flux and cross-section is typically very low, these experiments are huge. The mass of the inner detectors is, respectively, 32000, 1000, and 1000 tons for SuperK, SNO and KamLAND. Even with these massive targets, the interaction rate is only of the order of a few events per day. Therefore, extreme care must be taken to suppress background from radioactive decays of radon, uranium and thorium contamination. Hence, these experiments require extreme purity of the filling liquids. Special care needs also to be taken to reduce radioactivity from the materials used in the supporting structures, the PMTs and the read-out electronics. Both SuperK and SNO require the signal to have a (visible) energy above 5 MeV, because the background from radioactivity is mostly at energies below 5 MeV. The detectors are placed deep underground to reduce the background from atmospheric muons. The outer detectors shield the detection volume from neutrons and environmental radioactivity. They also act as a veto for incoming muons or charged particles from neutrino interactions in the surrounding rock.

Muon-induced spallation of oxygen nuclei (and carbon for KamLAND) is found to be one of the main backgrounds. These spallation products have anomalous neutron/proton ratios and therefore all kinds of decay modes, like β and inverse β -decay, gamma emission, and proton or neutron emission. For the neutron sensitive detectors, SNO and KamLAND, the neutron emitters are particularly problematic. However, these nuclei have typically very short half-lives and therefore decay close to their production point. As a passing muon can be detected and its path reconstructed, the background from muon-induced spallation can be reduced by removing events in time (a few millisecond) and spatial (a few meter for neutrons) coincidence with a passing muon track.


The main difference between the detectors is the liquid used in the inner detector, water in SuperK, heavy water in SNO, and liquid scintillator in KamLAND. Therefore, the different detectors are sensitive to different processes. SuperK and SNO use Cherenkov radiation to detect a neutrino interaction, while KamLAND uses scintillation light and can measure the total energy deposited in the inner detector better than the other two experiments.



KamLAND [92]



SNO [93]



Super-Kamiokande [94]

Figure 1.12: The experimental setups of SuperK, SNO and KamLAND characterized by an inner and outer detector monitored by large photo-multiplier tubes.

1.6.1 Super-Kamiokande

The SuperK detector is a 39.3 m diameter, 42 m high barrel of water. Around the inner detector, 33.8 m in diameter and 36.2 m high, a steel construction houses the PMTs pointing both inward and outward. The SuperK detector is sensitive to both solar and atmospheric neutrinos. It can identify the Cherenkov light produced by muons and electrons.

Atmospheric-neutrino detection

Atmospheric neutrinos from 10 MeV to 10 GeV can be detected via quasi-elastic (QE) and deep-inelastic charged-current reactions (DIS). In a QE interaction, the emerging lepton will generate Cherenkov radiation which is detected as a single ring of hit PMTs. The hit pattern depends on whether the lepton is a muon or an electron. For a muon, this ring

has sharp boundaries. For an electron or positron, on the other hand, the Cherenkov light is generated by electrons and positrons inside a quickly developing electro-magnetic shower and the hit pattern is generally more fuzzy. This difference is clearly visible, as shown in Figure 1.13 [101]. A likelihood analysis of the hit pattern is used to classify single-ring events as electron-like (e -like) or muon-like (μ -like).




Figure 1.13: Different single-ring hit patterns in SuperK, resulting from Cherenkov light from a single muon on the left and from an electron on the right. Each point represents a hit PMT.

Deep-inelastic interactions produce multiple rings from the different charged particles. However, a classification into electron and muon events can still be made based on the brightest ring (highest number of hits), normally associated with the leading lepton. The separation into ν_e and ν_μ events is however not so good as for single ring events because the corresponding likelihood distributions overlap. As a result, the e -like classification contains only 50% real ν_e charged-current interactions, while the μ -like classification is relatively pure. For these reasons, only the μ -like sample is used in the analysis for multi-ring events.

The atmospheric event sample in the oscillation analysis is split further in sub-GeV ($E < 1.33$ GeV) and multi-GeV ($E > 1.33$ GeV) samples. At low energy the direction of the lepton and the neutrino are not well correlated. Therefore, the incoming neutrino direction is not well determined and the effects of oscillation are smeared out.

Even though SuperK is very big, high-energy muons or events close to the detector edge, will escape from the detector. Events which deposit all their energy inside the detector, mostly electrons (showering) and low energy muons, are classified as fully contained. If the interaction point is inside the inner detector but some particle (usually a muon) leaves the volume and has hits in the outer detector, the event is classified as partially contained. As only a part of the energy is detected, the energy resolution is worse for partially contained events. After applying all selection criteria to reduce the atmospheric-muon background, the remaining partially-contained sample is mostly due to ν_μ interactions.

To increase the event sample, one can also detect neutrino events inside the rock surrounding the detector. These events are indistinguishable from atmospheric muons, but as no muon can penetrate the earth, muons which enter the detector from below the horizon can only be due to neutrino interactions. This event sample is called the upward-muon sample. Muon tracks of at least 7 m long in the inner detector with an upward direction are selected. A manual scan of the events is needed to remove the remaining atmospheric-muon background.

The backgrounds to the various event samples are different. The common background from internal radioactivity is removed by the visible energy cut at $E > 30$ MeV. Atmospheric muons are removed by not allowing hits in the outer detector. The remaining background in the fully-contained sample is below 1% and mostly from PMT flasher events and partially from neutrons and not-vetoed atmospheric muons. The neutrons generate an e -like signal via the interaction chain $p + n \rightarrow d + \gamma$, followed by Compton scattering $\gamma + e^- \rightarrow \gamma' + e^-$. In the partially-contained sample, the background is mainly due to not-vetoed atmospheric muons, but only at the 0.2% level. For the upward-muons near the horizon, horizontal atmospheric muons can be reconstructed accidentally as upward muons or appear to go upward due to multiple-Coulomb scattering. As the rock overburden has two weaker shielding regions, the first analysis bin with $-0.1 < \cos \theta < 0$ is corrected for the remaining atmospheric-muon background.

Solar-neutrino detection

Solar neutrinos are detected via the strongly forward-peaked neutrino–electron scattering interaction which is mainly sensitive to ν_e (see section 1.2.4). While the high-energy ν_e interactions from atmospheric neutrinos are easily recognized from their Cherenkov ring, the solar-neutrino interactions are much more difficult to extract from the background, even though their flux is much higher. The low-energy solar neutrinos give a maximum electron recoil energy of about 18 MeV. These electrons do not form clear rings of hit PMTs, as the number of detected Cherenkov photons is only about 7 hits per MeV. Furthermore, the background at low energies is high, especially below 6.5 MeV where most natural radioactivity takes place. The event trigger requires a certain number of hit PMTs in a 200 ns time window and therefore effectively acts as an energy threshold. Due to the rapidly rising background at lower energies, the energy threshold is determined by the maximum allowed trigger rate. The trigger threshold has been gradually decreased between 1996 (6.5 MeV) and 2000 (4.5 MeV) with the installation of more and faster computers. These computers are used online to select events whose vertex is reconstructed inside the fiducial volume. Because of the energy threshold, SuperK is only sensitive to a very small fraction of the solar-neutrino flux, namely the high-energy ${}^8\text{B}$ neutrinos. The endpoint energy is about 14 MeV for these neutrinos. The standard-solar-model prediction (BP2004) for the ${}^8\text{B}$ neutrino flux is $5.79 \cdot 10^6 \nu \text{ cm}^{-2} \text{ s}^{-1}$ [95], which corresponds to an expected detection rate in SuperK of 325.6 events per day.

The vertex position is reconstructed from the PMT timing information with a resolution of about 200 cm. The recoil electrons from neutrino–electron interactions create only a very short track of maximum a few centimeters. The direction of this track can not be determined from its begin and end-point. Instead, the direction of the electron is measured from the ring-pattern of the Cherenkov light with about 25° accuracy. The electron energy is correlated to the number of PMT hits corrected for water transparency and light propagation. The energy resolution runs from 18% at 5 MeV to 12% at 20 MeV.

A linear electron accelerator has been used to inject 5 to 16 MeV electrons at precise locations to calibrate vertex, direction and energy resolution. Stopping muons and ^{16}N β -decays are used to calibrate the absolute energy scale.

Except for the already mentioned background from radon, the main background below 6.5 MeV comes from radioactivity in the surrounding rock and from the detector material. This background is suppressed by requiring a minimum distance between the vertex and the detector wall. Above 6.5 MeV, the main background is from muon-induced spallation, already discussed at the beginning of this section. A remaining small background is due to beta-gamma decays of ^{16}N created by μ^- capture on ^{16}O .

1.6.2 SNO

The inner detector of the SNO experiment consists of a spherical, 12 m diameter, acrylic vessel filled with 1000 ton pure heavy water (D_2O). It is suspended in the outer detector which is a 22 m diameter and 34 m high cylinder filled with normal water (H_2O). The SNO detector is sensitive to three interaction processes of solar neutrinos: $\nu_x + e^- \rightarrow \nu_x + e^-$ (ES), $\nu_e + d \rightarrow 2p + e^-$ (CC), and $\nu_x + d \rightarrow p + n$ (NC). The charged-current ν_e interaction on deuterium has an energy threshold of 1.442 MeV. The neutral-current disintegration of the deuteron is independent of neutrino flavour and makes a measurement of the total flux possible. This measurement is only sensitive to the ^8B and hep neutrinos which have an energy spectrum extending above the 2.224 MeV dissociation threshold.

The electrons from the CC and ES interactions are detected by the Cherenkov light they generate in the water. The neutron from the NC reaction is captured by a nucleus with high-energy photons carrying away the released binding energy. In the first phase of the SNO experiment, the inner detector was filled with pure heavy water; in the second phase, 2000 kg of salt (NaCl) was added to the heavy water. The addition of salt enhances the NC detection efficiency, because of the much larger cross-section for neutron capture on ^{35}Cl compared to ^2H . Therefore, the neutron is captured closer to its production point and has less chance of leaving the detector. The photons from the neutron capture are detected indirectly by the Cherenkov light of an atomic electron released by Compton scattering. The signals of the three reactions are therefore not that different from each other as all are detected from the Cherenkov light emitted by electrons. The SNO experiment measures four parameters for each event: the total deposited energy (T_{eff}), the angle with the sun (θ_{\odot}), the isotropy of the distribution of light in the 4π solid angle (β_{14}), and the radial position of the event inside the detector (ρ). The three different processes and backgrounds have (slightly) different distributions of these four variables. A global fit to the measured distributions of the four variables extracts the relative contributions from ES, NC, CC and background interactions as function of energy. Figure 1.14 shows these distributions for the three different signals and the main background from external neutrons [114].

The ρ distribution allows one to subtract the background from external neutrons entering the inner detector. The ES interaction maintains the directionality. The CC reaction produces an electron with an energy highly correlated with that of the neutrino and is therefore sensitive to deviations from the expected energy spectrum. The capture of neutrons from the NC reaction on ^{35}Cl releases 8.6 MeV energy in on average 2.5 photons, compared to the single 6.25 MeV photon from the capture on ^2H . Multiple photons give rise to a more isotropic distribution of light from multiple Compton electrons. This higher isotropy shifts the β_{14} distribution for the NC reaction to lower values with respect

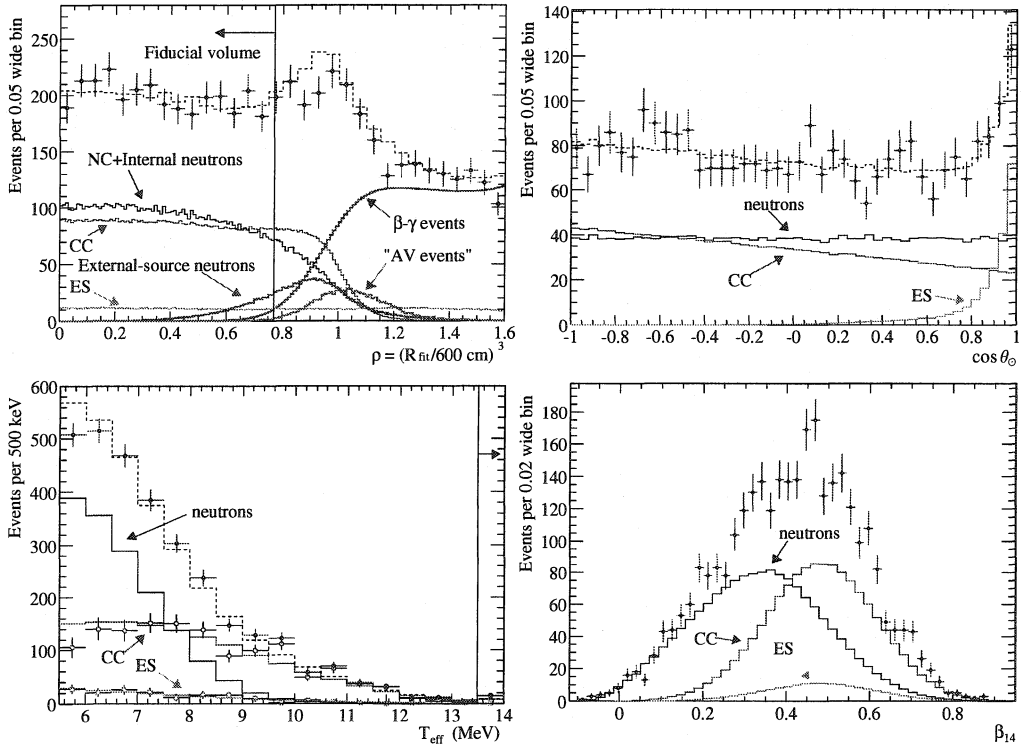


Figure 1.14: The contributions of the three different physics signals to the distributions of the four experimental variables in SNO. The solid lines are from Monte-Carlo calculations and the dashed line is the sum of all contributions. The crosses show the actual data. The fluxes derived from neutrino–electron scattering (ES), charged-current (CC) and neutral-current (NC) interactions of solar ^8B neutrinos are extracted from a global fit to the data for these four distributions. Some of the background is due to radioactive inclusions in the acrylic containment vessel, indicated by the abbreviation AV.

to the data with capture on ^2H and therefore enables a better separation between the NC and CC signals. An energy cut at 5.5 MeV removes most of the low-energy internal radioactivity from the thorium and uranium decay chains. The main background to the NC signal is photo-dissociation of deuterium. By keeping the radon contamination as low as possible, this background was kept small. The calculated background from photo-dissociation in the total sample of 4722 events is 102 events [115].

1.6.3 KamLAND

The KamLAND experiment utilizes a similar technique as was used to prove the existence of the neutrino by Cowan and Reines (section 1.1.1). This time, though, the nuclear reactors are much further away. The KamLAND experiment is situated in the cavern previously used for the Kamiokande experiment, more or less in the middle of Japan's main island, whereas most nuclear reactors in Japan are placed along the coastline. About 80% of the flux at the KamLAND site is from 26 reactors at distances between

134 km and 214 km with a weighted average of 180 km. This relatively narrow band of baseline distances makes it possible for KamLAND to study neutrino oscillation as function of energy.

The detector is a 13 m diameter balloon filled with liquid scintillator. The balloon is suspended in non-scintillating oil inside a 18 m diameter steel vessel, which itself is placed inside a cylindrical water-Cherenkov detector, the outer detector. The charged-current $\bar{\nu}_e$ interaction with protons, $\bar{\nu}_e + p \rightarrow e^+ + n$, generates direct scintillation light from the positron and its rapid annihilation (prompt energy deposit), followed by the delayed signal from neutron capture on hydrogen, giving a 2.2 MeV gamma. Using scintillator and high photo-cathode coverage of the inner detector makes a good measurement of the positron energy and the neutron capture gamma energy possible. The energy resolution is about $\Delta E/E \approx 7\%/\sqrt{E/\text{MeV}}$. As the recoil energy of the neutron is small, the measured positron energy is correlated with the neutrino energy.

The energy spectrum for the neutrinos from the neutron-rich fission products in the nuclear reactors goes up to about 10 MeV and peaks around $E_{\bar{\nu}_e} = 4$ MeV. Anti-neutrinos from ^{238}U and ^{232}Th decays in the earth's core have $E_{\bar{\nu}_e} < 2.49$ MeV and are suppressed by selecting events which have $2.6 \text{ MeV} < E < 8.5 \text{ MeV}$ together with a time and spatial coincidence for the neutron capture.⁴ The delayed signal must occur within a time-interval of $0.5 \mu\text{s}$ to $1000 \mu\text{s}$, take place within a distance of 2 m of the primary interaction point, and have an energy compatible with neutron capture on ^1H ($1.8 \text{ MeV} < E < 2.6 \text{ MeV}$).

The delayed coincidence technique with the applied energy window removes almost all background. The largest background is from muon-induced nuclear spallation and subsequent decays by positron and neutron emission (≈ 3000 events per day). A veto of 2 ms after a passing muon eliminates most of this background, except for the slow neutron emitters, ^9Li and ^8He , which are suppressed by applying a 2 second time and 3 m spatial window around a passing muon.

1.7 Status of oscillation research

1.7.1 Atmospheric neutrinos

The discrepancies in measured ratios for atmospheric neutrinos between the water-Cherenkov and iron-calorimeter experiments, have been investigated by the Soudan experiment, running between 1990 and 1999. The Soudan detector consisted of an iron calorimeter surrounded by a veto shield. It found $R_{\text{atm}} = 0.64 \pm 0.11(\text{stat}) \pm 0.06(\text{sys})$ [96], similar to the KamII and IMB results. It also showed that neutron-induced background could not explain the difference between the water-Cherenkov and iron-calorimeter detectors [97]. A different analysis for upward-going muons with $E > 1 \text{ GeV}$ gave much less stringent limits for $\nu_\mu \rightarrow \nu_\tau$ oscillation [98] and is in contradiction with the previous IMB result for upward muons [85].

The latest results on atmospheric neutrinos come from the SuperK and MACRO experiments. SuperK is described in section 1.6.1. MACRO is a large underground iron calorimeter. The results of MACRO are in agreement with the SuperK results [99]. The MINOS experiment has published preliminary results for its atmospheric-neutrino measurement.

⁴The low background in the detector and a special selection of low-energy events has enabled KamLAND to detect these so-called geo-neutrinos [92].

The magnetic field of the MINOS detector makes it possible to separate atmospheric ν_μ and $\bar{\nu}_\mu$ [100]. As the earlier mentioned discrepancies (see section 1.4.3) have now been clarified and the superK experiment yielded the most accurate results, this section only discusses the latest superK results of Ref. 101.

Zenith angle dependence

In accordance with earlier atmospheric-neutrino measurements, the simplest analysis gives the result as the data over Monte-Carlo ratio for the ratio $(\nu_\mu + \bar{\nu}_\mu)/(\nu_e + \bar{\nu}_e)$. This ratio yields for the latest data 0.67 ± 0.1 using a weighted average over all event categories. In the next step, one looks at the dependence of the muon and electron-neutrino flux on the baseline. The baseline is determined by the incoming neutrino direction and the earth's diameter. The main result is the up/down ratio (up $\equiv \cos \theta < -0.2$, down $\equiv \cos \theta > +0.2$) which turns out to be approximately 1 for low-energy μ -like events and for low and high-energy e -like events. The up/down ratio for high-energy μ -like events is, however, only about a half. At high energy the flux is believed to be isotropic and so this ratio should be 1 if every neutrino entering the earth is also leaving it, independent of any production model. A value below 1, clearly indicates that neutrinos are disappearing on their way through the earth.

In 1489 days of acquiring data, the superK experiment has collected enough events, 4099 e -like, 5436 μ -like and 2259 upward muons, to perform a more detailed analysis based on energy and baseline. At high neutrino energies, there is good correlation between the neutrino and lepton direction. Below 400 MeV/c the correlation between the neutrino and lepton direction is such that any effects depending on the baseline will be washed out. The results of the analysis as function of zenith angle for the different event categories are shown in Figure 1.15. The data show a clear deficit with respect to the Monte-Carlo expectation for μ -like events at higher energies and longer baselines. The number of e -like events are compatible with the Monte-Carlo, which suggest that the disappearance of ν_μ is not creating any ν_e . The data can be fitted very well assuming $\nu_\mu \rightarrow \nu_\tau$ oscillation, indicated by the solid lines, as discussed later.

The number of events is large enough to make a sub-selection of events which each have a small uncertainty in the L/E ratio. This sub-set is used to study the L/E dependence of the ν_μ flux, which should show the characteristic oscillation pattern. To see the maximal oscillation dip requires a good measurement of both L and E . The sub-set is taken from the ν_μ charged-current events in the fully and partially-contained event samples. The selection requires that $\Delta(L/E) < 70\%$ and is based on the Monte-Carlo calculation of $\Delta(L/E)$ for a given zenith angle and energy. In general, events with $|\cos \theta| < 0.1$ have large ΔL and events with $E < 1$ GeV have large $\Delta \theta$ and therefore also large ΔL . In the partially-contained sample, muons that stop in the outer detector have good ΔE , while escaping muons have bad ΔE . In general, any event with $E > 50$ GeV will leave the detector and is not selected.

For the selected 40% of fully-contained and 70% of partially-contained events, one can plot the ratio of the number of measured versus expected events as function of L/E , as is depicted in Figure 1.16 [102]. The characteristic dip, corresponding to the first oscillation maximum, occurs around $L/E = 500$ km/GeV. The data can be fitted well with two-flavour $\nu_\mu \rightarrow \nu_\tau$ oscillation, but disfavour at 3.4σ and 3.8σ the alternative explanations of neutrino decay [103] and decoherence [104].

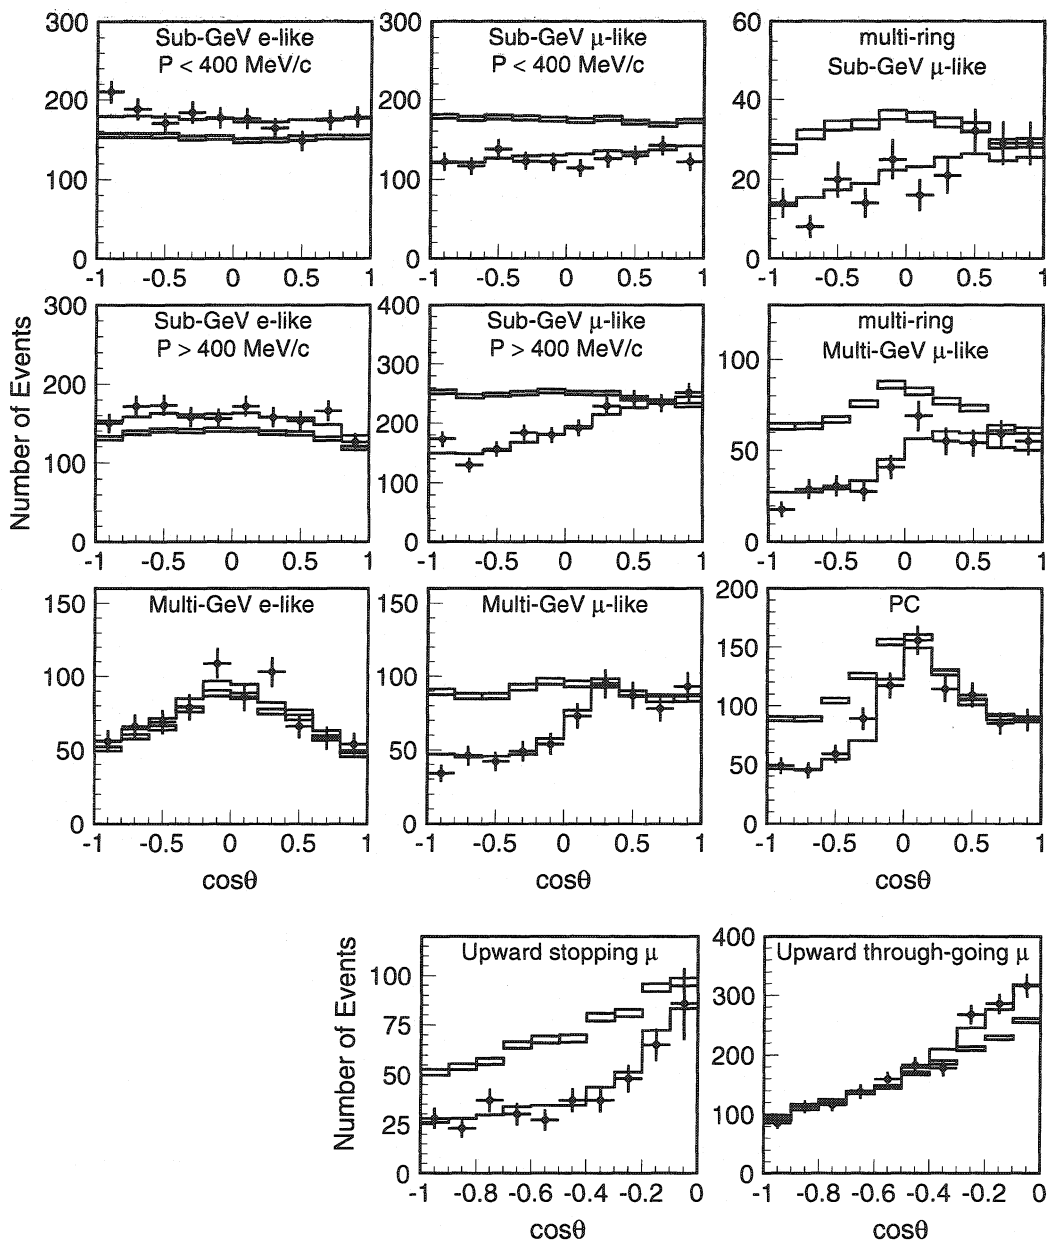
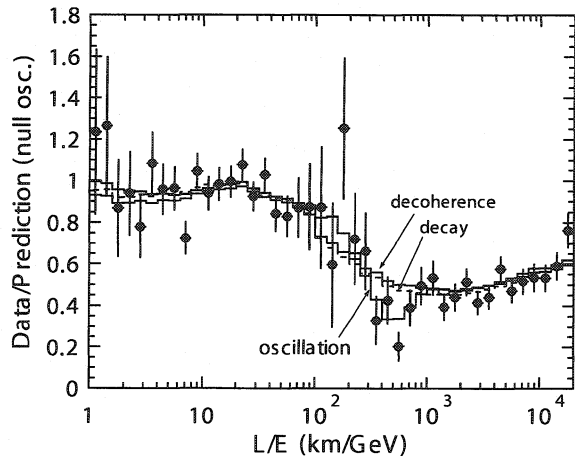


Figure 1.15: The zenith angle distribution for the different event categories of Super-Kamiokande. The e -like events are in accordance with the Monte-Carlo, while the μ -like events clearly deviate from the Monte-Carlo depending on the flight length given by $\cos\theta$. The crosses show the data, the boxes show the Monte-Carlo expectation without oscillation and the solid lines the best-fit expectations for $\nu_\mu \rightarrow \nu_\tau$ oscillation with $\sin^2(2\theta) = 1$ and $\Delta m^2 = 2.1 \cdot 10^{-3} \text{ eV}^2$. The definition of the different event classes can be found in Ref. 101.

Figure 1.16: The data over Monte-Carlo ratio of atmospheric ν_μ (μ -like events) interaction rate as function of L/E for a sub-sample of events with small uncertainty of L/E . The crosses show the measured rate and the solid line the best fit for $\nu_\mu \rightarrow \nu_\tau$ oscillation. Also shown are fits for neutrino decoherence and decay.



Oscillation parameters

From the above results, one can conclude that the oscillation is not ν_μ to ν_e , as there is no surplus of upward electron neutrinos. The natural assumption is then that ν_μ oscillates to ν_τ . Interactions of ν_τ are largely invisible in SuperK due to a charged-current threshold of about 3.5 GeV, although some ν_τ interactions could end up in the multi-GeV e -like sample, which has been excluded from the oscillation analysis.

Extraction of the oscillation parameters involves a fit of the data versus Monte-Carlo, with Δm^2 and $\sin^2(2\theta)$ as parameters in the Monte-Carlo. The simulated events are classified in bins for ν_e or ν_μ interactions, energy, and flight length. Fits to all data and for the separate categories give favoured parameter regions which are shown in Figure 1.17 [101]. For the combined fit, the 90% confidence-level intervals for $\nu_\mu \rightarrow \nu_\tau$ oscillation parameters are $\sin^2(2\theta) > 0.92$ and $1.5 \cdot 10^{-3} < \Delta m^2 < 3.4 \cdot 10^{-3} \text{ eV}^2$. The separate analysis in terms of L/E gives similar results with slightly smaller uncertainty for Δm^2 : $\sin^2(2\theta) > 0.90$ and $1.9 \cdot 10^{-3} < \Delta m^2 < 3.0 \cdot 10^{-3} \text{ eV}^2$.

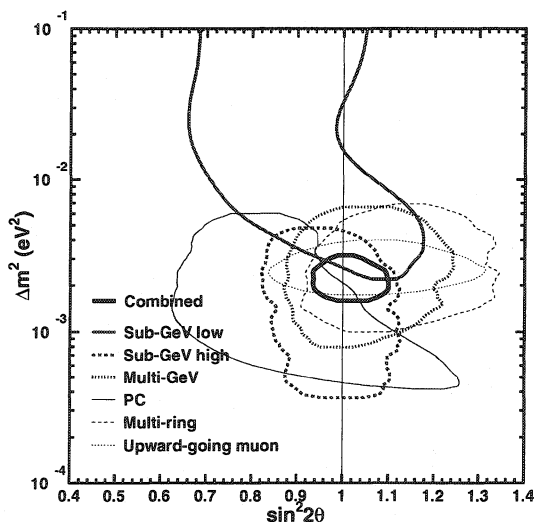


Figure 1.17: Allowed 90% confidence regions for $\nu_\mu \rightarrow \nu_\tau$ oscillation parameters extracted from different event categories of SuperK and for the combined analysis.

Alternatively, ν_μ could oscillate to a hypothetical sterile neutrino (ν_s). Sterile neutrinos have no charged-current or neutral-current interactions and therefore show up as pure ν_μ disappearance. However, an earlier analysis by SuperK has shown that oscillation to ν_s can be excluded at the 99% confidence level for the current oscillation parameters [105]. The difference due to the absence of neutral-current interactions for ν_s is exploited in two ways. First, in the low-energy, multi-ring, fully-contained sample, the e -like events contain about 30% hadron (pion) showers from neutral-current events. Part of these events would be absent if ν_μ would oscillate to ν_s instead of to ν_τ . Second, the earth's matter would cause a MSW effect (see section 1.3.2) for $\nu_\mu \rightarrow \nu_s$. This effect is invisible for $\nu_\mu \rightarrow \nu_\tau$ because they have identical neutral-current interactions. For $\nu_\mu \rightarrow \nu_s$, the oscillation parameters are changed by a factor 2 in the earth at $E_\nu = 15$ GeV and $\Delta m^2 = 3 \cdot 10^{-3}$. However, no effects are seen in the zenith angle distribution for the partially-contained sample with $E > 5$ GeV, nor for the high-energy upward-muon sample.

1.7.2 Solar and reactor neutrinos

The results from the radio-chemical solar-neutrino experiments, Homestake, SAGE, Gallex and GNO, have not changed very much since 1992 [106–109]. In the mean time, though, both Gallex and SAGE have performed calibrations with an intense ^{51}Cr (Gallex and SAGE), ^{71}As (Gallex) and ^{37}Ar (SAGE) neutrino sources. These calibrations showed that the reported detection efficiencies and thus the solar-neutrino rates are understood [110–113]. All the Gallium experiments report a solar-neutrino rate of about 70 ± 6 SNU, which is only 54% of the SSM prediction. The final Homestake result is 2.56 ± 0.23 SNU, which is only 0.3 of the SSM prediction. The latest result for the solar ν_e flux from SNO and SuperK are published in Refs. 114–116. These results are related via the *CPT* theorem with the KamLAND reactor $\bar{\nu}_e$ result [117,118], which will be described before the combined results for ν_e oscillation is given at the end of this section.

Before the measurements from SuperK, SNO and KamLAND, several solutions to the solar-neutrino disappearance, as measured by the radio-chemical and Kamiokande experiments (see section 1.4.2), were possible. These solutions defined four different regions in the $\sin^2(2\theta)$, Δm^2 plane for $\nu_e \rightarrow \nu_\mu/\nu_\tau$ oscillation:

Solution's name	Abbreviation	$\sin^2(2\theta)$	$\Delta m^2[\text{eV}^2]$
Large mixing angle	LMA (MSW)	$\approx 3/4$	$\approx 2 \cdot 10^{-5}$
Small mixing angle	SMA (MSW)	$\approx 6 \cdot 10^{-3}$	$\approx 5 \cdot 10^{-6}$
Low mass difference	LOW (MSW)	≈ 1	$\approx 8 \cdot 10^{-8}$
Vacuum oscillation	VAC	$\approx 3/4$	$\approx 10^{-10}$

In general MSW is important for $10^{-9} \lesssim \Delta m^2/E \lesssim 10^{-5} \text{ eV}^2/\text{MeV}$. For Δm^2 in the MSW range, the oscillation length is much smaller than the distance sun–earth and propagation of neutrino mass-eigenstates is therefore incoherent. The detection probability at earth is then given by:

$$P(\nu_e) = P(\nu_1) \langle \nu_1 | \nu_e \rangle + P(\nu_2) \langle \nu_2 | \nu_e \rangle$$

where $P(\nu_1)$ and $P(\nu_2)$ are the probabilities for a solar ν_e to leave the sun as ν_1 respectively ν_2 . For $\Delta m^2/E$ smaller than $10^{-9} \text{ eV}^2/\text{MeV}$ the distance sun–earth approaches the oscillation length and the propagation should be treated coherent. In the case of the

VAC solution, one reverts to the earlier survival-probability formula of equation (1.8):

$$P(\nu_e) = 1 - \sin^2(2\theta) \sin^2\left(\pi \frac{L}{L_{\text{osc}}}\right)$$

SuperK

Between May 1996 and July 2001, SuperK recorded about 35 million solar-neutrino triggers. The event selection reduces this with about a factor 120 to about 300,000 solar-neutrino candidate events. This sample still contains about a factor 10 more background than signal. The signal is extracted from the background using the directionality of the neutrino–electron scattering process. The cross-section peaks at forward angles and the solar-neutrino interactions are therefore correlated with the direction of the sun. The background on the other hand is largely isotropic. The signal contribution can therefore be extracted as the surface of the forward peak in the distribution of the measured event directions with respect to the sun. This distribution is shown in Figure 1.18 [116]. The peak contains $22404 \pm 226(\text{stat})_{-717}^{+784}(\text{syst})$ solar-neutrino events, which corresponds to a ^8B neutrino flux of $\phi_B = 2.35 \cdot 10^6 \pm 0.02(\text{stat}) \pm 0.08(\text{syst}) \text{ cm}^{-2}\text{s}^{-1}$. The measured flux is only 41 % of the SSM predicted flux.

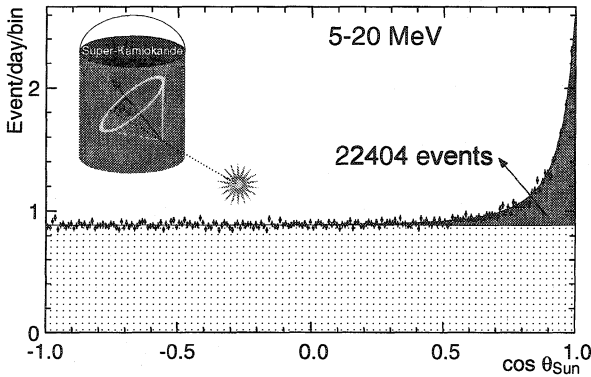


Figure 1.18: Number of solar-neutrino events as function of the angle with the sun's direction. The real solar-neutrino events show up as the forward peak on top of the uniform background.

Neutrino oscillation influences the SuperK measurements in three ways. First, the neutrino–electron scattering cross-section for ν_μ and ν_τ is only about $(1/6)^{\text{th}}$ of the cross-section for ν_e . Second, the neutrino energy spectrum is distorted because the ν_e survival probability, equation (1.8), is energy dependent. Changes in the spectrum translate directly to changes in the electron recoil-energy spectrum. Third, for the MSW solutions, the earth's matter influences the oscillation probability which would show up as a 24 hour modulation of the signal rate. For the VAC solution, on the other hand, the survival probability is baseline dependent which would show up as a seasonal fluctuation due to the eccentricity of the earth's orbit. With the large statistics of SuperK, it is possible to look for both energy dependence and time variations.

The SuperK data show no observable distortions from the expected shape of the ^8B neutrino spectrum. The seasonal variation is consistent with the expected $1/r^2$ behaviour at 69% confidence level. This excludes the VAC solution for solar-neutrino oscillation. The day versus night asymmetry was measured to be $[2.1 \pm 2.0(\text{stat})_{-1.2}^{+1.3}(\text{syst})] \%$. This value is both consistent with zero and with the expected value for part of the LMA solution. It does exclude some of the LMA region at low Δm^2 . A combined analysis

of the data as function of energy and zenith angle gives two allowed regions for the oscillation parameters at 95 % confidence level in the LOW and LMA regions, but excludes the SMA solution. The two allowed regions are depicted in Figure 1.19 [116]. From the χ^2 values, one can see that the best fit-point is in the LMA region at $\tan^2(\theta) = 0.52$ ($\Leftrightarrow \sin^2(2\theta) = 0.90$) and $\Delta m^2 = 7 \times 10^{-5} \text{eV}^2$.




Figure 1.19: Allowed $\nu_e \rightarrow \nu_{\mu\tau}$ oscillation parameter regions at 95 % confidence level from SuperK solar data. To the top and right are indicated the χ^2 values for the fit when the other parameter is set to give the lowest χ^2 (lines inside the regions).

SNO

The detection principle of the SNO solar-neutrino measurement has been explained in section 1.6.2. The results, covering the period up-to September 2003, can be found in Refs. 114,115. The SNO experiment has measured the charged-current, neutral-current and neutrino-electron scattering rate which have different contributions from ν_e and ν_μ or ν_τ interactions. For each interaction type, the measured flux and its uncertainty define bands of possible values in the $\nu_{\mu\tau}$ versus ν_e plane. The slopes of the bands are determined by the relative sensitivity of each process to ν_e and $\nu_{\mu\tau}$: charged-current $\propto \nu_e$; neutral-current $\propto \nu_e + \nu_{\mu\tau}$; and neutrino-electron scattering $\propto \nu_e + 0.16\nu_{\mu\tau}$. A combined fit to the four distributions of Figure 1.14 yields the relative amounts of charged-current, neutral-current, and neutrino-electron scattering rates. The ν_e and $\nu_{\mu\tau}$ fluxes derived from these numbers are shown in Figure 1.20 [115]. The measured charged-current rate gives a ^8B flux of:

$$\phi_B = (1.68 \pm 0.06(\text{stat})_{-0.09}^{+0.08}(\text{syst})) \times 10^6 \text{cm}^{-2}\text{s}^{-1} ,$$

which corresponds to only about $1/3^{\text{th}}$ of the expected SSM flux [119,120]. This result confirms the earlier radio-chemical experiments. However, the neutral-current measurement yields:

$$\phi_B = (4.94 \pm 0.21(\text{stat})_{-0.34}^{+0.38}(\text{syst})) \times 10^6 \text{cm}^{-2}\text{s}^{-1} .$$




Figure 1.20: Solar ^8B neutrino flux, split in ν_e and $\nu_{\mu\tau}$ contributions, as measured by SNO from the charged-current (CC), neutral-current (NC) and neutrino–electron scattering (ES) rates. The width of the bands represent 1σ errors, the slopes correspond to the sensitivity for each process to ν_e and $\nu_{\mu\tau}$. The total ^8B neutrino flux (neutral-current result at $\phi_e = 0$) is within the 1σ expectation of the SSM (dashed lines).

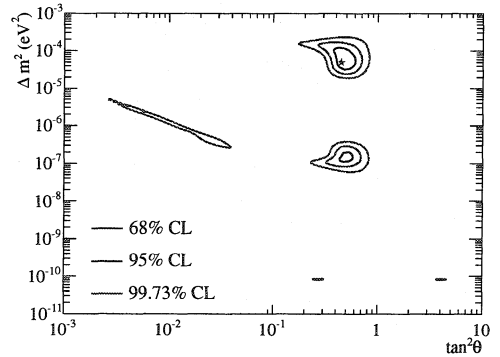


Figure 1.21: Allowed parameter regions at different confidence levels for solar-neutrino oscillation. The best fit point is indicated by the star and is in the LMA region. The LOW solution also gives 1σ fit probabilities. The SMA and VAC solutions are almost excluded except for some small regions at 3σ .

This value for the total neutrino flux lies within the range predicted by the SSM which has about 20% uncertainty.

The SNO neutral-current result provides evidence for the hypothesis that the ν_e from ^8B decay are transformed into active neutrino species (ν_μ, ν_τ). The neutrino–electron scattering measurement further limits the range of the ν_e versus $\nu_{\mu\tau}$ flux ratio, especially if the measurement from SuperK is used, as is shown in Figure 1.20.

A fit which compares the data to the calculated propagation (and matter effects) of solar neutrinos for a given set of $\nu_e \rightarrow \nu_{\mu\tau}$ oscillation parameters gives regions of allowed parameters. These allowed regions at different confidence levels are shown in Figure 1.21. The best fit point is in the LMA region at $\Delta m^2 = 5.0 \cdot 10^{-5} \text{ eV}^2$ and $\tan^2(\theta) = 0.45$ ($\sin^2(2\theta) = 0.86$). Also the LOW solution has a region which good χ^2 values. The SMA and VAC solutions are almost excluded as only some small 3σ regions remain.

KamLAND

Although KamLAND measures reactor $\bar{\nu}_e$, see section 1.6.3, its results are related to the solar-neutrino measurements via *CPT* invariance. The energy range of the reactor anti-neutrinos also matches the typical range of solar neutrinos. The last published result from the KamLAND experiment covers the data taken between March 2002 and January 2004. Based on the reactors power output, the expected number of events in the KamLAND experiment is 365.2 ± 23.7 . Instead, only 258 events were detected, which corresponds to $[65.8 \pm 4.4(\text{stat}) \pm 4.7(\text{syst})]$ % of the expected number [117]. This number falls squarely in

the range predicted by the LMA solar-neutrino solution ($L_{\text{osc}} \approx 130$ km); but is more than 4σ different from 1 which would be the case for the SMA, LOW and VAC solutions. The measured ratio and those from previous short-baseline reactor experiments are shown in Figure 1.22.





Figure 1.22: The ratio of measured over expected $\bar{\nu}_e$ neutrino fluxes for several reactor experiments [118]. The result of KamLAND at $L = 1.8 \cdot 10^5$ m shows the disappearance of $\bar{\nu}_e$ at long baselines. The result is consistent with the expectation for the solar LMA solution (shaded area). The best fit for $\bar{\nu}_e \rightarrow \bar{\nu}_x$ oscillation yields $\Delta m^2 = 5.5 \cdot 10^{-5} \text{ eV}^2$ and $\sin^2(2\theta) = 0.83$ and is indicated by the dotted line.

If neutrino oscillation is responsible for the $\bar{\nu}_e$ disappearance, KamLAND should be able to see a spectral distortion of the measured $\bar{\nu}_e$ flux for a given range of Δm^2 . Because about 80% of the flux has a baseline of around 180 km, oscillation would give a relatively clear sinusoidal shape as function of energy (E). Figure 1.23 shows the survival probability as a function of L/E , with L fixed to 180 km. In the figure the first maximum and first two minima of the oscillation pattern can be seen. From the maximum at $L/E = 32$ km/MeV one can directly derive Δm^2 which yields $\Delta m^2 = 7.7 \cdot 10^{-5} \text{ eV}^2$. The mixing angle determines the depth of the minima. From the figure one can estimate that $\sin^2(2\theta) \approx 0.7$. The limited statistics lead to large errors on the mixing angle. A precise fit of the KamLAND data yields $\Delta m^2 = 7.9^{+0.6}_{-0.5} \times 10^{-5} \text{ eV}^2$ and $\sin^2(2\theta) = 0.86^{+0.14}_{-0.30}$.

Figure 1.23: Dependence of the ratio of the measured neutrino-rate versus the expected rate in the KamLAND experiment as function of energy. In this analysis, all reactor neutrinos are assumed to come from a single reactor at $L_0 = 180$ km [117].



Combined solar results

One can take all solar-neutrino measurements together to limit the possible oscillation parameters further. In such an analysis, KamLAND's Δm^2 measurement limits the solution to the LMA region. KamLAND also has the best determination of Δm^2 from the spectral distortion in the L/E measurement. The mixing angle is best determined by SuperK and SNO. In Table 1.2 a summary of the results is given as well as the average of the combined results as reported in Refs. 115–117. The allowed parameter region for the combined result at different confidence levels is shown in Figure 1.24 [145].

	$\Delta m^2 [10^{-5} \text{ eV}^2]$	$\sin^2(2\theta)$	$\tan^2(\theta)$	$\phi_B [10^6 \text{ cm}^{-2} \text{ s}^{-1}]$
SNO	$5.0^{+6.2}_{-1.8}$	$0.86^{+0.06}_{-0.09}$	$0.45^{+0.11}_{-0.10}$	5.11
SuperK	$6.3^{+4.7}_{-2.1}$	$0.90^{+0.07}_{-0.08}$	$0.52^{+0.18}_{-0.12}$	-
	$7.6^{+4.4}_{-2.3}$	$0.90^{+0.07}_{-0.08}$	$0.52^{+0.18}_{-0.12}$	-
solar (combined)	$6.5^{+4.4}_{-2.3}$	$0.86^{+0.05}_{-0.07}$	$0.45^{+0.09}_{-0.08}$	5.06
KamLAND	$7.9^{+0.6}_{-0.5}$	$0.86^{+0.14}_{-0.30}$	$0.46^{+0.54}_{-0.26}$	-
solar + KamLAND	$8.0^{+0.6}_{-0.4}$	$0.86^{+0.05}_{-0.06}$	$0.45^{+0.09}_{-0.07}$	4.92

Table 1.2: Oscillation parameters from separate and combined results of solar ν_e experiments and the KamLAND experiment.

1.7.3 Accelerator experiments

The results from earlier accelerator neutrino experiments will not be mentioned here, as they are all superseded by the positive results from solar and atmospheric experiments. The discussion is limited to the confirmation of ν_μ disappearance by the K2K experiment and the unconfirmed positive oscillation result from LSND.

The K2K experiment is a combined near and far detector experiment in Japan [121]. It uses an accelerator ν_μ beam, generated at the KEK laboratory, pointing to the SuperK detector which is located at a distance of 250 km. The near detector is a smaller copy of the SuperK detector with additional fine-grained detectors. The near detector is used to accurately measure the neutrino flux. The short-spill timing of the beam allows to detect any beam related ν_μ interactions in SuperK with negligible background. The experiment reported a statistical significant deficit of ν_μ interactions: 107 events detected while 151 ± 11 were expected from the measured flux at the near detector [122]. This result is consistent with the oscillation parameters determined from atmospheric $\nu_\mu \rightarrow \nu_\tau$ oscillation.

One other experiment, with a near and a far detector, will also give results in the near future. The MINOS experiment in Soudan (USA), has just published its preliminary results. An observed disappearance of ν_μ [123] over the 735 km baseline from Fermilab's main injector confirmed the result of K2K. The oscillation parameters derived from these data are compatible with the SuperK atmospheric-neutrino measurements. The MINOS experiment has reached similar uncertainties as the SuperK experiment. With more data to come, the MINOS experiment will be able to determine the oscillation parameters with an uncertainty of less than 10%. Another long-baseline experiment, OPERA [124], will




Figure 1.24: Allowed oscillation parameter regions from all solar and KamLAND results combined. The overlap between these results limits the solution to the LMA region. The selected part of the KamLAND result (filled areas) constrains Δm^2 , while the solar experiments (filled and dashed lines) limit the mixing angle. The independent best-fit points are indicated by the star and triangle symbols. The plots on top and to the right show the χ^2 of the overall fit when the other parameter is optimized separately and indicate the uncertainty in the result.

become operational in 2007 at Gran-Sasso in Italy. This experiment plans to detect ν_τ events using emulsion in a high-energy ν_μ beam coming from CERN, 732 km away.

In the future, the planned T2K [125] experiment in Japan, will have a baseline of 295 km from an accelerator to be built at Tokai (near Tokyo) to the SuperK detector. The T2K experiment will make precision measurements of ν_μ disappearance and possibly ν_e appearance using a high-intensity neutrino beam. The neutrino energy is tuned to have the first oscillation maximum occurring at the SuperK distance of 295 km.

The reason some experiments have been looking for possible sterile neutrinos is the claim for a positive oscillation signal from the LSND experiment [126–128]. Between 1993 and 1998, LSND used the delayed neutron-signal coincidence technique to look for

$\bar{\nu}_e$ interactions ($\bar{\nu}_e + p \rightarrow e^- + n$) at a continuous, low-energy, neutrino source. The neutrinos are generated from decay-at-rest of π^+ and μ^+ stopped inside a beam dump. Negative pions and the few μ^- from decay in flight of π^- are mostly captured before decaying. Therefore, only ν_μ from π^+ decays and $\bar{\nu}_\mu$ and ν_e from μ^+ decays escape the beam dump. Hence, there are almost no $\bar{\nu}_e$ present in the beam ($\phi(\bar{\nu}_e)/\phi(\bar{\nu}_\mu) = 8 \cdot 10^{-4}$). The LSND detector was located 30 m behind this beam dump. The experiment reported an excess of 87.9 ± 23.2 $\bar{\nu}_e$ events with respect to the expected background of 30 ± 6 events. If interpreted as due to oscillation from $\bar{\nu}_\mu$, the oscillation probability would be $[0.264 \pm 0.067(\text{stat}) \pm 0.045(\text{sys})] \%$ [128]

However, most of the preferred oscillation region is excluded by other experiments, such as short-baseline reactor experiments and the KARMEN experiment. The KARMEN detector used a similar neutrino source as LSND, but with a pulsed beam. Their detector was placed perpendicular to the beam-dump and had a shorter baseline of 17 m. The KARMEN experiment did not see an excess of $\bar{\nu}_e$ events [129]. A combined analysis of the KARMEN and LSND data shows that only a small region with Δm^2 between 0.1 and 0.7 eV² and $\sin^2(2\theta) \approx 10^{-3}$ remains [130]. This region is incompatible with both the solar and atmospheric oscillation results. A fourth neutrino state would be needed to explain all oscillation results. This fourth neutrino state must be sterile because the number of light neutrinos coupling to the Z^0 is only three (see section 1.1.3). The MiniBoONE experiment should confirm or reject the LSND signal in the near future [131].

1.7.4 Cosmic neutrinos

The field of neutrino astronomy commenced with the detection by the Kamiokande-II and IMB experiments of neutrinos from the 1987 supernova in the Large Magellanic Cloud (distance $\approx 1.7 \cdot 10^{18}$ km) [132–135]. The cluster of detected events is shown in Figure 1.25 [134]. Nowadays, several collaborations are constructing enormous detectors to detect the very low fluxes of cosmic neutrinos using Cherenkov radiation. These detectors use large volumes of water deep in the sea (BAIKAL [136], ANTARES [137], NESTOR [138], and NEMO [139]) or a large volume of ice several kilometers down at the south pole (AMANDA [140] and its successor ICECUBE [141]). Other experiments and proposals exist to detect coherent radio-waves emitted by large electro-magnetic showers in either the south-pole ice, RICE [142] and ANITA [143], or from the moon, GLUE [144].




Figure 1.25: All events seen in the Kamiokande-II detector within a 45 second window around the time of supernova SN1987A. The height of each line represents the measured energy. Solid lines are electron-like events (left scale) and dashed lines atmospheric muons crossing the detector (right scale). A burst of electron neutrinos is seen at $t = 0$ s.

1.7.5 Three-flavour oscillation

The results of atmospheric and solar neutrino experiments can be fitted well in a theory with two-flavour neutrino oscillations dominated by small mass differences and high mixing angles. As there are three neutrino flavours, the results of all experiments should be analysed in a three-flavour oscillation theory. The results of all the experiments, except for LSND, can be fitted in such a model with the oscillations dominated by almost the same two small mass differences and mixing angles [145, 146]. However, for the three-flavour mixing matrix, one of the mixing angles remains unknown. Assuming a unitary mixing matrix for three neutrino flavours, the complete mixing matrix can be written as the product of three matrices:

$$U = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \times \begin{pmatrix} c_{13} & 0 & s_{23} e^{i\delta} \\ 0 & 1 & 0 \\ -s_{13} e^{i\delta} & 0 & c_{23} \end{pmatrix} \times \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix},$$

with $c_{\alpha\beta} = \cos \theta_{\alpha\beta}$ and $s_{\alpha\beta} = \sin \theta_{\alpha\beta}$. The first matrix dominates the oscillation of atmospheric neutrinos, the third the solar neutrino oscillation. The main unknown in this decomposition is the third mixing angle θ_{13} . At present, the upper limit for $\sin^2(\theta_{13})$ is 0.06 from the CHOOZ reactor experiment [147]. This limit is now supplemented by a three-flavour analysis of the SuperK atmospheric data, giving limits of $\sin^2(\theta_{13}) < 0.14$ and $\sin^2(\theta_{13}) < 0.27$ at 90% confidence level for the normal and inverted mass hierarchy, respectively (see section 1.8). The CP violating phase $e^{i\delta}$ is currently completely unknown and the chances of measuring it will depend very much on the magnitude of θ_{13} . More detailed discussions of the effects of three-flavour mixing and its experimental signatures can be found in, for example, Refs. 148–151.

1.8 Discussion and outlook

Although neutrino oscillations are now a well established fact and the mass differences and mixing angles are already relatively well constrained, there are many remaining questions regarding the intrinsic properties of neutrinos.

Concerning neutrino oscillations, the third mixing angle and the CP violating phase remain the main unknowns. New reactor experiments with two detectors could possibly observe the modulation of the $\bar{\nu}_e$ disappearance due to the θ_{13} angle. Another possibility is to observe the appearance of ν_e in a pure ν_μ beam. Currently, all results on neutrino oscillations for solar, atmospheric and reactor neutrinos are from disappearance experiments. The one exception is the neutral-current result of the SNO experiment which has measured the total solar-neutrino rate. However, also that result can not distinguish the weak eigenstates involved in the interactions. No experiment has as yet seen the appearance of the other neutrino states in oscillation.

Confirmation of the suggested oscillation channels can only be established by actually seeing the appearance of the other weak eigenstates due to oscillation. Such experiments are currently becoming operational at existing long-baseline neutrino beams (OPERA and MINOS). Other such experiments are being proposed using new neutrino beam concepts, like a muon storage ring or a beta-beam. In a muon storage ring [152], a beam of high-energy muons circulates inside a racetrack shaped ring. The decays of muons in the straight sections generates an intense neutrino beam. Such a storage ring is an ideal

source for appearance experiments as it produces only ν_μ and $\bar{\nu}_e$ neutrinos (or the inverse) with well-known energy spectra and with energies above threshold for charged-current reactions of either ν_μ or ν_τ . The proposed beta-beam of radioactive nuclei [153] with its pure ν_e or $\bar{\nu}_e$ flux is also a very clean neutrino source which makes detailed studies of neutrino oscillations possible.

The neutrino oscillation experiments have determined the two mass differences, but cannot determine the absolute mass scale. However, the lowest limit for any of the neutrino species is also a limit for the absolute mass scale. The current lowest limit is about 3 eV from the direct ν_e mass measurement. As the neutrino mass limits follow the same hierarchy in mass as the corresponding leptons, the most likely experiments to lower the absolute mass-scale limit are direct ν_e mass measurements. Currently a new tritium-decay experiment, called KATRIN, is being set up in Karlsruhe that will be able to measure down to $m_{\nu_e} \leq 0.2$ eV [154]. In the special case that neutrinos are Majorana particles (see below), the absolute mass could also be determined from the rate of neutrinoless double β -decay.




Figure 1.26: The two mass differences, measured for solar and atmospheric neutrinos, can be arranged in two different ways. Currently no experimental result can distinguish between the two. Also the absolute mass scale is currently unknown, although the direct ν_e mass measurements set an upper limit of about 3 eV.

The three neutrino masses can also be arranged in two different ways, called normal and inverted hierarchy, illustrated in Figure 1.26. Which particular hierarchy is right is yet unknown. If the absolute masses are close to the mass differences and the neutrino is a Majorana particle, neutrinoless double β -decay could determine the mass hierarchy. If, on the other hand, the absolute mass scale is high with respect to the larger Δm^2 , the neutrino mass-eigenstates are almost degenerate and the hierarchy can only be established by observing the sub-dominant oscillation in very accurate oscillation experiments.

There are two ways (or a linear combination of both) in which neutrino masses can be introduced in the Lagrangian of the standard model, known as a Dirac or Majorana mass-term. In any case, because the neutrinos have mass, right-handed neutrinos exist as well. If the neutrino is a Dirac particle, the right-handed neutrino state is generated by a Lorentz transformation of the left-handed helicity state. If the neutrino is instead a Majorana particle, the right-handed state corresponds to the anti-neutrino. This coupling between helicity and particle or anti-particle state is only allowed for neutrinos because they carry no conserved quantities. If the neutrino is a Majorana particle, an emitted

neutrino can interact as an anti-neutrino with a probability given by its mass. In this case neutrinoless double β -decay becomes possible for some nuclei. The Feynman diagram for neutrinoless double β -decay is depicted in Figure 1.27. The rate of this process can then be used to measure the absolute mass of the neutrino (if all the nuclear matrix elements can be calculated [155]). On the other hand, if neutrinoless double β -decay is not observed and the absolute mass scale is larger than 1 eV, one can conclude that neutrinos are Dirac particles.




Figure 1.27: Feynman diagram of neutrinoless double β -decay. In this process a nucleus emits two electrons but no neutrinos when two neutrons change into two protons. The neutrino emitted in the first β -decay is absorbed as an anti-neutrino in the second β -decay, which is only possible if the neutrino is a Majorana particle and has non-zero mass such that its helicity, indicated by the small arrows, is not a conserved property.

Chapter 2

The CHORUS experiment

The CHORUS experiment was built to detect $\nu_\mu \rightarrow \nu_\tau$ oscillation in an almost pure ν_μ beam. The detector was designed with one specific process in mind: locating and identifying a τ particle from a charged-current ν_τ interaction inside a large stack of nuclear emulsion plates. Emulsion is ideal for the detection of short-lived particles which is crucial to attain the design sensitivity for oscillation which requires the rejection of events due to ν_μ interactions to be better than 1 in 10^6 . The use of emulsion also permits detailed studies of events with similar length scales as τ decays, like charmed-particle production and decay.

The perfect detector does of course not exist and trade-offs between different detector choices need to be made. Sometimes a new technology allows improvements to be made while the detector is already running. One of those new technologies, a honeycomb tracker, was installed in the CHORUS experiment for the last one and a half years of data taking. Chapter 3 describes the development and performance of this detector.

This chapter explains the general layout and design of the CHORUS experiment. A detailed description of the full detector and its performance can be found in Ref. 156 and the details of several sub-detectors in Refs. 157–165. As Chapters 4 and 5 of this dissertation require a detailed understanding of the particularities of emulsion as a tracking detector, the emulsion target and the location and reconstruction of neutrino vertices inside it will be described in more detail.

2.1 Detection principle

The detection of $\nu_\mu \rightarrow \nu_\tau$ oscillation is based on the identification of charged-current ν_τ interactions in a ν_μ beam which does not contain any ν_τ . Any τ particle produced in a charged-current interaction of a ν_τ can therefore only be due to neutrino oscillation transforming a ν_μ in a ν_τ . The experiment therefore requires a good detection efficiency of τ particles while rejecting all processes that might mimic this signal. The τ identification is complicated, however, by its short lifetime and the fact that its decay produces neutrinos which leave the detector undetected. Hence, an accurate determination of the invariant mass from the decay products is impossible. In the CHORUS experiment, the τ identification uses both the short lifetime and the missing momentum.

2.1.1 Tau identification in emulsion

In its proposal [166, 167], the CHORUS experiment aims for the two easiest detectable τ -decay topologies where a τ^- decays into a single charged particle (see Figure 2.1). The τ^- can decay into a μ^- via $\tau^- \rightarrow \nu_\tau \mu^- \bar{\nu}_\mu$ with a branching ratio of $\text{Br}(\tau \rightarrow \mu) = 17.4\%$ or into a single charged meson (π^- or K^-) via $\tau^- \rightarrow \nu_\tau h^- n(\pi^0)$ with $\text{Br}(\tau \rightarrow h) = 49.5\%$. In both cases, the undetected neutrinos carry away part of the momentum and energy of the τ parent ($m_\tau = 1777 \text{ MeV}$). Therefore, the charged daughter particle has a different direction from its parent. In the experiment this is visible as a kink in the track made by the τ^- and its charged daughter. A kink angle of 50 mrad or more is measurable.




Figure 2.1: Decay modes of a τ^- produced in a charged-current ν_τ interaction into a single charged particle. Because of the momentum carried by the neutrino(s), the direction of the μ^- (a) or the negative meson (b) differs from the τ^- direction. In the experiment this is detectable as a kink in the apparent path of the daughter particle when it is followed back to the vertex. The average decay length is given by $\langle d \rangle = \langle \gamma \rangle c \tau_\tau \approx 87 \langle E/m_t \rangle \mu\text{m}$.

To detect the short flight path of any produced τ ($c\tau_\tau = 87 \mu\text{m}$), a detector is needed with a very good resolution in 3-D. To get enough events, the target must also have a large mass because of the very small interaction cross-section of neutrinos. Nuclear emulsions are both excellent targets and detectors, because they give 3-D track detection at sub-micron scale and are relatively dense due to the high silver content ($\rho = 3.815 \text{ g/cm}^3$). The disadvantage of emulsion is that it does not provide any time information; it records all ionizing tracks passing through it in the period between pouring and development.

In early emulsion experiments, the exposure time was usually short and the quantity of emulsion small, allowing human operators to scan all of the emulsion for interesting events. In the CHORUS experiment, the exposure time of two years is exceptionally long and the emulsion volume an order of magnitude larger than ever before. For CHORUS, it was no longer possible to scan all of the 1540 kg of emulsion for interesting events. There have been two runs with each 770 kg of emulsion; one spanning the years 1994 and 1995; the other in 1996 and 1997. The chosen solution was to use electronic tracking detectors to indicate where to look in the emulsion for a particular event. This is known as a hybrid emulsion-electronic detector. As the emulsion can only be examined under a microscope after it has been developed, the electronic data are also needed to separate the events recorded during the two years of data taking.

The limit on $\nu_\mu \rightarrow \nu_\tau$ oscillation [168] at the time of the CHORUS proposal was such that a maximum of 35 charged-current ν_τ interactions could be detected in a sample of $5 \cdot 10^5$ charged-current and $1.5 \cdot 10^5$ neutral-current ν_μ interactions. The total number of events ($6.5 \cdot 10^5$) exceeded the emulsion scanning capacity at that time by far. For the proposal, it was estimated that 40,000 events could be scanned, at maximum, during two years of analysis. Therefore additional detectors were needed for pre-selection of events with a higher probability of being due to a ν_τ interaction. For this pre-selection and to suppress background from charm decays, two magnetic spectrometers and a calorimeter were placed downstream of the emulsion target. The spectrometers measure the momentum and charge of particles leaving the target and the calorimeter measures the total energy in an event. The calorimeter also serves as passive muon filter.

The pre-selection was based on kinematic variables and would select one track in an event that is most likely the daughter of a τ particle. Following only that track back in the emulsion lowers the scanning load. The detection of the kink is then done in a single pass through the emulsion, avoiding the need to follow all tracks downstream from the interaction vertex. The kinematical selection requires a high reconstruction efficiency and good resolution of the kinematic quantities. Any inefficiency or wrongly measured variable lowers the maximum sensitivity to $\nu_\mu \rightarrow \nu_\tau$ oscillation. The kinematical cuts which were to be applied, are primarily based on the energy and momentum that is carried away by the neutrino(s) in the τ decay. In a neutrino charged-current interaction, the lepton's transverse momentum p_T balances the transverse momentum of the shower resulting from the nuclear breakup. In a ν_τ interaction part of that momentum is carried away by the neutrino(s) from the τ decay. The direction of the missing transverse momentum is typically opposite to the hadron shower direction. Finally, the energy spectrum of the muon or meson from τ decay is different from that induced by ν_μ events.

During the time the experiment was taking data and doing the analysis, automatic scanning microscopes have become much faster. The allowed scanning load has grown by more than a factor hundred. The increased scanning speed has made it possible to select all events with a reconstructed vertex inside the emulsion for scanning. Only the target trackers, located directly behind the emulsion (section 2.4), are used for this reconstruction, so any inefficiency due to wrong matching or identification in the downstream detectors is avoided. Only the muon spectrometer is used to preferentially select muon tracks for scanning. The vertex and possible decay topologies are reconstructed in the emulsion. This has the benefit that all tracks from the neutrino vertex can be checked for kinks. To increase the event location efficiency, an alternative track selection procedure has been applied. The selection of tracks for scanning is discussed in section 2.8, event

location and reconstruction in the emulsion in section 2.10. The data from the other detectors are used to measure momenta and identify particle type of event-related tracks reconstructed in the emulsion.

2.1.2 Background processes

There are at least 10,000 charged-current and 3,000 neutral-current ν_μ interactions for every charged-current ν_τ interaction. These numbers take into account a lower limit for oscillation at high Δm^2 of $4 \cdot 10^{-4}$, the difference in energy dependence of the neutrino cross-sections, and the energy spectrum of the neutrino beam. Multiplying this by the inverse of the branching ratio $\text{Br}(\tau^- \rightarrow \nu_\tau \mu^- \bar{\nu}_\mu) = 17.4\%$, the number of ν_μ events per charged-current ν_τ interaction becomes 75,000. Background events in the $\tau^- \rightarrow \mu^-$ channel should thus be suppressed by a factor 10^5 . For the $\tau^- \rightarrow h^-$ channel, the event has to be detected among the neutral-current interactions and those charged-current events where the primary muon is missed (about 10%). With $\text{Br}(\tau^- \rightarrow \nu_\tau h^- n(\pi^0)) = 49.5\%$, the $\tau^- \rightarrow h^-$ background should thus be suppressed by a factor of 10^4 . The probability to miss a primary μ^+ is higher than that for a μ^- . The probabilities to miss a primary e^- or e^+ are even higher. As background suppression is partly based on identifying the primary lepton, it is important to keep the relative flux of $\bar{\nu}_\mu$, ν_e and $\bar{\nu}_e$ in the neutrino beam as low as possible.

There are three processes which are identical to or mimic a τ decay and therefore contribute to the background for the $\nu_\mu \rightarrow \nu_\tau$ oscillation search:

1. Charged-current interactions of ν_τ contamination in the neutrino beam. The ν_τ 's originate from the decay of D_s and τ^- created in proton interactions with the primary target.
2. Decay of negatively-charged mesons close to their production vertex. These events are only a background if either the primary lepton in a charged-current interaction remains undetected or in neutral-current interaction where there is no primary lepton.
3. Elastic scattering of a muon or hadron on a nucleus with no visible recoil in the emulsion, known as white kinks.

Point 1 is identical to the oscillation signal and is therefore an irreducible background. The ratio ν_τ/ν_μ in the neutrino beam must therefore be as low as possible (see section 2.2). For points 2 and 3, the kink must be located first and the kink daughter identified. The efficiencies for kink detection in both real τ decays and these backgrounds are very similar, except for (small) differences in energy and momentum.

Regarding the background of point 2, the single-prong decays of $\pi^- \rightarrow \mu^-$ and $K^- \rightarrow \mu^-$ or π^- can be eliminated by requiring $p_T > 240 \text{ MeV}/c$ with respect to the kink parent's direction (see Table 2.1). As the decay $K^- \rightarrow \mu^- \bar{\nu}_\mu$ is close to this cut ($p_{T,\text{max}} = 236 \text{ MeV}$), the measurement uncertainties lead to a probability of about 10% to exceed this cut. As the lifetime of a K meson is much longer than that of a τ lepton, this background can be suppressed at the required level by restricting the flight length of the kink parent to be less than 3 mm. This background is only present in charged-current interactions where the primary lepton is not recognized, which leads to an additional reduction by a factor of about 10. For neutral-current interactions there is no primary

lepton, but the neutral-current cross-section is smaller than the charged-current cross-section and therefore this background is suppressed by the ratio of neutral-current to charged-current cross-sections.

The decay of the negative charmed meson, D^- , cannot be eliminated in this way, as it has similar flight length ($c\tau = 312 \mu\text{m}$) and mass ($m_{D^\pm} = 1869 \text{ MeV}$) as the τ lepton. One of the contributions to this background is from neutral-current interactions where a $c\bar{c}$ quark pair is produced. As the cross-section for this process is relatively small and the associated charm quark (in a D^+ , D^0 , or charmed baryon) can also be detected in the emulsion, this background is low. The production cross-section for a single c -quark in charged-current interactions of anti-neutrinos is typically 20 times larger, but these are efficiently rejected by the detection of the primary positive lepton. However, the chance of not identifying the primary μ^+ in a $\bar{\nu}_\mu$ charged-current interaction is still about 15%. The fraction of $\bar{\nu}_e$ in the beam is about a factor 10 smaller, but the probability for missing the primary e^+ is about 50%. The contribution to the background from $\bar{\nu}_e$ is therefore still about a third of that due to $\bar{\nu}_\mu$.

The cross-section for the white-kink background (point 3) was mostly unknown. The cross-section is normally described as the mean free path λ between white-kink scatters as function of p_T . To measure λ , two experiments have been done. One experiment was a test for a new neutrino-oscillation experiment [169], the other was dedicated to a measurement of the white-kink cross-section for pions [170].

In any case, the sensitivity of the emulsion (500 eV for rendering a grain developable) ensures that the probability of a scatter without visible recoil is very low. The energy transferred to the recoiling object depends on the p_T of the kink and consequently on the kink angle and the parent's momentum. The minimum kink angle of 50 mrad, the $p_T > 0.24 \text{ GeV}/c$ cut together with a minimum energy requirement for the daughter meson ensures that the energy transfer to the recoil is so large that the probability to miss it is sufficiently small. As the white-kink background is proportional to the total track length considered, the maximum decay length of 3 mm limits this background to an acceptable level in the $\tau^- \rightarrow h^-$ channel. For muons, high-angle scattering is less likely and therefore white kinks are not an important contribution to the background for the $\tau^- \rightarrow \mu^-$ channel.

2.2 Neutrino beam

The neutrino beam is generated by dumping the CERN super-proton-synchrotron (SPS) proton beam on a target. Most of the produced hadrons are π^\pm and K^\pm mesons which escape the thin rods of the target and can decay in flight. The decays of these secondary mesons generate the neutrino beam, consisting mostly of ν_μ and $\bar{\nu}_\mu$ neutrinos and a lower flux of ν_e and $\bar{\nu}_e$ neutrinos. The flux of ν_τ and $\bar{\nu}_\tau$ is almost negligible. Mainly muon neutrinos are produced, because of the preferential decays $\pi^+ \rightarrow \mu^+ \nu_\mu$ and $K^+ \rightarrow \mu^+ \nu_\mu$ (and their charge-conjugates). The decay to $e^+ \nu_e$ ($e^- \bar{\nu}_e$ for π^- and K^-) is suppressed by a factor $(m_e/m_\mu)^2$ due to the parity violating nature of the weak interaction. The lepton in the two-body decay of the spin-zero mesons has to be produced with the wrong helicity which favours the decay to the heavier muon. The flux of ν_e and $\bar{\nu}_e$ comes from the decays $K^+ \rightarrow e^+ \nu_e \pi^0$ and its charge-conjugate which have a branching ratio of 4.82% and $K_L^0 \rightarrow \pi^- e^+ \nu_e$ and its charge-conjugate which have a branching ratio of 38.81%. The ν_τ and $\bar{\nu}_\tau$ in the beam come from the decays of short-lived D_s mesons.

However, the production cross-section for D_s is much smaller than that for π and K mesons. The branching ratio for $D_s \rightarrow \tau \nu_\tau$ is quoted as $[6.4 \pm 1.5]\%$ [1]. The energy spectra of the ν_τ and $\bar{\nu}_\tau$ in the beam have two contributions, as also the τ from the D_s decay will quickly decay giving rise to a second τ -neutrino.

Because of the relativistic energies of the secondaries, the decay products are boosted forward. At higher energies, the neutrino beam will be more focused and more energetic. At the same time, however, the decay length $\gamma c\tau$ becomes longer. Consequently, a smaller fraction of the secondaries will decay in the available decay space. Therefore, one has to make a trade-off between energy and intensity of the neutrino beam. The kinematics of the two-body decay of the π^+ and K^+ to $\mu^+ \nu_\mu$ determine the energy spectrum of the neutrino beam. In the center of mass frame the muon and neutrino are emitted back-to-back and the available energy is balanced between the muon and the neutrino. In the lab frame, the momentum of the neutrino is therefore only dependent on the angle θ at which the neutrino is emitted with respect to the direction of the parent meson. The transverse momentum of the neutrino with respect to this axis is given by:

$$p_T = \frac{m_{\pi,K}^2 - m_\mu^2}{2m_{\pi,K}} \sin \theta \quad , \quad (2.1)$$

and the longitudinal momentum by

$$p_L = E \left(1 - \frac{m_\mu^2}{m_{\pi,K}^2} \right) \left(\frac{1}{2} + \frac{1}{2} \cos \theta \right) \quad . \quad (2.2)$$

Because the kaon has more mass than the pion, the neutrino carries (on average) more of the total energy for kaon decays. Therefore, the higher-energy neutrinos in the beam are mainly due to K decays. Table 2.1 gives a summary of the properties of the muon and the π and K mesons.

	mass [MeV]	τ [ns]	$\gamma c\tau$ [m/GeV]	Br($\rightarrow \mu \nu_\mu$) [%]	$p_{T,\max}$ [MeV]	$p_{L,\max}$ [p_{parent}]
μ	105.7	2197.0	6233.8			
π	139.6	26.0	55.9	99.99	30	0.427
K	493.7	12.4	7.52	63.43	236	0.944

Table 2.1: Properties of the muon and the π and K mesons. The last two columns give the maximum transverse and forward momentum transferred to the neutrino, calculated using equations (2.1) and (2.2).

For the CHORUS experiment, a hard neutrino spectrum is desired, because the charged-current cross-section for ν_τ has a high energy threshold due to the τ mass and increases rapidly above this energy. As was explained in section 2.1.2, anti-neutrinos contribute to the experimental background and should be suppressed. Also the direct ν_τ flux in the beam should be as low as possible. For the neutrino beam used by the CHORUS experiment, this irreducible background is calculated to be $4.1 \cdot 10^{-6} \nu_\tau$ charged-current interaction per ν_μ charged-current interaction [171]. For a total of $4 \cdot 10^{19}$ protons on target, this corresponds to 0.18 ν_τ interactions.

Figure 2.2 gives a schematic overview of the components of the neutrino-beam. The setup is described in full detail in Ref. 172. The 450 GeV protons from the SPS accelerator are extracted in two ‘fast-slow’ extraction spills and dumped on the target. Each spill lasts about 6 ms; long enough to separate multiple interactions in the neutrino target; short enough to make pulsed operation of the focusing magnets possible. The target consists of eleven 3 mm thin beryllium rods separated by 9 cm. Each rod is 10 cm long. Using thin rods minimizes secondary interactions of the hadrons produced by the proton-beryllium collisions and therefore optimizes the neutrino flux and spectrum.




Figure 2.2: Schematic layout of the neutrino beam setup. The angular acceptance at the CHORUS site is about 0.8 mrad.

The number of anti-neutrinos is minimized by sweeping out the negatively charged hadrons by the magnetic fields of the horn and reflector. These are two specially shaped, single-winding, magnetic lenses which generate a pulsed toroidal field that focuses positively charged particles and defocuses negatively charged particles [173]. The focusing increases the energy-weighted beam flux by about a factor 5 and suppressed the unwanted anti-neutrino flux by a factor of 2. The tapered collimator matches the secondary particle beam to the horn aperture. The collimator in front of the reflector absorbs defocused negative mesons before they decay. The two helium tubes limit absorption and scattering of the mesons before they enter the 289 m long vacuum decay tunnel. The remaining protons and hadrons are absorbed in the first few meters of iron shielding at the end of the decay tunnel. The muons need more shielding and several sections of earth, concrete and iron shielding are in between the decay tunnel and the experimental area. The muon-flux is measured in three gaps in the first iron shield. These measurements are used to monitor the beam shape and intensity.

The large mass of the calorimeter (section 2.6.3) has been used to measure the beam flux as function of radius and energy. The results of these measurements are reported in Ref. 174 and chapter 4 of Ref. 175. In Ref. 175, the discrepancies with earlier reported values [172] for the horn and reflector efficiencies and the beam Monte-Carlo simulation is discussed (see also Ref. 176). An earlier measurement of the neutrino-flux from the interaction rate in the calorimeter is shown in Figure 2.3. A normal-distribution fit to the data shows that the neutrino-beam has a RMS width of 0.7 m at the emulsion position in CHORUS. The estimated charged-current ν_μ event rate in the emulsion target is $2.1 \cdot 10^{-14}$ events per proton on target.




Figure 2.3: Position and width of the neutrino beam determined by neutrino interactions in the calorimeter.

2.3 Experimental setup overview

To point from the underground neutrino target to the experiments in the surface buildings, the neutrino beam has an upward slope of 42 mrad with respect to the horizontal. The coordinate system used in this dissertation is the one used in emulsion scanning where the x and y axis are in a vertical plane, with x vertical. The z -axis points downstream and lies in the horizontal plane.

The position of the neutrino interaction is referred to as the primary vertex. The tracks coming from the vertex are Lorentz boosted forward with respect to the direction of the incoming neutrino. These tracks will lie within a cone which has an opening angle determined by the exchanged transverse momentum and the neutrino energy. All detector components (except the veto trigger plane) are therefore located downstream of the target and increase in size to cover the same solid angle. The components are installed vertically, but shifted upward with respect to each other to follow the beam slope. Figure 2.4 shows an overview of all the detector components and their layout.

The emulsion is mounted in a rigid frame which also houses the trackers, their layout is presented in section 2.4.4. Emulsion is prone to fading (see section 2.9.2) where the latent image of the particle tracks fades away in time. The speed of this effect increases with temperature. The fading can be limited by keeping the emulsion cold. Therefore, the whole target region is kept at $5 \pm 0.5^\circ\text{C}$ during its two years of exposure. The constant temperature also increases alignment accuracy by limiting thermal expansion. Therefore, the target region is located inside a large refrigerated volume, known as the coolbox. Inside the coolbox are the emulsion target, its associated electronic trackers, and the hadron spectrometer.

Just downstream of the coolbox was a gap of 21 cm where originally a set of streamer tubes was placed. These tubes were later replaced by a honeycomb tracker described in Chapter 3. Downstream of these trackers is the calorimeter, consisting of a 1 meter thick block of instrumented lead which absorbs almost all hadrons. Only muons with a momentum higher than $1.6 \text{ GeV}/c$ pass through the calorimeter and enter the muon-spectrometer. The hadron spectrometer, calorimeter and muon spectrometer are briefly discussed in section 2.6.




Figure 2.4: Schematic layout of CHORUS detector and its sub-detectors.

2.4 Emulsion target and electronic tracking detectors

To predict the position in the emulsion of tracks from a neutrino interaction, electronic tracking detectors are used. These detectors are placed directly downstream of the emulsion target. In this section, the design of the emulsion target and the tracking detectors is discussed. The guiding principle is to locate a neutrino interaction in the emulsion accurately and efficiently.

2.4.1 Emulsion target considerations

An emulsion target is made out of separate emulsion plates. An emulsion plate cannot be thicker than about 1 mm as the lenses used for scanning the emulsion have a working distance of about 1 mm. During development, the chemical reducing solution must be able to diffuse into the emulsion layer which also limits the thickness of the layers. The plates used in the CHORUS experiment have about $350 \mu\text{m}$ of emulsion deposited on two sides of a $90 \mu\text{m}$ thick plastic (tri-acetate-cellulose) base. These plates are known as the target plates. The emulsion target is made out of stacks of such emulsion plates. The plate size is $72 \text{ cm} \times 36 \text{ cm}$ and is limited by the equipment needed to pour, dry, develop, and scan the plates.

Plate orientation

In a target, the emulsion plates can be put either perpendicular or longitudinal with respect to the beam direction, as shown in Figure 2.5a & c. The choice of orientation depends on how the emulsion scanned. One consideration is if events are reconstructed by a human operator or by an automatic scanning station. Another factor is how interesting events are located inside an emulsion stack. In the case of a hybrid detector like CHORUS, that means that track predictions need to be followed back to the vertex in the emulsion. As is indicated in Figure 2.5, the interesting tracks for CHORUS lie in a forward cone with respect to the beam direction.

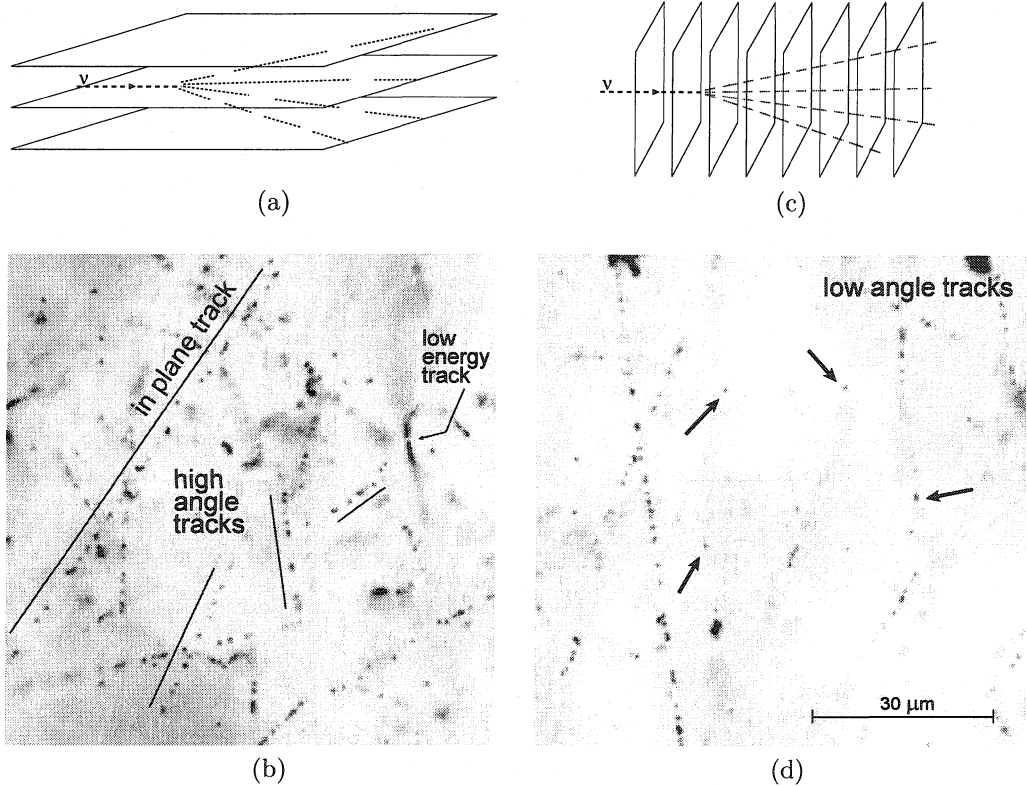


Figure 2.5: Difference in image for emulsion plates oriented longitudinal (a,b) or perpendicular (c,d) to the incoming beam. The emulsion images in (b) and (d) show only 240×240 pixels of the 1024×1024 pixel full image. The contrast in the images has been enhanced to let the grains stand out more.

If the plates are oriented longitudinally, the interesting tracks have large angles (θ_z) with respect to an axis perpendicular to the plate (z -axis). In a single microscope view of a piece of emulsion — typically covering an area of $150 \mu\text{m} \times 150 \mu\text{m}$ — these large-angle tracks are easy to recognize for a human operator. These tracks have several grains visible within the depth of field of the microscope which lie on a straight line, as is indicated in Figure 2.5b. If the depth inside the emulsion which is in focus is moved, grains will appear at one end of this line and disappear on the other, giving the impression of a traveling particle. These tracks are therefore easy to follow by a human operator. Once the track leaves the border of the view, the plate has to be shifted under the objective. Automatic track finding for these tracks is also relatively easy, because multiple hits of the track are visible in each view. To follow the track in the emulsion, the automatic track-finding needs to be done online, because the track leaves the microscope view fairly quickly which requires shifting the emulsion plate.

If the plates are oriented perpendicular, the interesting tracks have small angles θ_z . These tracks have typically only one grain per track visible inside the depth of field, as indicated by the arrows in Figure 2.5d. Visually, these tracks are identified by moving the depth in focus through the emulsion. The individual grains on a track will appear one after another. The position within the view of these grains will move slightly de-

pending on the track's angle. The closer a track's angle is to the z -axis, the smaller this movement and the more difficult it becomes for a human to spot the track in the background of randomly developed grains and the lines of the large-angle tracks. Using an automated scanning station, it is straightforward to take images at different depths inside the emulsion. One can then compute the correlation between grains at different depths and assign one grain per image to a particular track. Track finding can be done either online, for example in hardware, or offline from the analysis of a set of images.

For human scanning of an event in emulsion, it is easier if the plates are oriented longitudinally, because the tracks of interest lie in the plane of a microscope view. However, when a track crosses from one plate to another, it is difficult to locate the continuation of the track on the next plate because the alignment errors grow proportional with $\tan \theta_z$. Finding the prediction of a track requires a search at the outer edge of several emulsion plates; scanning through their full depth while looking for a track that matches in angle. Once the vertex is located, several plates need to be scanned at very different positions to reconstruct completely all the tracks from an event. Depending on a track's angle, a significant part of the track is missed where the particle crosses the plastic base of the plates.

If the plates are placed perpendicular to the beam, like in Figure 2.5c, the tracks cross all plates downstream of the interaction vertex at relatively small angles θ_z . Following a track upstream from plate to plate is done by scanning the upstream plate starting at the position where the track exited the downstream plate. An electronic prediction can be found by looking for an angular match in one or more microscope views around the predicted position on the most downstream plate. In the perpendicular orientation, all predicted tracks for a stack of plates can be located first on the most downstream plate, before the scanning of all found tracks on the next upstream plate. This step is then repeated until all interaction vertices have been located. Just downstream of a vertex, all vertex tracks will be in the same microscope view and can be reconstructed using the set of images already taken when following the predicted track to the vertex. Once all interaction vertices are located, all event related tracks can be followed in a second pass through all plates, now going downstream.

As several thousand events were subject to scanning using automatic scanning stations, the choice for a perpendicular orientation of the plates in CHORUS is obvious. Many discoveries using emulsion have used longitudinal exposures though, as can be clearly seen in the events that led to the discovery of the pion [187, 188] of which Figure 2.6 shows an example [177].

Stack thickness

A stack of target plates, called a module, is vacuum packed to preserve its water contents and to mechanically fix the relative positions of the plates. The neutrino beam has a radius of about 1.4 m at the site of the detector (see Figure 2.3). Therefore, 2×4 stacks of emulsion plates of 72 cm \times 36 cm are used to cover the beam cross-section. The eight modules are put in two rows of four modules with the long edge of the plates oriented vertically.

Given the density of emulsion, $\rho = 3.815 \text{ g/cm}^3$, the total stack thickness for 770 kg of emulsion would be 9.7 cm. With the radiation and interaction length of emulsion being $X_0 \approx 29 \text{ mm}$ and $\lambda \approx 35 \text{ cm}$, respectively, this thickness would represent roughly $3.4 X_0$ and 0.3λ . Absorption and showering in such a thick stack would lower the efficiency




Figure 2.6: Mosaic of micro-photographs of the decay chain $\pi \rightarrow \mu \rightarrow e$ from one of the first pion tracks seen in emulsion.

of reconstructing primary-vertex tracks downstream of the stack. The track parameters measured behind the stack are smeared due to multiple scattering [1]. Although, this smearing is not so important for finding the track in the emulsion, it does affect the vertex reconstruction accuracy. Another important consideration is that with a stack of about 140 plates, the scanning load would be high, because on average half of the number of plates needs to be scanned to follow a track back to the interaction vertex. For these reasons, the set of emulsion plates is split into four separate stacks, each containing 36 plates. Electronic trackers are inserted between these stacks to accurately predict the position of the primary-vertex tracks for each stack separately. A particle from the primary vertex now crosses on average only 18 plates before its track parameters are measured. The multiple scattering in the downstream stacks can be taken into account in the track fit. The average number of plates that needs to be scanned to locate the vertex in this configuration is also only 18 plates, instead of 72.

2.4.2 Interfacing emulsion and electronic tracking detectors

The design parameters of the electronic tracking detectors are mainly determined by the need to accurately locate a single track in the emulsion. The matching between tracks found in the emulsion (subscript 'e') and tracks reconstructed by the electronic tracking detectors (subscript 'p') is based on the χ^2 sum over four matching variables; the position and slope differences in x and y :

$$\chi^2 = \left(\frac{x_e - x_p}{\sigma_{xy}} \right)^2 + \left(\frac{y_e - y_p}{\sigma_{xy}} \right)^2 + \left(\frac{\theta x_e - \theta x_p}{\sigma_\theta} \right)^2 + \left(\frac{\theta y_e - \theta y_p}{\sigma_\theta} \right)^2 .$$

The χ^2 is mainly determined by the position resolution of the electronic tracking detectors and the angular resolution of the emulsion. If a 10% contamination of fake matches is allowed when scanning an area with sides of $3\sigma_{xy}$, then the matching should give less than 0.1 candidate for a random area of emulsion of this size averaged over the angular distribution of all tracks. The required resolutions are then determined by the track density in emulsion (ρ_{tracks}) and its angular distribution, i.e. :

$$\int_{3\sigma_{xy}} \rho_{\text{tracks}}(\theta_x, \theta_y) dx dy d\theta_x d\theta_y < 0.1 .$$

There are two ways to reach this goal: one, reduce the $3\sigma_{xy}$ volume; two, lower ρ_{tracks} in the emulsion. The latter can be achieved by reducing the exposure time. The slope resolution of emulsion is limited by distortion of the emulsion layers (section 2.9.2) to $\sigma_\theta \approx 15$ mrad for the target plates. A better slope resolution and a low track density can be achieved simultaneously by inserting special emulsion plates between the emulsion target and the electronic tracking detectors and exchanging them regularly. These plates are special in the sense that they use two $100\ \mu\text{m}$ thick emulsion layers on a $800\ \mu\text{m}$ thick plastic base. As distortion does not affect the position of a measured track at the emulsion–base interface ($\sigma \approx 0.5\ \mu\text{m}$), the slope of the track can be measured over the base with an accuracy of better than 1 mrad. This type of plate is known as interface plates.

Three of these interface plates are inserted in each stack. One, called special sheet (SS), is packed with the target plates and changed every year. Two others are placed between the emulsion stack and the first tracker plane and are called changeable sheets (CS). The changeable sheets are exchanged depending on the number of integrated tracks (beam-muons, X7 muon beam, and cosmic rays). Due to the increase in scanning power and the lower X7 intensity, the number of changeable sheet periods has been reduced during the experiment from 7 periods in 1994, 3 in 1995, 2 in 1996, to just 1 in 1997.

2.4.3 Tracking detector

The efficiency, resolution and two-track separation of the tracking detectors are important parameters for a reliable and accurate track match with the emulsion. The position resolution should be better than $160\ \mu\text{m}$ to limit the 3σ scanning area to $1\ \text{mm}^2$. The angular resolution should be comparable to that of the changeable sheets, i.e. $1 \dots 2$ mrad. As the tracking detector is close to the interaction vertex, the spacing between tracks and therefore the required two-track resolution is of the order of a millimeter. To limit absorption and reinteraction of hadrons before they reach the hadron spectrometer, the tracking detectors should also not interpose too much material. In the design of the tracking detectors, two other considerations for detectors were not important in the CHORUS experiment. The maximum detection rate is not an issue as the average event rate is less than 0.7 events per spill. Secondly, the occupancy in the detectors is low as the average number of primary tracks in a neutrino interaction is only 4.1 [178].

A good compromise between these requirements and building cost has been achieved using scintillating fibers. Plastic scintillators have relatively low mass (for solid-state detectors), are fast and efficient, but offer limited resolution as they are normally built in strips of several millimeters thick and several centimeters wide. Better resolution can be achieved using thin fibers, but at the cost of detection efficiency. High detection efficiency and good resolution has been achieved by stacking several layers of thin scintillating fibers. The fibers are read out individually using a CCD camera. Because the diameter of the individual fibers is small ($500\ \mu\text{m}$), a two-track separation at the level of about 1 mm is also achieved. The read-out of CCD cameras is normally slow with a read-out time of several milliseconds, but in CHORUS they can be used because the event rate is low. A similar optimization for low-rate and occupancy has been used several times in the experiment, for example in the honeycomb (section 3.4.1) and muon-spectrometer (section 2.6.4) drift-time measurements.

2.4.4 Target region experimental setup

Figure 2.7 shows the arrangement of the emulsion stacks and the electronic tracking detectors. The tracking detectors are referred to as the target trackers. The setup of the emulsion stacks and trackers consists of two identical sections, each containing two emulsion stacks and four tracking detectors. Each tracking detector consists of four rotated planes such that tracking of particles is possible in 3-D. The target trackers provide the missing time-resolution of the emulsion by uniquely matching a single track in the emulsion to a specific, electronically recorded, event. As discussed previously, good spatial and angular resolution and good two-track separation is crucial. The construction of the tracking planes is described below. The 3×4 tracking planes behind each pair of emulsion stacks provide the angular measurement and are sufficient to do stand-alone track reconstruction. The four planes between each pair of emulsion stacks are used to recover position accuracy for the upstream stack as the emulsion stack interposes about one radiation length of matter.




Figure 2.7: Layout of two emulsion stacks and the associated target trackers. This setup is identical to the setup of the other two emulsion stacks in the experiment.

The distance between the emulsion and the target trackers is a trade-off between two conflicting requirements, the two-track separation and the prediction accuracy. The changeable sheets, CS1 and CS2, are used to resolve this conflict by placing CS2 just 1 mm upstream of the first tracker plane. The CS1 plate is 14 mm further upstream and the actual emulsion stack another 38 mm. The two changeable sheets are mounted on a honeycomb panel which is traversed by 15 X-ray guns per emulsion module. A similar honeycomb panel with X-ray guns is placed between CS1 and the emulsion stack. The X-ray guns are brass cylinders with a ^{56}Fe X-ray source inside. The X-rays create a 1 mm diameter black dot on the surface of the two emulsion plates. These dots are used to determine the alignment between the interface sheets and the target trackers.

A detailed description of the target trackers can be found in Ref. 157. The target tracker consists of modules which contain four planes each. Each module contains one pair of horizontally and vertically oriented planes (XY) and one pair of rotated planes ($X^\pm Y^\pm$). The rotation angle is 8° and alters sign for successive modules. Each tracker plane is composed of seven layers of 2.3 m long scintillating plastic fibers with a diameter of $500 \mu\text{m}$. The far end of the fibers is coated with an aluminum mirror to increase the light-yield. The other end of the fibers is coupled to the camera. The light-output of a single fiber is too small to be detected directly by a CCD camera. Therefore, an opto-electronic image intensifier is inserted between the fibers and the camera. The image intensifier also demagnifies the image to match the fiber diameter to the CCD pixel size. The measured hit density is between 5 and 7 for a minimum-ionizing particle passing at 220 cm and 70 cm from the read-out end, respectively. The measured inefficiency of a plane is 0.2%. The disadvantage of the image intensifiers is that they need to be shielded from magnetic fields (even the earth's magnetic field) and therefore no magnetic (stray) fields are allowed in the target region.

The read-out of the CCD camera takes about 20 ms, but the CCD chip can store one image in a memory zone within $125 \mu\text{s}$. Using the memory zone, two events can be buffered in the CCD during the 6 ms beam spill. The buffered events are read out during the time between the spills. Because of the limited buffer capacity of the camera, the image recording needs to be delayed to allow for the application of a trigger signal. The image intensifier contains a multi-channel plate that can be electronically gated to expose the CCD only for triggered events. A fluorescence phosphor with a long decay time in the first stage of the image intensifier is used to delay the image. If the trigger enables recording of the event, the CCD captures about 30% of the light in a time window of $20 \mu\text{s}$ after the arrival of the scintillation light.




Figure 2.8: Position and angular resolution of the target tracker as measured by comparing scanning predictions to tracks found in the emulsion.

The track residual of a single target-tracker plane was measured to be around $180 \mu\text{m}$. The final resolution of the scanning predictions can be evaluated by comparing the predictions with the tracks found in the changeable sheet. The resulting distributions for muons, after alignment (section 4.6.3), are shown in Figure 2.8. The position resolution is $\sigma_{xy} \approx 190 \mu\text{m}$ and the angular resolution $\sigma_{\theta} \approx 2.3 \text{ mrad}$ after unfolding the 1 mrad emulsion resolution. Because of the asymmetric distribution of tracking planes around the emulsion stacks, these resolutions are different for each stack. The plots in Figure 2.8 are for the last stack which has the smallest amount of tracking planes behind it.

2.5 Trigger

The main purpose of the trigger system is to select primarily events due to neutrino interactions in the emulsion. For this, the emulsion target is surrounded by several scintillator planes, as shown in Figure 2.9. All planes are made out of two staggered planes of plastic scintillator strips. The trigger planes are coded as follows: E = emulsion, T = trigger, H = hodoscope, V = veto and A = anti-counter. Both the T and V plane provide accurate timing ($\approx 1 \text{ ns}$) by averaging the time of a detected hit on both sides of the scintillator strip (mean-time).




Figure 2.9: Schematic view of the trigger planes surrounding the emulsion target.

A coincidence of hits in the T and H planes indicates the presence of a charged particle that left the emulsion, while an anti-coincidence with the V-plane makes sure that no charged particle entered the emulsion. To avoid vetoing events due to back-scattered particles, the V-plane is put 2 m upstream of the emulsion. This gives a time difference of 13 ns between forward and backward going particles. The accurate timing of the V and T planes is used to distinguish between these two cases. Another important criterion in the trigger design is the high rate of beam related muons that must be efficiently vetoed which requires a high efficiency of the V-plane. In effect, the V-plane has an inefficiency of less than $1.5 \cdot 10^{-3}$.

As the mass of the material surrounding the target (metal supports, concrete & iron floors, shielding) is much larger than the target mass, many more interactions will take place around the target than in the target itself. The expected rate of neutrino interactions in the emulsion is about 0.34 events per spill at the maximum spill intensity

of $1.5 \cdot 10^{13}$ protons on target. The number of events (k) per spill is then (ignoring dead time) given by the Poisson distribution $P(k; \mu = 0.34)$. The CCD read-out of the target trackers limits the maximum number of events that can be recorded per beam spill to two. If the trigger would fire only for events in the emulsion, the recording efficiency is given by:

$$\varepsilon = \frac{1}{\mu} \left(P(1; \mu) + 2 \sum_{k=2}^{\infty} P(k; \mu) \right) ,$$

which yields for $\mu = 0.34$: $\varepsilon = 98.4\%$. If, instead, the total triggered mass with respect to the emulsion mass is larger by a factor f , then the average event rate μ' is $f \times \mu$. The recording efficiency will decrease as the average number of triggered emulsion events is now given by:

$$\langle k \rangle = \frac{1}{f} P(1; \mu') + \frac{2}{f} \sum_{k=2}^{\infty} P(k; \mu') .$$

The decrease in the recording efficiency $\varepsilon = \langle k \rangle / \mu$ for real emulsion events is shown in Figure 2.10 as a function of f . For a T +H + \bar{V} trigger $f = 6$ and the recording efficiency has dropped to about 72%. Most of the additional triggers are due to cosmic rays and neutrino interactions in the iron floor, the frame and read-out equipment of the target trackers, and the concrete floor in front of the experiment. Requiring an additional coincidence in the E-plane (installed in the 2nd year of data taking) improves the selection of emulsion events. Putting the A-plane in anti-coincidence removes events from the concrete floor. Cosmic rays and events from the iron floor are suppressed by requiring a hit combination in T and H consistent with a particle track with $|\tan \theta| < 0.2$ with respect to the neutrino beam. The size of the V-plane is such that any incoming cosmic ray not hitting the V-plane crosses the emulsion at a larger angle than this. The final trigger rate for emulsion events corresponds to a total mass of 1700 kg ($f = 2.2$). For typical spill intensities of $1 \cdot 10^{13}$, the expected efficiency is then 96.7%, which is accounted for as dead time of the detector. The acceptance of the trigger to neutrino interactions in the emulsion target has been estimated to be 99%.




Figure 2.10: Relative recording efficiency as function of the ratio of the total trigger mass and the emulsion mass due to the two-event limit of the CCD read-out.

Several other triggers are made for the sub-detectors and other physics. The details of these triggers and the hardware and software implementation of the trigger logic are described in Ref. 159.

2.6 Downstream detectors

The detectors downstream of the target region are used for particle identification and for energy and momentum measurements. Originally foreseen to measure the kinematic variables for a pre-selection of τ -decay candidates, in the final analysis they are mainly used to assign charge, momentum and energy to the tracks found in the emulsion.

2.6.1 Hadron spectrometer

The hadron spectrometer is placed directly behind the emulsion target and target trackers. It consists of three scintillating fiber trackers, called the diamond trackers, which are placed around a magnet. The purpose of this spectrometer is to measure the momenta of hadrons up to about $20 \text{ GeV}/c$. This spectrometer is also used to measure the momentum of muons of less than $2 \text{ GeV}/c$ that do not reach the muon spectrometer (section 2.6.4). The momenta of the hadrons must be known in order to suppress some of the background (section 2.1.2). All the components of this detector must be light to minimize multiple scattering and showering which would affect the momentum resolution and the energy measurement in the calorimeter located downstream. The spectrometer magnet must have a very low stray-field because of the image intensifiers used in the read-out of the both the target trackers and the tracking planes in the hadron spectrometer itself (image distortion). The depth must also not be too large in order to keep the lateral dimensions of the downstream detectors reasonable for the same solid angle.

A solution was found by using a superposition of toroidal magnetic fields. Toroidal fields have closed field lines and therefore the stray-field outside the windings is very weak. Another advantage of a toroidal field is that the B-field is perpendicular to the particle's direction which gives the largest bending power. The main disadvantage of a standard toroidal magnet is that the material of the windings is in the particle's path and that all the windings cross at the center of the magnet. A compromise was found by using very light aluminum windings (0.04 radiation length) and distributing the center windings. The center windings are spread out over six spokes of a hexagonal shaped magnet. The magnet, shown in Figure 2.11a, was specially designed for the CHORUS experiment and is described in more detail in Ref. 161. It has a depth of 0.75 m and consists of six equilateral triangular sections with 1.5 m wide sides. The field inside a triangle is homogeneous and the field lines are parallel to the triangle's outer edge. The field in the triangles is homogeneous because the number of windings per unit length contributing to the field at any point is constant. The magnet is pulsed synchronously with the beam spills and has a field strength of 0.12 Tesla. The field is oriented such that negative particles are focused. The overall current running along the windings creates a single winding running once around the whole magnet which creates a solenoidal stray-field. This was compensated for by winding the feeding wires once in the opposite direction along the outer rim of the magnet.

The target trackers (section 2.4.4) give an accurate measurement upstream of the magnet. The same detection technique used for these trackers has been used for the tracking detectors around the magnet. Upstream of the magnet there is one tracking module (DT1) and downstream there are two (DT2, DT3). Each module consists of two




Figure 2.11: Hexagonal magnet (a) and arrangement of surrounding tracking detectors (b) of the hadron spectrometer.

layers of three diamond-shaped paddles, with the second layer rotated by 60° , as shown in Figure 2.11b. Each paddle is made out of seven layers of scintillating fibers. Of the two paddles which cover a magnet triangle, one measures the coordinate parallel to the base of the triangle (perpendicular to the bending plane) and the other a coordinate rotated by 60° . In this configuration, the DT1 module upstream of the magnet provides an accurate measurement of the entry point of a track reconstructed in the target trackers. The two planes behind the magnet, DT2 and DT3, are oriented such that for each magnet triangle they provide two measurements in the bending plane and two coordinates rotated by respectively $+60^\circ$ and -60° with respect to the bending plane.

2.6.2 Streamer-tubes and honeycomb detector

The four measurements from DT2 and DT3 behind the magnet are not sufficient to perform stand-alone tracking. Four planes of streamer tubes, recovered from the CHARM II detector [179], were placed in the 21 cm gap between calorimeter and coolbox. These 1 cm by 1 cm streamer tubes give additional hits to aid the track finding. In 1995, at the beginning of the second year of data taking, two additional planes were added.

However, the streamer tubes had a limited resolution and small stereo angles (7°) between the planes. A new tracker which could be used for stand-alone 3-D track reconstruction was proposed to replace the streamer tube planes. This new tracker, the honeycomb tracker, was installed halfway the 1996 data-taking run. Its construction and read-out electronics are the subject of Chapter 3. The new honeycomb tracker turned out to be essential in determining the alignment of the diamond tracker paddles. This alignment was then applied to pre-honeycomb events to improve the momentum resolution of the hadron spectrometer for the data from 1995 and beginning of 1996.

2.6.3 Calorimeter

The energy measurement, originally needed for τ -decay candidate pre-selection, is done using the calorimeter. The calorimeter in CHORUS requires tracking capabilities, because a muon which passes through the calorimeter and is detected in the muon spectrometer, must be connected to the corresponding track in the target trackers. For this, streamer tubes were recovered from the CHARM II experiment and interspersed between planes of calorimeter modules.

The energy resolution depends mainly on the ratio of active to passive material. The active material is used to measure a fraction of the total energy. The passive material is used for the development of the shower. Any energy deposited in the passive material is not measured, but is assumed to be proportional to the energy deposited in the active material. Due to pair production in electro-magnetic interactions of electrons and positrons, the deposited energy can vary with small spatial dimensions. To have an accurate measurement of the energy in an electro-magnetic shower requires, therefore, fine-grained sampling inside the passive material.

The calorimeter is constructed from bar-shaped modules made out of lead with plastic scintillator as active material. Each module has a separate read-out channel (one photo-multiplier read out with an analogue-to-digital converter). The high-sampling rate is obtained by interspersing many small-diameter scintillating fibers inside the lead. All fibers from a single module are read out by a single photo-multiplier tube. The ratio of lead and scintillator mass is chosen such that differences in shower development for electrons and hadrons are compensated in the total energy measurement [180]. The size of individual modules of the calorimeter is adapted to the need to measure the energy deposited by individual particles.

The calorimeter contains three sections. Each section consists of several planes of the lead-scintillator modules placed perpendicular to the beam direction. The planes are oriented alternatively horizontally and vertically. Electrons and positrons deposit most of their energy in the first section. This section consists of four planes constructed from 4 cm wide and 4 cm deep bars. The second section contains five planes with bars of 8 cm wide and 8 cm deep. In the final section, the electro-magnetic component of the shower is so much reduced that fine sampling is no longer necessary. The bars in this section are constructed with alternating layers of lead and scintillator strips. These bars are 10 cm wide and 10 cm thick.

In total, the calorimeter represents 5.2 hadronic interaction lengths and 144 radiation lengths. Showers of 5 GeV hadrons are fully contained in 99% of the cases. Additional information about the calorimeter and its performance can be found in Refs. 163, 164.

2.6.4 Muon spectrometer

The calorimeter acts as a muon filter because the muons are the most likely to pass through, having no hadronic interactions and much smaller radiation losses than electrons. The muon spectrometer, placed downstream of the calorimeter, measures the charge and momentum of muons. The magnetic field necessary to measure the momentum and charge is generated in iron disks. At maximum magnetization, the field reaches 1.7 Tesla inside the iron disks using an electric current of 700 A. The toroidal field in the disks is oriented such that negative muons are focused.

The muon spectrometer consists of seven tracking sections interleaved by six magnets. Each magnet is made from twenty iron disks with scintillator planes interspersed. The fast scintillator planes are used to measure the muon arrival time and are used in the trigger. They are also used to measure the energy of hadron showers leaving the calorimeter. In this way, the first few modules of the spectrometer work as a tail catcher for the calorimeter. The magnet modules were recovered from the CDHS experiment [181]. The tracking sections consists of wire chambers also recovered from CDHS [182] and streamer tubes recovered from CHARM II [179]. The layout of one section is shown in Figure 2.12a.

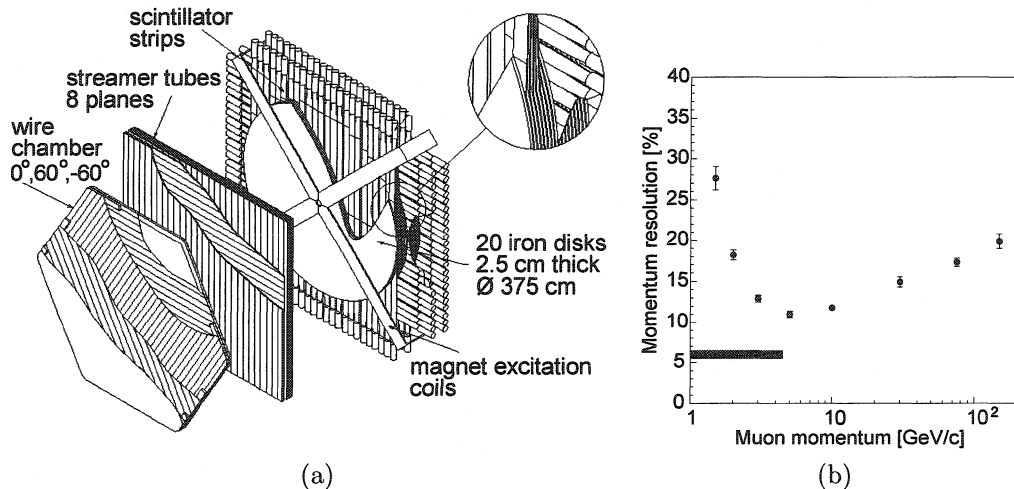


Figure 2.12: Exploded drawing of a single section of the muon spectrometer (a). The resolution of the muon spectrometer as function of muon momentum is plotted in (b). For muons with $p_\mu < 7 \text{ GeV}/c$ the momentum is determined by the range, indicated by the gray bar at 7%.

The wire chambers and streamer tubes are read out using time-to-digital converters (TDCs). As the occupancy in the spectrometer is low, the drift-time measurement in the wire-chamber is multiplexed with four wires connected to a single TDC channel. For the streamer-tube planes, each tube is read out digitally and 4×8 tubes are grouped together to a single TDC channel with 10 ns accuracy. The four groups of eight cells are time-multiplexed by delaying the signals of each group. To get a better drift-time measurement in the streamer tubes, all 352 tubes in a plane are also wired together to a single 1 ns resolution TDC channel.

The muon momentum is measured from the bending in the six magnet sections. The resolution $\Delta p_\mu/p_\mu$ as function of the momentum p_μ is shown in Figure 2.12b. The resolution is limited to about 12% due to multiple-Coulomb scattering inside the iron of the magnets. At momenta above 10 GeV/c, the resolution of the tracking sections also contributes to the resolution. A better momentum measurement can be achieved for muons that stop inside the spectrometer from their range [183]. For stopping muons ($p_\mu < 7 \text{ GeV}/c$), the momentum resolution is about 7%, indicated by the gray bar in Figure 2.12b.

2.7 Online monitoring

The hardware and software of the CHORUS data-acquisition system is extensively described in Ref. 160. Here, only the monitoring and visualization system is described. This system reads histograms created by online tasks and displays them on request of the user. A graphical user interface (GUI) was written using Tcl/Tk [184] and the Tix extension [185]. This interface has a connection with a C++ program that reads and displays histograms using the CERN library HIGZ.

Tix uses a kind of class for the widgets it adds to Tk. Using a similar design, the GUI could be written in an object-oriented fashion, even in a scripting language like Tcl/Tk. A class concept was designed for Tcl scripts which makes it possible to define classes and methods in Tcl and to instantiate objects. These objects become Tcl commands that invoke the methods given to them. The Tcl interpreter made it possible to do this in plain Tcl code. Using these Tcl classes, common classes for displaying parts of the user interface can be reused.

The object-oriented C++ program contains a class hierarchy for all the different histogram types of the CHORUS histogram libraries and also encapsulates the CERN library HBOOK histograms. A bridge class for the HIGZ histogram plotting toolkit was created which reorganizes the page and histogram layout options of HIGZ. These page and histogram layout classes define all values, distances, and sizes of items (like headers, labels, axis values, etc.) in a consistent top-to-bottom and left-to-right order and automatically resize the histogram zones to fit everything in the window. Displaying histograms is then done by specifying a page-layout (containing title, legends, etc.) and histogram layout (containing labels, histogram header, etc.) to the abstract plot method in the histogram classes. The plot method can be called with different layout objects. Layout objects were defined, for example, to display histograms with all axis and values attached and to create a compact display of many histograms using the same scale and only one set of axis labels.

A tool was developed to map C++ classes to corresponding Tcl classes. This tool was used to make the histogram class hierarchy, layout classes, and HIGZ interface class directly available as Tcl classes and objects. Using these classes, all histogram manipulations, layout, and plotting could be done directly from a Tcl script. For example, the title of a HIGZ page or the axis label of a histogram become simple variables of a Tcl object and can be manipulated using standard Tcl commands. The final trick, necessary to create a fully interactive histogram monitoring task, was to capture the main HIGZ window inside a Tk window. This allows (Tk detected) mouse events on the HIGZ display to become active events. Having achieved that, it was then easy to read histograms and display them in all kinds of configurations from a Tcl/Tk script. Histograms of different runs can be added to look at averages, overlaid in a single plot to spot differences between runs, or automatically compared to a reference histogram. With a mouse-click, a histogram can be selected for individual display, its scale changed or bin values extracted, for example to locate the exact dead and hot channels in some detector.

Figure 2.13 shows the GUI for selecting a set of histograms and runs to display. The pull-down menus on top select between the main categories of histograms for all the sub-detectors and some derived values calculated online. The left-panel in the window shows a detailed layout of histogram sub-categories which can be selected using buttons, radio-buttons and active graphical elements. In the right panel all runs accessible from disk are shown. One or more runs, including (if present) the currently active run, can

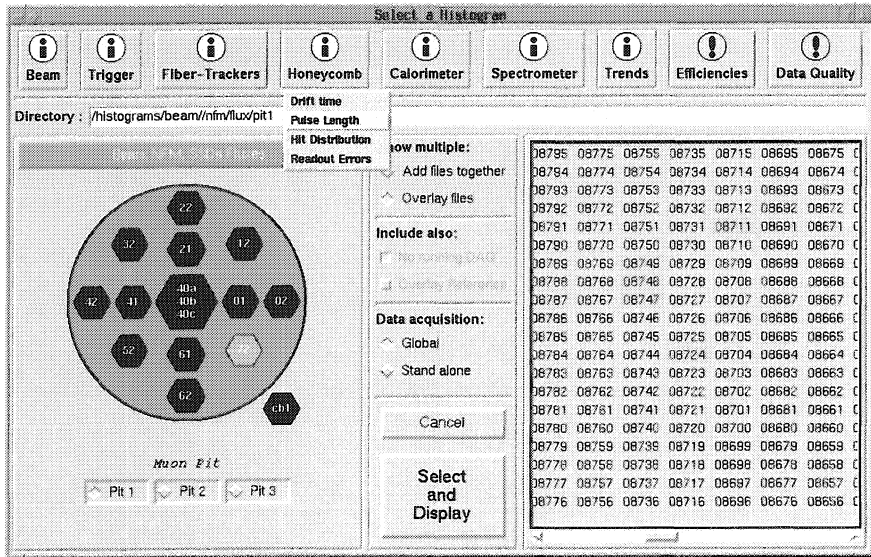


Figure 2.13: The graphical user interface for online monitoring. This screen-shot shows the configuration used to select histograms for the muon counters in one of the beam-line pits. In this case, the left-side of the GUI selects between different counters in the muon pits, the right-side selects one or more run numbers to be displayed.




Figure 2.14: Histograms of honeycomb drift time distributions (see also Chapter 3) in the online-monitor main window. A left mouse click pops up a position sensitive menu with, for example, the 'zoom' submenu to manipulate the display options of one or more histograms. A separate 'Bin Contents' window at the bottom shows the contents of the bin the mouse cursor is over.

be selected in this window and then displayed, either added together or all overlaid in a single plot. These kind of options are enabled or disabled with the buttons in the middle column. Clicking on the ‘Select and Display’ button brings up the main window showing the histograms for the runs selected. The histograms shown can then be manipulated using the mouse as described above. A screen-shot of the displayed histograms and the mouse activated menu is given in Figure 2.14, in this case for honeycomb drift time histograms (see also Chapter 3).

2.8 Track reconstruction and scanning predictions

The offline analysis program (CHANT) is used to reconstruct events from the data [186]. One of its main tasks is to select one track for each event which should be followed down into the emulsion to locate the vertex. These tracks are referred to as the scanning predictions or scan-back track. The scanning predictions are selected from the tracks reconstructed in the target tracker (section 2.4.4). First a vertex position is calculated from the target-tracker tracks. If the vertex lies inside an emulsion stack within the uncertainty, a scanning prediction is generated from the tracks attached to the vertex.

The track finding starts with clustering the CCD images recorded by the target tracker. These clusters correspond to the center positions of track segments in the seven fiber layers. In the second step, all combinations of hits in each module of four planes are considered. Track segments inside a module should have a hit in at least three out of the four planes. The track segments from the four modules in a target section are then combined into tracks using a minimum spanning tree. Connected branches in this tree yield an initial estimate for a common vertex. If the event is located in the first two stacks, then the tracks in the upstream target section are combined with tracks and hits in the downstream section.

Each found track is fitted taking into account multiple scattering in the various materials between the hits. The track fit yields a χ^2 -probability $P(\chi^2)$. In this calculation, the hit resolution is assumed to be $\sigma = 180 \mu\text{m}$. Two effects cause the $P(\chi^2)$ distribution to deviate from the expected uniform distribution. First, the momentum of the track can be unknown. In this case, a momentum of $1 \text{ GeV}/c$ is assumed. Depending on the real momentum of the particle, this can over or under estimate the multiple scattering contribution to the hit resolution. Therefore, a cut on $P(\chi^2)$ introduces a bias on momentum, selecting preferably high momentum tracks. Second, the hit-residual distribution has tails with a RMS width of about $750 \mu\text{m}$. These tails are probably due to cross-talk in the input window of the image intensifier or backscattering in the multi-channel plate. This leads to reconstructed tracks with very low χ^2 -probability. However, these tracks have been found in the emulsion.

A vertex is defined as the point of closest approach between tracks. If a vertex is reconstructed in one of the emulsion stacks, the event is marked for scanning. There can be several vertices in an event, in which case a quality value is assigned to each vertex. The quality value is defined such that it is highest for the vertex that is most likely the primary vertex. It takes into account the presence of a muon track, the number of tracks attached to the vertex and the vertex position (more upstream higher value).


The increase in scanning power has made it possible to remove the selection of the most likely τ -decay daughter as the scanning prediction. Instead, the track most likely to be found in the emulsion is taken as prediction. If the event contains a vertex with a primary muon, the muon track is always chosen as prediction. Muons are identified independently

in the muon spectrometer (section 2.6.4). A matching is attempted between the target-tracker tracks and the spectrometer muon tracks in order to identify primary muons. If there is a muon in the spectrometer but no target-tracker track can be matched to it with good probability, the same selection procedure as for events without a muon is used. For these events, one track is selected from those attached to the primary vertex with the additional requirement that it has a fit probability $P(\chi^2) > 10\%$. Because of the non-uniformity of the $P(\chi^2)$ distribution, the $P(\chi^2)$ cut is lowered to 10^{-4} if no track in the event has an acceptable $P(\chi^2)$ value. In all cases where there are several candidate tracks, track isolation is used as a tie-breaker. The selected tracks are then refitted to give the impact point on the most downstream plate of the emulsion stack containing the vertex. Because of the distribution of emulsion and target-tracker planes, the various emulsion stacks show different reconstruction and event-location efficiencies.

2.9 Emulsion and scanning techniques

Nuclear emulsion has a rich history of scientific discoveries, starting in 1896 with the accidental discovery of radioactivity by Becquerel (see Figure 2.15). In 1947, Powell and his collaborators discovered the pion in emulsion plates exposed to cosmic rays [187,188]. This discovery was quickly followed by other discoveries of elementary particles with emulsion, for example, the charged kaons, the Λ^0 [189] and $\bar{\Lambda}^0$ [190], and the Σ^+ [191,192]. In hind-sight, also the Ω^- was first seen in emulsion [193]. Emulsion experiments also played a crucial role in elucidating the parity-violating decays of the charged kaons, the so called $\tau-\theta$ puzzle [194–197]. The power of emulsion is maybe best illustrated by a rocket experiment in which a few postage-stamp sized emulsion pieces gave detailed information about the components and spectra of the earth's Van Allen radiation belts [198]. For over a decade, emulsion remained one of the prime tools in experimental particle physics. After that, emulsion was largely replaced by bubble chambers and electronic detectors which eliminated the need for laborious manual scanning. Only for very few experiments did the benefit from emulsion's unrivaled spatial resolution outweigh the drawbacks of manual scanning.

Figure 2.15: The original photographic emulsion plate which led to the discovery of radioactivity by Becquerel. The plate, wrapped in opaque paper, was left near uranium salts and became dark even though it was never exposed to light.



2.9.1 Hybrid experiments and automatic scanning

In the last three decades, the drawbacks of the manual scanning of emulsion have been remedied, largely due to two interrelated developments. First, the combination of nuclear emulsion with electronic detectors bridges the gap between the micrometer position resolution of emulsion and the nanosecond time resolution of electronic detectors. Such hybrid experiments offer the best of both worlds: real-time measurements in the electronic detectors and high-resolution track and vertex information in the emulsion. The idea of hybrid experiments has led to the revival of emulsion. Second, the development

of automatic scanning techniques opens up the possibility to analyze several hundred-thousand events within a reasonable amount of time.

With scanning predictions from hybrid experiments, only a minute fraction of the emulsion volume has to be inspected. In 1965 an experiment along those lines was performed using a neutrino beam at CERN and seven neutrino interactions were found [199, 200]. The first large-scale hybrid experiment was WA17 at CERN, exposing 31.5 ℓ of emulsion in front of BEBC, the Big European Bubble Chamber. The WA17 experiment located 169 neutrino charged-current interactions, including eight candidates of a charmed-particle [201].

In the middle of the eighties, a breakthrough came with the development of the first semi-automatic scanning techniques. With the plates oriented perpendicular to the beam, the emulsion was used as a tracking device rather than as a visual display. These techniques were pioneered by the E531 experiment [202] at FNAL which exposed 58.6 ℓ of emulsion to a neutrino beam. To evaluate the advantages and drawbacks of different scanning techniques, part of the emulsion was oriented parallel to the beam, part of it perpendicular. The E531 experiment can be considered the direct predecessor of CHORUS. With the continuing work of the Nagoya university in Japan, fully automatic scanning became possible in the beginning of the nineties [203]. Using these techniques, the E653 experiment at FNAL performed a study of hadronically produced heavy flavour states using a 71 ℓ emulsion target, exposed to an 800 GeV proton beam and a 600 GeV pion beam [204]. In total, more than 50,000 interactions were located; the proton exposure yielded 146 identified charm events [205] and the pion exposure yielded nine beauty pairs [206]. A review of hybrid experiments, focusing in particular on neutrino experiments, is given in Ref. 207.

The development towards more massive emulsion targets in which a larger number of events could be located with the aid of automatic scanning have led to the CHORUS experiment. To this date, the 400 ℓ of emulsion used in this experiment remains the largest emulsion target ever built. Also its two year exposure is exceptionally long. A detailed description of the pouring and processing facilities used for CHORUS can be found in [208]. In 2007, OPERA will start operating needing at least 12,000 ℓ of emulsion.

The following quote from the beginning of Barkas's first book on emulsion — considered as the standard texts on nuclear emulsion [209,210] — puts all of the above in perspective:

“Nevertheless it should be pointed out that emulsion has no great advantage when one is looking for a particle of predicted properties. It is of greatest use for discovering utterly new things, the anomalous behavior of which often can be recognized from a single event.”

A recent example of how telling even a single event may be is the observation by the CHORUS experiment of neutrino-induced diffractive production of D_s^* followed by the decay chain $D_s^* \rightarrow D_s^+ \gamma$, $D_s^+ \rightarrow \tau^+ \nu_\tau$ and $\tau^+ \rightarrow \mu^+ \nu_\mu \bar{\nu}_\tau$, reported in Ref. 211.

In the following sections, the typical properties of emulsion will be discussed, followed by a description of the scanning microscopes and the automatic track finding hardware.

2.9.2 Emulsion characteristics

Probably the best short-description of emulsion as a particle detector is given by Barkas, to quote from Ref. 209:

“Silver-halide emulsion of the type used for registering the tracks of charged particles consists of about equal parts by volume of 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. [...] With a microscope the paths of charged particles that 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.”

Latent image formation

Nuclear emulsions are made from a suspension of tiny silver-halide crystals, usually AgBr, with typical crystal sizes in the order of a few tenths of a micron. The composition of nuclear emulsion and the physical processes underlying the formation of the image and its development are largely identical to those for photographic emulsion. The chief differences are the larger thickness of the plates — up to several hundred μm as opposed to few μm in the case of photographic film — and a higher density of crystals. Photons or a passing charged particle can excite electrons from the valence band to the conduction band, leaving behind positively charged holes. The electrons are free to move through the crystal and can get trapped at internal lattice defects or special surface sites formed by colloidal silver or silver-sulfide. There they reduce silver ions to neutral silver atoms which form little centers of so called photolytic silver, referred to as the latent image. During development, these silver specks, which can be as small as three silver atoms, serve as growth centers for larger silver crystals, referred to as grains.

Normal photographic emulsion is not sensitive enough to detect minimum ionizing particles. A through-going charged particle is not very efficient at rendering a crystal developable. A particle creates many electron-hole pairs in a crystal in a very short time which greatly enhances the recombination probability. Furthermore, photons from ionization have energies far above the excitation threshold of 2.5 eV for AgBr. Whereas about 30 eV, deposited by a few visible light photons, suffice to render a silver-halide crystal developable, a particle must deposit between 500 eV and 3 keV to achieve the same. Especially for small crystals, a relatively small fraction of the energy loss effectively contributes to the formation of the latent image. Addition of several inorganic (like sulfur or gold) or organic molecules (dyes) sensitizes the emulsion so that minimum-ionizing particles can be seen. The processes responsible for this are still not well understood and some ingredients are kept secret. Sensitization probably works by creating additional internal lattice defects or surface sites where silver specks can form easier and are stabilized. The sensitivity of emulsion is usually measured in the number of developable grains per 100 μm track length. The world record for minimum-ionizing particles is about 70 grains/100 μm , the CHORUS emulsion has about 30 grains/100 μm .

The binding potential for latent image specks, especially if they are composed of only a few atoms, is relatively shallow and, with time, the latent image of a particle track gradually fades when the silver atoms re-oxidize. This process, known as fading, is strongly temperature dependent and enhanced by the presence of oxidizers in the emulsion. Fading can therefore be reduced by lowering the temperature and deploying the emulsion in a protective atmosphere of nitrogen or in vacuum. Tests on the CHORUS emulsion showed that the rate of fading is negligible for an exposure of up to two years if the emulsion is kept at 5°C and vacuum packed. The CHORUS emulsion was monitored during its exposure using small samples developed at intervals, to check for fading and fog (explained below).

Gelatin matrix

The value of gelatin as the emulsion matrix stems from a unique combination of properties. It encapsulates the silver-halide crystals, allowing them to be dispersed and to produce a uniform distribution of crystals. The components of gelatin, mostly organic molecules containing sulfur compounds, are photographically sensitizing. The gelatin has an exceptional mechanical strength, but does not prevent the penetration of watery solutions needed for pouring and development. As gelatin binds strongly to silver and silver-halide crystals, it keeps both the undeveloped silver-halide crystals and the silver grains in place, even though the emulsion can swell by as much as a factor of ten during development. The cross-linked network of long molecules in the gelatin ensures that the crystals or grains return to their original position when the emulsion is dried, except for some distortions which are discussed below. Unfortunately, gelatin is strongly hygroscopic, causing it to swell with increasing humidity and to shrink when drying. Therefore, emulsion plates need to be kept in a controlled environment with stable humidity.

Development

The latent image is converted into a stable image in the development process. Development works by reducing the silver-halide crystals and depositing silver on the latent image specks. The silver-halide crystals that are not part of the latent image are dissociated into their respective ions which are evacuated from the emulsion in the fixing stage. The various procedures are similar to those used in photographic development, but the larger thickness of nuclear emulsion plates leads to greatly increased soaking times. The development process is based on a delicate balance between different reaction rates, the formation of metallic silver aided by the presence of latent image centers and reduction and dissolving of silver-halide.

Because of the immense number of crystals present and the differential nature of the development process, some random crystals will become developed as well, leading to a background of track unrelated grains, known as fog. Various effects lead to a higher fog density at the surface of the emulsion. During development, the surface layer of the emulsion will develop more rapidly than the bulk, causing a built up of fog at the surface. The surface could have been exposed to light, which does not penetrate more than a few micron in the undeveloped emulsion. The surface is also subject to mechanical abrasion. Both processes cause excessive development leading to black spots and streaks. These two processes are actually exploited to print patterns on the emulsion. Using light, reference points are printed on the emulsion with a carefully aligned mask placed on the undeveloped emulsion. The stack and plate numbers are written with a blunt, non-writing, pencil and they show up after development. To be able to scan a developed plate under a microscope, the dark surface layer is removed without removing also the reference marks.

Resolution, shrinkage and distortion


The power of emulsion is its ability to visualize the paths of particles in 3-D with high resolution. The silver spheres that are produced during development tend to be larger, by up to a factor of two, than the corresponding silver-halide crystals. This growth is largely uniform. The resolution is therefore determined by the crystal diameter in the undeveloped emulsion, which ranges from $0.05\ \mu\text{m}$ to $0.5\ \mu\text{m}$. This resolution should not be taken at face value, because the intrinsic resolution is only meaningful to the extent

that measurements can be made in the same reference frame as where the emulsion has been exposed. Shrinkage as well as distortions will affect the practical resolution.

Shrinkage results from the removal of the undeveloped crystals during development. The silver-halide contents is about 50 % in volume of the exposed emulsion whereas only a very tiny fraction contributes to the latent image. During development, the undeveloped silver-halide crystals are removed and therefore the thickness of the emulsion layer is much reduced after development. The lateral area is unaffected because it is determined by the size of the mechanical support plate. The shrinkage factor, defined as the emulsion thickness at the time of exposure divided by the thickness at the time of scanning, is typically about two. However, the shrinkage factor varies with the water content of the emulsion. Because of the hygroscopic nature of gelatin, the shrinkage may vary with the scanning conditions. For instance, during scanning the top surface is typically insulated by a thin layer of immersion oil whereas the bottom surface is kept under vacuum leading to dessication. As a consequence, the shrinkage factor may vary along the depth.

Although the gelatin network is exceptionally strong, it is also very flexible and thus free to distort. This actually happens when tension built up during pouring and drying is released during development. During processing, the emulsion swells in the development baths. These tensions can then relax, causing the emulsion to distort. Additional distortions appear because the emulsion tends to swell in all directions whereas the surface mounted on the plastic base is constrained to retain its dimensions. The fact that particle tracks in developed emulsion do not appear as straight lines is generally referred to as distortion, regardless of its origin.

Figure 2.16: Effects of distortion and shrinkage on one side of an emulsion plate. The mechanical support plate is located at the bottom (dashed area). With respect to the original track (dashed line), the distorted track has a parabolic shape (thick solid line).



The most common form of distortion results in a parabolic shape of tracks. This kind of distortion results from shearing stress in the emulsion layer which is proportional with the distance to the base. Because the emulsion is fixed at the base, the stress force there is zero. The release of a stress that increases linearly with the distance to the base leads to a displacement which is quadratical with the depth in the emulsion. Because the outer emulsion surface is not constrained, it will have more or less the same shear down to a certain depth. The net result is that the position at the support plate is unchanged, whereas the slope of the track at the outer surface is free of distortion. Figure 2.16 schematically represents the combined effect of shrinkage and this type of distortion. Because the base position is unaltered, accurate angular track information can be obtained from double-sided emulsion with a relatively thick base. This idea was first suggested in Ref. 212 and has been used for the interface sheets as discussed in section 2.4.2.

2.9.3 Scanning microscopes

The layout of a microscope for emulsion scanning is depicted in Figure 2.17. The microscope consists of a movable stage on which the emulsion plate is mounted, a high numerical-aperture but low power objective, and a light source. The stage movement is in the plane perpendicular to the optical axis. The stage is used to position a selected part of the emulsion precisely under the objective. The objective creates a magnified image of the emulsion with a limited field of view. The size of the view depends on the magnification and the size of the image sensor. It is typically around $150\ \mu\text{m}$. Because of the magnification the depth of field is small (typically a few microns), meaning that only a thin slice of the emulsion is in focus. Moving the objective along the optical axis brings slices at different depths in focus. Normally, emulsion microscopes use trans-illumination where a condenser lens focuses the light through the emulsion. With this kind of illumination the grains in the image show up as tiny dark dots on a more or less uniform light-gray background. These dots are the shadows of the corresponding silver spheres that lie within the depth of field. The gray background reflects the fact that gelatin and the plastic base are transparent in most of the optical spectrum and that the grains out of focus cast very dispersed shadows. As was shown in Figure 2.5, tracks parallel to the emulsion surface are seen within a field of view as an aligned set of dots. For tracks perpendicular to the surface, typically only one or two grains lie within the depth of field. Moving the focus up and down, subsequent grains appear one after another and one can picture the track in three dimensions.




Figure 2.17: Schematic view of a microscope for emulsion scanning. The data flow between the various components involved in automated track finding is indicated by the arrows. A photograph of the real setup in the CERN scanning laboratory is shown in Figure 4.2.

2.9.4 Automatic track recognition

Figure 2.17 also indicates the data-flow in the automatic scanning. In an automatic scanning station, the stage and objective movement is controlled by a computer, the human peering through the objective is replaced by a digital camera, and the images are processed by computers. To reconstruct tracks, a set of images at a fixed position

but at different depths inside the emulsion is taken. A pattern recognition algorithm is applied to this tomographic image set to reconstruct the tracks. The results are usually stored in some kind of database. The Nagoya FKEN laboratory has developed hardware to perform the pattern recognition using FPGAs and processors. The original hardware was built specifically to look for a track at a given angle and was therefore called track-selector [203, 213]. The hardware takes a tomographic set of 16 images up to a depth of $100\ \mu\text{m}$ in the emulsion. These images are preprocessed with a digital filter to enhance the contrast between grains and background and a threshold cut is applied to binarize the images. The track finding algorithm can be cursorily summarized as shifting and summing. First, each of the binarized images is shifted by the reverse of the displacement corresponding to the predicted direction. Then, the summed image is obtained by counting for each pixel the number of shifted images in which this pixel is above threshold. If a track is present, it will show up in the sum as a distinct peak above a flat background. The peak's height reflects the average grain density. Because of uncertainties in the predicted track direction, the hardware repeats the shifting and summing for each point in a grid of angles that covers a given angular acceptance.

The track selector can also be used to look for all tracks in the emulsion within a certain (wide) slope range by simply trying out all shifts. For a view of $150\ \mu\text{m} \times 120\ \mu\text{m}$ and slope acceptance of $400\ \text{mrad} \times 400\ \text{mrad}$, the early version of the track selector would take more than 30 seconds. The latest version used by CHORUS takes only 0.3 seconds. This increase in speed was achieved using a 120 Hz CCD camera and parallel processing hardware. The track-finding efficiency is found to be above 98% for track slopes less than $400\ \text{mrad}$ [214].

2.10 Reconstructing tracks and vertices in emulsion

In this section, the way neutrino interactions are located and reconstructed in the emulsion is explained. The required plate-to-plate alignment for this is discussed first, followed by an explanation of the scan-back procedure. Finally, the reconstruction of the event, including secondary vertices, using the net-scan technique is described. The net-scan technique was originally developed for the DONUT experiment (see section 1.1.3) and then used for analysis of the CHORUS emulsion. More details about the vertex location and reconstruction in the emulsion can be found in Refs. 214–216.

2.10.1 Alignment

The scanning predictions must be located inside the emulsion which requires a precise alignment between the emulsion interface sheets and the target trackers, see section 2.4.4. The 1 mm diameter black dots left by the X-ray guns (section 2.4.4), shown in Figure 2.18, provide an initial alignment between the target trackers and the three interface plates. The position of the X-ray guns is known from a survey. The initial alignment is refined by scanning a set of about 200 well reconstructed predictions. For this initial set, single isolated muons with a slope of about $200\ \text{mrad}$ are selected. The slope requirement is necessary for determining the longitudinal alignment. Because of the uncertainty of the initial alignment of about 0.5 mm, an emulsion area of $1.5\ \text{mm} \times 1.5\ \text{mm}$ around each prediction is scanned. A new alignment can then be made by minimizing the distances between found tracks and their predictions. The parameters of the alignment are a




Figure 2.18: Photograph showing lower half of an interface and a target sheet. The ruler in the background is 50 cm long. The target sheets are much darker because of their thicker and longer-exposed emulsion layers. They are also flexible enough to be rolled up as demonstrated with the target sheet at the bottom of the picture. The interface sheet is rigid enough to stand by itself. On the interface sheet (an upstream changeable sheet) the X-ray marks for alignment are enlarged in the inset.

perpendicular and longitudinal shift of the interface sheet (either CS2 or CS1) and a 2×2 matrix representing rotations and linear deformations. This new alignment has a typical residual of about $300 \mu\text{m}$. All predictions are then scanned using this alignment and a scanning area of $1 \text{ mm} \times 1 \text{ mm}$. The same procedure is applied to determine the alignment between the three interface sheets. For this inter-plate alignment, the typical uncertainty is about $150 \mu\text{m}$. The scanning area is therefore reduced to $500 \mu\text{m} \times 500 \mu\text{m}$ on the special sheet.

The emulsion target is made out of stacks of plates. Following tracks from plate to plate in a stack requires precise alignment between the plates. This alignment is different for the two years in an exposure because the stack has been repacked after the first year for changing the special sheet. Furthermore, the alignment has to take into account that the target plates are flexible and deform easily because of their very thin base ($90 \mu\text{m}$). This flexibility is illustrated in the photograph of Figure 2.18.

The special sheet contains two sets of reference points. The X-ray marks which connect it with CS1 and a set of printed reference points, called fiducials, which are also

printed on each target plate. The fiducials are a grid of $50\ \mu\text{m}$ diameter dots and dots with a $200\ \mu\text{m}$ diameter ring, spaced $19.5\ \text{mm}$ apart. Before development, the mask with fiducials is aligned with each special and target sheet in one corner and along two axes, the sheet is flattened by vacuum and then exposed using a flash of light. Because of the flexibility of the target emulsion sheets, the accuracy of this alignment is about $1\ \text{mm}$.

To get a better alignment between the sheets, six to eight track maps are scanned on each sheet. A track map is an area of $3\ \text{mm} \times 3\ \text{mm}$ centered in a square of 4 fiducials. The emulsion is then scanned to reconstruct all tracks within this area. Matching the extrapolated tracks between the maps, yields a shift of the fiducials between each plate. For each position on the plate, a local alignment is obtained from the shifts of the three nearest maps. This alignment has a typical uncertainty of about $20\ \mu\text{m}$.

Pattern recognition software on the scanning microscopes is used to locate the fiducials and to reconstruct their central position. The fiducials are often damaged such that parts of the ring or the dot is missing. Furthermore, the radii of the dot and ring vary with the light intensity and the applied gray-level cut. A picture of several fiducials with the reconstructed central position indicated by white markers is shown in Figure 2.19. The pattern-recognition algorithm uses an edge-detect filter followed by a minimum-spanning-tree to collect points belonging to the same edge. The central position is then determined from calculating circle solutions to maximally spaced sets of three points on the same edge. As can be seen from Figure 2.19, the algorithm is capable of locating the fiducial marks reliably, also when part of the fiducial is missing or when edges are connected by black streaks on the emulsion.

2.10.2 Interaction location by scan-back

Emulsion data taking for a single module starts with the search for all scanning predictions within an area of $1\ \text{mm}^2$ centered around each prediction on one or both of the changeable sheets. Emulsion tracks are selected as candidates to be followed further upstream on the basis of an improved alignment. This alignment is obtained by finding the best match of the full set of pairs consisting of a prediction and the track found in the emulsion. The process of following tracks upstream through the stack of emulsion plates is drawn schematically in Figure 2.20 and called scan-back. The introduction of the interface sheets helps in matching predictions to emulsion tracks and reduces the pick-up of accidental background because of their low track density. On the interface-sheets (CS and SS), the most selective parameter for candidate matching is the direction, because of the accurate base-slope measurement (section 2.4.2). However, this implies that four track segments must be found on the two interface-sheets, namely the upstream and downstream side of the CS and the SS.

The found tracks are followed upstream from one plate to the next, where the scanned area reduces as the position resolution from the measured emulsion tracks and the plate alignment improves. Within the stack of target plates, the scanning area is reduced to a square of $50\ \mu\text{m} \times 50\ \mu\text{m}$. Due to distortion, the direction resolution of the target plates is only approximately $10\ \text{mrad}$. The matching inside the stack of target plates is therefore dominated by the position resolution. After the precise alignment between candidates and predictions from the previous plate, the position resolution is about 5 to $10\ \mu\text{m}$. Within the candidate-selection area of typically $30\ \mu\text{m} \times 30\ \mu\text{m}$, there are usually no other (unrelated) candidates. In case a second candidate is found, both tracks are followed.




Figure 2.19: Images of a typical single dot fiducial (a) and dot-with-ring fiducial (b). Images (c) and (d) show fiducials which are properly reconstructed, even though a large part is missing or a big black streak is present.

In the scan-back of the target plates, the images are taken in the $100\ \mu\text{m}$ upstream surface of the plate. If the track is not found in this upstream layer, something must have happened inside that plate. The vertex plate is defined as the first (most downstream) of two consecutive plates where the scan-back track is not found. The requirement for two misses suppresses scanning inefficiencies quadratically. The event-location process is described in detail in Refs. 214,215. The name ‘vertex plate’ only indicates that the scan-back track disappeared in that plate. The most likely reason is indeed that a neutrino interaction took place in that plate. However, the scan-back track could also have changed direction or originate from a secondary interaction in that plate.




Figure 2.20: Illustration of the scan-back procedure in which a track is followed upstream from plate to plate. Only the most upstream $100\ \mu\text{m}$ thick layer of emulsion is scanned around the predicted position. Track candidates found in the emulsion are indicated with thick black boxes. The scan-back procedure stops when a predicted track is not found in two consecutive plates, indicated by the empty boxes and crosses. The first plate where the track was not found is considered the vertex plate. The predicted track either originates in this plate or changes direction inside it. In the net-scan procedure (section 2.10.3), volumes in eight plates around the disappearance point of the track are scanned to find all event-related tracks. The light-gray areas indicate these volumes. Each volume is $1500\ \mu\text{m}$ wide, $100\ \mu\text{m}$ thick, and centered on the disappearance point of the scan-back track, corrected for the beam slope.

2.10.3 Vertex reconstruction with net-scan

Once the vertex plate is identified, a detailed analysis of the emulsion volume around the vertex position is made. In this procedure, known as net-scan, track segments within a given volume and with a given angular acceptance are reconstructed. The scanning is done on several consecutive plates around the vertex plate. In the CHORUS experiment, the net-scan volume spans eight emulsion plates, corresponding to a length of 6.3 mm along the beam direction. It is 1.5 mm wide in each transverse direction. The eight plates contain the vertex plate itself, the plate immediately upstream, and the six plates downstream of the vertex plate. As the track-selector hardware can only scan layers of $100\ \mu\text{m}$ thickness, the upstream surfaces are scanned. The plate upstream of the vertex acts as a veto for passing through tracks. The six plates downstream of the vertex act as decay volume and are used to detect the tracks of the decay daughters. The scanning area on the vertex plate is centered on the extrapolated transverse position of the scan-back track. On the other plates, the area is shifted with the neutrino beam slope (42 mrad). The angular acceptance for tracking corresponds to a cone with a half-opening angle of 400 mrad.

The first step of the vertex reconstruction is the selection of only those segments belonging to the neutrino interaction out of the large number of track segments found. A coarse plate-to-plate alignment is determined by comparing the pattern of segments in a plate with the corresponding pattern in the next upstream plate. With this coarse alignment each segment found on one plate is extrapolated to the next plate where a matching segment is looked for within a cone of 20 mrad and an area of about $4\ \mu\text{m}^2$. The size of the matching area corresponds to 3σ of the alignment resolution. To recover scanning inefficiencies, if no track is found on the next plate, the matching is tried one plate further upstream. A second and more accurate inter-plate alignment is performed using tracks passing through the entire volume after the connection of all matched segments. These tracks are mainly coming from muons associated with the neutrino beam or charged-particle beams upstream of the experimental area. After this alignment, the distribution of the residual of the segment positions with respect to the fitted track has a RMS width of about $0.45\ \mu\text{m}$. At this stage, typically about 400 tracks remain in the volume. The majority of these are tracks of low-momentum particles (mainly Compton electrons and δ -rays) with momentum less than $100\ \text{MeV}/c$. These background tracks are rejected with a criterion based on the χ^2 of a straight-line fit to the track segments. The final step is the rejection of all tracks not originating from the scanning volume. After this filtering, the average number of tracks originating in the scan volume is about 40.

Figure 2.21 shows an example of net-scan data and the result of a simplified vertex search. It shows how a vertex can be located among all the track segments by simply requiring a small impact parameter. The real net-scan analysis is more sophisticated and also uses segments reconstructed on a single plate, for example a short-lived particle decaying in the next plate.

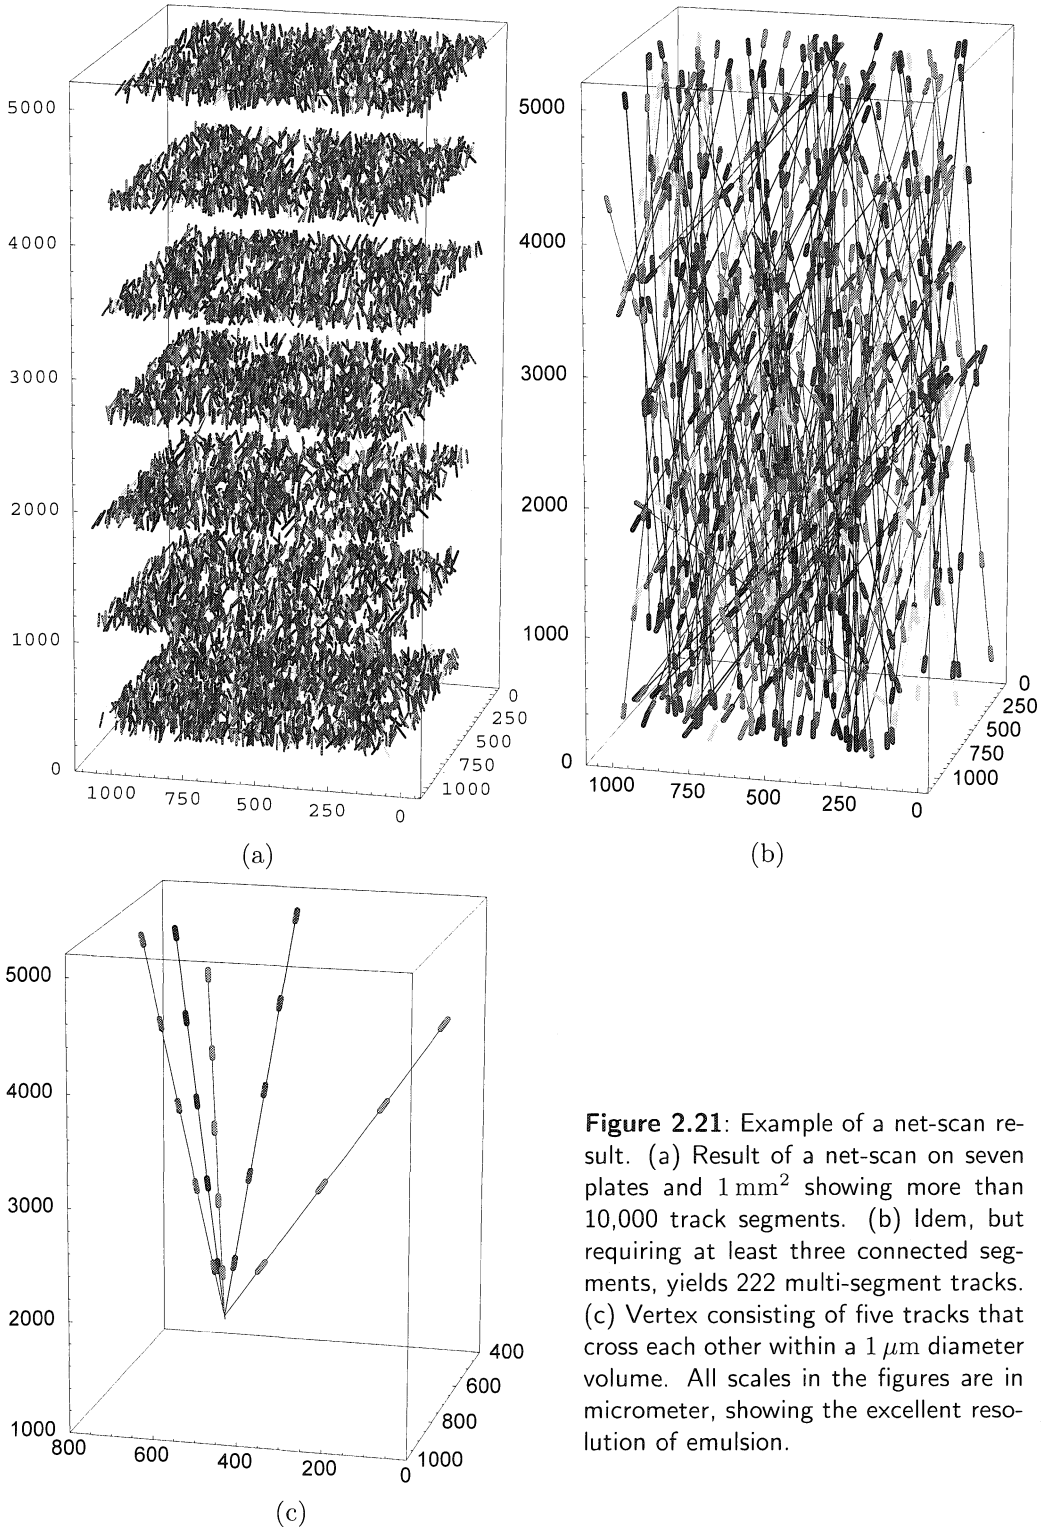


Figure 2.21: Example of a net-scan result. (a) Result of a net-scan on seven plates and 1 mm^2 showing more than 10,000 track segments. (b) Idem, but requiring at least three connected segments, yields 222 multi-segment tracks. (c) Vertex consisting of five tracks that cross each other within a $1\ \mu\text{m}$ diameter volume. All scales in the figures are in micrometer, showing the excellent resolution of emulsion.

Chapter 3

Honeycomb tracker

Several of the selections to suppress background to the $\nu_\mu \rightarrow \nu_\tau$ oscillation signal (section 2.1.2) require an accurate measurement of a particle's momentum and charge. For example, to suppress background from kaon decays in kink candidates, the charge and transverse momentum of the kink-daughter needs to be measured accurately. The performance of the hadron spectrometer is therefore a crucial parameter in determining the sensitivity of the experiment to $\nu_\mu \rightarrow \nu_\tau$ oscillation. The foreseen configuration of the hadron spectrometer with only four planes of fiber trackers behind the hexagonal magnet (section 2.6.1) is unable to do independent tracking behind the magnet. This makes the alignment and the pattern recognition in these four planes difficult.

A new technique to build quickly a large number of straw tubes had been developed at NIKHEF [217, 218]. As the equipment to build these new honeycomb straw-tube planes was available, it was proposed to replace the six planes of streamer tubes behind the hadron spectrometer with eighteen planes of honeycomb straw-tubes. The improvement in the resolution of the hadron spectrometer achieved with this new detector enables, among others, the charm production study described in Chapter 5.

This chapter describes the design and construction of these honeycomb straw-tube planes. The read-out electronics are described in detail, followed by a description of the pattern recognition algorithm used to find tracks in the honeycomb detector. Finally, the performance of the detector and its use within the whole CHORUS detector is given.

3.1 Motivation and requirements

To reconstruct multiple 3-D tracks, the tracks need to be detected in at least three projections. If a detector with two read-out projections is traversed by two particles, each projection will yield two hits corresponding to two lines in the other projection. The four lines have four intersections and it is impossible to tell which two belong to the actual particle crossings. With a third, different, projection this ambiguity can be resolved. In practice, the projections have a finite width due to the resolution of the detector. Especially in high multiplicity events, the resolution of the detector and the angular difference between projections become important to separate the real crossing points from fake ones. Furthermore, in a tracking detector it is beneficial to have several detector planes per projection, both for detection efficiency as well as for track reconstruction.

The low number of tracking planes behind the hadron-spectrometer magnet (section 2.6.1) required that the alignment and track finding was done with tracks reconstructed by the target tracker upstream of the magnet. These tracks cross the magnet first and therefore the alignment accuracy is degraded. Moreover, as most tracks have small angles with respect to the beam direction, longitudinal shifts in the alignment of the DT2 and DT3 planes of the diamond tracker are difficult to measure. As a result, the momentum resolution of the hadron spectrometer was worse than expected.

To improve this, additional detectors behind the magnet are necessary. However, the available space was only 21 cm. Originally, there were four low-resolution (3 mm) streamer-tube planes in this gap, later extended to six. These planes had only 7° stereo angles. Combined with a relatively low efficiency and poor resolution, they were not sufficient to support independent tracking behind the magnet. Neither was their resolution good enough to improve the alignment of the diamond tracker paddles.

A new tracker that could independently find 3-D track segments in the 21 cm gap could improve the performance of the hadron spectrometer considerably. The positions and slopes of track segments found in such a tracker could be used to validate hit combinations in the diamond tracker planes, reducing fake tracks. Tracks reconstructed in such a detector could also be used to determine the alignment of the DT2 and DT3 paddles, using for example large-angle cosmic-ray tracks. Furthermore, in the downstream direction, the track segments could be used to determine the energy deposited in the calorimeter on a track-by-track basis, which allows to distinguish between electrons, muons and hadrons. Such a tracker would also add an intermediate segment to connect muon tracks reconstructed in the muon spectrometer with a target tracker track.

A design for a new tracker was subject to the following restrictions and requirements. As already said, it should fit in the 21 cm gap and support 3-dimensional track reconstruction. Furthermore, its acceptance should cover the whole hexagonal magnet surface. Because of the physical limit due to multiple scattering, there is no need to improve upon the diamond tracker's resolution. So a $200\ \mu\text{m}$ resolution was taken as the design aim. Even though the average primary track multiplicity is only about 4.1 tracks per event [178], all detector elements must be read out individually, because for some events the showering has started upstream of the gap. As the sole purpose of this detector is tracking, no energy measurement is required. However, some information on a hit's pulse height is useful to distinguish between noise and signal hits. The tracker should be kept preferably as light as possible to minimize additional showering. Finally, the read-out system has to fit in the CHORUS data-acquisition system. This required that the read-out can buffer at least 16 events per beam spill and that all events can be read out in less than 500 ms.

These criteria could be fulfilled with a honeycomb tracker: a, lightweight, straw-tube, gas detector, read-out with drift-time measurement to reach the required resolution. This tracker was built at NIKHEF starting halfway 1995 and installed in CHORUS in August 1996, half-way the 1996 neutrino-beam run. This chapter describes the construction, read-out, and performance of this honeycomb tracker. The detection principles of gas-avalanche detectors are briefly recalled in section 3.2. The design and mechanical construction of the honeycomb tracker are described in section 3.3. The read-out electronics are presented in section 3.4, the read-out protocol in section 3.5, and the tracking code in section 3.6. The chapter ends with a discussion on the honeycomb tracker performance and some ideas for future improvements.

3.2 Detection principle

The principles of gas-avalanche detectors, used by the honeycomb detectors, will be shortly recapitulated in this section. For a more extensive description of gas-avalanche based detection and amplification of ionizing particle tracks see Ref. 219 or any textbook on detectors for high energy physics, for example Refs. 220, 221.

A honeycomb tracker consists of a set of tubes made from folded conductive foils with a thin central wire. The wire is kept at a positive high voltage relative to the tube wall. Each tube, shaped as a hexagonal cell, functions as a single wire drift chamber. A charged particle traversing a cell ionizes gas molecules along its path, liberating electrons. The electrons, under influence of the radial electric field, drift toward the central wire. The ions drift toward the cell walls which are normally kept at ground potential. Near the wire, the electric field $E_r = r^{-1} \cdot V / \ln(r_{\text{cell}}/r_{\text{wire}})$ is high enough for the electrons to gain enough energy between collisions with gas molecules to produce secondary electrons. These secondary electrons are also accelerated towards the wire, producing more electrons. The electrons produced in this avalanche are collected by the wire in a few nanoseconds. The ion-cloud of the avalanche is left behind and drifts towards the cathode. The signal on the wire has two contributions. A fast negative current pulse from the collected electrons and an induced signal from the moving ion cloud [1, page 177]. The wire current can be detected electronically, usually with a trans-impedance amplifier. The gas amplification itself can be increased by raising the voltage to enlarge the size of the avalanche region. The single atom gasses (especially the noble gasses) are most suitable to get high amplification, because they lack rotational and vibrational excitation modes. Therefore, they have a longer mean free path. However, excited atoms in the avalanche radiate UV-photons that can cause ionization far from the primary avalanche. These secondary ionizations can lead to electrical breakdown of the gas causing discharges (like in a fluorescent tube lamp). Therefore, a quench gas is usually added to absorb these photons. Typical quench gasses are CO_2 and some organic molecular gasses.

Gas-avalanche detectors can be operated in different modes, depending on the voltage, gas composition and wire thickness used. For relatively low voltages, in the so-called proportional mode, the amplification process is localized near the wire and the current pulse is proportional to the initial ionization. Because the avalanche region is small and localized, the temporary insensitivity is limited to tens of nanoseconds at the position of the avalanche only.

For high voltages and low quench gas concentrations, the proportional mode gives way to the Geiger-Müller mode. In this mode, the avalanche spreads along the wire

due to secondary avalanches caused by photo-electrons. In the Geiger-Müller mode, a constant amplitude pulse (independent of the amount of primary ionization) can be detected without any additional electronic amplification. However, the detector is dead for several milliseconds due to the space charge of the ions produced in the avalanche that shields the electric field near the wire.

For high voltages combined with high quench gas concentrations, a so-called streamer can occur [222]. In this case the avalanche grows also in the direction towards the cathode. The electric field due to the charge separation in the avalanche in combination with the field of the wire induces secondary avalanches of photo-electrons near the top of the primary avalanche. In the weaker field, further away from the wire, the streamer stops. Large pulses are produced in this way, but the wire is only dead in a limited region around the avalanche position. This mode of operation is usually achieved using thicker wires (100 μm) such that the field gradient is smaller.

The position resolution of a gas-detector cell can be improved by measuring the time it takes for the primary electrons to drift into the avalanche region. This drift time is measured with respect to some other fast detector, usually scintillators. The position resolution is limited by diffusion of the drifting electrons and by the resolution of the drift-time measurement. The drift time is a function of the radial distance of the (closest) primary ionization to the wire and the drift velocity. The drift velocity is mainly determined by the gas mixture (also its pressure and temperature) and the electric field. The single-atom noble gasses have high drift velocities, while adding a molecular gas with rotation and vibration modes lowers the drift velocity. Scattering of the electrons leads to diffusion of the primary ionization cloud, decreasing the achievable position resolution. Scattering becomes more important for longer drift distances unless a magnetic field is applied to bind the electrons to field lines, like in a time-projection chamber.

Although the resolution of straw tubes is usually a bit worse compared to multi-wire drift chambers, they have some significant advantages. Because each wire is enclosed in its own (grounded) cathode, cross-talk between wires and pick-up of electro-magnetic signals from the environment are largely eliminated. In general, drift tubes are more reliable than wire chambers, because the tube prohibits a breaking wire from creating short circuits with other wires. The honeycomb technology has two additional advantages over standard aluminum straw tubes. First, they are easier to manufacture because the wires are inserted in half open cells and do not have to be pulled through a tube. Second, the plastic cells contain less material and with lower Z than aluminum tubes.

3.3 Design and mechanical construction

The moulds and the folding machine available in the NIKHEF workshop for making honeycomb straw-tubes limited the basic building block to flat rectangular planes with 1 to 236 consecutive tubes of 100 cm length and 1.1 cm diameter. With respect to previously built honeycomb straw-tubes [217, 218], the type of cell wall material was changed from copper-sputtered mylar to conductive plastic. The easiest way to cover the full area of the hexagonal magnet is to stack three modules of 3 m \times 2.6 m, rotated around the beam axis by -60° , 0° and 60° . The rotated modules yield three projections needed for 3-D track reconstruction. The stereo angle of 60° is optimal considering the effects of the position resolution. Given the gap space, the design consists of three rectangular modules with six planes per module. The six planes per orientation yield enough hits to find two-dimensional track segments on a module-by-module basis.

3.3.1 Monolayer

The basic element of the honeycomb tracker is a $100\text{ cm} \times 274\text{ cm}$ monolayer consisting of 216 cells of 100 cm length. Such a monolayer is constructed out of a 100 cm wide and $75\text{ }\mu\text{m}$ thick foil of conductive pocalon plastic (polycarbonate mixed with graphite). A computer-controlled machine makes a sequence of folds in the foil. The distance between folds is such that when the foil is bent alternately twice at $+60^\circ$ and twice at -60° , it gets the bottom (or top) structure of a series of hexagonal cells. The distance between the folds is programmable. For the CHORUS honeycomb tracker and the moulds available, the cells have a diameter of 11.54 mm with a 1.13 mm spacing between the cells. The folded foil is depicted in Figure 3.1a. After folding, the foil is vacuum-sucked onto the bottom mould.

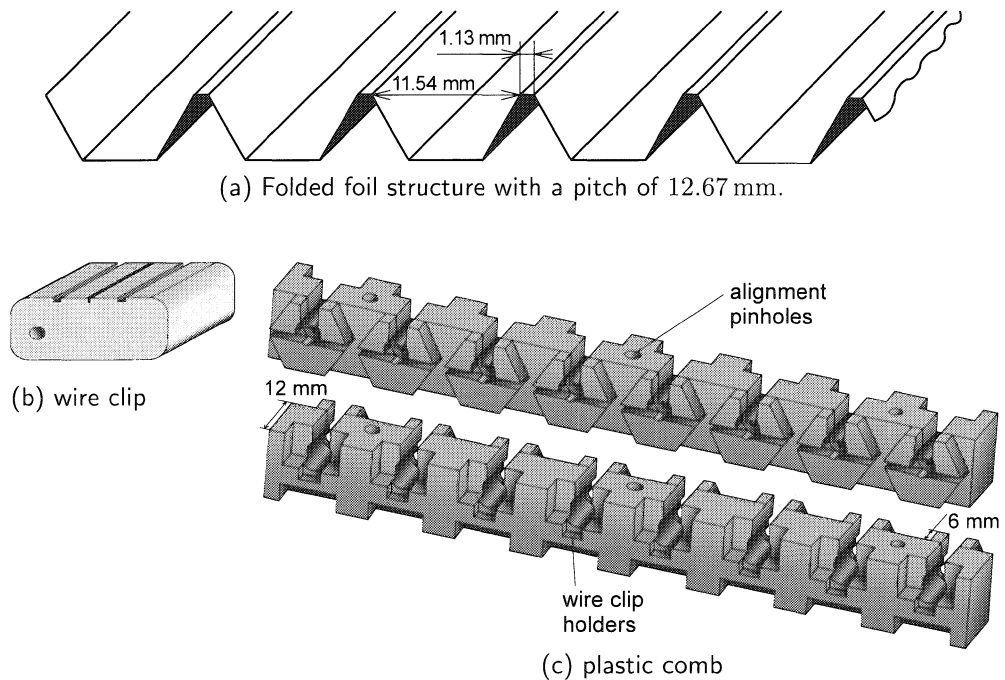


Figure 3.1: Components of a single honeycomb monolayer: (a) the folded foil which gives the characteristic hexagonal structure of the honeycomb cells; (b) a wire clip in which the wires are crimped; (c) the plastic combs in which the wire clips are placed and to which the foils are glued.

Plastic combs, shown in Figure 3.1b, are glued at both sides of the foil. These combs center the wires and give mechanical stability to the whole structure. The precision in their position, which also defines the precision of the wire centering, is better than $20\text{ }\mu\text{m}$. This precision is achieved by placing the combs on the dowel pins of the mould. Small wire clips, made of a copper-tellurium alloy (depicted in Figure 3.1c), are glued inside the combs to fix the wires. These clips also provide the electrical connection with the wires and between monolayers (see next section). The assembly procedure is illustrated in Figure 3.2.




Figure 3.2: Components and construction of the bottom half of single honeycomb monolayer. After the wires have been put in the open bottom-half, another mould holding the top foil is placed on top and glued and point-welded together to close the honeycomb cells.

A $30\ \mu\text{m}$ thick gold-plated tungsten wire is first wound with the required tension onto a frame. The frame is placed over the mould such that the wires fall into the central grooves of the wire clips. The clips are then crimped with a special tool, fixing the wire.

A second, identical foil, is sucked onto the top mould which is then placed on top of the bottom mould. The top foil is also glued to the combs and the top and bottom foil are point-welded together in the $1.13\ \text{mm}$ spacing between cells to ensure electrical contact. Finally, a single flat pocalon sheet is glued, partly with conductive silver glue, on top of the monolayer to improve the mechanical stability. This flat foil also ensures good electrical contact between stacked monolayers.

The structure resulting from stacking several monolayers forms a regular honeycomb structure, see Figure 3.3, and is stiff enough to resist the pull caused by the wire tension.




Figure 3.3: When several monolayers are stacked, the resulting structure resembles a honeycomb.

3.3.2 Honeycomb module

To cover the hexagonal magnet, a module with a width of at least 3 m and a height of 2.6 m is needed. Three monolayers are therefore combined to make one large plane of 308 cm × 274 cm by mounting them side-by-side along the long edge. The wires are electrically connected between monolayers using pins placed in holes in the wire clips. The holes contain a small contact spring to ensure good electrical contact. Electrical contact between the pocalon cell walls of the adjacent monolayers is made by placing copper-tape strips every 30 cm over the full width of the plane.

Six of these planes are subsequently glued on top of each other, partly with conductive glue. Each plane is shifted by a few millimeters in order to detect particles traversing the small dead space between the honeycomb cells and to resolve left/right ambiguities inherent to drift tubes. The staggering of the layers is shown in Figure 3.4. The shifts are optimized to give a minimal average loss of hits for particles with a uniform angular distribution between -30° and $+30^\circ$ with respect to the normal of the detector plane. The expected number of hits versus angle and position is plotted in Figure 3.5. The periodical structure contains six bands of about 1.2 mm wide due to the dead-space between cells. The apparent widths of these bands reflect the angular range covered by the dead-space which depends on the distance to the position reference plane (taken to be in the middle of the six planes). There are two small regions with only two expected hits for each column of cells. Both these regions are at angles larger than 20° .




Figure 3.4: Staggering of six honeycomb planes inside a module.




Figure 3.5: Density plot of the expected number of hits versus track angle and position. The dashed lines indicate the radius r of a cell and the dead-space d between cells. The position reference plane is in the middle of the six planes.

The six planes of one tracker module are enclosed in a 1 mm thick aluminum gas-tight box. The electrical and mechanical construction is schematically drawn in Figure 3.6. Printed circuit boards (PCB) to distribute the high voltage are connected with 12 mm pins to the wire clips on one end of the wires. At the other end, two sets of interconnected PCBs carry the signal from the wires via AC-coupling capacitors to two rows of eighteen 64 pin connectors. The first set of these PCBs is also connected to the wire clips with 12 mm pins. These wire-connection boards are glued gas-tight to the aluminum gas-enclosure box. The second set of PCBs at the read-out end consists of six layers of high-voltage separation boards which contain the high-voltage decoupling capacitors and the connectors for the read-out cards. These boards are mounted in the same plane as the wires. Each set of three boards is interconnected to route the signals of 24 wires from three honeycomb planes to two connectors mounted on the outside boards. The 12 mm distance between the cells and the PCBs is needed to ensure sufficient gas flow through all honeycomb cells. The gas in- and out-let are placed diagonally opposite of each other. Two 5 m long aluminum bars are fixed to the sides of a module. They are used later to build the three modules in the hexagonal configuration. At the read-out side of a module, a set of six copper bars mounted between the two side-bars, supplies power to the read-out cards (see section 3.4.2).




Figure 3.6: Printed circuit boards used for mechanical and electrical connections of the honeycomb wires to high-voltage supply and read-out cards.

3.3.3 Honeycomb tracker

The complete honeycomb tracker is built from three modules, put together at angles of 60° , 0° and 120° . The final layout, including the support structure, is shown in the photograph of Figure 3.7. The power supplies and cooling ventilators for the read-out electronics were placed on top of the electronics huts, close to the tracker.

A 1:1 mixture of Ar and CO_2 is used as drift and avalanche gas. This mixture combines a slow drift velocity of $2 \text{ cm}/\mu\text{s}$ with a reasonable working voltage of about 2200 V. Both Ar and CO_2 have the advantage that they are not flammable and not toxic. Small leaks can therefore be tolerated and no additional safety measures were needed. These two gasses were also readily available at the detector and therefore the honeycomb tracker could easily be incorporated in the existing CHORUS gas system.




Figure 3.7: Full honeycomb tracker during its installation in the CHORUS detector. The 5 m long side-bars connect the three modules under 60° stereo angles. The power-bus and power-distribution cables are visible on the top of each module. The gleaming copper shielding was an earlier attempt to avoid amplifier oscillations by shortening the ground path. The shielding was later completely removed after recurrent problems of bad electrical connection with the aluminum cover sheets. [Photo courtesy of J. Visschers.]

3.3.4 Prototype measurements

Before construction of the honeycomb tracker started, it was decided to build a small prototype of a few cells. Three monolayers of eight cells were built and then connected to form 3 m long tubes. This prototype was put in a gas tight plastic enclosure and it was flushed for a couple of days with a 4:1 or 1:1 Ar:CO₂ mixture. Measurements on this small prototype were performed to check signal propagation, signal shape, cross-talk, and working voltages. The full results can be found in Ref. 223; only the results that influenced the operation of the final tracker are repeated here. The prototype used the same front-end amplifier and discriminator as was used in the final honeycomb tracker (section 3.4.2).

A test was performed on signal propagation through a cell in which the coaxial structure of wire and cell is disrupted by the interconnection between monolayers. There was also some worry about possible signal attenuation caused by the finite conductivity of the pocalon foil, $R_{\square} \approx 120 \Omega$, and the wire, $R_{\text{wire}} \approx 75 \Omega \text{m}^{-1}$. The pulses formed by the single photo-electron of a 5.9 keV X-ray ^{55}Fe source were recorded at positions along the wire in different monolayers. No significant signal loss could be observed, supported by the idea that the wavelength of the signal is much longer than the distance over which the coaxial structure of wire and cell is disturbed.

Some cross-talk between wires was expected because the shielding of the signal by the pocalon cell wall is not perfect due to its finite conductivity. A measurement with the ^{55}Fe source was made with the wire voltage adjusted to give 600 mV pulses (60 % of maximum amplifier output voltage), while simultaneously monitoring the signal on an adjacent wire. An inverse signal of less than 10 mV was seen on the adjacent wire. Because the cross-talk has the opposite polarity, the discriminator behind the amplifier will reject these pulses. For the expected low occupancy and rate in the CHORUS experiment, this cross-talk can safely be neglected.

To determine the working voltage, the count rate for the ^{55}Fe source was measured at different voltage settings for both gas mixtures. The count rate for a ^{106}Ru β -ray source was also measured for the 4:1 Ar:CO₂ mixture. This electron source has an energy spectrum with a 3.5 MeV endpoint. Electrons of different energies produce different numbers of primary ionizations in the cell which leads to signal pulses of different strength. The results are plotted in Figure 3.8. For the ^{55}Fe measurements, a clear plateau can be seen. The onset of the plateau is about 250 V higher for the 1:1 mixture. This can be understood as less Ar leads to a smaller gas amplification. For the β -ray source the plateau is less clear, because at higher gas amplification (higher voltage), smaller signals will be amplified such that they cross the discriminator threshold. However, above 1850 V some saturation occurs. Extrapolating to the 1:1 mixture, one expects a working voltage above 2100 V.




Figure 3.8: Count rate versus wire voltage for 4:1 and 1:1 Ar:CO₂ gas mixtures for a ^{55}Fe X-ray source and a ^{106}Ru β -ray source.

3.4 Data-acquisition and read-out electronics

To measure the drift time in a honeycomb cell, the signal arrival time is usually measured with respect to a trigger signal. This measurement is normally done using a time to digital converter (TDC). The pulse height of a hit can give information about the total ionization and can be used to distinguish between noise and signal hits. If a discriminator is used to digitize the analogue pulse, the pulse length, defined as the time over threshold, is a measure of the pulse height. To measure the pulse length with a TDC, both the rising and the falling edge must be timed.

There are several ways of measuring the time difference between a start and stop signal. In an analogue TDC [224], a capacitor is charged by a constant current in the period between the start and the stop signal. The voltage is then measured after the stop signal. These TDCs have usually very good time resolution and a large dynamic range. In general, however, these TDCs have no multiple-hit capacity and long dead times due to the analogue-digital converter needed to convert the voltage to a digital signal. In a digital TDC, an on-chip delay line is used to interpolate between transitions of an $\mathcal{O}(10)$ MHz clock to reach nanosecond resolution [225]. The stop signal latches the clock count and the active delay-line stage in a time-stamp register. With multiple registers, a multi-hit TDC can be constructed. Since frequency generators can be made very stable, these TDCs have also good time stability.

For the honeycomb tracker in the CHORUS experiment, the use of TDC modules (VME or CAMAC) was not possible due to high costs of cabling and the space needed in read-out crates. An alternative would be to use monolithic TDC chips which can be used in custom-designed electronics and could be mounted directly on the honeycomb modules. As the read-out is much closer to the wires there is less chance of picking up noise from the environment and the cost can be reduced by replacing cables with PCBs. As the average number of tracks per event in the CHORUS experiment is low, one could consider reducing the cost by multiplexing wires in time. By delaying the signals by more than the maximum possible drift time, hits on different wires can be combined in a single TDC channel. This requires multi-hit capability of the TDC. Not only the occupancy is low in the CHORUS experiment, also the event rate is low due to the small neutrino interaction cross-section. Therefore gasses with a low drift velocity can be used. With low drift velocity, the time resolution required for a given space resolution is reduced. A low time-resolution makes it possible to build TDC circuits with discrete components. The 1:1 Ar:CO₂ mixture has a drift velocity (v_{drift}) of about 2 cm/ μ s, which yields for a design goal of 200 μ m position resolution (Δr) a minimum required time resolution of

$$\Delta t < \frac{\Delta r}{v_{\text{drift}}} = \frac{200 \cdot 10^{-6}}{2 \cdot 10^4} = 10 \text{ ns} ,$$

something which can be achieved with discrete logic components.

3.4.1 The bit-stream principle

The first idea was to build a TDC with a programmable gate-array (FPGA) using the time-stamp technique. The value of an $\mathcal{O}(100)$ MHz counter could be latched into a FIFO for each wire hit. However, a design that can handle both signal edges and multiple hits is still quite complex. Instead an approach was adopted that can be characterized by the

term bit-stream. The binary over-threshold signal is sampled at discrete intervals using a central clock and stored in a memory. The stored bit-stream contains the complete history of the over-threshold signal during a certain time period before the stop signal. The time period depends on the size of the memory available. The principle is essentially the same as that used in a digital oscilloscope. A literature survey showed that a similar principle was already applied successfully before [226–228]. In those implementations, fast ECL memories were used. Using ECL memories it is possible to go down to 5 ns sampling time. However, the depth of these memories is small and they consume a large amount of power compared to TTL and CMOS technology.

For the CHORUS experiment other memory chips were considered. Using just four synchronous dual-port memories with a 36 bits wide data bus and 20 ns access time, 72 channels could be integrated on a single read-out board. The four chips each store a different part of the 200 MHz bit-stream. Several events can be stored in the 32 KWord of memory available per card. Switching between event buffers can take as little as one clock-cycle and therefore the dead time after a trigger is practically zero. In the configuration used for the honeycomb tracker, the available 32 KWord is divided into 64 buffers with 256 bits for each of the 72 channels. The trigger signal freezes the last 1280 ns of the over-threshold signal in these 256 bits buffers. The delay between the time of the event and the trigger must be set such that the drift time and pulse length of event related hits are fully contained in this time window. This is achieved by delaying the trigger signal with a simple delay line.

3.4.2 Chambercards

The read-out of the honeycomb tracker is implemented in electronic boards mounted directly on the honeycomb modules. These so-called chambercards are plugged into the two connectors depicted in Figure 3.6. Each chambercard contains the amplifiers, discriminators, memories and control-logic to digitize and store the signals from 72 wires. One honeycomb module is equipped with 18 chambercards which are connected in a pipeline for the read-out (see section 3.5). The layout of a chambercard is schematically depicted in Figure 3.9.

Analogue signal processing

The signal from a wire is amplified using a trans-impedance amplifier chip, specially designed for wire chambers.¹ The amplification factor is about $25 \text{ mV}/\mu\text{A}$. The signal is digitized using discriminators² with an adjustable threshold. A small PCB contains one amplifier and one discriminator chip and contains four channels. It fits in a 30 pin SIMM-socket which facilitates replacement of broken amplifiers. The design is based on an earlier design with a small modification of the threshold circuit [229]. When the input signal goes over threshold, a negative feedback signal lowers the threshold by about 15 mV. The threshold is raised again by 15 mV when the signal goes under threshold. This threshold shift prohibits oscillation of the output when the signal crosses the threshold slowly.

¹LeCroy TRA 402S

²LeCroy MVL 407S

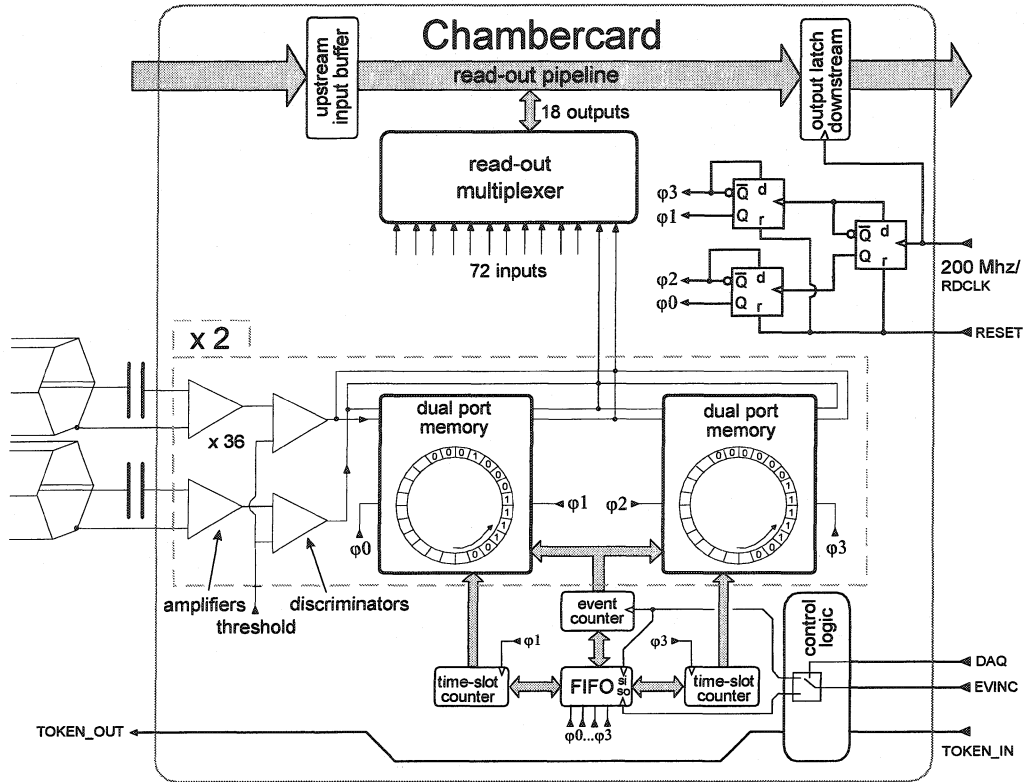


Figure 3.9: Schematic diagram of a chambercard which digitizes and stores the history of 72 wires. In total, there are four memory chips on each card. The four phases ($\phi_0 \dots \phi_3$) correspond to a sampling period of 5 ns.

Sampling clock

The binary output of the discriminators, which corresponds to the over-threshold of the analogue signal is sampled every 5 ns, timed by a central 200 MHz clock implemented on the clockcard (see section 3.4.3). Because all chambercards derive their sampling clock from the same master clock and with the same delay, all channels of a module have the same time offset with respect to the stop signal (t_0). During read-out, the same clock input is used to generate read-out cycles. Using the same clock input for both sampling and read-out simplifies the clock circuitry on the chambercards.

Circular buffers

The bit-stream of the over-threshold signals is stored in a 2^n bits circular buffer. The value of n can be chosen between 6 and 10. The buffers are implemented using two dual-port synchronous memories.³ The two ports of the 8KWord deep and 36 bits wide memories can be written independently, as long as both ports access a different 2KWord block defined by the highest two address bits. The memories are read and written synchronously at a maximum speed of 50 MHz. They have an internal pipeline

³Quality Semiconductor 75836

of two stages which implies that the delay for read access is actually two clock cycles. By splitting the 200 MHz bit-stream of the wires in four phases of 50 MHz ($\varphi_0 \dots \varphi_3$), it is possible to store the bit-stream in two memories. The over-threshold signals of 36 wires are distributed to both the left and the right port of both memory chips. The first memory clocks in the samples at the phases φ_0 and φ_1 , the second memory at φ_2 and φ_3 . A continuous block of memory for each phase (port and chip) serves as a 2^{n-2} circular buffer. There are thus 2^{14-n} of such blocks for each memory port. As the memories do not have automatic address generation, external counters are needed to generate the addresses. The memories latch the address and data to be written so no external latches are needed to hold the address and data inputs stable during a cycle. This simplifies the data path and address generation logic.

In the CHORUS experiment the chambercards were configured with $n = 8$, which gives 64 buffers of 256 bits deep, equivalent to 1280 ns at 200 MHz. The reconstruction of the time-over-threshold signal for the last 1280 ns before the event trigger is therefore possible. The drift time of hits is determined by the sample number times 5 ns minus t_0 for a $0 \rightarrow 1$ transition in the bit-stream. The pulse length of a hit on a wire can be deduced from the number of samples between the $0 \rightarrow 1$ and the subsequent $1 \rightarrow 0$ transition.

Address generation and event FIFO

The four phases are derived from the central clock with a simple divider chain of three flip-flops. The normal and inverted output of the first flip-flop drive the clock inputs of the others. The four outputs (normal and inverted) of the second stage correspond to the four phases. These phase clocks are used to clock the memories and counters.

To generate the pointers for the circular buffers, there are three different address generators implemented as counters in three programmable logic chips.⁴ The use of programmable logic and careful routing of signals between the counters allows one to trade buffer depth versus the number of buffers. In a test experiment with a small honeycomb tracker [169], the chambercards were set up to have 128 buffers of 128 bits deep (640 ns). For each memory chip there is a circular-buffer pointer, called the timeslot counter, which is common to both phases and incremented by the later phase. A chambercard contains two of these circular-buffer circuits. The corresponding memory chips in each circuit share the same timeslot counter. A single event counter, which selects between buffers, is used for all four memory chips.

The timeslot counters determine the address in the circular buffers where the next write will take place and they generate the lower bits of the address ($A_0 \dots A_5$ for $n = 8$). They are implemented as $(n-2)$ bits Gray counters⁵ to limit noise due to many bits toggling at the same time. The timeslot counters wrap around after reaching their maximum count. The $(14-n)$ bits event counter generates the high address bits ($A_6 \dots A_{11}$ for $n = 8$) of all memories. The highest address bit (A_{12}) of the memories is hardwired to 0 and 1 for the left and right port, respectively. This makes sure that the left and right port write to different 2 KWord blocks of the memory. The event counter selects which of the 2^{14-n} independent circular buffers is used for writing. It is incremented by the trigger signal. The trigger input on the central clockcard first halts the clock to stop writing samples in the buffers and freeze the time-slot counter and clock-phases. The clockcard

⁴AMD MACH215A

⁵Gray counters are binary counters in which only 1 bit changes at every count

then generates an event-increment pulse (EVINC). This pulse causes the chambercards to store the current values of the phase-clocks, timeslot counter and event counter in a FIFO. These counters point at that moment to the next address to be written in the last circular buffer which is equivalent to the oldest timeslot in the buffer. After that, the event counter is incremented and writing continuous in the next buffer.

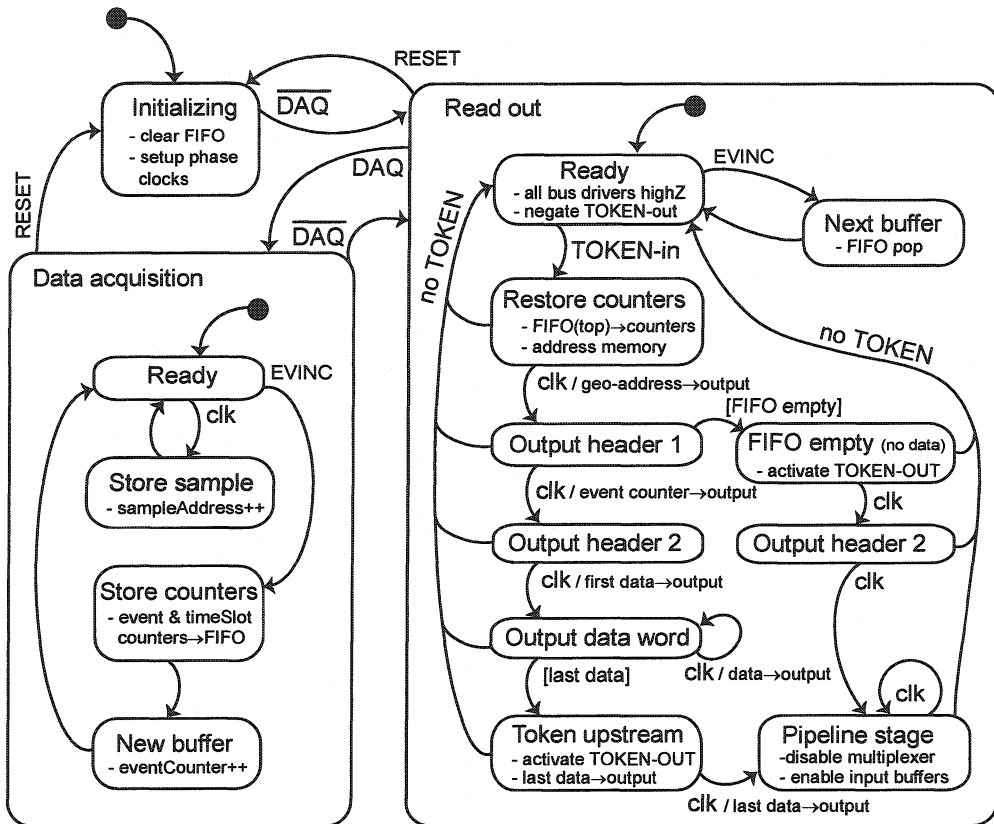


Figure 3.10: State diagram of chambercard logic, consisting of a sequence of states for taking samples in data-acquisition mode (left) and externally controlled states in read-out mode (right).

Control logic and read-out bus

The working mode of the chambercards is controlled by several control signals on the read-out bus. Here, only the three main control signals are discussed. All control logic for a card is contained inside one programmable logic chip.⁶ The control logic implements the state-machines, shown in Figure 3.10, for the data-acquisition and read-out mode. The mode is selected by the DAQ signal. The EVINC signal is used to select the next circular buffer. Its operation in data-acquisition mode is described above. In read-out mode, the next contents of the FIFO are reloaded into the address generators. In read-out mode, the control logic watches the TOKEN_IN input. Detection of an active TOKEN_IN

⁶AMD MACH230A

input triggers the card to become the active read-out bus driver. The control logic also configures the multiplexer to select the next 18 bits of data from the 72 memory outputs in each read-out cycle. Once all data of an event have been read out, it activates the `TOKEN_OUT` signal and configures the card as a clocked pipeline stage. The complete read-out system is described in section 3.5.

Printed circuit board

The amplifier cards were placed as close as possible to the wire connectors to keep the distance traveled by the small wire signals as short as possible. The lower quarter of a chambercard contains all the amplifier cards and has a separate analogue-ground layer. The analogue ground is connected to the signal ground (cell walls) via the aluminum gas box and the pocalon structure. It was found that the aluminum connection led to grounding problems, discussed in section 3.7.2. To avoid pick-up of noise from the digital logic, the high-speed clock circuit and counters were placed on the other side of the PCB. All digital logic is connected to a separate digital-ground layer which was coupled with $10\ \Omega$ resistors to the analogue-ground layer.

3.4.3 Clockcard

A central clockcard per honeycomb module controls the timing of the data-acquisition and the read-out. The clockcard distributes a central clock signal to all the chambercards using balanced ECL signals over eighteen 3.5 m long, shielded, twisted-pair cables. Figure 3.11 shows the state diagram and transitions for the clockcard logic. In data-acquisition mode, the central clock is running at 200 MHz (5 ns period) from which the chambercards derive the four 50 MHz phase clocks. A single trigger input on the clockcard is used to freeze the contents of the buffers on all chambercards. When the trigger input on the clockcard is asserted, the central clock is stopped for 520 ns during which a 200 ns `EVINC` pulse is sent to the chambercards to select the next buffer for storing samples. After the 520 ns delay, the clockcard automatically restarts the 200 MHz clock and the chambercards continue data acquisition. The dead time caused by this delay is less than the overall dead time of the experiment and hence does not affect the data-taking efficiency.

Because the trigger signal and the sampling clock are asynchronous, the clockcard will stop the sampling clock three clock-cycles after the trigger input became asserted. A series of three flip-flops ensures stable operation with respect to the asynchronous trigger input. The timing of the trigger signal with respect to the last clock cycle is measured with a 1 ns resolution TDC. The TDC is started by the trigger and stopped by the `TDC_STOP` output signal generated by the last edge of the stopping clock.

In read-out mode, the clockcard passes control of the `EVINC` signal and the clock drivers to the read-out controller. The read-out controller generates pulses on the `RDCLK` input of the clockcard. These pulses are forwarded to the chambercards to clock data during the read-out. As the transition from data-acquisition to read-out mode is also asynchronous with respect to the sampling clock, this signal is routed from the read-out controller to the clockcard first. The clockcard will detect the transition to read-out mode and will stop the sampling clock before forwarding the transition to the chambercards. Vice versa, the sampling clock will only be started 200 ns after the transition to data-acquisition mode.

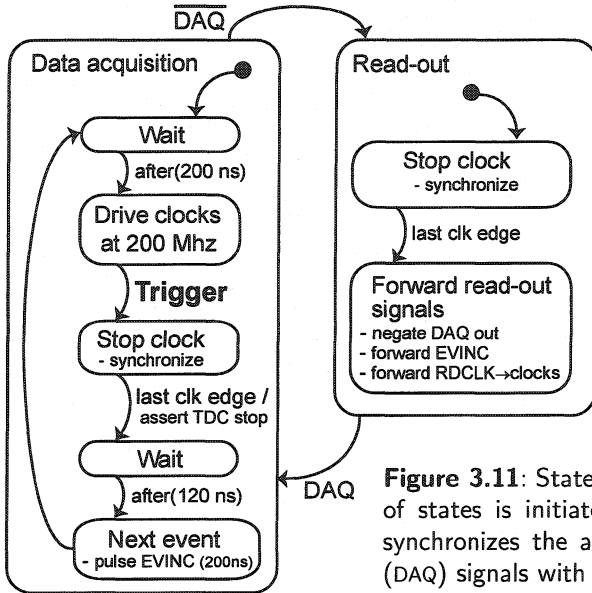


Figure 3.11: State diagram of the clockcard. A sequence of states is initiated by a trigger. The clockcard also synchronizes the asynchronous trigger and mode-switch (DAQ) signals with the central clock.

3.5 Read-out system

The complete read-out system for a single honeycomb module is schematically drawn in Figure 3.12. Instead of reading out each chambercard individually, all chambercards are read out using a single bus, with the cards forming a clocked pipeline (like a bucket brigade). A latch at the output stage of each card is used to hold the output data stable during read-out cycles. The eighteen cards on a module are connected using 34-way flat-cables. The number of read-out cables from the detector to the central counting room is therefore kept to one per module. The interconnections between chambercards facilitates replacing or removing a faulty chambercard.

The read-out pipeline is controlled by a token passing protocol. All cards of which the `TOKEN_IN` input is not active are quiescent. Only one card owns the token (`TOKEN_IN` asserted, `TOKEN_OUT` negated) and drives the 24 data lines. All cards downstream of this card have passed the token already upstream. These cards transfer the data synchronously from card to card at each read-out clock pulse. After the active chambercard has put all its data on the bus (for one event), it activates the `TOKEN_OUT` signal to pass the token to the upstream card. The proper functioning of this pipeline relies on equal timing of the clock signal in all cards which is assured by distributing the read-out clock via the clockcard.

A piggy-back on the most downstream chambercard, called the transfer-card, converts the signals on the read-out bus to balanced ECL in order to transfer the data over longer distances. This piggyback is connected with a single 64 wire shielded cable to a VME I/O-module with an on-board CPU.⁷ This VME-module, called RIO (for RISC I/O module), is the read-out controller for a honeycomb module. The RIO has a MIPS R3000 processor which runs a custom-made program that is downloaded over the VME interface. It is programmed to do the honeycomb read-out and to compress the data. The read-out

⁷Creative Electronic Systems S.A. RIO 8260

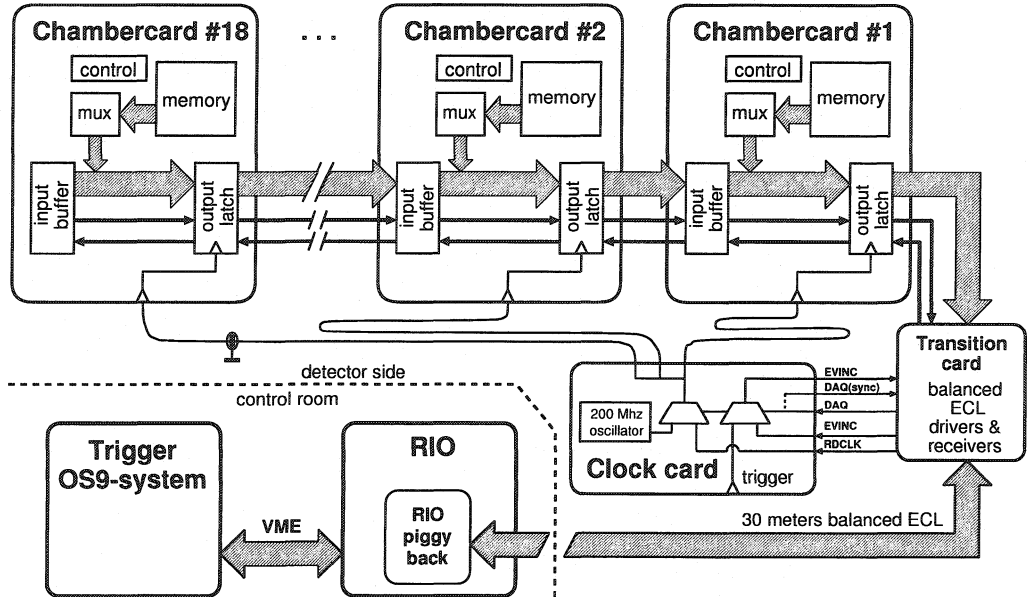


Figure 3.12: The read-out system for one honeycomb module.

protocol, the program, and the data compression are discussed in sections 3.5.1, 3.5.2, and 3.5.3, respectively.

The RIO has an interface card on the I/O-bus connector which contains an 8 bits output register and a 24 bits input register. All control signals are generated using the output register. The DAQ and EVINC signal are synchronized by the clockcard hardware (section 3.4.3). Any read instruction from the input register creates a pulse on the RDCLK signal which serves as the read-out clock. The data belonging to the read-out clock pulse N are available at the RIO input approximately 400 ns later due to the propagation delay over the 30 m long cable (round-way trip). As a read cycle of the RIO processor normally takes 120 ns, instead of letting the processor wait, the data for clock pulse N are fetched by the processor at the next read from the input register. This read then generates the read-out clock pulse $N + 1$. The RIO must therefore not read data faster than the cable delay (≈ 2 MHz). A dummy read followed by a 500 ns delay is necessary to fetch the first read-out data. The 500 ns delay in between read cycles is used for the data compression.

The read-out of the honeycomb tracker forms a sub-system of the CHORUS trigger data-acquisition system [160]. As explained in section 3.4.2, the chambercards can be either in data-acquisition or in read-out mode. Requests for transitions between these modes are sent by the trigger data-acquisition program. During a neutrino spill the detector is in data-acquisition mode. After the spill, a read-out request is sent to the sub-systems telling them how many events have been triggered and need to be read out. When all sub-systems have been read out, they are switched back to data-acquisition mode. For the honeycomb tracker read-out, all actions in data-acquisition mode are implemented in hardware. In read-out mode, all signals are generated by software. The trigger data-acquisition program assembles the data produced by the RIOs for the three honeycomb modules and sends them to the central event builder. The event builder assembles all data into a single structure, distributes the data to other tasks (e.g. histogramming) and writes them to disk.

3.5.1 Read-out protocol

The read-out program which runs on the RIO controls the transitions between the states of Figure 3.10 implemented in hardware on the chambercards. The transitions between the read-out states are controlled by three signals, TOKEN_IN, EVINC, and CLK. The TOKEN_IN signal of the first chambercard is directly connected to a bit in the RIO's output register. The EVINC signal is controlled by another bit in the output register, but routed via the clockcard, because in data-acquisition mode the clockcard generates this signal. The CLK signal is generated by a read of the RIO's input register which sends a RDCLK pulse to the clockcard which distributes it to all chambercards.

The RIO software starts by asserting TOKEN_IN on the first chambercard to reload the counters and phase-clocks with the address of the first sample in the first buffer. The memories require two clock cycles before the data are available at their output (internal pipeline). These two empty data cycles are used to transfer some additional internal-state information of the chambercards to the RIO. The RIO generates the first CLK pulse with a dummy read from the input register and then fetches the first word of the first chambercard from the input register 500 ns later. This first word is a card identifier; it contains a logical address of the card in the following bit pattern:

23	16	15	14	8	7	0
not used		FIFO empty	not used		logical address	

The value of the address is set on a card with rotary switches and is coded according to the position of the card inside a tracker module. It is used in the read-out program to identify which card is actually being read and in the tracking code to recover the wire corresponding to each hit. The RDCLK pulse for capturing the first word clocks out a second status word from the chambercard. This second word contains the current value of the event counter (section 3.4.2):

23	16	15	14	8	7	0
not used		FIFO empty	not used		event counter	

The value of the event counter is used to check the event sequence by comparing it to its previous value if more than one event is read. If the FIFO empty bit is 0 in both words then the current card contains data. If it is 1 then the card contains no data (or all events were already read); in this case the token has been passed on directly to the next card. The first two words are referred to as the header.

If a card contains data, the next 1024 CLK cycles, generated by read instructions from the input register, will transfer all the data of the 256 timeslots of the 72 wires (assuming configuration with $n = 8$). These data are transferred as four groups of 18 wires. The multiplexer on a chambercard connects the internal data bus to the read-out bus and selects the correct wire group. The multiplexer is programmed such that all timeslots of one wire group are transferred sequentially, before the data of the next wire group is transferred. The 24 bits read from a chambercard contain six extra status bits from the phase-clocks and timeslot-counters:

23	20	19	18	17	0
timeslot counters		phase bits		wire sample bits	

The extra bits from the two timeslot counters are used to detect counter or transfer sequence errors. To each phase belong four bits of one of the two timeslot counters. In four consecutive data words, all phases are present and the complete bit pattern of the two timeslot counters can be reconstructed. The sequence of phases and the values of the timeslot counters can then be checked for consistency. The full sequence of header and data words is shown in Figure 3.13.

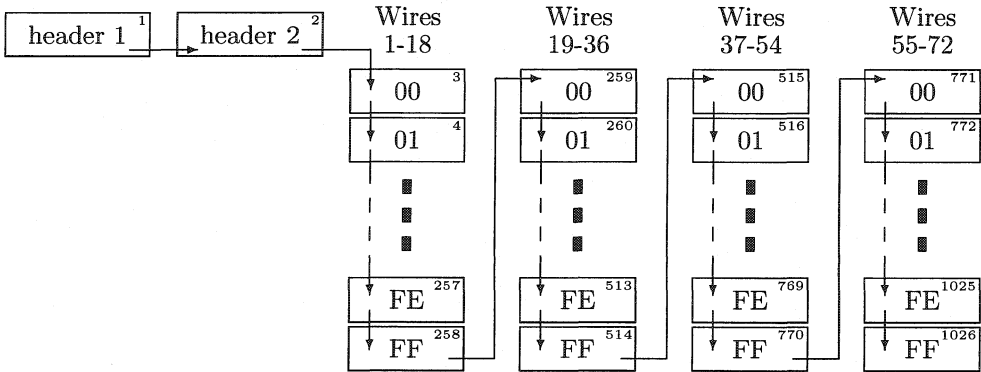


Figure 3.13: Data transfer sequence of one chambercard: after two header words, all 256 timeslots of the first group of wires are transferred, followed by the timeslots for the other three wire groups. The small number in the upper right corner indicates the transfer sequence number ($1 + 4 \times 256 = 1026$). The large number inside the boxes indicate the timeslot in hexadecimal code.

One memory access time after CLK pulse 1025 ($2+1024-1$), the last data word will be on the internal read-out bus of the active chambercard. At this time the `TOKEN_OUT` output of the card is asserted, activating the next chambercard in the pipeline. On CLK pulse 1026 the last data word of the card is latched in its output buffer and immediately afterward the input buffer is enabled. The upstream card's output latch now contains that card's first header-word. With the next clock pulses, the data from the upstream card are passed on to the downstream card. This timing of passing the token ensures a seamless transfer of data.

The RIO will continue reading headers and data from all cards until the data of all cards for one event have been read. With all data of one event read, the RIO negates the `TOKEN` signal and can then either re-read the current event or pulse `EVINC` and start reading the next event. The first option is used in the case of a recoverable read-out error, for example due to a noise spike on the `TOKEN` line (later resolved by a spike filter in the control logic). In case of read-out errors, the RIO can negate `TOKEN` at anytime and bring all cards back to the initial state of the read-out mode. It can then retry the read-out.

3.5.2 Read-out program

To add the honeycomb modules to the read-out of the trigger sub-system of the CHORUS data-acquisition system, the trigger data-acquisition program had to be extended. The trigger data-acquisition program was written in C++ and runs on a VME-master module.

The actual read-out of the hardware is done by a program running on the RIO, which was written in C. A small part was written in assembly to do fast read-out and data compression. The trigger data-acquisition program will be referred to as the master, the RIO's main program and the client task connected with it as the slave.

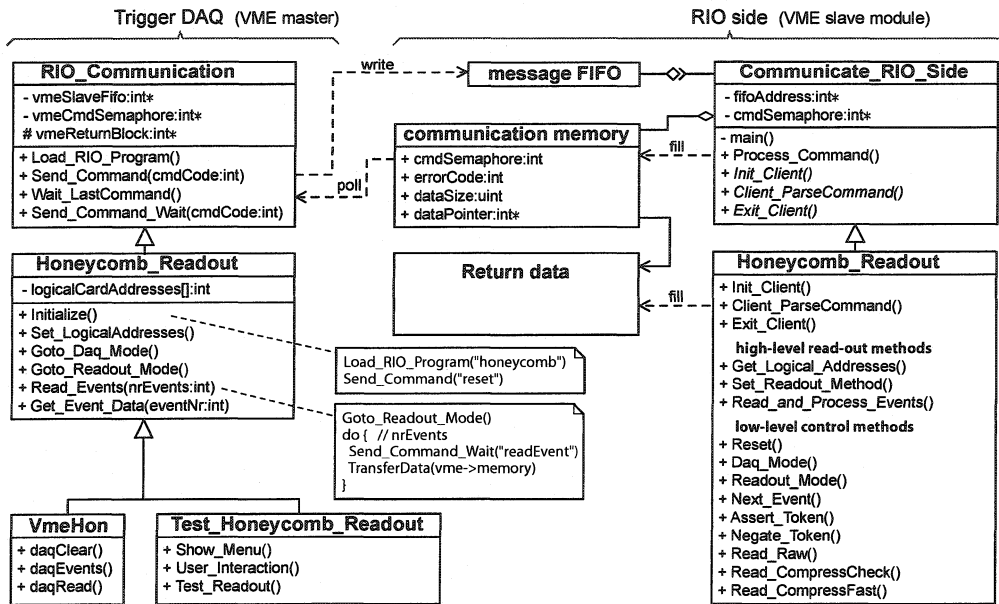


Figure 3.14: Read-out and communication class diagram of the honeycomb module read-out program.

The class diagram of the read-out programs is depicted in Figure 3.14. The communication between the master and the slave is based on a simple send-command, wait-for-completion polling algorithm. The master writes via the VME bus a command to a message FIFO on the RIO. The command will be interpreted by the slave's main communication program and executed by its client application. Upon reception of the command, a VME-readable semaphore is set to the value of the command code. When the client finishes processing of the command, some VME-readable memory words will contain the command-completion value and the length and address of any command-specific data. The slave then inverts the command semaphore to signal to the master that it has finished processing the command. The master can then retrieve the returned data over the VME bus.

The basic communication (FIFO, command semaphore, result passing) is done by the main communication program on the slave side using the `Communicate_RIO_Side` class and on the master side using the `RIO_Communication` class. The latter class also downloads and starts the slave's program via the VME bus. On the slave side, the honeycomb read-out client is attached to the main program via three function pointers. The function `Init_Client` is called at startup of the main slave program and the function `Client_ParseCommand` is called for each command received. When the slave's program is told to quit by the master, it calls the `Exit_Client` function.

On the master side, the honeycomb read-out class, `Honeycomb_Readout`, is derived from the communication class. It implements all the methods needed to control and read out the chambercards. Most of the methods of this class just send a command to the honeycomb read-out client on the slave side. During the read-out, the raw or compressed honeycomb data are transferred over the VME bus from RIO internal memory to the master's memory. The `VmeHon` class is a bridge class that encapsulates the `Honeycomb_Readout` class. Its definition is conform to the instrument class interface as expected by the trigger data-acquisition program. This class just calls the `Honeycomb_Readout` class methods and puts a standard data header in front of the data.

Several read-out methods are provided by the read-out client on the slave side. The main read-out routine loops over all cards, checks the header words and then calls one of three possible data reading routines. The simplest routine just reads the data, checks the continuous counting of the phase- and timeslot-counters in the six extra bits and puts the raw data directly into the output buffer. The other two routines perform data compression described in the next section. Of these, one routine also checks the consistency of the data sequence. The other is written in assembly and performs the compression in only 320 ns per word. This routine is used when the read-out must be fast. In other cases, the check-and-compress routine is used to detect and diagnose any malfunctioning of the system.

A stand-alone program was developed which uses the `Test_Honeycomb_Readout` class. This class allows the user to make menu-driven tests of the read-out pipeline and protocol. This test program can also analyze the data to check for broken memories or other hardware problems on the chambercards.

3.5.3 Data compression

The data can be compressed online by a simple algorithm. For each group of 18 wires the data bits of the first timeslot are stored. The next sample is compared to this value, if there is any change in one or more of the bits the new value of the 18 data bits will be stored in the output buffer. The 32 bits word added to the output also contains the current timeslot, wire group and card number (modulo 16):

31 card number	28	27 wire group	26	25 timeslot number	18	17	new wire status bits	0
----------------------	----	---------------------	----	--------------------------	----	----	-------------------------	---

In this way, all changes in the 18 data bits are stored for all 256 timeslots. For typical events, this reduces the data volume to less than 2%, on average, of the original size of 40.5 Kb. For the fast routine, the compression time is completely absorbed in the minimum read-out time of 9.2 ms per event dictated by the maximum read-out speed of 2 MHz.

3.6 Tracking

A tracking routine was developed [230] to use the honeycomb data in the CHORUS reconstruction and analysis program [186]. In this, the hits in a module are used to reconstruct track segments. The segments in the different modules are then combined into 3-D tracks, which can then be used by the general reconstruction and analysis.

3.6.1 Drift-time calibration

To be able to use the drift time in the tracking, the correlation between drift time t and distance r needs to be known first. This r - t relation is extracted from the data with some assumptions about the drift velocity. For a 1:1 mixture of Ar and CO₂, the drift velocity v depends strongly on the electric field strength. The radial electric field is

$$E(r) = \frac{V_{\text{anode}}}{\ln r_{\text{cell}}/r_{\text{wire}}} \cdot \frac{1}{r} ,$$

with r_{cell} (r_{wire}) the radius of the cell (wire) and V_{anode} the anode high voltage. With the wire chamber simulation program *garfield* [231], the field configuration and expected drift velocity or drift time can be calculated for a honeycomb cell and a given gas mixture. Contour lines for the electric field strength are shown in Figure 3.15a, while Figure 3.15b shows the expected drift time. From the simulation it can be seen that in the corners of the honeycomb cells the electric field is low and the drift time becomes long.




Figure 3.15: Electric field contours for wire at 2250 V (a) and drift time contours (b) in a honeycomb cell as calculated by the *garfield* wire chamber simulation program. Note that 1/3th of the maximum drift is spent in the corners of the cells where the field is low and not radial anymore.

To calculate the r - t relation, one could make the first-order approximation that the drift velocity v is proportional to $E(r)$. One then finds the following r - t relation:

$$t(r) = \int_{r_{\text{wire}}}^r \frac{dr'}{v(r')} = \frac{\ln r_{\text{cell}}/r_{\text{wire}}}{c V_{\text{anode}}} \int_{r_{\text{wire}}}^r r' dr' \Rightarrow t(r) - t_0 = \frac{\ln r_{\text{cell}}/r_{\text{wire}}}{c V_{\text{anode}}} r^2 , \quad (3.1)$$

where c is the proportionality constant. Taking the inverse of equation (3.1), the dependency of the wire distance on drift time is $r \propto 1/\sqrt{t - t_0}$. A more general dependency $v(E) \propto E^{c_2/(1-c_2)}$ between drift velocity and electric field leads to a power law for the r - t relation of the form:

$$r(t) = r_{\text{cell}} \left(\frac{t - t_0}{t_{\text{max}}} \right)^{-c_2} . \quad (3.2)$$

Under the assumption of proportionality of v and E , equation (3.2) is identical to equation (3.1) with $c_2 = 1/2$. The coefficients t_0 , t_{\max} and c_2 are derived from a set of tracks by minimizing the sum of the χ^2 of tracks fitted using equation (3.2) for the $r - t$ relation. The coefficients are calculated for each plane individually since they depend strongly on the high-voltage setting which is not the same for all planes. Typical values are $t_{\max} = 300$ ns and $c_2 = 0.44$. This value of c_2 is close to the initial assumption of $v \propto E$. The coefficient t_0 is determined by the trigger delay time. For hits in the corners of the honeycomb cells which have long drift times, the $r - t$ relation of equation (3.2) is not adequate and the hits are simply assigned the maximum drift distance. The tracks used for the drift-time calibration are isolated muon tracks. The same set of tracks is also used to align the three modules of the honeycomb tracker internally and to align the tracker with the rest of the CHORUS detector.

3.6.2 Track finding per module

After decompressing the raw data, each hit that falls inside the event-related drift-time window yields a circle around the wire. The procedure to look for track segments in a module starts with looking for the highest-quality segments which have 6 out of 6 planes hit. The procedure is then repeated for finding lower-quality segments which have only n_{\min} hits, where n_{\min} is 5, 4 and 3. Hits in a segment are marked as used and cannot be used for segments requiring less hits. The procedure combines all unused hits in two planes which can lead to a segment with n_{\min} hits. For $n_{\min} = 6$ only the hits in planes 1 and 6 are considered, while for $n_{\min} = 4$ all hit combinations in the planes 1-6, 1-5, 1-4, 2-6, 2-5, and 3-6 are tried. The combinations are also restricted to those hit pairs that define segments with a slope of less than 45° with the normal to the planes. Each combination of two hits has four possible tangents to the two circles around the corresponding wires. At least one more hit is needed to resolve which of those 4 tangents could be a particle track. Therefore n_{\min} must be at least 3.

Each of the four tangents defines a possible track path. The unused hits that are closer than $1000 \mu\text{m}$ to such a track path in any of the other planes are collected. If the number of collected hits on a segment is larger than n_{\min} , a straight line is fitted through the hits and a quality value is assigned to the segment. The quality value is a function of the resolution, cell efficiency, and track isolation. The precise definition of the quality value is given in Table 3.1. The segment is then inserted in a list. If there are already segments in this list which have two hits on the same side of a wire in common with the newly found segment, only the best segment is kept.

Since in a particle shower many cells in a region are hit, the number of segments could become very large. Therefore the length of this list is limited to 100 entries and if it gets full the segment with the lowest quality is removed. When all segments with n_{\min} hits are inserted in the list, the list is pruned for segments sharing hits. The list is first sorted according to the quality value and the hits from the best segment are marked as used. This procedure then continues with the lower-quality segments where the hits are marked if the segment does not include any used hits, otherwise the segment is removed.

The results of this track-finding algorithm are shown in Figure 3.16. From the figure it is clear that the single-cell efficiency is high as 18 out of 18 cells have a hit. The resolution of the individual planes is discussed in section 3.7.1.

The quality value for each segment is calculated using the following formula:

$$Q_{\text{segment}} = +0.8 \cdot \sum_{i=1}^{n_{\text{hits}}} K_{\text{isol},i} + 0.5 \cdot n_{\text{gap}} + \\ - \left(\sqrt{\chi^2} - 1 \right) - |\theta_{\text{segment}}| + \\ - 0.5 \cdot n_{\text{nohit}} - n_{\text{late}} - 0.25 \cdot n_{\text{early}} \quad ,$$

where:

- n_{hits} is the number of hits found within the 2000 μm wide road defining the segment.
 K_{isol} is the isolation factor for each hit. It has the value 1 if any of the two neighbouring cells in the plane are also hit else it has the value 1.5. The isolation is used to create a preference for non-shower tracks.
 n_{gap} is the number of planes where the segment fit crosses the 1.13 mm gap in between cells. The coefficient is positive because missing two cells can actually be considered a hit in the region between them with a resolution of $1.13/\sqrt{12}$ mm.
 χ^2 is the normalized sum of the squared residuals divided by the degrees of freedom of the fit: $\chi^2 = (n_{\text{hits}} - 2)^{-1} \sum (\Delta r / \sigma_{\text{hit}})^2$
 θ_{segment} is the angle between the track segment and the normal to the plane measured in radians. This value is subtracted to give preference to more forward going tracks.
 n_{nohit} is the number of planes without a hit where one was expected, which counts as a penalty. Therefore this number is subtracted from the quality value.

If a cell was hit, it is still possible the hit was not included in the segment, either because it is outside the acceptance road or because it is marked as used. If a used hit is closer than $4\sigma_r$ to the segment, there is no penalty for not including it. If it is far from the segment there are two cases, which have different penalties:


- n_{late} is the number of these hits that have drift distance longer than expected from the segment's position. This case is relatively badly penalized. If a particle really passed closer to the wire, its hit should have been detected at a shorter drift time.
 n_{early} is the number of these hits that have a drift distance shorter than expected from the segment's position. This case is less heavily penalized, as the segment's expected hit could overlap with this earlier hit.

Table 3.1: Calculation of quality value for a segment.

3.6.3 Reconstructing 3-D tracks

The track segments from the three modules are combined to make 3-D tracks. A similar algorithm as in segment finding is used. Each combination of two segments from two orientations is combined with the remaining orientation. If there is a matching segment or at least two matching hits in the third orientation, the combined hits from all three orientations form a 3-D track candidate. If the 3-D track reconstructed from the initial two segments crosses the inactive zone between monolayers (section 3.3.2), the combination is also accepted as a valid track candidate.

Figure 3.16: A set of segments in a module showing the drift circles around the wires and the fitted track segments tangent to the hits.



The hits on the track candidate are then refitted to find the four track parameters (two positions and two slopes). Because the coordinate along the wire in an orientation is now known from the other orientations, the drift distance is first corrected for the signal propagation time along the wire. The track candidate is then assigned a quality value and put in a list of 3-D candidates. If this list gets full, the lowest quality candidate is removed. The definition of the quality value is shown in Table 3.2. In exactly the same way as for segments in a single module, the list with track candidates is sorted and then pruned from overlapping segments in a module.

The 3-D tracks can be extrapolated to other detectors and refitted taking into account all available data. The main use of the honeycomb tracks is the momentum fit using the hexagonal magnet.

The quality value for 3-D track candidates is calculated as follows:

$$\begin{aligned}
 Q_{3-D} = & +0.8 \cdot \sum_{i=1}^{n_{\text{hits}}} K_{\text{isol},i} + 0.5 \cdot n_{\text{gap}} - 18 \cdot \left(\sqrt{\chi^2} - 1 \right) + \\
 & - 0.5 \cdot n_{\text{nohit}} - 0.25 \cdot n_{\text{early}} - n_{\text{late}} + 0.8 \cdot n_{\text{used}} \\
 & + 6 \cdot (n_{\text{segments}} + c_{\text{dead}} - 3)
 \end{aligned}$$

The definitions of the first six terms are identical to the ones in Table 3.1, but with different weight factors. The additional terms take into account the orientations participating in the track. The various terms are weighted such that the plane and module contributions can be added. The new terms with respect to Table 3.1 have the following meaning:

- n_{used} is the number of hits which are within $\pm 4\sigma_{\text{hit}}$ around the track path, but excluded from a segment because they were already used in another segment.
- n_{segments} counts the number of modules where the track has segments.
- c_{dead} is a correction factor for tracks that cross (near to) one of the dead zones between monolayers. If a track hits the center of a dead zone, 0.5 is added. If a track passes within 5 mm of a dead zone, 0.25 is added. And if a track misses a module completely, 0.25 is added.

Table 3.2: Calculation of quality value for 3-D track candidates.

3.7 Performance and discussion

3.7.1 Resolution

The resolution of the individual honeycomb tracker planes has been determined from the hit residuals with respect to reconstructed 3-D single muon tracks. The widths of the distributions of the residuals are shown in Figure 3.17 for each plane. Most planes have a resolution of about $200\ \mu\text{m}$. The last two planes of the first module had many wires removed and are operated at a lower high-voltage setting.





Figure 3.17: Track residual distribution per honeycomb plane [230].

3.7.2 Read-out electronics

The design of the read-out electronics has shown to offer several advantages over standard TDCs, namely: the (almost) dead-time free, multi event, buffering; low cabling cost; and excellent time stability. The bit-stream design proved also to be very resilient against failures. The full reconstruction of the signal on a single wire made it possible to find signal hits in the case of a broken memory chip or an oscillating amplifier.

A broken memory generates a random pattern for all addresses. Due to the phase splitting of the bit-stream, half of the phases are stored in the other memory chip. By only looking at the two phases stored in the functioning memory chip, the signal on the wire could be reconstructed with a time resolution of $10/\sqrt{12}$ ns instead of $5/\sqrt{12}$ ns. Also oscillations of the amplifiers could be filtered. Most oscillations were due to the aluminum cover plate near the chambercards, see Figure 3.18. The connection to ground for this cover plate was about 3 m long. The cover plate could therefore act as an antenna for frequencies higher than 100 MHz, retransmitting noise generated by the digital logic. This noise was picked up by the amplifiers and caused them to oscillate. The oscillations generated a single one bit every three to four timeslots. By only accepting hits longer than four timeslots, the real hits on the wires could be filtered from the about 70 noise hits. The oscillations were eventually suppressed by removing the cover plate.

Figure 3.18: Cut-away view of ground layers on chambercard and aluminum shielding and ground connection paths of honeycomb module.



3.7.3 Conclusion

The honeycomb tracker has achieved its original aims; it makes stand-alone reconstruction of 3-D tracks possible with good resolution. These tracks have been used to align the diamond tracker paddles with respect to each-other and the other sub-detectors. Before applying the alignment determined with the honeycomb tracker, the momentum resolution of the hadron-spectrometer was found to be:

$$\frac{\Delta p}{p} = \sqrt{0.35^2 + (0.038p/1 \text{ GeV})^2} .$$

After applying the new alignment, the momentum resolution was found to be [232]:

$$\frac{\Delta p}{p} = \sqrt{0.35^2 + (0.025p/1 \text{ GeV})^2} .$$

In this, the first term is due to multiple scattering and the second incorporates the spatial resolution and alignment effects. The introduction of the honeycomb tracker has led to a significant improvement in the momentum determination of hadrons. This improvement could even be extended to the period before the installation of the honeycomb tracker. The alignment obtained with the honeycomb data in the 1997 run was in good agreement with that of the 1996 run. As the alignment between these different periods was consistent, the same alignment of the diamond tracker paddles was applied to the data of the 1995 run.

Chapter 4

Track finding in emulsion

In Chapter 2 the setup of the CHORUS experiment has been introduced, including a hardware-based method to scan emulsion. This chapter describes the development of a new approach undertaken at CERN. The principles in this approach were to use off-the-shelf hardware and to use software as much as possible. This required the development of track-finding code for the particular case of emulsion images. Electronic detectors usually yield information in one or two projections; emulsion yields 3-D position information. Electronic detectors typically measure only a few track points; emulsion tracks typically contain many hits. The difficulty with emulsion is that these hits are buried in a large number of background hits.

An algorithm that could efficiently find tracks in this high noise environment was developed. Although originally written for track finding in emulsion, the algorithm and its tools could be used in more general applications and have, therefore, been implemented as an object-oriented C++ toolkit. Part of this chapter is a copy of a published paper describing this toolkit [233]. In this chapter, the algorithms and implementations of the track-finding classes and the container classes developed for fast searching in multi-dimensional spaces are presented. The track-finding efficiency, estimated using a Monte-Carlo simulating, is also presented. The expected performance of the algorithm has been investigated. The tracking code was originally designed to reconstruct all tracks. However, in the scan-back stage of event location, a track-selector like approach (section 2.9.4) is sufficient and faster. This was also implemented in software and is described in section 4.5. Finally, the application to real emulsion data is presented.

4.1 Introduction

Automatic emulsion scanning with computer-controlled microscope stages and digital read-out and processing of emulsion images was pioneered by the the FKEN laboratory in Nagoya (Japan). As is described in section 2.9, the Nagoya approach to emulsion scanning is based on a brute force, hardware based, track-finding system which examines a fixed set of 16 images. Originally, only a track with known slope could be located automatically. With the development of ever faster hardware, this restriction disappeared because the hardware could simply check for many slopes.

When one examines the emulsion-scanning strategies used in CHORUS in detail (section 2.10), three different stages can be distinguished: scan-back, net-scan and eye-scan. These stages differ in the area and thickness scanned and whether all tracks or only a single track is being searched for. During scan-back (section 2.10.2), a single track is looked for and the area scanned is large on the interface sheets but very small on the target sheets. During net-scan (section 2.10.3), all tracks are reconstructed and the area is large. In both stages, it is sufficient to examine only a thin slice of one emulsion surface to find all interesting tracks. The exception is scan-back on the interface sheets where both surfaces need to be scanned.

The net-scan procedure has a short-coming which becomes apparent for events with a secondary vertex or kink. Net-scan is comparable to electronic tracking detectors in the sense that tracks and vertices are reconstructed from a few measurements along the paths of the tracks. The complete particle track or the actual vertex is not seen. From the net-scan data alone, it is impossible to tell if a secondary vertex was caused by a decay or an interaction. The net-scan procedure can also not distinguish between the decay of a charged or neutral short-lived particle if the decaying particle does not cross the upstream surface of at least one emulsion plate. These limitations re-introduced human-eye scanning in the emulsion analysis. The advantage of net-scan is that now only a small sample of events needs to be scanned at a well known location in the emulsion and with a partially known topology. During such eye-scanning, one to several plates are examined through their full thickness and the tracks and vertices of interest (some of which are already known) are inspected, measured and registered in a computer readable format. Eye-scan corresponds thus to a scanning stage with small to medium areas but full thickness.

During the development of the scanning and track-finding hardware in Nagoya, the optics and the limitation to 16 images have never changed. So even though the scanning speed has increased several orders of magnitude (including the CCD camera speed), the optics still limit the field of view to about $150\ \mu\text{m} \times 150\ \mu\text{m}$ and the hardwired 16 image limit restricts the scanning to emulsion slices of around $100\ \mu\text{m}$ thick. Historically of course, the track-selector was designed for doing scan-back only; it was the increased scanning speed which led to the development of the net-scan procedure. Net-scan is probably close to the best that can be done for full automatic event reconstruction given the limitations of the hardware and a given time frame determining the time that can be spent on each event.

Within the CERN scanning laboratory, the idea took root to redevelop automatic scanning techniques, keeping the ideas that had already been developed while avoiding known limitations and human eye-scanning as much as possible. The guiding principles in these developments were to use up-to-date instrumentation and off-the-shelf electronic com-

ponents wherever possible and implement as much as possible all pattern recognition in software. Using off-the-shelf components and implementing pattern recognition in software, allows one to profit directly from Moore's law,¹ while avoiding the long development time and relative inflexibility of home-built designated hardware.

The main subject of this chapter is the software developed for track-finding in a set of emulsion images. Another pattern-recognition problem improved in the CERN developments was the location of reference points on the emulsion plates. This has already been addressed in section 2.10.1. In the next sections, some of the new instrumentation developed for the CERN scanning laboratory will be briefly described, before returning to the main subject with a discussion of the characteristics of emulsion data and the constraints these place on a tracking algorithm.

4.1.1 Microscope optics and stages

Normal microscope optics are designed to render an accurate image of an object to the smallest detail possible usually using white light. On the contrary, for the reconstruction and measurement of charged particle tracks in emulsion, the shapes of the grains are not important; the only relevant parameter for each grain is its position. For emulsion scanning, the optical system should yield an image of sufficiently high contrast such that grains can readily be identified and of sufficiently high resolution, both transversely and axially, such that their position can be accurately determined. The typical grain size in the CHORUS emulsion is $0.8\ \mu\text{m}$. Given a typical dimension of pixels on image sensors of around $10\ \mu\text{m}$, the transverse resolution dictates a magnification of around $40\times$ to have about three pixels per grain. The depth of field, the size of the axial region that is in focus, should be below $2\ \mu\text{m}$ such that the z position of individual grains can be determined with reasonable accuracy. In order to scan large areas, the field of view should be as large as possible, while the field curvature should be minimal such that imaged slices in the emulsion are basically flat planes. The free working distance, the distance between the first lens surface in the system and the object in focus, has to be more than 1 mm to be able to scan the full-thickness of an emulsion plate.

In the 1970s, Tiyoda designed an objective specifically intended for emulsion work, on request from and in collaboration with Nagoya University. This $50\times$ oil-immersion lens represented a compromise between automatic scanning and comfort for eye-scanning. It was designed with a numerical aperture (NA) of 0.85 using green light at 550 nm. A higher NA or shorter wavelength would give a better resolution, but it also decreases the contrast making grain recognition more difficult for the human eye. Its field of view is free from distortions up to a diameter of about $200\ \mu\text{m}$, which is about the maximum a (trained) human can quickly oversee. The practical depth of field for grain recognition is about $2.6\ \mu\text{m}$.

A comprehensive study of the optics required for emulsion scanning [234], showed that a larger field of view could be achieved with a new optical design purely intended for automatic scanning. In collaboration with industry [235], a new optical system was developed with as goals a field of view of $500\ \mu\text{m}$ diameter and a depth of field of $1.5\ \mu\text{m}$. The different refractive index of the various types of emulsion plates required that the

¹Moore's law, posed in 1965, states that the number of transistors on integrated circuits doubles every 18 months which is accompanied by a similar increase in processor speeds. For various takes on this not-so-constant law, see <http://www.answers.com/topic/moore-s-law>.

optics could be tuned to deal with these differences. These specifications were realized with an oil-immersion objective with a NA of 1.05 using a blue light source at 436 nm (g-line of a mercury-vapour arc-lamp). It can accommodate a variable refractive index between the object and the front surface of the objective lens, within the range $1.49 < n < 1.54$, by moving a group of lenses inside the objective, which contains a total of eleven lenses. The magnification is selectable from 28 \times , 40 \times , 60 \times , and 80 \times by exchanging an adapter tube. The high NA and short wavelength ensure good resolution in both transverse and axial directions, even at the minimum magnification of 28 \times . With a typical one square centimeter image sensor, the actual field of view is $350\ \mu\text{m} \times 350\ \mu\text{m}$, seven times larger than the $150\ \mu\text{m} \times 120\ \mu\text{m}$ field of view of the Tiyoda lens and CCD system used in Nagoya. The field curvature is less than $1\ \mu\text{m}$ up to the very edge of the field of view. The practical depth of field for grain recognition is about $1.2\ \mu\text{m}$, more than two times better than the Tiyoda objective. A more extensive description of the optical system can be found in Refs. 215,236.




Figure 4.1: Mega-pixel camera view of a piece of emulsion with the 50 \times Tiyoda objective (top, width about $220\ \mu\text{m}$) and the 40 \times Jenoptik objective (bottom, width about $280\ \mu\text{m}$). The optical axis is located at the right-hand side of the images and indicated by the small black crosses. The Tiyoda objective suffers from clear radial distortion at longer radial distance from the optical axis (left-side of image), while the Jenoptik objective shows no such imaging artefacts. From the images, one can see that the Tiyoda objective has a higher contrast (due to lower NA, higher magnification, and longer wavelength) making it much easier for the human eye to spot the grains. The black line and area in the images are camera defects.

Due to the smaller depth of field of the new optics more independent images can be taken inside an emulsion layer. For the same reason, the z resolution of the grain positions is also increased. For scanning $100\ \mu\text{m}$ thick slices of emulsion, typically 20 to 30 independently imaged planes inside the emulsion are taken, called layers. The larger field of view reduces the number of views to be taken when scanning large surface areas. For small areas, the processing can be restricted to the central area of the image. As there is no hard limit on the number of images that can be taken, the eye-scan stage can be replaced by simply taking images through the whole depth of the emulsion. With a $3\ \mu\text{m}$ layer spacing, this gives between 100 and 120 images per side of the target emulsion plates.

Even though the new optics allow more of the emulsion to be viewed per operation, the amount of data collected in a single view is still only a very tiny fraction of the total information on a single plate. A single pixel in the view of the microscope covers

a volume of about $0.35 \mu\text{m} \times 0.35 \mu\text{m} \times 3 \mu\text{m}$ of the emulsion. Considering that for scanning purposes its value can be given by a single bit as simply black or white, each target plate contains about 250 terabits of data. The 2304 plates in CHORUS contain thus 570 petabits of data. This corresponds to 344 years of continuous black and white TV images of 1024×1024 pixels at 50 Hz frame rate. From this amount of data, the need for a hybrid detector is clear as it is impossible to scan all the emulsion, see section 2.9.1. Even just scanning the predictions per emulsion module, the grain position data generated by scan-back and net-scan is several terabytes per module, with a typical burst data rate of about 2 megabyte per second. These data volumes and rates require some thoughtful design of computing and storage infrastructure. For example, interface-plate scan-back and net-scan data is normally processed offline on a cluster of computers, while scan-back in the target sheets can be handled online.

Another, straightforward, development was the introduction of a bigger microscope stage with a stroke of $40 \text{ cm} \times 80 \text{ cm}$ that accepts one complete emulsion sheet and that can handle the much heavier objective. Some minor technical upgrades are the use of an immersion-oil containment device and the introduction of plate holders that facilitate quick exchange of the plate on the microscope stage. Figure 4.2 is a photograph of one of the CERN microscopes in its latest configuration with the new optics and a new faster CMOS camera.




Figure 4.2: Photograph of an automated scanning microscope system at CERN. The photograph shows several of the components discussed in the main text.

In the CERN scanning setup, the ‘image processing’ and ‘track finding’ units in Figure 2.17 have been replaced by a digital signal processor (DSP) board and tracking software. Modifications have also been made to the ‘offline analysis’ unit in this figure, which are discussed in section 4.6.2. The DSP applies a digital filter to the images to recognize the grains. The processed images are then transferred to the control computer. On the control computer, a fast clustering algorithm reconstructs the grain positions from the pixels in the processed images. These grain positions are either stored in the database for later processing or fed directly to track-finding software running online. The tracks found online are also stored in the database. In the CERN setup, the ‘database’ block in Figure 2.17 is an object-oriented database. The data model is comprised of classes that store predictions, acquisition parameters, compressed grain data, several types of tracks, reference points, and alignment data with many references between them. The database is both read and written by the online scanning program and by the offline analysis tools. A detailed description of the instrumentation of the CERN scanning laboratory can be found in Ref. 215.

4.1.2 Tracking input characteristics

The track-finding’s job is to reconstruct particle tracks out of the grains in a set of tomographic images. The input to the track-finding consists of processed grain data which just contains the 3-dimensional grain positions, referred to as hits. The largest part of the scanning results consists of 20 to 30 layers, where each layer contains about 4000 hits. The xy -resolution (in the plane of a layer) of the hit coordinates is of the order of $0.5\ \mu\text{m}$. The z resolution is defined by the layer spacing and is about $4\ \mu\text{m}$. In this sense, emulsion data are not much different from a multi-layer 2-D electronic detector, like silicon pixel layers, although with much better resolution and layer spacing. The typical emulsion thickness used for tracking is about $100\ \mu\text{m}$ in which a track has about 30 high-resolution 3-D hits (for CHORUS emulsion).

Track reconstruction would be straightforward if the 30 track hits were not hidden in about 1200 other background hits. A typical volume of CHORUS emulsion data on which track finding needs to be performed, contains of the order of 10^5 grains. Of these, only about 2500 belong to interesting tracks. The noise consists mainly of randomly developed grains (fog) and low-energy tracks. Distortion of the emulsion implies that tracks can only be considered straight on a scale of about $20\ \mu\text{m}$, which complicates the track finding. Fog and distortion have been explained in section 2.9.2.

4.1.3 Algorithm restrictions and requirements

Due to distortion, the track’s direction changes gradually over a distance of around $20\ \mu\text{m}$. Position correlations between track hits are therefore only well defined for a sequence of about 5 to 10 hits. This leads naturally to a track-finding algorithm which looks only at hits in close proximity to hits already considered as part of a track. The large total number of hits limits the time a algorithm can spend on investigating each hit. Therefore fast algorithms are required for retrieval of hits by position and for acceptance calculations. The close-range relationships and fast look-up can be achieved by constructing a connection network of links between neighbouring hits. Building this network, however, still requires finding all hits in the neighbourhood of each hit. To speed up this operation a set of multi-dimensional search tools were developed. These tools are based on the extension of a binary-search tree to multi-dimensional space. These tools, implemented as ordering containers in any dimensional space, are described in section 4.2.

The track-finding algorithm uses such a 3-D ordering container for creating the links network of close hits. The network is searched for patterns consistent with particle tracks. Conceptually, the method is based on selecting the best path of connected hits in a tree of all track-compatible paths from a certain starting point. The actual implementation follows more closely a depth-first search algorithm from graph theory [237]. The algorithm is described in detail in section 4.3.

4.1.4 Toolkit abstraction

The implementation of the algorithm is general enough that it can be used for any situation where points, not necessarily 3-D, have close-range correlations. Setting up the links network and searching it do not require any direct knowledge of the hit or track model. The algorithm only requires yes or no decisions for hit acceptance and a way of comparing track candidates. In the C++ programming language, these decisions are easily isolated by calling them as abstract methods of a class representing a track segment. The track-finding algorithm can therefore be implemented as an object-oriented toolkit. The user has to implement the concrete class for doing the acceptance and comparison calculations. In the implementation of the decisions, the calculation time can be balanced with the tracking input characteristics. This allows one to tune for a particular background condition or to tune the track-finding efficiency by considering more paths through the links network.

The toolkit is currently used in two applications: in CHORUS it is used to reconstruct tracks in emulsion; in HARP [238, 239] it is used to reconstruct bent tracks in the magnetic field of a time-projection chamber. These two applications use the same tracking toolkit, but a different implementation of the hit acceptance class. In CHORUS, the implementation is tuned to be efficient in an environment with a large number of noise hits. Because of the redundancy of track hits, high hit-to-track assignment efficiency is not required and therefore strict cuts are applied to avoid including noise hits. In HARP, the implementation takes into account the track curvature due to the magnetic field.

4.2 Multi-dimensional ordering containers

In general, a tracking algorithm in k -dimensions (abbreviated to k -D) requires a k -D look-up mechanism to search for other hits in a certain range near a given hit. A simple and fast algorithm for retrieving elements in a given range in 1-D is described first. In section 4.2.2, the properties of binary-search trees and their extension to more dimensions is presented. These trees are used to construct containers for ordering elements in k -dimensional space. The range look-up algorithm can be extended to k -D space using these containers. The implementation of the k -D containers is described in section 4.2.4. In section 4.2.5 a summary of the performance with respect to a simplistic approach is presented.

4.2.1 Find-in-range algorithm

Finding all elements in a set \mathcal{P} of unique² numerical values which lie within a given range can be done fast if the elements are sorted. For an ordered set \mathcal{S} with elements p_i ,

²In practice identical values can be included.

$i \in [1, n]$, which has the property that

$$\forall i \Rightarrow p_{i-1} < p_i \quad , \quad (4.1)$$

it follows that

$$\forall k > 1 \Rightarrow |p_i - p_{i\pm 1}| < |p_i - p_{i\pm k}| \quad , \quad (4.2)$$

where $|p_i - p_j|$ represents the distance between elements p_i and p_j . Equation (4.2) states that the element with smallest distance to some element p_i is one of its neighbours in the sorted set. To find all elements in a given range can then be done by locating the first element larger than the lower bound of the range using a binary search, which runs with an upper limit of $\log n$ in time. One then takes the following elements in the set as long as they are below the upper bound of the range. The time for sorting the set \mathcal{P} has an upper limit of $n \log n$. Because the tracking requires a range search for each hit, the sorting time amortized over all searches is of the order of $\log n$.

This algorithm cannot be extended directly to more than one dimension, because the distance operator in equation (4.2) is not valid for vectors. There exist strict weak ordering operators defined on the set of k -D points that can be used in equation (4.1). However, none of these will leave equation (4.2) valid if the absolute difference is interpreted as a distance. The underlying reason is that there exists no mapping of a k -D space to a space with less dimensions that also maps distances. To make equation (4.2) valid for vectors, one would have to order them in a Voronoi tessellation [240–242], where each point occupies a volume whose borders are determined by the closest points around it. The time needed by the fastest algorithm to build a Voronoi tessellation is proportional to $n \log n$ for 2-D space and proportional to $n^{\lceil k/2 \rceil}$ for $k > 2$ [243].

4.2.2 Search trees

In a multi-dimensional space another range-finding algorithm is required, because of the impossibility to satisfy equation (4.2). The sorted sequence of equation (4.1) can be realized as a binary-search tree [237, 244]. In a binary-search tree, each node contains a value and has a left and right branch to sub-trees. The left sub-tree contains all smaller values than that of the parent, the right sub-tree all larger values. A node with no branches is called a leaf. The value stored in the node is usually associated with other data and is therefore often called a key. The time for a key search has an upper limit of h , where h is the height of the tree (number of levels). In balanced trees, the leaves are at almost equal height $h \propto \lg n$, with $\lg n \equiv \log_2 n$. Algorithms exist to build balanced trees in a time with an upper limit of $n \lg n$. Values can also be efficiently retrieved in sorted order from a tree by a walk through its nodes.

Although sorting in multiple dimensions is not possible, the concept of splitting a range can be extended to more dimensions. The equivalent of splitting a 1-D range into sub-ranges at some key value, is splitting a cube into 2^3 sub-cubes in 3-D. Each sub-cube is then the root of a 3-D sub-tree for an octant of the space around the parent's 3-D key value. This kind of trees are generally known as k -ary-trees or K d-trees. Here, the space covered by a sub-tree for multi-dimensional trees is referred to as a sub-volume, independently of the dimension of the space.

A balanced k -D tree with n keys has a height proportional to $\log_m n$, where $m = 2^k$ is the dimensionality multiplicity. Balancing operations (see Ref. 237) rotate nodes, as shown in Figure 4.3a, to ensure that the sub-trees of each node are approximately of

the same height. In 1-D trees this rotation involves three sub-trees. Such a rotation is needed, at maximum, twice per insert of a key in a tree. The rotations can be done in constant time because the shaded area in Figure 4.3a that moves over to node 1 in the rotation, corresponds exactly to the sub-tree γ of the new top-node 2. Rotations in more than 1-D require a complete restructuring of a large part of the tree and can therefore not be done in constant time. As can be seen in Figure 4.3b, in two (or more) dimensions, a rotation involves partial areas of the old top-node 1 which have to be mixed with existing sub-trees of the new top-node 2. Only the sub-trees α , β , and γ are unaffected by the rotation.




Figure 4.3: (a) Rotating the nodes 1 and 2 in a 1-D tree requires moving the shaded area. As this area is just the left sub-tree of node 2, the complete sub-tree γ can be set as the right sub-tree of node 1. The sub-trees indicated by α and β are not affected by the rotation. (b) In a two-dimensional tree the nodes 1 and 2 cannot be easily exchanged as the shaded area, which would become sub-areas of node 2, overlap with parts of sub-areas of node 1. Nodes 4, 5, and 6 need to be redivided as would all of their children.

A tree in which the points are inserted in random order has on average a height with an upper limit of $\log_m n$. However, in track-finding applications, the hits are usually sorted (by detector layer, row and column coordinates) and a tree could, in the worst case, degenerate to a linear sequence which has a look-up time of order n . A simple solution to avoid this kind of unbalance exists when the range of keys is known beforehand and the keys are more or less uniformly spread over the range. In this case one can build a binary tree in which all keys are stored in the leaves and internal nodes split their range through the middle. A node controls the range it covers and is in one of the following three states: it is either empty; it holds a key somewhere in its range; or it holds the branches to the nodes below it which each cover half its range. Because all keys are stored in the leaves, the key look-up time becomes proportional to $\log_m n$ which is still of the same order as for a normal tree.

The principle of inserting keys in this kind of tree is shown in Figure 4.4a. Inserting a new key starts at the root and it goes down the branches until it reaches an empty or an occupied node at the bottom of the tree. In the first case, the sub-volume represented by the empty node gets the new key assigned to it. If an occupied node is encountered, that node's range is split into equal size sub-volumes — half segment, quarter area, octant, etc. — and the current node key is moved to its sub-volume inside the split node. Next,




Figure 4.4: Structure of one, two, and three-dimensional fixed-range trees. The algorithm for inserting new keys into the sparse tree is illustrated for the 1-D case. When inserting a new key (open square) a decision into which sub-volume the key belongs is made at each level. If a sub-volume is a leaf and already occupied (black dots), that volume is split and both the new and the current key are moved to their own sub-volumes. This process is repeated until an empty leaf is reached. The numbers in (a) indicate the steps taken by the algorithm. Figures (b) and (c): structure of fixed-range trees for two and three dimensions, respectively.

the new key is placed in its corresponding sub-volume. This process is repeated if the two keys are close together and end up in the same sub-volume, as shown in Figure 4.4a at steps 3 and 4. An empty tree consists of just the empty top-node which controls the total range spanned by the tree.

There are two disadvantages of trees with a fixed range. First, prior knowledge of the range of keys to be inserted is required. In the type of application described in this work, this is not a problem. The maximum range of hit coordinates is known a priori and almost always limited by physical constraints, like the size of the detector. The second disadvantage is that the amount of memory needed to store a tree can become prohibitive. There are more internal nodes than keys when the tree is fully developed down to its smallest spacing between keys. A fully developed k -dimensional tree with height h (the root of the tree has $h = 1$), has $(m^h - 1)/(m - 1)$ nodes and ends in m^{h-1} leaves. However, only nodes actually used need to be created. The key-insert algorithm, described above, only creates those nodes which have occupied sub-volumes.

The algorithm to find elements within a given range (section 4.2.1) requires a binary search for the lower value of the search range. This search can now be replaced by a tree search which runs in the same time. The next step is to take ordered elements which in a 1-D tree can be done by a walk through the tree's nodes. However, no such walk exists for multi-dimensional trees and therefore the second part of the algorithm has to be adapted as well. Finding all elements within a certain range is done by traversing the tree structure down the branches. Any internal node whose range overlaps the requested range is searched recursively. Any key in the nodes traversed or in the leaves of the fixed-range tree lies in a range which overlaps with the search range. A final verification is needed to check whether that key actually lies inside the search range and, in that case, to add it to the list of found keys.

4.2.3 Hash table

A competitor of binary-search trees is a hash table [237,244]. With proper tuning of the hash function and the number of hash bins, hash tables have constant insert and key look-up time. In k -D space the total volume can be divided in sub-volumes as a regular k -D array of bins, each bin containing the points which lie inside it. This is equivalent to a hashing algorithm with a simple linear hash function to convert key coordinates to bin numbers. For a range look-up to work with a hash table, the choice of hash function is in any case limited, because the hash-function has to preserve the order of the elements. For the key-coordinates to bin mapping applied here, that requirement is fulfilled.

Normally in hash-tables the number of bins is larger then the expected number of keys to be stored in the table, such that the average number of keys per bin is less than one. In a standard approach, multiple keys in the same bin are often stored as a linked list. In the type of application here, having more bins than points would slow down a range-search as many empty bins inside the search volume would have to be examined as well. The find-in-range algorithm for a hash container requires the selection of keys from the bins that overlap the search range. For bins completely inside the range all keys can be taken immediately. For bins overlapping the border of the range, selecting keys is a linear search, but now for a much smaller number of entries $n_{\text{bin}} \approx n/N_{\text{bins}}$, with N_{bins} the number of bins. A search using a hash table will therefore be faster than a k -D tree search if $c_{\text{linear}} \times n_{\text{bin}} < c_{\text{tree}} \times \log_m n$, where c_{tree} and c_{linear} are the time constants for a tree search and a linear search respectively. The number of bins needed for hashing to be faster than a tree is then given by

$$N_{\text{bins}} > n \cdot \frac{c_{\text{linear}}}{c_{\text{tree}} \times \log_m n} . \quad (4.3)$$

In this calculation the overhead caused by many empty bins is ignored and the inequality of equation (4.3) is only an indication. Equation (4.3) grows almost linear with n . The constant c_{linear} is normally smallest when the keys to compare are stored sequentially in memory, which is not the case for a linked list. The keys in a bin can be stored sequentially in memory by ordering an array of keys by bin number. A hash-bin then points to a sub-range of the ordered array. However, inserting elements is now no longer a constant time operation, as for standard hash tables. The time taken for the sorting the keys array by bin number, amortized over n look-ups, is limited to $\lg n$. In the inequality of equation (4.3), one should also take into account that the requested search-range can span several bins due to overlap with the bin boundaries, even when the search-range is smaller than a bin-volume. Therefore, c_{linear} should be replaced by $c_{\text{hash}} = c_{\text{linear}} \times f_m$. The value of the multiplicity factor f_m can be anywhere from close to one, if the search volume is much smaller than the bin volume, to several times 3^k , if the search volume is much bigger.

In conclusion, if enough information about the input data and the search-ranges is known and the condition of equation (4.3) is fulfilled, this kind of k -D hash table can be faster than the k -D search tree. A comparison of the relative timing between the fixed-range binary-tree implementation and a hash table implementation is given in section 4.2.5. The k -D hash table is also known as a bucketting container and is used for example in many of the fast k -D Voronoi-tesselation algorithms [245].

4.2.4 Implementation

Both the normal and fixed-range tree have been implemented in C++. Because the types to store vary, the trees are designed as template classes. The classes follow the C++ Standard Template Library (STL) conventions [246,247] and are implemented as container adaptors on top of an STL `vector` class. Like the STL `map` class, the implementations differentiate between the objects to store, called elements or values, and the key to sort those objects with. Keys can have up to eight separate dimensions.³ The classes provide two interfaces to the user to access the data. One is the standard STL-vector interface for linear access using iterators and indexing. The other accesses the data as an ordered set in k dimensions using the keys and are used to find a given key or to look up all elements in a given k -D volume. Figure 4.5 gives a unified modelling language (UML) diagram of the classes and methods.

Like the STL containers, the k -space container classes have different behaviour but (almost) identical interfaces. Which type of container to use depends on the type of application. The k -D tree class is called `map` as it behaves identical to the STL `map` class. The times taken for both the insert and find operations on this class have an upper limit of $\log_m n$. Different from the 1-D STL map, which is normally implemented using a balanced red-black tree, the worst-case timing for these operations for the k -space map is order n . The `map_fixed_range` class guarantees an insert and find time proportional to $\log_m n$, but can only be used if the range of keys is known beforehand. Hash tables are not implemented in the standard template library. A simple hash container in k -D space has been implemented. If the range and the number of keys as well as the typical volume of a search range is known beforehand, then the hash container class can be faster than the map classes as explained in section 4.2.3.

In the STL ordered-container classes, the ordering operator is given as a template parameter. For the multi-dimensional containers, this ordering operator is replaced by a key-traits class. The methods of this class are used for all key operations. A default key-traits implementation is provided that works for simple key classes (identical coordinate types accessible via index operator). The k -space containers have been optimized for speed. This optimization implies that there is no checking of the input parameters or key values.

Map containers

The map classes have four template parameters: the type of the elements to store, represented by class `value_t`; the type to sort on, represented by `key_t`; the dimension of the space (which gives the number of used coordinates in `key_t`); and a key-traits class which lets the map compare and modify key objects, represented by `key_traits_t`. All operations the map performs on key objects are handled by static methods in `key_traits_t`. The map classes therefore require no knowledge of the coordinates of `key_t`. The only requirement on the `key_t` class is that its objects can be constructed with the copy-constructor (using C++ placement-`new`); no default constructor or assignment operator for `key_t` is required. The only requirement for the `value_t` class is that it must be storable in an STL `vector`.

³The maximum dimension of keys can easily be extended to more than 8.

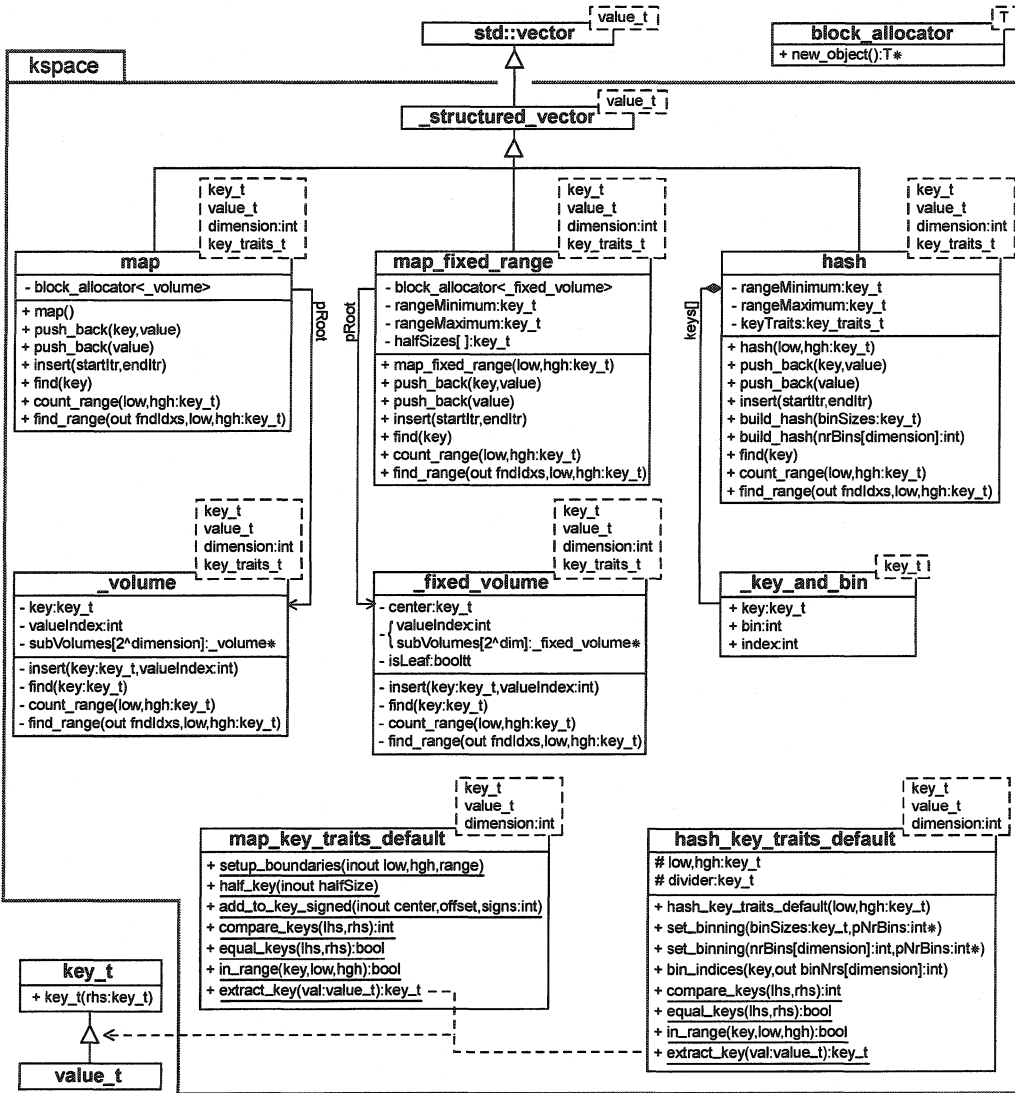


Figure 4.5: UML diagram of the *k*-space container classes and the node and helper classes. The `value_t` class represents the elements stored in the container, the `key_t` class represents the multi-dimensional points which are the keys on which the ordering is based.

The `_structured_vector` base class inherits all methods from the STL vector class from which it derives. However, values can only be inserted from one of the derived *k*-space container classes. Elements are inserted in the containers using the `insert` and `push_back` methods. They add a copy of the value to the underlying vector and update the tree's node structure for the associated key by calling the `insert` method on the root node. In general all methods using the underlying tree structure forward the call to the root node where the actual recursion over the tree's nodes is done. Often, the key is a sub-class (as indicated in Figure 4.5) or member of the `value_t` class. In this case, a

set of values can be inserted in the maps in one operation by extracting the key of each value using the `extract_key` method of `key_traits_t`. The `push_back` method follows the algorithm described in section 4.2.2 with one additional step. For the standard tree structure, at every node where the comparison of the new key and the node's key yields zero (equivalent to the new key is **not less-than**), the two keys are checked for equality. For the fixed-range tree, this equality check is done once and only if the key descends to an occupied leaf. If the keys are equal, the new element is either discarded or an exception is raised, depending on how the container object was configured.

The `find` method can be used to retrieve the value associated with the specified key. The `find_range` method retrieves all elements for which the keys lie within a volume specified by the inclusive lower boundary and the exclusive upper boundary. It returns an array with the indices of the elements within the search range. These indices can then be used to retrieve the values through the vector interface of the map.

Map node and leaf class

The private classes `_volume` and `_fixed_volume` represent both the nodes and leaves of the tree. The maps allocate blocks of these objects, using the `block_allocator` helper object. In the standard map, `_volume` contains the key, an index to the associated value in the map's underlying vector, and a set of 2^k child-node pointers stored in the `subVolumes[]` array. The `subVolumes[]` array is indexed using a bit-coded comparison of key coordinates. For the `map_fixed_range` class, the node and leaf are objects of class `_fixed_volume`. An internal node just holds a set of child-node pointers and the data member `center` contains the key value for the center of the volume spanned by the node. In a leaf node, indicated by the `isLeaf` data member, the `center` member is the key associated with the value and the corresponding index is stored in place of the child pointers. All methods in `_volume` and `_fixed_volume` call themselves recursively on all existing sub-volumes that contain part of the requested search key or range. For efficiency reasons, the `insert` method is actually implemented as a loop.

Key traits

The `key_traits_t` template argument contains a set of static methods required for the key operations of the map implementations. The `map` class uses only the key comparison methods: `compare_keys`, `equal_keys` and `in_range`. The `compare_keys` method returns an integer with the results of individual coordinate comparisons shifted into the corresponding return bits. The methods `equal_keys` and `in_range`, on the other hand, return a boolean which is the logical AND of all coordinate compares.

For the `map_fixed_range` class, one needs to calculate ranges and centers of the keys. The range given to a fixed-range map's constructor is first passed through the `key_traits::setup_boundaries` method. This makes it possible for a traits implementation to, for example, adjust integer-type boundaries to be a power of two or to set string boundaries to the first character of the strings. The map calculates centers of nodes by adding or subtracting half the parent-node's range from each coordinate. The half-ranges for each level are calculated using the `key_traits::half_key` method and are cached by the map. The center key for a sub-volume is obtained by adding the half-range of the current level to the parent's center correcting the sign for each coordinate. The signs are defined by the corresponding bit in the value returned by the `compare_keys` method.

A default implementation for the `key_traits_t` template member is provided by the `map_key_traits_default` class. This class is also set as the default template argument. The default implementation will work for `key_t` classes with simple numerical coordinates. It requires access to the `key_t` coordinates using the index operator `[]`. Most of its methods are implemented as template-meta-programs that iterate the operation over all used key coordinates [248–250]. The default implementation of `extract_key` is just a C++ `static_cast` of `value_t` to `key_t` which works if `value_t` is indeed derived from `key_t`. This method is only used when a key and its associated value are combined in a single object as in the `push_back(value_t)` and `insert` methods of the maps.

If the coordinates in the key class are not accessible, are not numeric (e.g. strings), or require more complicated operations for the calculations of the centers in a fixed-range map, the user must provide a specific implementation for the `key_traits_t` class. One can either derive from the default traits class and override the methods that need to be changed or implement all methods in the specific traits class.

Hash container

The `hash` class has the same template parameters as the map classes, but its default key-traits class is different. The `key_traits_t` class takes the role of the hash function in standard 1-D hash tables and is responsible for the mapping of keys to bin numbers. The `bin_indices` method in the `key_traits_t` class is the actual hash function that maps `key_t` coordinates to a k -D array of bin numbers. In order to use the `find_range` method, the hash function should preserve the order of the keys which requires an ordering of the key's coordinates. If only the `find` method is used, no restrictions are imposed on the hash function. The default implementation of the hash key-traits is `hash_key_traits_default`. The `bin_indices` method in this default divides the range of each key's coordinate in equal sized bins.

The `hash` class requires a slightly different setup before being used, because, in this implementation, the hash structure can only be built once the number of bins to use is specified. For this, the `build_hash` methods are used. These methods assign a bin number to all inserted keys using the `bin_indices` method of the `key_traits_t` class. The key and the bin number, associated with an element, are stored in an array of the helper class `_key_and_bin`. The k -D bin numbers are collapsed into a single value $b = \bar{b} \cdot \bar{d}$, where \bar{b} is the k -D vector with bin numbers given by the hash function to each coordinate and $d_j = \prod_{i=1}^j n_{i-1}$ with n_i the number of bins for each dimension and $n_0 = 1$. The array of `_key_and_bin` objects can then be sorted according to the bin value, such that all keys in the same hash bin are stored sequentially in memory. The `find_range` method in class `hash` uses a recursive template-meta program to loop over the k -D range of bins that overlap the search volume and in each bin performs a linear search of the keys to select the ones inside its search range.

4.2.5 Timing performance

Because the `map_fixed_range` class requires no tuning and has an guaranteed time behaviour, it is used as the default range look-up container for the tracking algorithm. In this section, its range look-up time is compared to a very simplistic linear-search algorithm. This simplistic algorithm is still useful, because the timing for an optimized hash-table can be calculated from the linear search time.

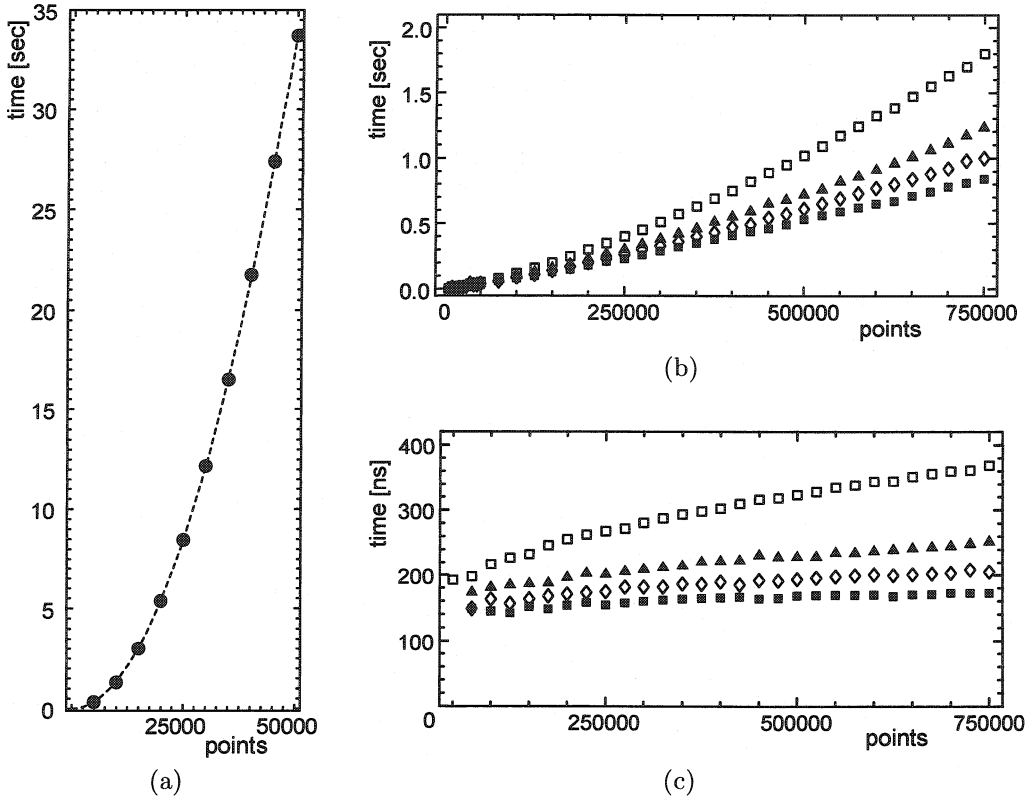


Figure 4.6: Time needed to find neighbours for all points versus the number of points: (a) using a simple linear search; the dashed curve indicates the expected n^2 time behaviour; (b) using a 3-D fixed-range map. The vertical line indicates the point range spanned by figure (a). Results in figure (c) are obtained from (b) after dividing the data by the expected $n \cdot \log_8 n$ behaviour. The dependence on the search volume is indicated by the different symbols. The length of the sides of the cubic search volume is: $\blacksquare = 1$, $\diamond = 2^4$, $\blacktriangle = 2^5$, and $\square = 2^6$, corresponds to a cube with sides of 2^{14} . For (a), the result is independent of the search volume.

The time it takes to traverse the fixed-range tree in a search of the map is proportional to $\log_m n$. In the link-building step of the tracking algorithm (see section 4.3.1), all points close to a given point need to be found. This operation is done for each point, which leads to an additional factor n (the number of points) in the total time needed for the link-building step. Figure 4.6b shows the timing results for a 3-D fixed_range map containing randomly distributed points on a 1.8 GHz CPU.⁴ In Figure 4.6a, the time needed when using a straightforward linear search algorithm is given. A fit to the data (dashed line) with $t = c_{\text{linear}} \cdot n^2$ yields $c_{\text{linear}} = 13.5$ ns. The search time for the map depends on the size of the search volume, indicated by different symbols. This can also be seen in Figure 4.6c which shows the same search time after dividing the data by the theoretical n -dependence of $n \cdot \log_8 n$. The tree's search-time coefficient c_{tree} depends slightly on n because a larger search volume (or equivalently a higher point density) leads to more

⁴all values in this paper are for a 32 bits AMD Athlon 2500+ CPU at 1.8 GHz.

points inside the search volume. In that case, the probability increases that the search volume spans multiple nodes in the tree, which all have to be searched. From Figure 4.6c one can determine that a single search in the tree scales as $t = c_{\text{tree}}(v) \cdot \log_8 n$, where c_{tree} depends slightly on the size of the search volume v and is about 170 ns for the smallest search volume.

The one-time overhead to build the underlying tree structure of the map should also be considered. To create the tree, each point in the map has to be inserted in a tree containing the already inserted points. For each insert, the tree has to be traversed to find the place where to insert the new point. The total time needed for n points is thus proportional to $\sum_{i=1}^n \log_m i \propto \ln(n!)$. In Figure 4.7, the tree building time is shown as a function of the number of points for a 3-D fixed-range map. A fit yields a building time of $t_{\text{build}} = 46.5 \ln(n!)$ ns. As $\ln(n!)$ is difficult to calculate for large n , an approximation using an upper limit for $\ln(n!)$ is $t_{\text{build}} = -26.6 + 46.9 \cdot 10^{-6} n \ln(n)$ ms, valid for the range of points in Figure 4.7.




Figure 4.7: Building time for a fixed-range map as function of the number of 3-D points to insert. The dashed line is a fit to the data yielding $t_{\text{build}} = 46.5 \ln(n!)$ ns.

From the results found above, one can calculate the required number of bins for a hashing container using equation (4.3). For typical volumes and search-ranges involved in the track finding for emulsion data of the CHORUS experiment, the search-range bin multiplicity $f_m \approx 3^3 \times 0.6$. Substituting the values found above for $n = 75000$ gives $n_{\text{bin}} \approx 4$ and $N_{\text{bins}} \approx 17900$. Using 2^{15} hash bins, the k -space hash container is a factor 1.21 (for search box size = 1) to 1.55 (box size = 64) faster than the fixed-range map used by default. The dependence on the size of the search range shows that tuning of the number of bins is important.

4.3 Track-finding algorithm

As described in the introduction, the track-finding algorithm has been designed for an environment with a large number of hits and for tracks which can only be considered as straight lines on the scale of a few hits. The hits of a track show therefore only local correlations. The large number of background hits rules out tracking algorithms that do complex calculations or examine a large fraction of all possible combinations of hits. In the following, a collection of hits that are part of a possible track is referred to as a segment. A fully-grown segment is called a track candidate. In the implementation of the algorithm, a custom-made class is required to make all decisions about accepting hits in a segment. This class, containing all the hit and track acceptance criteria, is known as the criterion.

4.3.1 Concept

Linked hits network

The algorithm first builds up the network of links. A criterion defines a cuboid volume, relative to the position of a hit, to be searched for other hits. The k -D map containing all hits is used to find the neighbouring hits in this volume for every hit. The map's `find_range` method dictates the use of a cuboid region. In an application, however, the link acceptance region for track hits is not necessarily rectangular. For example, the hit acceptance region based on the extension of a segment with constant uncertainties for both position and direction has the shape of a topped-cone. The algorithm therefore applies a criterion to select acceptable links from all the links formed by the base hit and the hits found in the acceptance volume around it. The implementation of this criterion allows the user to define any arbitrary acceptance volume. The track-finding algorithm is in general isotropic, but can be restricted according to the experimental conditions. Any restrictions that are applied when building the links network also limit the tracks that can be found. For example, an angular restriction in the link-acceptance region limits the solid angle of the tracks that can be found.

The hits and links correspond to the vertices and edges of a large graph. If only forward links are accepted, this connection graph is a directed acyclic graph. The connection graph links each hit to the other hits which may belong to the same track. As a result, the look-up of all possible hits that might be added to a segment is very efficient.

Segment growing

The graph of linked hits, built in the previous step, is searched with a modified depth-first algorithm [237] for paths compatible with tracks. All hits are taken as possible starting points for segments. A criterion is applied to select the hits that should be used as starting points. All links attached to a selected starting hit form the initial set of segments containing just two hits. Each segment is then expanded recursively by adding hits linked from the last hit in the segment. For this, a criterion decides which new hits are accepted. In an application, this criterion should accept hits that correspond to a topology consistent with its particular track model. The growing of a segment stops when none of the links from the last hit are accepted. The segment then becomes a track candidate. A segment splits into multiple new segments whenever there is more than one accepted hit. Each new segment is also followed until it stops. To do this, the algorithm backtracks to previous hits that have multiple accepted links and then follows these. The algorithm behaves therefore as a depth-first graph search, except that stepping to already visited vertices (hits) is not disabled.

Each track candidate, formed this way for a selected starting hit, traces a path through the links network. All the track candidates share the same starting hit, indicated by **S** in Figure 4.8. No hit is ever exclusive to a single segment or track candidate. In fact, many track candidates are identical, apart from a single hit. In Figure 4.8, the growing procedure can be imagined as moving from left to right through the links network, creating track candidates for every path that is compatible with a given track model.

The algorithm selects the best track candidate from all track candidates for each starting point. The result of the comparison between track candidates is decided by another criterion. Comparisons are only meaningful if similar entities are compared, like track candidates or segments spanning the same range of hits. Because the algorithm




Figure 4.8: Topological representation of the recursive segment growing tree. The black dots represent the hits. The links between hits are indicated by the numbered lines. The numbering restarts at each hit. The starting point of the track search is indicated by **S**. The thick line is the final track candidate. Dashed links indicate links that are rejected by the hit-acceptance criterion using the preceding segment. Decisions which branch to retain have to be taken in points **A**, **B**, **C**, and **S**.

behaves as a depth-first search, comparisons are only made between track candidates which are complete segments. In practice, selecting the best track candidate is done on the fly whenever there are several accepted links for a segment. A first track candidate is created by following either the first or the best accepted link at every step. Backtracking along these hits, that track candidate can be compared with others following the other branches. At each hit that has more than one accepted link, the other accepted links are grown as well. The track candidate with the current best branch can be compared with a track candidate taking a new branch. Only the best of the two is kept at every step. Effectively, this amounts to a branch decision at each hit. In Figure 4.8, this branch-pruning procedure can be imagined as keeping the track candidate with the best right-hand side from the decision point and moving right-to-left back to the starting point **S**. The example in Figure 4.8 corresponds to the case in which the track candidate with the most hits is considered the best track. If two (or more) track candidates contain the same number of hits, the one that corresponds best to a straight line is chosen.

For every starting point after segment growing, there is therefore one remaining track candidate. A final criterion is used to pass a judgement about its validity as a possible track. If accepted, the track candidate is stored in a list of found tracks. After processing all starting points, one is left with a list of possibly overlapping tracks. The track candidates from different starting points can share hits. A criterion is used to decide which of the overlapping track candidates to keep.

Limiting combinatorics

The link-following algorithm, as described above, considers all possible combinations of linked hits in the network and therefore always finds the best track candidate. A determination of the tracking time (section 4.3.3) shows that the algorithm scales approximately as $\sum_k^{k_{\max}} (c_l n)^k$, where k_{\max} is determined by the typical segment length and the volume fraction $c_l = v_l/V$ is the size of the link acceptance region v_l divided by the total volume V . The product $c_l n$ corresponds to the expectation value for the number of links per hit.

As long as this value is reasonably low for track-unrelated hits, the tracking time remains polynomial in the number of hits n . Unfortunately, the link-following algorithm can suffer from an inverse combinatorics problem. On a track, the number of possible segments could be very large, such that the tracking time becomes too long for all practical purposes. The problem is that the acceptance criteria for links and hits have to accommodate for hit-finding inefficiency. In the CHORUS experiment, the hit finding efficiency in each emulsion image is about 86%. The segment-growing algorithm must therefore be able to cross one or more layers with no hit on a track. The link-acceptance region must thus span several layers. In that case, hits belonging to the same track can have both direct and indirect links pointing to them. A hit that can be reached following two or more short links, can also be reached by a single link. Links numbered 1 and 2 which connect hit **C** to **A** in Figure 4.8 are an example. Due to these longer links, the basic algorithm follows the same set of links at the tail of a track very often. The segments created in these steps are usually identical apart for one hit. In case a hit on a track can be reached either directly or via one intermediate hit (the link acceptance region spans two layers) and assuming that a track has hits on all layers, a hit i is visited from its predecessor and from the hit before that. In this case, the number of visits $HV_2(i)$ to hit i on a track is given by:

$$HV_2(i) = HV_2(i-1) + HV_2(i-2) \quad . \quad (4.4)$$

The subscript 2 indicates that the next two layers can be reached via links. The first two hits are only visited once, independently of the number of layers links can span, so that

$$HV_2(1) = HV_2(2) = 1 \quad . \quad (4.5)$$

Equations (4.4) and (4.5) correspond to the definition of the Fibonacci series. The total number of hits visited, $THV_2(N)$, for a track with N hits is then given by:

$$THV_2(N) = \sum_{i=1}^N HV_2(i) = HV_2(N+2) - 1 \quad . \quad (4.6)$$

This sum can become quickly large. For a track with $N = 25$ hits, the total combinatorics of the algorithm is $THV_2(25) = 196,417$. The result will be even higher when links spanning two or more intermediate hits also exist. The sum in equation (4.4) then gets extended with more previous terms (known as the tribonacci and tetranacci series [251]). For links connecting the next four layers $THV_4(25) = 3,919,944$.

In a standard depth-first search of a graph [237], the above problem of large combinatorics doesn't occur because vertices (hits) are marked as visited. If a search path reaches a visited vertex it stops. The same strategy cannot be applied in a segment-growing procedure because the first visit to a hit does not necessarily correspond to the best path. A similar strategy can be used, however, if the correct path can be identified. One can then mark the hits on that path and disable future visits. One case in which the correct path is known is when a complete path is found and taken as a track candidate, but then it is too late as all the combinations have already been tried. However, because of the recursive nature of the segment-growing algorithm, this solution can be applied by marking the hits at the tail of a segment when the number of hits in the tail exceeds a given value t . The value t is chosen such that a tail containing more than t hits is probably the correct tail of the segment. The tail mark is only valid for a single starting point,

so it does not affect segments grown from other starting points. With this tail marking, the total number of hits visited is now limited to $THV'_2(N, t) = N - t + THV_2(t)$. In the example above with $N = 25$ and using $t = 7$ one obtains $THV'_2(25, 7) = 51$. With a link acceptance region spanning 4 layers one finds $THV'_4(25, 7) = 78$.

Unfortunately this marking strategy creates a dependence on the order in which the links are followed. The set of hits which will be included in a track is now affected by the link order. All possible paths at the tail of some segment are examined, but when the algorithm backtracks to earlier hits, continuations of other paths are blocked by the earlier segment that included sufficient hits. The optimization of the beginning of the segments is thus effectively disabled. To control which hits are included from the start of a segment, the links have to be followed in a certain order. There are two solutions to this problem. The first solution is called initial link ordering. When building the links network, the accepted links are sorted according to a value defined by a criterion. This value determines the order in which links are followed during segment growing. For example, it can favour short links such that the first segment built contains the largest number of hits. The second solution is called followed link ordering. In this solution, the segments are sorted each time a hit is added according to how good the new hit fits in the segment. If there is more than one accepted link, the best new segment is grown first. This solution yields better results as the most likely segment continuation is tried first. However, the first solution sorts the links only once and should therefore be faster. A simulation showed that the time gained in following the best link first compensates for the time spent in sorting the accepted links (section 4.4).

The tail marking solves the combinatorial problem in the link-following algorithm. A related problem exists with the starting points. Each hit is tried as starting point for a segment and therefore the same track is found again for all hits on a track. To avoid this, the hits in an already found and accepted track candidate can be marked. Hits marked in this way can then be skipped. Again, the restriction creates an ordering dependence, now on the order in which starting points are processed. Therefore, a criterion is used to sort the hits a priori to determine the starting order.

4.3.2 Implementation

The implementation is organised as a source-code library containing the tracking toolkit. The user must provide the hit, track, and criterion classes for a specific application. Specific optimizations for each environment are dealt with by the corresponding criterion implementation. The user must also put code in its application program that feeds the hits to the track-finding procedure and converts the found track candidates to tracks. A simple Monte-Carlo simulation program, used to assess the efficiency of the tracking algorithm, can act as a template for such a main program. This program is described in section 4.4.

Two header files specify which input hit class, coordinate class, and criterion class to use. The criterion type specification has as default the abstract base class `VCriterion`. By replacing this default with a concrete class, the overhead of calling virtual functions can be eliminated. Using the abstract `VCriterion` instead, a criterion can be selected at run-time. A test showed that the time overhead for calling abstract methods is less than one percent (see Figure 4.11b). Another header file defines a set of compile-time flags. These flags specify which of the link ordering options, described before, to use. A class diagram of the implementation is shown in Figure 4.9.

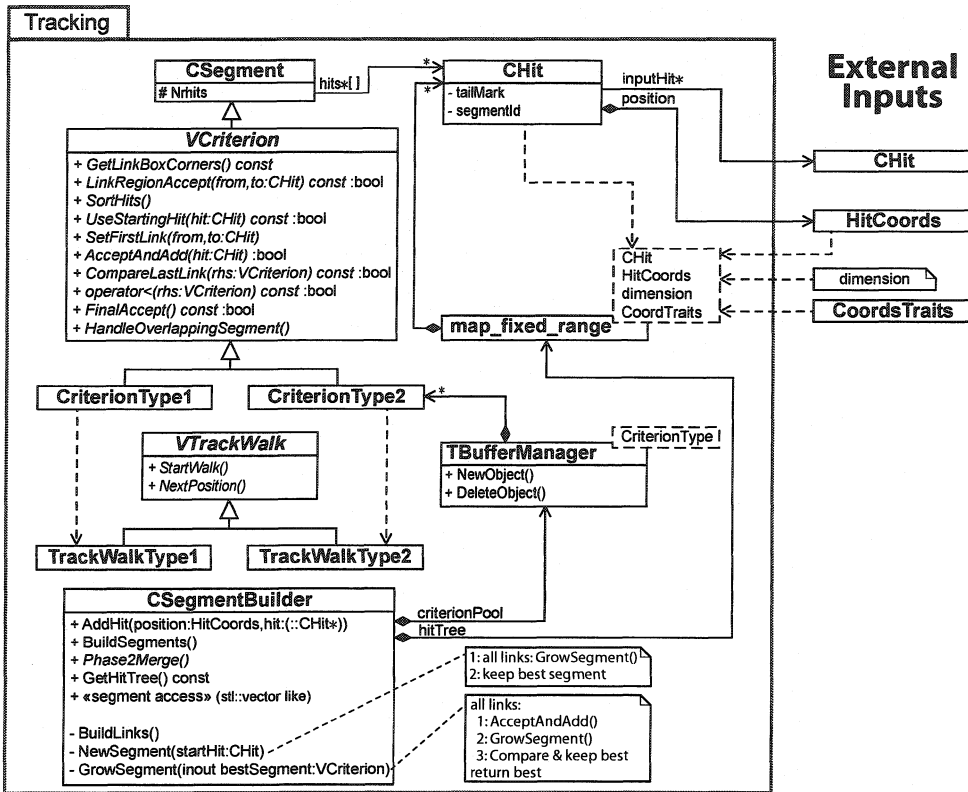


Figure 4.9: Class diagram (UML) of the classes involved in the tracking. All classes are in the 'Tracking' namespace. At compile time the user must provide the dimension of the hits, the input hit class (CHit), the coordinates class (HitCoords), and the traits for working with objects of this last class (CoordTraits). The last two classes correspond to the template parameters key_t and key_traits_t, respectively, of the k-D fixed-range map (section 4.2.4). The HitCoords class corresponds to the position of a hit and it is a data member of the internal hit class (Tracking::CHit). An array of these internal hits forms a segment and is the type of object included in a criterion object to make acceptance decisions. The user must supply concrete implementations of the criterion class, indicated by CriterionType1 and CriterionType2 and their corresponding track-walk implementations TrackWalkType1 and TrackWalkType2.

The tracking classes use an internal hit class which points to a user-defined hit object. The user-defined hits are not directly used by the tracking code; the pointer can be used in the main program to locate hits assigned to tracks. The input to the tracking code consists only of the hit coordinates represented by class HitCoords. In the tracking code, only the map's key_traits_t operates on this type of objects which is represented by the CoordTraits class (see section 4.2.4). The coordinate class can be replaced, at compile time, by any coordinate class chosen by the criterion implementer, as long as a corresponding CoordTraits and the dimension is defined as well. For most simple coordinate classes, the map's default traits class can probably be used.

The class `VCriterion` encapsulates the segment class `CSegment`, which holds a fixed-size array of pointers to hits to avoid time-consuming heap allocations. Removing an automatic heap allocation in the inner segment-growing code in a previous version, reduced the tracking time by about a factor three. The assignment operator in `CSegment` is optimized to only copy the used part of the array. All criteria, mentioned in section 4.3.1, appear as abstract methods in `VCriterion`.

The actual segment building and bookkeeping of track candidates is performed by an object of class `CSegmentBuilder`. The structure of its recursive `GrowSegment` method is shown in pseudo-code in Figure 4.10. The actual code contains several optimizations to avoid loop overhead, to undo tail marking and to reuse discarded segment objects. The method starts with a pointer to the current segment. The procedure will replace this pointer with the new segment it has built from the given one.

The lines 6-12 in Figure 4.10 iterate over all links of the last hit in the current segment. The criterion is asked to add the destination hit of each link to a copy of the input segment in line 10. If this hit is accepted, the extended segment is stored in a buffer in line 11. The followed-link-ordering solution to the tail recursion problem (section 4.3.1) changes the type of this buffer from a FIFO to a priority queue. The priority queue is implemented using a binary heap. If no hit was accepted, the segment has reached its end and the method returns in line 14. If there are accepted links, the current segment is replaced with the top of the buffer contained stored segments in lines 18 and 19. If there is more than one accepted link, the first segment is recursively grown in line 22. If the tail from the current hit to the end of this now complete segment contains more hits than the tail-mark limit, the tail is marked. The iteration in lines 25-33 pops the other stored segments one-by-one from the buffer and grows each one using a recursive call to `GrowSegment`. If one of these fully-grown segments is better than the current one, they are swapped in lines 29 to 32, which also update any tail marking. The worse segment is immediately discarded at line 33. At the end of the iteration, any tail mark is removed in line 35 before the method returns. The infinite loop in line 3 and the if-statement in line 21 handle the case that only one link is accepted. In that case a recursive call can be avoided and the current segment is simply replaced by the one accepted and the loop repeats.

The `GrowSegment` method creates and discards many criterion objects. Using the C++ `new` and `delete` calls for individually allocating these objects is not efficient. Instead, the template class `TBufferManager` is used which allocates a pool of objects. Individual objects can quickly be taken from and returned to this pool. The pool works as a LIFO which has the additional advantage of improving data cache efficiency.

Found tracks are represented by criterion objects that pass the criterion's method `FinalAccept`. They are kept by the segment building class. In the end, the list of criterion objects retained are the tracks found by the procedure. The main program, which knows the exact type of the criterion object, can access this list and process the tracks further. Standard operations would include fitting a track model through the hits and adding possible missed hits close to the fitted track. For this, the map containing all hits can be used and it is therefore accessible via `CSegmentBuilder`. An implementation of such code is available as the `Phase2Merge` method in `CSegmentBuilder`. This method uses the abstract interface class `VTrackWalk` to walk along the track contained in a criterion object. The criterions instantiate the corresponding concrete implementation of the track-walk interface. The `Phase2Merge` implementation looks for hits within a certain

```

1 void GrowSegment(VCriterion*& pCurSegment)
2   VCriterion* pNewSegment = 0
3   loop: // Loop if only one link is accepted => no recursive call needed!
4     // — STORE NEW CRITIONS FOR ACCEPTED LINKS —
5     AcceptedLinkBuffer acceptedLinks // Type is link-ordering dependent
6     iterate: pCurSegment->lastHit->links =>
7       if: pNewSegment == 0 =>
8         pNewSegment = criterionPool.NewObject()
9         *pNewSegment = *pCurSegment
10      if: pNewSegment->AcceptAndAdd(link->toHit) =>
11        acceptedLinks.push(pNewSegment)
12        pNewSegment = 0
13      nrAccepted = acceptedLinks.size()
14      if: nrAccepted == 0 => done. // NO LINK ACCEPTED => DONE
15
16     // — START ITERATION ACCEPTED LINKS WITH FIRST ONE —
17     currentLength = pCurSegment->nrHits // to calculate tail length new segments
18     criterionPool.DeleteObject(pCurSegment)
19     pCurSegment = acceptedLinks.pop()
20     // — MORE ACCEPTED LINKS => LOOP AND KEEP ONLY BEST —
21     if: nrAccepted > 1 =>
22       GrowSegment(pCurSegment) // RECURSIVE CALL
23       if: pCurSegment->nrHits - curLength > tailMarkLength =>
24         MarkSegmentTail(pCurSegment)
25       while: acceptedLinks.not_empty() =>
26         pNewSegment = acceptedLinks.pop()
27         GrowSegment(pNewSegment) // RECURSIVE CALL
28         if: *pNewSegment > *pCurSegment =>
29           UndoSegmentTailMark(pCurSegment)
30           swap(pNewSegment, pCurSegment)
31           if: pCurSegment->nrHits - curLength > tailMarkLength =>
32             MarkSegmentTail(pCurSegment)
33           criterionPool.DeleteObject(pNewSegment) // Delete discarded segment
34         // — UNDO TAIL MARK BEFORE RETURNING —
35         UndoSegmentTailMark(pCurSegment)
36       done.
37     else: =>
38       // continue grow segment infinite loop for 1 accepted link

```

Figure 4.10: Pseudo-code of the recursive segment growing code. The indentation marks blocks of code to be executed in a loop or if-statement.

distance of the (extension of) path of the track candidate. These hits are then added to the track candidate. The `Phase2Merge` method also checks if several track candidates share the same hits. The implementation of the criterion's `HandleOverlappingSegments` method decides if these track candidates should be merged. Running the merge step after looking for hits in the extensions of track candidates reduces the number of split tracks. Tracks can split into several found track candidates due to gaps of several missing hits on a track path.

4.3.3 Tracking time

In this section a determination of the running time of the tracking algorithm is given. The time needed to do track finding depends of course on the complexity of the criterion implementation and on the tracking environment. In this section, the cylindrical acceptance criterion implementation, described in section 4.4.2, is used. The results have been obtained using the Monte-Carlo simulation program described in section 4.4. The simulation includes noise and realistic track hit distributions. The setup generates 80 tracks in a volume of $512 \mu\text{m} \times 512 \mu\text{m} \times 100 \mu\text{m}$ with 25 evenly-spaced z -layers. The default tracking and criterion parameters include a link-span maximum of three layers, a tail-mark limit of $t = 7$, and both link-ordering options enabled (section 4.3.1). When followed link ordering is enabled, the initial link ordering option is used to determine the order in which seeds are created from each starting hit. The segment builder was configured to call the cylindrical-acceptance criterion's methods directly. Naturally, the results are processor dependent.

Coordinate type

In Figure 4.11, the tracking time versus the number of hits is given as function of the coordinate type. It can be seen from Figure 4.11a that building the fixed-range map and the links network is slightly faster when using a 16 bits integer as the coordinate type. This is mainly due to the four times larger size of a 64 bits floating-point number. The segment growing time is shown in Figure 4.11b. The segment growing is fastest when using floating point types for the coordinates. This is probably due to scheduling of instructions in the floating-point unit, which leaves the processor free to execute other instructions.




Figure 4.11: Tracking time as a function of the number of hits for different types of hit coordinates. Figure (a) gives the total time needed to sort the hits, build the fixed-range map, and build the links network. Figure (b) shows the segment-search time. The basic coordinate type for the different symbols is: ● = double (64 bits); ○ = double and virtual criterion methods; □ = float (32 bits); and ▲ = integer (16 bits).

Link ordering and tail marking length

As described in section 4.3.1, ordering the links helps in finding the best segments and is necessary when tail marking is enabled. Figure 4.12 shows how the tracking time depends on the tail recursion limiting parameter and the link ordering options. These parameters influence also the efficiency and one has therefore to optimize. The optimization depends on the particular environment where the tracking is applied. A simple test showed that using a priority queue to sort the accepted links takes on average about 160 ns instead of 60 ns for a non-sorting ring buffer. This time is more than compensated for by finding better segments first, as can be seen from Figure 4.12. The time dependence can be fitted with a constant term and a function of the tail-marking length t . A good fit is obtained with the function $c_0 + c_1 t + c_2 F(t)$, where $F(t)$ is the t^{th} Fibonacci number.




Figure 4.12: Segment search time for 50,000 hits as function of the recursion-limiting tail-marking parameter. Upper curve, \diamond symbol: links are ordered after being accepted for segment extension. Middle curve, \bullet symbol, links are sorted at build time. Lower curve, \triangle symbol: links are sorted at build time and after being accepted for segment extension. Note that each point in this figure has a different tracking efficiency. The dashed lines and legend in the figure are fits (see text) as function of the tail-marking parameter t .

Final segment length

In the following, a derivation of the running time of the segment-growing algorithm is presented. Reasonable estimates can be made in the case of uniformly distributed background hits and a criterion which uses fixed-size hit-acceptance volumes v_ℓ . In this case the expected number of accepted links per hit can be derived. The probability that a hit falls in the link acceptance volume v_ℓ is then given by v_ℓ/V , with V the total volume. The probability $P(k)$ that a volume v_ℓ contains k hits if n hits are distributed in V is given by the binomial distribution:

$$P(k) = \binom{n}{k} \left(\frac{v_\ell}{V}\right)^k \left(1 - \frac{v_\ell}{V}\right)^{n-k} \quad (4.7)$$

The last factor in equation (4.7) can be written as $e^{n \cdot \ln(1 - v_\ell/V)}$. For $v_\ell/V \ll 1$, the logarithm can be approximated to first order as $-v_\ell/V$. With the definition of the

volume fraction $c_\ell = -v_\ell/V$, the last factor in equation (4.7) is approximately equal to $e^{-c_\ell n}$. As n/V is just the density ρ , $c_\ell n$ is just the average number of hits in the volume v_ℓ . Using some approximations valid for large n , equation (4.7) turns into the Poisson distribution with mean $c_\ell n$:

$$\begin{aligned} P(k; c_\ell n) &= (c_\ell n)^k e^{-c_\ell n}/k! \\ \Rightarrow P(0) &= e^{-c_\ell n} \\ P(> 0) &= 1 - e^{-c_\ell n} \\ \langle k \rangle &= c_\ell n \end{aligned} \tag{4.8}$$

The parameters c_ℓ and n in the definition of P are dropped in the following equations unless they are different.

The track-finding algorithm uses all hits which have links as starting points for growing segments. If the criterion rejects a fixed fraction $1 - c_0$ as starting points then the number of starting points n_0 is given, according to equation (4.8), by

$$n_0 = c_0 n (1 - e^{-c_\ell n}) \tag{4.9}$$

The procedure yields one track candidate for each starting point by keeping the best of all segments grown from that starting point. The number of track candidates of a certain length $s_k(n)$ as function of n is shown in Figure 4.13. The sum over all k for some fixed n satisfies equation (4.9) with $c_0 = 0.72$ which is the expected value for this setup where 7 out of 25 layers not used. For increasing k , the curve $s_k(n)$ peaks at higher n as the probability to get k hits in a segment first increases with n before the probability to include the $(k + 1)^{\text{th}}$ hit becomes dominant.




Figure 4.13: Number of track candidates of a given size versus the number of background hits. Each size peaks at a different number of hits as the probability to include k hits in a segment grows with n until the probability to include a $k + 1^{\text{th}}$ hit becomes higher. The curve for segments with more than 7 hits (tail-mark length) is scaled by a factor $1/4$ (scale on right-hand side).

The probability to get a segment with two hits is given by the probability to have at least one link from the starting point and no links from any of the second hits. If the volume fraction for accepting a third hit is c_3 and there are k second hits, then the probability that none of these k hits accepts a third hit is given by $P(0; c_3 n)^k$ provided

that volumes do not overlap. The number of segments which include exactly two hits is then:

$$\begin{aligned} P(1 \rightarrow 2) &= \sum_{k=1}^{\infty} P(k) k e^{-k c_3 n} \\ &= c_\ell n e^{-(c_3 + c_\ell - c_\ell e^{-c_3 n})n} \end{aligned} \quad (4.10)$$

The volumes are very likely to overlap and therefore not independent. In practice, the formula $s_k(n) = c_0 n^k e^{-n(c_1 - c_2 \exp[-c_3 n])}$, which has the same form as equation (4.10) but groups coefficients, fits the curves of Figure 4.13 reasonably well.

Segment search timing

For a uniform background, the number of links per hit should simply be given by $\rho v_\ell = c_\ell n$, which is consistent with the expectation value of the Poisson distribution. If the time spent in the criterion's acceptance calculation is dominant, then the time for tracking n hits can be estimated as follows. For the n starting hits, the algorithm creates segment seeds with a fraction a_2 of the $c_\ell n$ linked hits. The fraction a_2 is less than one if marked hits (segment growing is actually a depth-first search) or hits already on other tracks are excluded. For each of these seeds of size two, there are $c_\ell n$ new links to check of which a fraction a_3 yields segments of size three. In total there are then $a_2 a_3 \times c_\ell^2 n^3$ segments of size three. For each of these segments there are again $c_\ell n$ linked hits, and so on. The number of acceptance calculations is then given by:

$$\begin{aligned} N_{\text{accept}} &= \sum_{k=2}^{k_{\text{max}}} c_\ell^{-1} \left(\prod_{j=2}^k a_j \right) (c_\ell n)^k \\ &= \sum_{k=2}^{k_{\text{max}}} (c_k n)^k \end{aligned} \quad (4.11)$$

In the second equation, all coefficients for the term with n^k are absorbed in a single constant c_k . The upper limit k_{max} of the sum is determined by the point where the number of accepted hits decreases rapidly. From equation (4.11) one can deduce that if $c_k n > 1$ for all k , the algorithm's running time and the number of segments will diverge, because for all segments there will always be more than one hit to extend it with. In practice this boundary condition is higher due to the effect of tail marking.

Figure 4.14 shows the segment-search time for n background hits with no tracks. The tracking time increases rapidly for $n > 10^5$, which is also the point where the number of segments with more than 7 hits increases rapidly, as can be seen in Figure 4.13. Equation (4.11) suggests that the tracking-time can be fitted by a polynomial of order k_{max} . A 5th degree polynomial seems indeed to fit the tracking time reasonably well. The residuals show that the fit is less good above 125,000 hits where the tail-marking effect becomes important and where exclusion of already assigned track hits lowers the number of starting points.

4.4 Tracking efficiency for simulated data

In order to estimate the performance of the tracking algorithm, a simple Monte-Carlo simulation program has been used. In this, track and background hits are generated that can be used by the tracking code. The simulation is a simple model of the real CHORUS emulsion data and does not include any physical models of particle propagation and hit formation in emulsion, emulsion distortion, or emulsion scanning. The track and noise




Figure 4.14: Tracking time as function of number of background hits. The dashed line is the result of a fit to the points with the function $\sum_{k=1}^{k=5} (c_k n)^k$:

$$\text{time}(n) = (n/826)^1 - (n/2581)^2 + (n/5731)^3 - (n/11511)^4 + (n/20675)^5.$$

hit-generators use parameters determined from measured tracks and scanned emulsion. Without distortion, tracks are perfectly straight and easily found. A simple distortion model was therefore introduced to test the tracking in real conditions.

The program uses the hit generators described in the next section. It uses an instance of `CSegmentBuilder` and (one of) its associated criteria (see section 4.4.2) to do the tracking. The actual analysis is done by one or more analysis classes that can be plugged into the simulation program. The analysis classes use a general framework for creating, filling and saving histograms or n-tuples to disk. The histograms or n-tuples can be displayed or further analyzed using Mathematica [252]. The results of two analyses are presented in section 4.4.4. The first analysis yields statistics of tracking time and the number of fake and found tracks. In the second analysis, tracking and hit-assignment efficiencies are quantified. Other analysis plug-ins were used to histogram the differences between the track parameters of the Monte-Carlo and the reconstructed tracks or to write all tracks to n-tuples for more detailed studies.

4.4.1 Hit generators

The generators use an occupation grid for generating hits. Each bit in the grid represents a small volume in 3-D space. The grid is used to avoid generating multiple hits with similar coordinates and takes into account that hits cannot overlap in the emulsion. The hits, like in real data, can still be spaced arbitrarily close, because the position of a hit inside a grid cell is random. The occupation-grid's cell size reflects the scanning microscope's resolving power for neighbouring grains (xy). The spacing in depth (z) is set either to a fixed layer spacing or to the microscope's depth of field, depending on whether the generators are setup for fixed layers or for real 3-D hit positions. When hits are generated for fixed z -layers, the generated 3-D hit positions are projected onto the layer. For hits on a track, the slope of the track is thus not taken into account, in accordance with the grains seen by the microscope. Any grain detected in the depth of field of the microscope is assigned a z -coordinate at the center of that layer.

Background generator

The background generator is rather straightforward. It generates uniformly distributed random positions for a requested number of hits. The hits are marked as background to distinguish them from hits belonging to tracks. Track-correlated background hits (like delta-rays) can be important in the real data. However, this is not simulated.

Track generators

The track generator is chosen at run-time. The simplest track generator creates straight tracks. An extended version, used in the results presented in section 4.4.4, distorts the tracks using a simple distortion model explained below. The basic track generator is set up with the following parameters: fixed layers or free 3-D coordinates; a variable number of tracks originating from a common vertex; the range in z where tracks can start and end; a minimum track length in case a track leaves the volume; and a distribution for the track slope (θ) with respect to the z -axis. Another set of parameters determines the distribution of hits along a track: the hit residual, defined as the sigma of a normal distribution modelling the perpendicular distance to the track; the probability of having a hit on a layer in the case of discrete layers, or the parameters of a Poisson distribution for free hit positions. In the CHORUS experiment, the number of hits per unit track length for a fixed number of layers is accurately described by a binomial distribution. This binomial distribution is the cumulative effect of, among others: intrinsic emulsion sensitivity, blind spots in the emulsion, microscope depth of field, layer spacing, image filtering, and grain detection. The use of a constant hit probability per layer is justified because it also leads to a binomial distribution for a fixed number of layers. When using free hit positions, the response of a uniform emulsion to a traversing particle is modelled. The model uses a fixed probability per unit length of having a developed emulsion grain. This also leads to a Poisson distribution for the number of grains on a fixed length of track.

For the distorted-track generator, only the most common form of distortion is modelled. One end of the tracks (at $z = 0$) is shifted by a random distance δ from its original starting point. The shift $\delta(z)$ is a quadratic function of z , such that the original track's direction is preserved at $z = 0$, $d\delta(z)/dz = 0|_{z=0}$, and the track's position at the other end is fixed, $\delta(z_{\max}) = 0$.

In the results presented here, no vertices were generated. Hits close to the vertex can usually be assigned to multiple tracks, an effect which requires a study by itself, but is not relevant to the track-finding efficiency discussed here. The other settings were chosen to reflect the typical CHORUS emulsion scanning values: 25 fixed layers of $4 \mu\text{m}$ thickness; tracks enter from the bottom at $z = 0$ and exit via the top or sides of the volume; minimum track length for detection is 14 layers; and a $\pm 300 \text{ mrad}$ wide uniform distribution of track slopes. The hit residual is set to $0.38 \mu\text{m}$ and the probability of detecting a hit on a layer is set to 75 %, lower than that of the real data which has 86 % hit efficiency per layer. The goal of this study is to estimate the track-finding efficiency for recognizable tracks, therefore tracks which have not enough hits to pass the final acceptance criteria of the criterion class (see section 4.4.2) are discarded. Distortion parameters were set according to a normal distribution with a mean shift of $0 \mu\text{m}$ and a sigma of $5 \mu\text{m}$ for all tracks in a single set. Each track has an additional random shift with a sigma of $0.63 \mu\text{m}$, 1.7 times the assumed track residual.

The track hits are merged with the generated background. In order to compare reconstructed tracks on a hit-by-hit basis with the generated tracks, the track to which each hit belongs is recorded. In the analysis of reconstructed tracks, the Monte-Carlo track with the largest number of hits in a reconstructed track is considered to be the matching Monte-Carlo track. A reconstructed track with only background hits is considered a fake track.

4.4.2 Acceptance criteria

As discussed in section 4.3, the criterion implementation has to make all decisions whether to accept hits in a segment or not. As it is called very often, its execution must be fast. Using vectors for the track's parameters, one can avoid time-consuming calculations involving trigonometric functions. The points \bar{p} on a straight line, parametrized by a single parameter λ , can be represented by:

$$\bar{p}(\lambda) = \bar{b} + \lambda \bar{r} \quad , \quad (4.12)$$

with a base vector \bar{b} and a direction vector \bar{r} . In this model there are two free parameters, \bar{b} can be any point on the line and the length of \bar{r} is not fixed. Using vector products (dot and cross-products), all calculations for distances and angles consist of fast multiplications and additions as shown in Figure 4.15. The division is optimized out as a multiplication with $1/\bar{r} \cdot \bar{r}$. This vector model has as advantage that it is applicable in any dimension and that it is isotropic (no preferred direction).




Figure 4.15: Straight track representation using vectors in 3-D, but applicable in any dimension. The vector equation for calculating \bar{d} from the straight-line parameters, \bar{b} and \bar{r} , and the position of the hit, \bar{q} , are shown in the figure.

As discussed in section 4.1.3, tracks can be considered straight only for short sections of their path. To get the local direction for the prediction, the criterion uses only the last k of n hits in a segment. The criterion uses two running averages:

$$\begin{aligned} \bar{S}_{\text{low}} &= \sum_{i=n-k+1}^{n-k/2} \bar{p}_i \quad , \quad \text{and} \\ \bar{S}_{\text{up}} &= \sum_{i=n-k/2+1}^n \bar{p}_i \quad , \end{aligned} \quad (4.13)$$

with \bar{p}_i the position of the i^{th} hit. The direction is estimated using the difference of the upper and lower half of the end of the segment: $\bar{r} = \bar{S}_{\text{up}} - \bar{S}_{\text{low}}$. The base vector is calculated from the average position of the last k hits in the segment: $\bar{b} = k^{-1} (\bar{S}_{\text{low}} + \bar{S}_{\text{up}})$. The running averages can be updated quickly for each hit added to the segment by simply subtracting the hit that has been dropped and adding the new hit. If a segment has

fewer than k hits the sums in equation (4.13) run over $\lfloor n/2 \rfloor$. If n is odd, the middle hit is not used for the direction estimate but is added to determine the average position \bar{b} .

In the design of a hit-acceptance criterion there are typically two quantities that determine how well a hit matches a prediction obtained from previous hits, namely the transverse distance to the prediction ($|\bar{d}|$) and the accuracy of the prediction. Two criteria were implemented. The first uses a cylindrical acceptance region for new hits using only the transverse distance. The second takes the uncertainty of the direction into account by defining an acceptance region with the shape of a topped-cone. The cut on $|\bar{d}|$ is then a linear function of the extrapolation length of the prediction $\lambda = |\bar{r}'| / |\bar{r}|$ (see Figure 4.15).

If all hits have the same position uncertainty, a prediction obtained with fewer hits has a larger uncertainty. Assuming the direction vector of the tracks runs mainly in the positive z -direction, the track slope with respect to the z -axis is given by:

$$\begin{aligned} \tan \theta_z &= \Delta \bar{x}y / \Delta z \\ &= \left(\sum_{up} \bar{x}y_i - \sum_{low} \bar{x}y_i \right) / \left(\sum_{up} z_i - \sum_{low} z_i \right) , \end{aligned} \quad (4.14)$$

where the sums run over $n/2$ hits. If the uncertainty in the xy -coordinate σ_{xy} is independent of the uncertainty for the z -coordinate σ_z , propagation of errors in equation (4.14) yields:

$$\sigma(\tan \theta_z) = \frac{2}{\sqrt{n} \Delta z} \sqrt{\sigma_{xy}^2 + \sigma_z^2 \tan^2 \theta_z} . \quad (4.15)$$

Assuming that the hits are evenly spaced by the fixed layers, $z_i = i \cdot \Delta l$ with Δl the layer spacing, the sums for Δz in equation (4.14) give $\Delta z = n \cdot \Delta l / 2$. This leads to the following dependence of $\sigma(\tan \theta_z)$ on n :

$$\sigma(\tan \theta_z) = \frac{4}{n \sqrt{n}} \frac{\sigma_{xy}}{\Delta l} \sqrt{1 + \frac{\tan^2 \theta_z}{12} \left(\frac{\Delta l}{\sigma_{xy}} \right)^2} , \quad (4.16)$$

using $\sigma_z = \Delta l / \sqrt{12}$ as the uncertainty for hits distributed uniformly within the thickness of a layer. The uncertainty for the prediction's direction decreases therefore as $n^{-3/2}$. Both criteria take this effect into account by letting the cut applied to $|\bar{d}|$ vary with n as long as $n < k$.

The criteria are tuned by histogramming $|\bar{d}|$ as function of λ and n for segments created directly from the Monte-Carlo tracks. The cuts are determined by requiring a certain fraction of these distributions to be contained. In the results given in section 4.4.4, a containment of 95% was used. The number of hits k used to generate the predictions was set to 6. The value k must be chosen such that the direction of the last piece of the segment is determined sufficiently accurate with respect to the position resolution of the hits, provided that other effects (distortion, magnetic field) do not affect the resolution. The cone-opening angles thus found are as expected from equation (4.16).

The track-walk algorithm used in both criteria first calculates track positions on all internal layers between its first and last hit, before extrapolating. Any additional hits compatible with the track are used to recalculate track parameters for the next step. The implementation of the `HandleOverlappingSegments` method in the criterion is based on a comparison of the track parameters. Only if two track candidates have similar position

and direction are they compared on a hit by hit basis. If the two track candidates share most of their hits, only the best of the two is kept. Two track candidates are merged if they share a significant part of their hits.

4.4.3 Efficiency criteria

The main objective of the simulation is the determination of the tracking efficiency: how many of the generated tracks are reconstructed. The efficiency depends on the parameters of the track-finding algorithm. These parameters also influence the tracking time. The simulation is used to choose the trade-off between efficiency and tracking time. The reconstruction efficiency ε is given by the ratio of the number of found tracks and generated tracks. The Gaussian approximation of the binomial distribution for large N is used to determine the uncertainty on the calculated efficiency from $\sigma_\varepsilon = \sqrt{\varepsilon(1-\varepsilon)/N}$, with N the number of generated tracks. For small N and when ε approaches its limits 0 and 1, the Gaussian approximation is no longer valid.

Other quantities can be studied in the simulation, for example: reconstructed versus real track parameters, number of background hits assigned accidentally to real tracks, hit residuals for reconstructed tracks, etc. The results presented here include the hit-to-track assignment efficiency and the number of background hits accidentally included. The hit-to-track assignment efficiency is given by the ratio of the generated track hits assigned to a reconstructed track and the total number of hits in the simulated track.

4.4.4 Results and discussion

The tracking code has been applied to simulated track and background hits in 27 different configurations for the two criteria implementations. Each configuration has been run 125 times with different sets of 80 tracks and a different background of 50,000 hits every five runs. As already discussed in section 4.3.1, the combination of initial link ordering and followed link ordering gives the best results and is also the fastest. The largest effects of disabling one or both of the link-ordering options are to lower the hit-to-track assignment efficiency and to increase tracking time.

Table 4.1 gives the results with both options enabled as function of the tail-marking length and the number of layers that links can span. The ‘found tracks/run’ entries list the average and RMS of the number of found tracks that are genuine. It can be higher than the number of generated tracks because a single track can be found several times, for example when it gets split in several pieces. The average multiplicity for split tracks (so excluding the ones found only once) is given in the row marked ‘multiplicity’. The third part of the table lists the results as they are after running the tracking algorithm on the input data. The phase-II hit-look-up code (section 4.3.2), contains a simple algorithm to merge overlapping or split tracks. The results after application of this step are reported in the fourth part at the bottom of the table.

The reconstruction efficiency is derived from the total number of tracks not found from the 10,000 generated tracks. From the table it can be seen that to increase the track-reconstruction efficiency, one has to increase the number of layers that links can span. For three layers, the reconstruction efficiency is already close to 100%. The tracks that are not found are mainly due to a combination of several gaps with missing hits and a background hit in or near the gap that pulls the segment away from the actual track. From the table it can be seen that the more layers that links can span, the higher the

average multiplicity becomes. At first glance, the opposite is expected as the tracking can bridge longer gaps with no hits. The reason for the higher multiplicity is that overlapping tracks are allowed by the algorithm. If links span more layers, more small segments are found that consist of background hits. Some of these background segments are then connected to a real track and acquire enough hits for final acceptance as a track candidate. With increasing size of the link-acceptance region, the total number of links and therefore the tracking time increases. The number of track hits included in a track candidate ('hit-assignment efficiency') also increases with a bigger acceptance region. However, the amount of background hits within the link-acceptance region increases proportional to the region's size, which leads to more fake tracks. Both for three and four acceptance layers, the tracking efficiency is about 99.5%. Except for the tracking time due to the number of links, there is only a small 2% difference in hit-assignment efficiency between three and four acceptance layers. The lower hit-assignment efficiency is because track segments are shorter and therefore less well determined when a smaller number of layers can be spanned by links.

The tail-marking length does not affect the reconstruction or hit-assignment efficiency very much, but does have a large effect on the tracking time. A tail-marking length $t = 7$ is 25% slower than $t = 3$ but yields slightly better determined track parameters (not shown), while $t = 11$ increases tracking time by a factor of 2 with almost no gain. The reason for the small effect of this parameter is the link-ordering applied during segment growing. If followed link ordering is disabled, this parameter becomes more important.

The last part of the table shows the efficiencies after picking up additional hits close to the reconstructed tracks and merging overlapping track candidates. Adding close hits improves the hit-assignment efficiency to almost 100%, at the cost of a bit more included background. The description at the end of section 4.4.2 explains how the criterion decides to remove or merge overlapping track candidates. The merging of track candidates leads, however, to a small ($\approx 0.2\%$) loss in the number of found tracks. A closer investigation showed that the few real tracks that are lost actually get merged with another close-by real track.

Another set of simulation runs showed that the reconstruction efficiency remains at almost 100% for all background levels between 50,000 and 250,000 hits (highest value tried). The number of fake tracks, though, increases from about 13 per run at 50,000 background hits to 26,500 per run at 250,000 background hits. Therefore, one should either restrict the hit acceptance criteria or remove the fake tracks by examining the χ^2 of the hit residuals for higher background levels.

Table 4.2 reports the results for the cone-shaped acceptance criterion. Due to a square root in its acceptance calculations it turns out to be a factor 1.6 slower for the largest number of links (4 accepted layers). Because its hit-acceptance region is more restricted, this criterion creates about a factor 3 fewer fake tracks and also includes slightly fewer background hits in the tracks. The cone-shaped acceptance criterion rejects hits on the tails of the hit-residual distribution ($|\bar{d}| > 3\sigma_{\text{hit}}$), whereas the cylindrical acceptance region is largely oversized ($5.5\sigma_{\text{hit}}$ for short prediction extensions) and accepts these outliers. As a result, the efficiency using the cone-shaped acceptance criterion is slightly smaller (0.2%). The hit-assignment efficiency is reduced by 2% in the results from the segment-growing procedure, but 0.2% higher after close hit look-up. This increase is due to a better determination of the track parameters because there are fewer background hits included.

Cylindrical acceptance criterion									
Tracking configuration values									
Accepted layers		2		3		7		11	
Tail-marking limit		3		3		7		11	
Tracking input parameters and input processing time									
Number of runs	125								
Background hits/run	50000								
Tracks/run	80								
Hit probability/layer	75%								
Track hits/run	1490 ± 19								
hit sorting time [ms]	32								
Tree build time [ms]	15								
Link building time [ms]	75 ± 5		27, 876 ± 184		134 ± 5		73, 981 ± 349		213 ± 5
Number of links									152, 233 ± 630
Segment searching results									
Segment search time [ms]	17 ± 5	19 ± 3	33 ± 6	92 ± 4	115 ± 6	245 ± 26	349 ± 8	440 ± 11	792 ± 59
Found tracks/run	87.8 ± 3.9	88.2 ± 3.9	87.7 ± 3.8	101.5 ± 6.5	101.6 ± 6.1	98.3 ± 5.7	108.0 ± 7.8	108.1 ± 8.7	101.2 ± 7.6
Fake tracks/run	1.8 ± 1.4	1.9 ± 1.3	1.9 ± 1.3	11.3 ± 2.9	13.6 ± 3.8	13.6 ± 3.7	11.2 ± 3.4	14.8 ± 4.1	14.9 ± 4.0
Multiplicity found tracks	2.1 ± 0.3	2.1 ± 0.3	2.1 ± 0.3	2.4 ± 0.8	2.5 ± 0.8	2.4 ± 0.7	2.7 ± 1.0	2.8 ± 1.1	2.7 ± 1.0
Background hits/track	0.86 ± 0.98	0.85 ± 0.98	0.83 ± 0.97	0.94 ± 1.02	0.92 ± 1.03	0.91 ± 1.01	0.91 ± 0.98	0.88 ± 0.98	0.86 ± 0.97
Hit-assignment efficiency [%]	78.32 ± 0.10	78.71 ± 0.10	78.96 ± 0.10	90.20 ± 0.07	90.56 ± 0.07	90.81 ± 0.07	92.45 ± 0.06	92.79 ± 0.06	93.06 ± 0.06
Total tracks not found	431	415	413	50	47	49	48	47	50
Reconstruction efficiency [%]	95.69 ± 0.20	95.85 ± 0.20	95.87 ± 0.20	99.50 ± 0.07	99.53 ± 0.07	99.51 ± 0.07	99.52 ± 0.07	99.53 ± 0.07	99.50 ± 0.07
After phase 2 hit pickup and track candidate merge									
Phase 2 time [ms]	7 ± 4	7 ± 4	7 ± 5	8 ± 4	9 ± 3	8 ± 4	8 ± 4	9 ± 3	9 ± 3
Found tracks/run	76.3 ± 2.0	76.5 ± 1.9	76.5 ± 1.9	79.9 ± 1.2	80.1 ± 1.2	80.0 ± 1.1	80.5 ± 1.2	80.8 ± 1.5	80.4 ± 1.3
Background hits/track	3.56 ± 1.88	3.56 ± 1.88	3.56 ± 1.88	3.55 ± 1.90	3.55 ± 1.89	3.54 ± 1.89	3.49 ± 1.88	3.49 ± 1.88	3.49 ± 1.88
Hit-assignment efficiency [%]	98.10 ± 0.03	98.15 ± 0.03	98.18 ± 0.03	98.82 ± 0.03	98.85 ± 0.02	98.85 ± 0.02	99.17 ± 0.02	99.20 ± 0.02	99.25 ± 0.02
Total tracks not found	469	446	449	81	72	71	69	67	70
Reconstruction efficiency [%]	95.31 ± 0.21	95.54 ± 0.21	95.51 ± 0.21	99.19 ± 0.09	99.28 ± 0.08	99.29 ± 0.08	99.31 ± 0.08	99.33 ± 0.08	99.30 ± 0.08

Table 4.1. Tracking time and efficiency as function of the tracking parameters indicated in the first two rows. The second part of the table gives the tracking input parameters and setup times. In the third part, the track candidates are counted as they are found by the segment builder. The fourth part shows the efficiencies for the case when hits are recovered and (identical) tracks merged. The errors, except for efficiencies, correspond to the RMS values of the underlying distributions. For efficiencies the error corresponds to the statistical uncertainty on its value (see text). An explanation of the different rows is given in the main text.

Cone-shaped acceptance criterion												
Tracking configuration values												
Accepted layers	2		3		4							
Tail-marking limit	3	7	11	3	7	11	3	7	11	3	7	11
Tracking input parameters and input processing time												
Number of runs	125											
Background hits/run	50000											
Tracks/run	80											
Hit probability/layer	75%											
Track hits/run	1490 ± 19											
hit sorting time [ms]	32											
Tree build time [ms]	16											
Link building time [ms]	75 ± 5		135 ± 5		214 ± 5							
Number of links	27,876 ± 184		73,981 ± 349		152,283 ± 630							
Segment searching results												
Segment search time [ms]	15 ± 5	18 ± 4	28 ± 5	99 ± 4	124 ± 6	262 ± 24	581 ± 12	748 ± 17	1301 ± 79			
Round tracks/run	84.7 ± 3.6	84.4 ± 3.7	84.1 ± 3.6	96.1 ± 5.4	94.5 ± 4.9	92.1 ± 4.3	99.8 ± 6.5	98.4 ± 6.0	92.8 ± 5.2			
Fake tracks/run	0.5 ± 0.8	0.5 ± 0.9	0.5 ± 0.9	4.6 ± 1.6	5.2 ± 2.1	5.3 ± 2.0	4.3 ± 1.8	5.8 ± 2.1	5.9 ± 2.2			
Multiplicity found tracks	2.1 ± 0.2	2.1 ± 0.2	2.1 ± 0.3	2.4 ± 0.7	2.5 ± 0.7	2.4 ± 0.6	2.7 ± 1.0	2.8 ± 1.0	2.7 ± 0.9			
Background hits/track	0.62 ± 0.81	0.62 ± 0.81	0.61 ± 0.81	0.74 ± 0.89	0.73 ± 0.89	0.71 ± 0.88	0.71 ± 0.87	0.70 ± 0.86	0.67 ± 0.84			
Hit-assignment efficiency [%]	75.61 ± 0.10	76.18 ± 0.10	76.48 ± 0.10	87.87 ± 0.08	88.33 ± 0.07	88.70 ± 0.07	90.14 ± 0.07	90.58 ± 0.07	90.99 ± 0.07			
Total tracks not found	586											
Reconstruction efficiency [%]	94.14 ± 0.23	94.45 ± 0.23	94.39 ± 0.23	99.26 ± 0.09	99.28 ± 0.08	99.24 ± 0.09	99.41 ± 0.08	99.27 ± 0.09	99.28 ± 0.08			
After phase 2 hit pickup and track candidate merge												
Phase 2 time [ms]	7 ± 5	6 ± 5	7 ± 5	7 ± 4	8 ± 4	7 ± 4	8 ± 4	8 ± 4	8 ± 4			
Round tracks/run	75.0 ± 2.1	75.2 ± 2.1	75.2 ± 2.1	79.4 ± 1.0	79.5 ± 0.9	79.4 ± 0.9	79.8 ± 1.1	79.9 ± 1.2	79.7 ± 1.0			
Background hits/track	3.47 ± 1.86	3.47 ± 1.86	3.47 ± 1.86	3.48 ± 1.88	3.48 ± 1.87	3.48 ± 1.87	3.44 ± 1.86	3.45 ± 1.86	3.44 ± 1.86			
Hit-assignment efficiency [%]	98.51 ± 0.03	98.54 ± 0.03	98.56 ± 0.03	98.99 ± 0.02	99.09 ± 0.02	99.11 ± 0.02	99.35 ± 0.02	99.40 ± 0.02	99.42 ± 0.02			
Total tracks not found	631											
Reconstruction efficiency [%]	93.69 ± 0.24	94.02 ± 0.24	94.01 ± 0.24	98.95 ± 0.10	99.01 ± 0.10	99.02 ± 0.10	99.16 ± 0.09	99.05 ± 0.10	99.10 ± 0.09			

Table 4.2: As table 4.1, but for a cone-shaped acceptance region.

4.5 The track trigger

The tracking algorithm and its implementation is by design completely isotropic and was developed to find all tracks in a set of emulsion images. As was discussed in the introduction (section 4.1), during scan-back of the target sheets only a small volume of about $50\ \mu\text{m} \times 50\ \mu\text{m} \times 100\ \mu\text{m}$ of emulsion on the upstream surface of a plate is scanned. The only question concerning this volume is whether a track with the predicted slope is present or not. The tracking algorithm, described in this chapter, does not use this additional information. On the contrary, the track selector hardware, described in section 2.9.4, is exquisitely tuned for this job. A similar approach as the track selector applies in hardware can also be implemented in software. The slope information of the predicted track is then used to speed up the track-finding algorithm. The basic idea is that the direction of the track prediction defines positions for acceptable grains on subsequent layers. A similar summing algorithm, as applied in hardware to the pixels of the images by the track selector, is now applied to the grains. The sum serves as a track trigger. The trigger defines a small region of interest in which the full tracking algorithm is then applied. The tracking step verifies the existence of a track, assigns grains to it, and calculates the track parameters.

4.5.1 Concept

The steps taken in the track-trigger algorithm are depicted in Figure 4.16. All grains are given an initial sum value of one. Starting with the first layer, the algorithm iterates over all grains in a layer. It looks up all grains in the next layer within a certain acceptance region. The position of this region is given by the grain position in the previous layer shifted by the direction of the track prediction. The angular acceptance determines the dimensions of the acceptance region. The sum values of all found grains are set to the sum value of the grain in the previous layer incremented by one. However, this is only done if that value is higher than the current value of a found grain. Such a case happens at area E in Figure 4.16. If there are no grains in the acceptance region, a virtual grain is created at the predicted position without incrementing the sum value. This is necessary because a track does not need to have a grain on each layer. To maintain the angular acceptance, the acceptance area in the next layer for these virtual grains is enlarged (see area D in Figure 4.16). The whole procedure is then repeated for the next layers. Throughout this procedure, the information on the position of grains in preceding layers is discarded. As a result, small curvatures due to distortion can be accommodated.

After summing all layers this way, the grains on the last layer with a sum value above threshold indicate the position of a possible track candidate. The tracking algorithm is then applied on the grains inside a rhombohedral volume. The position of the grain in the last layer that triggered the volume is the center of the top rectangle. The lengths of the sides are given by the angular acceptance. The slope of the volume is given by the predicted-track slope. Because this volume contains only a small number of grains, typically 30 to 100 of which about 20 belong to the track, the tracking procedure runs very fast.

4.5.2 Implementation

The implementation of the track-trigger algorithm is straightforward. The grains for each layer are stored in 2-D fixed-range maps. Iterating over all grains in a layer, the look-up of grains in the acceptance region in the next layer is done using these maps. Virtual grains are stored in a temporary linear array. The procedure is then repeated




Figure 4.16: Illustration of the grain-summing procedure for the track trigger. For each grain, the predicted-track slope and the angular acceptance define a region in the next layer. The sum value (indicated next to the grains) for the one (A) or more (B) grains in that region is incremented. If no grains are present in the predicted region (C), a virtual grain (open circles) is created at the predicted position. To keep the same angular acceptance, the acceptance region for virtual grains is increased (D). If a grain falls in multiple acceptance regions, the highest value is used for the new sum value (E).

for the next layer. Processing of each layer starts with the virtual grains first. This way, the same array can be reused for storing any new virtual grains, because each virtual grain can create at maximum one new one. If a grain's sum value is too low to pass the trigger threshold taking into account the remaining number of layers, it is immediately discarded.

After processing the one but last layer, the sum values of the virtual grains for the last layer and of all grains on the last layer are compared against a threshold. If the sum value is above threshold, an object is created to hold all the grains in the rhombohedral acceptance volume and the parameters of a possible track. The maps for each layer are used to find the grains inside this volume and insert them in object. Finally, the track-finding algorithm is run over each object. If a track is found which matches the prediction, the track parameters are included in the object, otherwise the object is deleted.

4.6 Application in CHORUS emulsion scanning

The tracking code described in this chapter has been applied to the emulsion data taken by the CHORUS experiment at CERN. The tracking code could have been applied directly online, but in most cases it is not sufficiently fast on the microscope control systems. Typically, a tomographic image set of 25 images of $300\ \mu\text{m} \times 300\ \mu\text{m}$ is acquired every two seconds. The tracking time for such an image set on the online computer⁵ is somewhere between several seconds to half a minute, depending on the grain density. Applying the track-finding algorithm online would therefore introduce considerable overhead and would leave the microscope idle. Two solutions have been adopted to eliminate this overhead.

First, for scanning areas that consists of one or more full microscope views, the scanning system writes the grain data to an object-oriented database. In the emulsion

⁵Dual-processor 500 MHz Pentium-III.

scanning at CERN, the grain data is stored for all scan-back on the interface sheets, for all net-scan, and for full vertex-reconstruction data. These grain data are then processed offline, while the data taking of the emulsion continues online. For processing the grain data offline, a main program was developed. This program reads the grain data from the database, applies the tracking algorithm, and writes the found tracks to the same database. The track data consist of the track parameters and the grains in the small volume around it. The parameters used by the tracking algorithm are also stored in the database. In normal operation, many instances of this program run on a computer farm, consisting of standard PCs running the Linux operating system. Each instance processes a different event and the distribution of the events is controlled by an external program. This program communicates with the tracking programs via the network using a message dispatching server [253].

The second solution is to increase the track-finding speed by using the angular information available in scan-back. The track trigger (section 4.5) is used for scan-back in the target sheets, where an area of only $60\ \mu\text{m} \times 60\ \mu\text{m}$ is scanned. The track-finding efficiency during scan-back of the target sheets is also less important, because two consecutive missed plates are required for scan-back to stop for a prediction (see section 2.10.2). Therefore, in the tuning of the track-trigger procedure, efficiency can be traded in for speed. Typical tracking time using the track trigger on target sheet data is about 0.5 seconds. The tracks found online during target-sheet scan-back are also written to the database. The grain data around the track is stored as well.

The tracking results stored in the database are available for further analysis. Because much of the grain data is stored in the database, many options are available for analysis. One can, for example, refit tracks, rerun track finding, or look for hot pixels in the data and eliminate fake tracks based on the grain positions. A typical analysis task needed for CHORUS is the matching of predictions to found tracks and the generation of new scanning predictions for the upstream emulsion plate. Because of the new approach to the scanning of the interface sheets, a more efficient candidate matching procedure could be applied which is described in section 4.6.2. The matching involves determination of the precise alignment for which a new strategy has been developed. This local-alignment strategy is discussed in section 4.6.3.

4.6.1 Tracking configuration

The main difficulty in applying the tracking code to real emulsion data lies in the large variability of the actual data. Emulsion differs from plate to plate and even within one plate. Effects like shrinkage and distortion, grain size, and fog density can change from one microscope view to another (section 2.9.2). The physical grain size and the fog density depend on the development process. Some of these variations arise from the scanning conditions, like transparency of the emulsion layer or optics tuning. Shrinkage and distortions are unavoidably different for each plate and depend also on the relative water content of the plate when scanned. Part of this variability is already dealt with while scanning, for example by changing the illumination or exposure time, and part, like contrast and grain size, is dealt with by the grain recognition algorithm (section 4.1.1). As the input to the tracking code has large variations in the number of grains per view, it is difficult to predict the tracking time required, which excluded running the tracking online. The variation in distortion and shrinkage implies that the hit-acceptance criteria have to be insensitive to these parameters.

To deal with these variations, the criterion class implemented for emulsion tracking uses simple but fast acceptance calculations, relatively large acceptance criteria, and a small number of consecutive hits for making the prediction. The acceptance criterion has been designed based on real emulsion data. The criterion implementation is identical to the cylindrical criterion described in section 4.4.2. To accommodate the distortions, only the last six grains in a segment are used to define the track prediction. A track segment of six consecutive grains has a typical length of about $20\ \mu\text{m}$ and is therefore unlikely to be affected much by distortions. The acceptance radii for additional grains depend on the number of grains already in the segment and are set as follows: $1.6\ \mu\text{m}$ for a 2 grain segment, $1.4\ \mu\text{m}$ for 3 grains and $1.2\ \mu\text{m}$ for 4 or more grains. Although this last value is large with respect to the track residual (5σ), the grain density in the emulsion is sufficiently low that not too many background grains are picked up.

The other settings of the criterion and tracking code were determined by the characteristics of the events in CHORUS. Event-related tracks lie in a forward cone (relative to the direction of the incident beam) with $400\ \text{mrad}$ half-opening angle. As the layer-to-layer uncertainty of the grain positions corresponds to about $100\ \text{mrad}$, the link acceptance region is therefore limited to a forward cone with $500\ \text{mrad}$ half-opening angle. The grain-detection efficiency per layer has an average value of about 86% . In order to have high tracking efficiency, the link-acceptance region has been set to span four layers. The tail-marking length parameter was set to 7 hits. The track-walk code used in the close-hit look-up procedure, picks up all grains in a radius of $1.5\ \mu\text{m}$ around the track candidate, but only grains closer than $0.6\ \mu\text{m}$ are used in a new track fit. Track candidates which pass a final selection criterion based on the number of grains on the track and the grain density are then handed off to the application using the tracking code.

4.6.2 Prediction matching on the interface sheets

During scan-back, the candidates for a prediction are selected using a χ^2 -probability cut, $P(\chi^2) > p$, based on four quantities: Δx , Δy , $\Delta\theta_X$, $\Delta\theta_Y$. The exact cut value p is chosen depending on the amount of background present. A procedure that derives the real efficiency from the number of multiple candidates shows a clear plateau below some value for p . A value slightly below the onset of this plateau is then used for the cut value p . The covariance matrix needed to calculate $P(\chi^2)$ is derived from the data. In a first approximation, the covariance matrix is taken to be diagonal and defined by the sigmas extracted from Gaussian fits to the distributions of the residuals. In an iterated procedure, the covariance matrix is used to select candidates and then recalculated from the correlations between predictions and candidates.

Originally in scan-back on the changeable and special sheet, the matching required four segments to be found that are matched with the base measurement of both emulsion sheets, indicated by the check marks in Figure 2.20. The base measurement is required for an accurate slope measurement to reduce pick-up of accidental background (section 2.4.2). With this selection, the efficiency for matching a track to the prediction includes a factor

$$\varepsilon_f^4 \varepsilon_{\text{base}}^2, \quad (4.17)$$

where ε_f is the probability of finding a track segment on one side of the emulsion and $\varepsilon_{\text{base}}$ the efficiency of connecting two segments over the plastic base. Inefficiencies lowering the value of ε_f are very costly due to the power of four of ε_f .

In the scanning approach used at CERN, the changeable sheet and special sheet are scanned independently. As the tracking algorithm does not look for the scan-back prediction in particular, many tracks are reconstructed on both sides of the plastic base. These additional tracks make some extra options for matching the found tracks to the prediction available. With sufficient tracks found on both emulsion layers and matched over the base, the distortion and shrinkage of the emulsion layers can be corrected using the base-slope measurements. This reduces the uncertainty in the measured direction of track segments found in each emulsion layer to about 4 mrad. These single-sided track segments can then be matched independently to a prediction. At CERN, this has been used to match candidates requiring only **three out of four** segments. A prediction has to match at least one track with a base measurement on either the changeable or special sheet. In addition, that matching track, extended to the other plate, must have either a matching base track or one matching segment on one emulsion side. The efficiency for matching predictions to track segments in the changeable and special sheets then contains the factor:

$$\varepsilon_f^4 \varepsilon_{\text{base}}^2 + 4\varepsilon_f^3 (1 - \varepsilon_f) \varepsilon_{\text{base}} \quad . \quad (4.18)$$

Assuming $\varepsilon_{\text{base}}$ to be 1 and setting conservatively $\varepsilon = 0.95$ for the CERN system, the value of equation (4.18) is 98.6%. This should be compared to equation (4.17) which yields 92.2% for the quoted track-selector efficiency of $\varepsilon_f = 0.98$.

There is however a complication in the distortion and shrinkage correction procedure. The distortion varies from one small piece of emulsion to the next. The typical scale over which the distortion changes is about a centimeter, but sometimes it can change rapidly within less than 100 μm . If such an abrupt change happens within the area scanned for a prediction, a global distortion correction for the whole area cannot be applied. To overcome this problem, a local distortion correction is applied instead. First the shrinkage factor, which changes over much larger scales than distortion, is determined from all tracks in the scanned area. The distortion for every segment is then determined from about 20 nearby tracks with a base measurement. The base slope of these tracks is used to estimate the local distortion at the position of the segment under consideration. A 2-D map is used to look up the nearby tracks. This map contains the positions of all tracks with a base measurement.

4.6.3 Local-alignment procedure

A similar problem with local variations was also found to exist in the alignment of the interface sheets. Although the interface sheets are quite rigid, a global alignment, calculated from minimizing the distances between all found predictions and their candidates, showed residuals which depend on the position and slope. Further investigation showed that these variations were probably due to deformation of the interface sheets at boundaries of the honeycomb support structure on which they were mounted.

These kind of local deformations were corrected for by applying a local-alignment correction. This is done efficiently using again a 2-D map. The map contains a set of track prediction and candidate pairs, indexed by position. A local-alignment correction can then be applied to any track prediction. This correction is calculated from a number of prediction-candidate pairs (usually about 50, depending on the track density) in the vicinity of the prediction being corrected. Because the alignment correction is local (typically several square centimeters), it takes into account the remaining deformations

of the sheets. This local-alignment procedure improves the overall resolution and reduces the variation in the residuals. By reducing position and slope dependent residuals, it also reduces the unknown bias in the candidate selection based on a χ^2 probability cut (section 4.6.2) from the position and slope dependence of the covariance matrix.



4.6.4 Tracking results

Some of the results obtained with the tracking code applied to emulsion data are shown in the following. In Figure 4.17 the slope distributions of all tracks found on a special sheet are shown. The peak from beam muons at $\theta_x = 42$ mrad is clearly visible. The typical track density on the special sheet is more than a factor two lower than that for the target sheets as they are exposed during only seven months before development. The target sheets are exposed for two times this period. In addition, background is accumulated during the five months that they are stored underground between two runs. From this figure, the necessary accuracy for making predictions can be deduced (section 2.4.2).


The result of running the track reconstruction on a set of images containing a neutrino-interaction vertex is shown in Figure 4.18. The emulsion slice containing the interaction vertex is displayed in Figure 4.18a and shows the typical star of black tracks due to the break-up of the nucleus. The tracks of interest are almost perpendicular to this picture and are only visible as single dots. The tracks reconstructed in the full image set are overlaid in the image. The interaction vertex can be distinguished from the four tracks that point to the same location. Some unrelated tracks can be seen as well. Close to the vertex, the four vertex tracks are not reconstructed, because they disappear in the black region caused by the nuclear break-up. A 3-D representation of the tracks and grains attributed to these tracks close to the vertex is given in Figure 4.18b. One track is only reconstructed further away from the vertex and not visible in this close-up of the vertex.

Figure 4.17: [top opposite page] Track density as function of track direction in a single special sheet of one of the corner modules of the emulsion stack. Data is based on 1.42 million tracks in 6156 scanned events with a total scanned area of 130 cm^2 . Three peaks of particle directions are clearly visible. The highest peak is from beam-related muons at $\theta_x \approx 40$ mrad, $\theta_y \approx 0$ mrad. The two other peaks with $\theta_x \approx -100$ mrad are due to background from beams in the adjacent experimental hall.

Figure 4.18: [bottom opposite page] A neutrino interaction vertex in the emulsion. (a) The emulsion layer containing the vertex is shown inside the emulsion player program which allows to view a movie of the tomographic image set. The black tracks are the result of the nuclear breakup. Overlaid on the image are all reconstructed tracks. The four coloured tracks come from the vertex, pink tracks are event unrelated tracks. (b) Shows a close-up of the vertex in 3-D, where one track is not visible anymore as it is only reconstructed further away from the vertex.



(a)



(b)

4.7 Conclusion and discussion

The multi-dimensional container classes have proved useful in other applications. For instance, to determine the initial alignment of target sheets, one has to connect tracks from track maps on one plate to the other plate (see section 2.10.1). With several thousand tracks per track map, the look-up of a possible match becomes time consuming, because trying to match each track with all the other tracks has quadratic complexity. Using a 4-D map, the look-up for matching tracks can be done fast. Furthermore, the look-up automatically matches both position and slope of the track within a limited 4-D volume, as the map is indexed by both the position and slope of the tracks. The slope matching reduces the number of fake matches. The maximum number of matches as function of the position offset determines the alignment. Trying out a grid of position offsets yields an initial estimate for the alignment. The set of matched tracks with this offset are then used to fit a refined alignment which also includes possible small rotations and a longitudinal shift.

If one examines the track-finding algorithm, it is clear that the principles on which the track finding is based require the following two conditions:

1. hits that belong together have close-range relationships,
2. a set of hits defines a volume that is defined by the close-range relationships of condition 1.

If examined in a more abstract way, condition 1 is actually not a strict requirement of the algorithm. The track-finding algorithm, described in section 4.3.1, can be used in any situation where a volume can be defined in which other hits of a track can lie with respect to some track hit. As long as that volume is small with respect to the total volume and the volume is continuous in its coordinates, the k -space containers can be used to build up a hit-connection graph. The notion of hits and tracks can also be replaced by other abstractions. The algorithm and its implementation as a C++ toolkit can be used in any environment where nodes in a graph need to be connected according to some acceptance and rejection criteria. As long as the above conditions are fulfilled, the tracking framework can be applied for such cases in any dimensional space.

Separating the acceptance criteria, which involve the actual hit and track model, from the track-finding algorithm has allowed to create a flexible toolkit. As discussed in section 4.3.2, all environment-specific characteristics — like geometry, track propagation, and magnetic fields — are handled by the criterion class implementation which the user has to provide. The toolkit can therefore easily be adapted to other applications. For example, the code has been used to combine five-dimensional track segments (3-D position and 2-D slope) from multiple emulsion sheets into single tracks.

For many track-finding applications, more specific algorithms, for example Kalman filtering [254, 255]), could be more efficient and faster. In general, tracking algorithms need to look up hits in regions of space and the k -space containers can be used to do this efficiently. The track-finding algorithm, described in this chapter, performs very well when applied to the CHORUS emulsion data. It has been shown that the tracking finds almost all tracks despite the high background. Its application has enabled improvements in the vertex-location efficiency in scan-back. If applied to data taken over the full thickness of the emulsion, it can also reconstruct all tracks and vertices in an event. This could be used to reduce the need for human confirmation of secondary vertices after the net-scan procedure. The algorithm has also been applied successfully to the reconstruction of tracks in the HARP time-projection chamber.

Chapter 5

Measurement of D^{*+} production

In this chapter, a result published by the CHORUS collaboration [256] is presented. The large amount of emulsion exposed to the high-intensity ν_μ beam at CERN makes it possible to study charm-meson production and decay with relatively high statistics and excellent resolution. About 100,000 charged-current ν_μ interactions have been found in the emulsion. Among these events, charmed-meson decays have been identified from the presence of a secondary vertex. Part of these charmed mesons are produced as a D^{*+} resonance, which immediately decays strongly to a pion and a D meson. In the CHORUS experiment, one can identify a D^0 meson from its decay in the emulsion and charged pions can be identified and measured in the hadron spectrometer. Hence, a measurement of the D^{*+} production cross-section can be made.

The measurement relies on an accurate determination of the charge and momentum of pions. Therefore, the upgrade of the hadron spectrometer with the honeycomb detector, described in Chapter 3, played an important role in this measurement. The techniques to find and reconstruct events in the emulsion have been discussed in Chapter 2. The exact layout of the emulsion stacks can also be found in that chapter. Possible upgrades to the event location and reconstruction techniques, although not used for this measurement, have been given in Chapter 4.

In this chapter, the phenomenology of charmed-meson production in neutrino interactions is given first. The experimental procedure to extract charmed-meson decays in the emulsion is described next, followed by the selection of events due to D^{*+} decays. After a discussion of the background and selection efficiency derived from a Monte-Carlo simulation, the number of D^{*+} events is given. From this number, the ratio of charmed vector meson (D^*) to scalar meson production (D^+, D^0) and the D^{*+} production cross-section with respect to the total charged-current ν_μ cross-section are derived.

5.1 Introduction

The measurement of the D^{*+} production cross-section and the comparison with the total D^0 production cross-section give some insight into the charm-production mechanism. In particular, the ratio of charmed vector meson and scalar meson production in deep-inelastic scattering can be obtained. Several experimental groups have performed a study of D^{*+} production in neutrino charged-current interactions [257–259]. In these experiments the identification of charmed particles relied mainly on the reconstruction of the invariant mass of the assumed decay products.

In emulsion experiments, charm production can be observed without the need to reconstruct the invariant mass. The tracking in emulsion has enough spatial resolution to clearly separate the charm-decay vertex from the primary neutrino-interaction vertex. In hybrid experiments, combining the emulsion technique with electronic detectors, the high spatial resolution at the neutrino and decay vertices can be combined with kinematical measurements of the outgoing particles. This technique has been applied in a neutrino beam by the E531 experiment at FNAL. However, the number of events accumulated in the 25 kg of emulsion was limited.

Recent improvements in automatic emulsion scanning systems made it possible to measure several orders of magnitude larger volumes of emulsion. In the CHORUS experiment, the high intensity neutrino beam from the CERN SPS (section 2.2) was used in conjunction with an 800 kg emulsion target. The net-scan technique was used to examine a large volume of emulsion around neutrino interactions and a high statistics sample of charm decays has been collected.

Contrary to the previous experiments which use the reaction $D^{*+} \rightarrow D^0 \pi^+$ to tag the D^0 , in the CHORUS experiment the D^0 is directly recognized by its decay topology, independently of the presence of the D^{*+} parent. Therefore, the D^{*+} sample is obtained through the recognition of the above reaction as a subset of an already very pure sample of D^0 events. This offers the unique opportunity to measure the D^{*+} production rate relative to the D^0 production cross-section without the need for an external normalization.

5.2 Charm-quark production and fragmentation

An extensive description of the theory of charm-quark production and fragmentation and a discussion of the previous experimental results can be found in Ref. 215. Here a brief overview of the phenomenology of charm-quark production and fragmentation is given.

The theoretical description of neutrino–nucleon scattering is rather complex because the nucleon is not an elementary point-like particle. The complexity arises, not from the theoretically well-understood weak interaction, but from the substructure of the nucleon and from nucleon–nucleon interactions inside a nucleus. The scattering cross-section is normally described in terms of structure functions which describe the internal structure of the nucleon. The structure functions represent quark and gluon-density distributions inside a nucleon.

The interactions between quarks and gluons is described by a theory known as quantum-chromo-dynamics (QCD). It is a non-Abelian gauge theory which means that the force-carrying bosons, the gluons, carry colour charges themselves. Because the glu-

ons have colour charges, they can interact with each-other, a process which is thought to be responsible for the confinement of quarks in colour-neutral hadrons. In interactions with high four-momentum transfer, the quark masses can usually be neglected, and the structure functions and cross-sections for these processes can be calculated using perturbation theory. On the contrary, in charm production and at the typical four-momentum transfer $Q^2 = 10 \text{ GeV}^2/c^2$ for the CHORUS experiment, the mass of the charm quark $m_c \approx 1.5 \text{ GeV}/c^2$ can no longer be neglected. As the charm-quark mass is similar to the nucleon mass, one can also no longer neglect the nucleon mass in the cross-section calculations.

Charged-current neutrino interactions are mediated by W^\pm bosons and are sensitive to the quark-flavour distributions. Neutrino–nucleon scattering is therefore an essential tool in determining the individual quark-density functions and related structure functions. As a case in hand, charm production by neutrino–nucleon interactions is especially sensitive to the strange-sea distribution in the nucleon. Even though the down valence-quark distribution is about 10 times bigger in a proton, the $d \rightarrow c$ transition is Cabibbo suppressed. As a result, the $d \rightarrow c$ and $s \rightarrow c$ processes give about equal contributions to the cross-section. The diagrams for these processes are given in Figure 5.1. From the difference of the charged-current neutrino and anti-neutrino scattering cross-sections with a charmed quark in the final state, one can extract the valence down-quark density function. The quark and anti-quark contributions from the down and strange sea are expected to be equal and cancel in the difference.




Figure 5.1: Schematic diagrams of charmed meson production by neutrino–nucleon interactions. The valence diagram is Cabibbo suppressed. The sea contribution, indicated by a gluon splitting in an $s\bar{s}$ pair, is Cabibbo favoured and similar to the valence contribution at the typical energies in CHORUS.

After the production of a charm quark, one still has to understand how the charm quark ends up in the hadronic final state. This process is known as hadronization or fragmentation. From the number of available spin states, one expects to produce three times as many vector mesons than pseudo-scalar mesons. In the case of charm production, the vector mesons are the D^{*+} and the D^{*0} ; the pseudo-scalar mesons are the D^+ and the D^0 . Furthermore, under isospin invariance and assuming similar mass for the up and down quark, the probability that the charm quark combines with an anti-up or anti-down quark is equal. The production cross-sections for D^0 and D^+ (D^{*+} and D^{*0}) are therefore equal. The D^{*0} resonance decays strongly to $D^0 \pi^0$ or electro-magnetically to $D^0 \gamma$, because the decay to $D^+ \pi^-$ is energetically forbidden. The D^{*+} resonance decays for about $2/3$ to $D^+ \pi^+$ and for about $1/3$ to $D^+ \pi^0$. One then finds that the expected

production ratio of D^+ to D^0 in the final state is 1:3. This fragmentation is schematically indicated on the right-hand side of the diagrams of Figure 5.1.

One should remember that in an experiment only some quantities of the final state are accessible. In neutrino experiment at a wide-spectrum neutrino beam, not even the initial state is known. However, in CHORUS the final-state muon momentum and the total hadronic energy is measured. In a charged-current ν_μ interaction, these two variables do determine the initial neutrino momentum.

Most charm-production cross-section measurements in neutrino experiments rely on the muonic decay of the charmed hadrons, as the individual charm particles cannot be identified due to their short flight length. These events, known as dimuon events, are characterized by two opposite-sign muons in the detector coming from the same initial vertex. Experiments based on this signature basically group all charmed hadrons for which an average charm to muon branching ration, $B_{c \rightarrow \mu}$, is observed. This average branching ratio is given by the sum over all types of charmed hadrons

$$B_{c \rightarrow \mu} = \sum_t f_{h_t} \cdot B(h_t \rightarrow \mu) \quad ,$$

where f_{h_t} is defined as the probability that a charmed quark hadronizes into a charmed hadron h_t , with the obvious normalization $\sum_t f_{h_t} = 1$. The semi-leptonic branching ratios entering into the definition of $B_{c \rightarrow \mu}$ have been determined to a reasonable degree of accuracy in $e^+ e^-$ experiments, both at the $c\bar{c}$ threshold and at higher energies. Unfortunately, this is not true for the production fractions f_{h_t} . Furthermore, f_{h_t} depends on energy, mostly due to the energy dependence of the Λ_c^+ production cross-section.

By their very nature, the f_{h_t} can only be measured in those neutrino experiments that are capable of identifying the charmed hadron. So far, only the E531 experiment [202] has conducted such an experiment. To tag and identify charmed particles produced in neutrino interactions, it used an emulsion target followed by a series of electronic detectors providing both kinematic reconstruction and particle identification. From a study of 122 charmed-particle decays among 3855 located neutrino interactions, the E531 collaboration has extracted the production fraction for each of the charmed hadrons [260]. The experiment also measured the dependence on the kinematic variables [261]. This sample of 122 events continues to be the only reference for the production fractions f_{h_t} and has been reanalyzed in Ref. 262. The results of all dimuon experiments rely heavily on this single measurement.

The E531 experiment was set up specially for charm physics studies, CHORUS on the other hand has many more events in the emulsion but is unable to reliably identify the type of charmed hadron because it has no particle-identification detectors and lacks accuracy in the hadron spectrometer. However, CHORUS can perfectly distinguish charged charmed hadrons from neutral ones and can make a sample selection that is essentially background free. Statistically, the f_{h_t} fractions can be extracted from the distributions of flight length and energies. On an event by event basis, a particle-identification probability can be assigned from the flight length and decay topology measured in the emulsion.

The charm detection in CHORUS relies on the net-scan technique. The primary neutrino vertex in a ν_μ charged-current interaction can be efficiently located by tracking back the muon through the emulsion plates. With the net-scan technique, all tracks (up to 5 mm) downstream from the neutrino vertex can be reconstructed and secondary vertices extracted. Secondary vertices with even numbers of tracks and no nuclear remnant

tracks, can only be due to decays of short-lived neutral particles. At the energy scale of CHORUS, only the D^0 contributes to this signal. Secondary vertices with an incoming track connected to the primary vertex, an odd number of outgoing tracks, and again no interaction are most likely from D^\pm decays.

5.3 Vertex reconstruction and charm sample selection

As was described in section 2.8, tracks found in the target tracker are used to predict the position of neutrino interactions in the emulsion. In particular, for charged-current ν_μ interactions, the primary muon is used to predict with high precision the point where it traversed the changeable sheets. The muon is identified by its range in the calorimeter and the muon spectrometer. The layout of the emulsion target and trackers is described in section 2.4. The approximate neutrino interaction point is located by following the prediction upstream from plate to plate in the scan-back procedure. In order to select charm-decay candidates, the net-scan procedure is used to reconstruct all high-energy tracks originating in a $1.5\text{ mm} \times 1.5\text{ mm} \times 6.3\text{ mm}$ volume around the located neutrino interaction. The scan-back and net-scan procedures are explained in section 2.10. As is described there, in the net-scan procedure tracks are connected over multiple plates and only tracks originating from within the net-scan volume are kept. The reconstruction algorithm then tries to associate those tracks to common vertices. A track is attached to a vertex if the distance of the vertex point to the reconstructed track (referred to as the impact parameter) is less than $10\ \mu\text{m}$. At the end of the vertex-assignment procedure, one defines a primary vertex and possibly one or more secondary vertices. In the following, tracks origination from secondary vertices are called daughter tracks.

Candidate events for ν_μ charged-current interactions with a charm decay are then automatically selected from the set of reconstructed events. An event is defined as a candidate for a charged-current ν_μ interaction if the primary muon track, identified in the electronic detectors, is connected to the neutrino-interaction vertex and is found in more than one emulsion plate in the net-scan data. Decay topologies are then selected with the following criteria. At least one of the tracks connected to a secondary vertex is detected in more than one plate and the direction measured in the emulsion matches that of a track reconstructed in the target trackers. The parent angle, in the case of a neutral particle decay deduced from the line connecting the primary and secondary vertex, should be within 400 mrad from the beam direction. To avoid including too many fake secondary vertices due to multiple scattering of low-energy tracks from the primary vertex, the impact parameter to the primary vertex of at least one of the daughter tracks must be larger than a certain value which is determined on the basis of the angular resolution. The angular resolution depends on the track angle (θ) with respect to the beam as $\sigma_\theta(\theta) = \sqrt{3^2 + (19.4 \cdot \tan \theta)^2}\text{ mrad}$. In order to remove random association, the impact parameter is also required to be smaller than a value depending on the distance covered by the track extrapolation to the vertex. The typical tolerance is $130\ \mu\text{m}$. The flight length of the parent candidate should be longer than $25\ \mu\text{m}$.

Out of a sample of 93,807 scanned and analysed neutrino-induced charged-current events, these criteria select 2752 events as having a decay topology. The automatic scanning only measures the most upstream $100\ \mu\text{m}$ parts of the tracks on successive emulsion plates. It, therefore, does not see the actual vertices and can, for example, not distinguish between a neutral or charged short-lived particle if it does not cross a plate

boundary. Nor can it distinguish between decay and interaction vertices. Therefore, the selected events are visually inspected to confirm the presence of a decay. A secondary vertex is accepted as a decay if the number of charged particles is consistent with charge conservation. In addition, no other activity (Auger-electron or visible recoil) should be observed. In this analysis, only those events are used in which the secondary vertex is consistent with the decay of a neutral particle. Thus the selection and identification of the D^0 sample used in this analysis is based on the decay topology of the D^0 alone.

Decay candidates are selected from all charged-current events:

Located charged-current events	93,807
Selected for visual inspection	2752

Visual inspection of the decay candidates rejected 739 events because:

Decay topologies with flight length $< 25 \mu\text{m}$	3
Topologies with kink angle $< 50 \text{ mrad}$	11
Secondary interactions	278
Electron-positron pairs	95
Overlay neutrino interactions	44
Uncorrelated (overlay) secondary vertices	21
Passing-through tracks	128
All tracks from primary vertex	142
δ -rays	2
Other	15
<hr/> Total rejected by visual inspection	<hr/> 739

The remainder are charm-decay candidates with the following signatures:

Charged charm: 1-prong decay (kink)	452
Charged charm: 3-prong decay	491
Charged charm: 5-prong decay	22
<hr/> Neutral charm : 2-prong decay vertex (V2)	<hr/> 819
Neutral charm : 4-prong decay vertex (V4)	226
Neutral charm : 6-prong decay vertex (V6)	3
<hr/> Total charm candidates	<hr/> 2013

Table 5.1: Number of located and reconstructed events in the emulsion from the charged-current data sample and the charm sub-sample.

The results of the visual inspection are summarized in Table 5.1. The purity of the automatic selection is $[73.1 \pm 0.8]\%$. The rejected sample consists of secondary hadronic interactions, δ -rays or gamma conversions, overlay neutrino interactions, and of low-momentum tracks which, due to multiple scattering, appear as tracks with a large impact parameter. The remainder consists either of false vertices, being reconstructed using one or more background tracks, or of vertices with a parent track not connected to the primary (passing-through tracks).

5.4 Event selection of $D^{*+} \rightarrow D^0 \pi^+$

The identification of D^{*+} is based on its decay into D^0 and π^+ :

$$D^{*+} \rightarrow D^0 \pi^+ . \quad (5.1)$$

The reconstruction of the invariant mass of the D^{*+} would require to measure momenta of at least three charged hadrons, namely the π^+ from the D^{*+} decay and two charged decay daughters of the D^0 for the most favourable decay mode, $D^0 \rightarrow K^- \pi^+$ with a branching ratio of $[3.80 \pm 0.09]\%$. Such a strategy would result in a very small event-selection efficiency and has therefore not been applied.

The procedure applied here relies on the high purity of the sample of D^0 events identified by the decay topology in emulsion and by the low Q -value of the D^{*+} decay. The Q -value corresponds to a maximum p_T of 39 MeV/ c . The π^+ from the D^{*+} decay has a relatively low momentum (< 4 GeV/ c) and, as a consequence of the low Q -value, a small angle with respect to the D^0 direction. In this two-body decay, the transverse momentum spectrum of the π^+ with respect to the D^0 direction shows a Jacobian peak below 40 MeV/ c . The kinematical quantities needed for this approach are the direction of the π^+ , the momentum of the positive pion, and the direction of the D^0 .

5.4.1 D^0 secondary vertex selection

The selection starts from the sample of D^0 events, characterized by the decay topology into two or four charged particles as is described above and in Ref. 263. In total, the sample contains 1048 D^0 events of which 1045 decayed into two and four prong modes. For each particle track, recognized in the emulsion as originating from the primary vertex, a charge selection is made and the transverse momentum, p_T , with respect to the direction of the D^0 is measured.

In the CHORUS experiment, the track reconstruction in the emulsion target determines the slope of the charged particles with high precision (order 1 mrad). The direction of the D^0 is determined from the measurement of the positions of the primary neutrino interaction vertex and the decay vertex. The precision of this determination depends on the decay length. For this measurement of the D^{*+} cross-section, only decays with a flight length of more than 100 μm are used to ensure good resolution.

5.4.2 Primary π^+ selection

The momentum measurement of the particles reconstructed in the emulsion is performed by extrapolating the tracks from the emulsion to the target tracker and to the hadron spectrometer (section 2.6.1). In order to select the tracks with a reliable momentum determination, strict requirements are applied on the measurements in the hadron spectrometer. Only tracks with a hit in all six paddles of the diamond tracker or with track segments in the honeycomb tracker (Chapter 3) and a smaller number of hits in the diamond tracker were retained. The 488 selected events with a D^0 (see section 5.4.4), had 1116 hadron tracks reconstructed at the primary vertex. Of these, the momentum of 377 particles was measured fulfilled the criteria just given. The inefficiency can be attributed in about equal parts to secondary interactions in the target material and to the limited geometrical acceptance of the hadron spectrometer. The resolution of the hadron spectrometer is about 35% for the low momentum range relevant for this study.

5.4.3 Monte-Carlo simulation

In order to study the potential separation power between signal and background, the p_T distribution was studied using a simulation of the events originating in the four stacks of the emulsion target separately. Figure 5.2 shows the p_T distribution of the simulated D^0 candidates (including events where the D^0 was produced from the decay of the D^{*+}), selected by the same selection criteria as used for the data. The solid curve is the prediction for combinations of D^0 's with positively-charged hadrons including the expected signal. The dashed curve is the sample of D^0 's which do not originate from the reaction $D^{*+} \rightarrow D^0 \pi^+$. The data for the four stacks are based on the same simulated sample and thus predict the relative weight of the signal to be observed in the four stacks. A good separation between signal and background is possible only for interactions originating in the two downstream stacks (stack 3 and 4). Therefore, only the events originating in these two stacks were considered in the analysis. The signal-to-background ratio is most favourable in the p_T region from 10 MeV/c to 50 MeV/c. This region is used as signal region in the analysis.

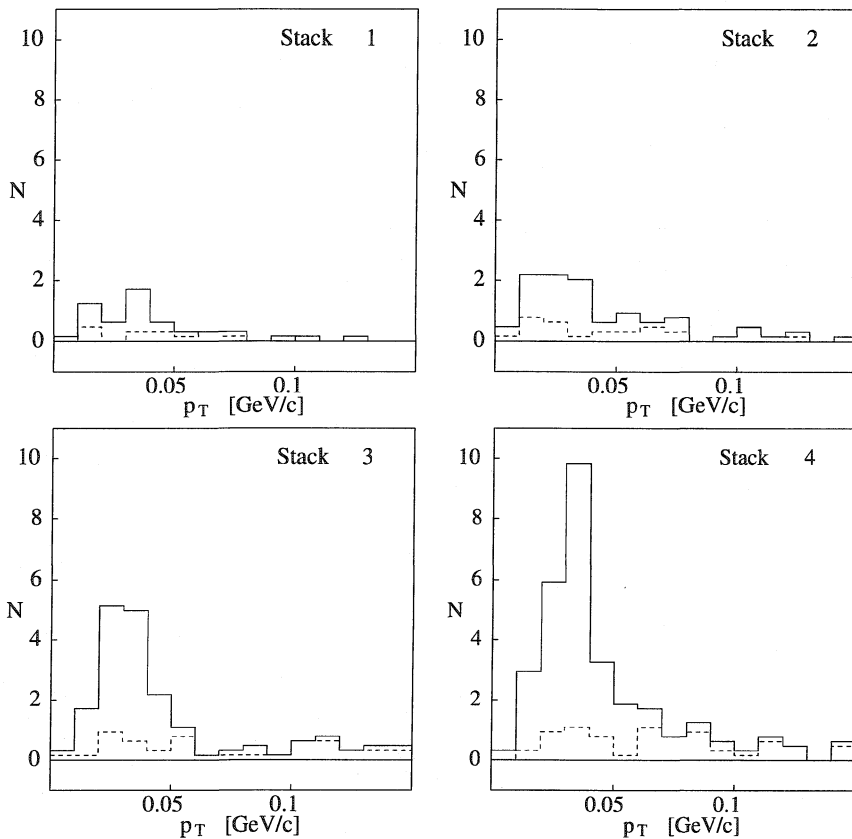


Figure 5.2: Number of charged hadrons (N) from the primary vertex as a function of the transverse-momentum (p_T) with respect to the D^0 direction as obtained from a simulation. The four panels show the results of the simulation for the four emulsion stacks, stack 1 being the most upstream stack. For an explanation of the curves see the main text.

The loss of efficiency and signal-to-noise ratio for the upstream stacks are easy to understand. They are the result of the low average momentum ($\approx 1 \text{ GeV}/c$) of the π^+ in reaction (5.1). Multiple scattering plays a large role together with re-interactions in the downstream stacks. Therefore the matching of the track seen in the emulsion with the hits in the hadron spectrometer has a low efficiency when a large amount of material has been traversed by the particle.

5.4.4 Signal extraction

The candidate events were searched within the sample of 488 D^0 events in stacks 3 and 4 with a D^0 flight length l_f exceeding $100 \mu\text{m}$. After removing tracks identified as muons, 1116 tracks originating from the primary vertex were found in these events. To obtain an enriched sample of candidate events within this sample, the following set of kinematical criteria was applied. Since the typical momentum of the π^+ from D^{*+} decay is smaller than $4 \text{ GeV}/c$, hadrons with a momentum greater than this value were not considered. In addition, tracks having a measured momentum smaller than $400 \text{ MeV}/c$ were rejected. These tracks are more likely to come from secondary hadrons and from random coincidences of hits in the tracking system of the hadron spectrometer. The angle between the hadron and the D^0 was required to be smaller than 60 mrad . The simulation showed that the signal does not populate the region with larger angles. For the remaining tracks, a separation between positive and negative charges was made and the p_T spectra were further analysed. The result of the event and track selection is given in Table 5.2.

Number of events with secondary D^0 and associated hadron tracks:

Stacks:	1 and 2	3 and 4
Visually confirmed D^0 decays	521	527
D^0 decays selected for analysis ($l_f \geq 100 \mu\text{m}$)		488
Hadron tracks at the primary vertex		1116
Hadrons from primary vertex with measured p		377

Selection of known-momentum hadron tracks in $D^{*+} \rightarrow D^0 \pi^+$ signal region:

Charge:	positive	negative
Hadrons from primary vertex with measured p	248	129
Within angular range ($\theta(\pi, D^0) < 60 \text{ mrad}$)	62	26
Within momentum window ($0.4 < p < 4 \text{ GeV}/c$)	47	12
In p_T signal region ($10 < p_T < 50 \text{ MeV}/c$)	27	4

Table 5.2: Number of events and tracks used in the analysis.

Figure 5.3 shows the p_T distributions of positively charged hadrons and negatively charged hadrons originating from the primary vertex in the D^0 data sample and in a sample of charged-charm production events. In the figure, the points with error bars show the number of candidate events with the statistical error. The dashed curve shows the background prediction. In the distribution for the D^0 with positive hadrons, the solid

curve represents the expected shape of the candidate events (including the background simulation) normalized to the observed events. In the other panels, the solid curve shows the prediction for the shape normalized to the total number of D^0 's and charged charm events, respectively. An excess which can be attributed to the signal is visible in the range between 10 MeV/c to 50 MeV/c in the p_T distribution of positive hadrons in the D^0 sample, while in the p_T distribution of negative hadrons no such excess is seen. In addition, if one makes the same comparison between the p_T spectra of positive and negative hadrons with respect to charged charm particles found in the same stacks and with the same kinematical cuts, no such signal can be seen. This behaviour is a clear indication of D^{*+} decays.




Figure 5.3: The number of charged hadron tracks (N) from the primary vertex as function of the transverse-momentum (p_T). The panel in the left (right) upper corner shows the p_T of positively (negatively) charged hadrons with respect to the D^0 direction. The panel in the left (right) lower corner shows the p_T for positively (negatively) charged hadrons with respect to direction of positively charged charm particles. The meaning of the different points and curves is explained in the text.

5.5 Efficiency and background evaluation

The detection efficiencies and expected background contamination were evaluated with a detailed Monte-Carlo simulation of the experiment based on the GEANT3 simulation framework [264]. The GEANT3 framework was also used for the modelling of the neutrino beam. A large number of deep-inelastic neutrino interactions were generated according to the beam spectrum by the JETTA generator [265] developed from LEPTO [266] and JETSET [267]. The simulated response of the detectors was processed through the standard chain of reconstruction programs. The efficiency of the event location and reconstruction in the emulsion was obtained by a detailed simulation of the efficiency and resolution of the net-scan procedure. In order to simulate realistic conditions of background grains in the emulsion, the data of empty volumes, i.e. net-scan volumes where no neutrino interaction is present, were overlaid on the simulated events.

5.5.1 Background evaluation

The main source of background is due to π^+ tracks from the primary vertex in events with a D^0 which have a p_T with respect to the D^0 in the signal region as defined in section 5.4. The number of background events in the signal region was estimated using D^0 events from the Monte-Carlo simulation, excluding the events coming from D^{*+} production. The simulated p_T distribution in the $\pi^+ + D^0$ channel was then normalized by comparing the number of events observed in the region $p_T < 0.1$ GeV/ c in the $\pi^- + D^0$ channel with the data in that channel. The total background was found to be $4.9 \pm 1.6(\text{stat}) \pm 1.0(\text{syst})$. The statistical error comes mainly from the normalization with respect to the events with a π^- (or other negative hadron). The systematic error is an estimate of the accuracy of the simulation.

Within the statistical accuracy, the simulated background agrees well with the events observed in the combination of a D^0 with a negative hadron, and for events with charged charm particles combined with positive and negative hadrons, respectively. Only in the channel with a D^0 and a positive hadron, is a clear excess of events above the predicted background observed. The p_T distribution of the excess is compatible with the simulated detector response.

The background to the sample of D^0 events coming from neutral strange particle decays such as Λ^0 and K_S^0 is negligible due to the much longer decay length of these particles. Their number has been evaluated for the four emulsion stacks and found to be 11.5 ± 1.9 Λ^0 's and 25.1 ± 2.9 K_S^0 's in the full D^0 sample, respectively [268]. Since the calculation of the background coming from randomly associated positive hadrons from the primary vertex in combination with a D^0 has been normalized to the number of negative hadrons with the same kinematic properties, the random associations of a neutral strange particle with a π^+ have been taken into account implicitly. Their number can be estimated to give approximately $0.3 \cdot 10^{-3}$ events in the signal region. A potentially larger background is caused by K^* decays into a K_S^0 and a π^+ . With the assumption that all K_S^0 's originate from K^* decay, 0.3 ± 0.1 events would pass the selection criteria. This estimate can thus be regarded as an upper limit for this background.

5.5.2 Detection efficiency

Since only the ratio of D^{*+} events with respect to the D^0 events is considered, only the efficiency associated with the detection of the π^+ has to be considered. Any energy-dependent effect in distinguishing the D^0 from D^{*+} decay from other D^0 's, which would spoil the cancellation of the efficiency is very small and is included in the simulation. The efficiencies in locating events of the two types in the emulsion are expected to be equal to a high accuracy.

The important issues related to the π^+ detection are the efficiency in attaching the π^+ to the primary vertex in the reconstruction, the efficiency in measuring its momentum, and the ability of the Monte-Carlo chain to simulate the p and p_T resolution correctly. The combined systematic uncertainty in the product of the efficiency to find the π^+ and in the efficiency in measuring its momentum was estimated to be 7%. Such a value is obtained by observing the differences in the fraction of tracks with a measured momentum predicted by the simulation and observed in the data for different track samples. The samples compared were the tracks from primary neutrino vertices with those from D^0 decays. The efficiency in finding the π^+ track at the primary vertex is well reproduced by the simulation, and the uncertainty in this calculation contributes to a smaller extent to the systematic error.

A study of the effects governing the p_T resolution shows that the most critical ingredient is the measurement of the angle of the D^0 . From a comparison of measurements using the vertex reconstruction with the tracks found in the net-scan procedure and the vertex measurement in the manual scanning process, an angular resolution of 10 mrad is deduced at short flight paths and 5 mrad at long flight lengths. Using these numbers a 5% uncertainty in the efficiency of the cut in p_T is evaluated.

The overall uncertainty in the efficiency is 11%, which includes the systematic uncertainty in the efficiency in measuring the π^+ momentum, the uncertainty in the effect of the p_T cut and the statistical error of the simulation.

5.6 Results and conclusion

There are 27 events with a positive hadron in the signal region in the D^0 sample. Using the evaluation of the background as described in section 5.5.1 amounting to 4.9 ± 1.9 , a signal of 22.1 ± 5.5 events is obtained.

The most direct measurement which can be obtained is the ratio of D^{*+} and D^0 production in charged-current neutrino interactions. This ratio can be expressed as:

$$\frac{\sigma(D^{*+})}{\sigma(D^0)} = \frac{N(D^{*+} \rightarrow D^0 \pi^+)}{N(D^0)} \cdot \frac{\epsilon(D^0)}{\epsilon(D^{*+} \rightarrow D^0 \pi^+)} \cdot \frac{1}{B(D^{*+} \rightarrow D^0 \pi^+)},$$

where $N(D^{*+} \rightarrow D^0 \pi^+)$ and $N(D^0)$ are the number of D^{*+} and D^0 events observed, respectively, and $\epsilon(D^{*+} \rightarrow D^0 \pi^+)$ and $\epsilon(D^0)$ their relative detection efficiencies. The ratio of these efficiencies is 0.176 ± 0.020 . The branching ratio $B(D^{*+} \rightarrow D^0 \pi^+)$ is 0.677 ± 0.005 . Substituting the numerical values, one obtains:

$$\frac{\sigma(D^{*+})}{\sigma(D^0)} = 0.38 \pm 0.09(\text{stat}) \pm 0.05(\text{syst}). \quad (5.2)$$

Under the assumption that the D^{*0} and D^{*+} production rates are equal and recalling that the D^{*0} always decays into a D^0 , it can be concluded that most D^0 's in neutrino interactions are produced through the decay of a D^* :

$$\sigma(D^* \rightarrow D^0) / \sigma(D^0) = 0.63 \pm 0.17 \quad .$$

Following Ref. 269 and defining R_2 as the fraction of D^0 's coming from D^{*+} decays, one obtains $R_2 = 0.25 \pm 0.10$ where the statistical and systematic errors are combined. With the definition $f_V = V / (P + V)$, the ratio of the vector D meson production and the sum of vector and pseudoscalar production of D mesons as introduced in Ref. 269, one finds $f_V = 0.51 \pm 0.18$. Within the precision, this determination is consistent with the value for the ratio of vector meson and pseudoscalar meson production (three to one) expected from simple spin arguments and with more precise measurements in $e^+ e^-$, πN and γN experiments [270].

The rate of D^{*+} meson production relative to the neutrino charged-current interaction cross-section can be obtained by combining the result in equation (5.2) with the measurement $\sigma(D^0) / \sigma(CC) = [2.69 \pm 0.18 \pm 0.13] \%$ obtained using the same D^0 sample [263].¹ One then obtains:

$$\frac{\sigma(D^{*+})}{\sigma(CC)} = [1.02 \pm 0.25(\text{stat}) \pm 0.15(\text{syst})] \% \quad .$$

The NOMAD experiment [259], operating in the same neutrino beam as CHORUS, has reported the rate of D^{*+} production per charged-current neutrino interaction to be $[0.79 \pm 0.17(\text{stat}) \pm 0.10(\text{syst})] \%$. The BEBC bubble chamber data [257] were reanalysed combining several datasets using a neutrino beam with energies similar to CHORUS. This group reported the rate of D^{*+} production per charged-current neutrino interaction to be $[1.22 \pm 0.25] \%$. Our result is consistent with these measurements. At the higher energies of the Tevatron beam a value of $[5.6 \pm 1.8] \%$ has been found [258].

¹The D^0 production rate had been obtained using only the four-prong decay topology for which practically all decay modes are available in the literature. This procedure minimizes the systematic error due to the unmeasured decay modes in the two-prong topology and the decay modes into final states with neutral particles only.

Bibliography

- [1] "Review of particle physics"
Particle Data Group; S. Eidelman *et al.* Phys. Lett., **B592**, p.1 (2004)
- [2] "Letter to the physical society of Tübingen", W. Pauli (1930)
reproduced in: "The idea of the neutrino"
L.M. Brown Phys. Today, p.23, September (1978)
- [3] "Versuch einer Theorie der β -Strahlen I"
E. Fermi Zeit. für Phys., **88**, p.161 (1934)
- [4] "A proposed experiment to detect the free neutrino"
F. Reines and C.L. Cowan Jr. Phys. Rev., **90**, p.492 (1953)
- [5] "Detection of the free neutrino"
F. Reines and C.L. Cowan Jr. Phys. Rev., **92**, p.830 (1953)
- [6] "Detection of the free neutrino: a confirmation"
C.L. Cowan Jr. *et al.* Science, **124**, p.103 (1956)
- [7] C.D. Anderson and S.H. Neddermeyer
in papers and discussions of the international conference on physics, London, 1934
- [8] "Über die Eigenschaften der durchdringenden Korpuskularstrahlung im Meeresniveau"
B. Rossi Zeit. für Phys., **82**, p.151 (1933)
- [9] "Cloud chamber observations of cosmic rays at 4300 meters elevation and near sea-level"
C.D. Anderson and S.H. Neddermeyer Phys. Rev., **50-4**, p.263 (1936)
- [10] "The absorption of cosmic-ray electrons"
J.C. Street *et al.* Phys. Rev., **47-12**, p.891 (1935)
- [11] "New evidence for the existence of a particle of mass intermediate between the proton and electron" ; J.C. Street and E.C. Stevenson Phys. Rev., **52**, p.1003 (1937)
- [12] "Search for improbable meson decays"
S. Lokanathan and J. Steinberger Phys. Rev., **98**, p.240 (1955)
Proceedings of the American Physical Society, Chicago, November 1954
- [13] "Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos"
G. Danby *et al.* Phys. Rev. Lett., **9**, p.36 (1962)
- [14] "Evidence for anomalous lepton production in $e^+ - e^-$ annihilation"
M.L. Perl *et al.* Phys. Rev. Lett., **35-22**, p.1489 (1975)
- [15] "Properties of anomalous $e\mu$ events produced in e^+e^- annihilation"
M.L. Perl *et al.* Phys. Lett., **B63**, p.466 (1976)
- [16] "Neutrino counting"
G. Barbiellini *et al.*; in "Z Physics at LEP1: vol. 1: Standard Physics", p.129
G. Altarelli *et al.*(eds); CERN (Geneva), 1989 CERN/89-08
- [17] "Measurement of the decay $\tau^- \rightarrow \rho^- \nu_\tau$ "
ARGUS collaboration; H. Albrecht *et al.* Zeit. für Phys., **C56**, p.339 (1992)
- [18] "A combination of preliminary electroweak measurements and constraints on the standard model, 1997" ; D. Abbaneo *et al.*, CERN report: PPE/97-154

- [19] "Direct measurement of leptonic coupling asymmetries with polarized Z bosons"
SLD collaboration; K. Abe *et al.* Phys. Rev. Lett., **79-5**, p.804 (1997)
- [20] "Observation of tau neutrino interactions"
DONUT collaboration; K.Kodama *et al.* Phys. Lett., **B504**, p.218 (2001)
- [21] "Quarks & Leptons"
F. Halzen and A.D. Martin; Wiley (New York), 1984 ISBN: 0-471-81187-4
- [22] "Question of parity conservation in weak interactions"
T.D. Lee and C.N. Yang Phys. Rev., **104**, p.254 (1956)
- [23] "Experimental test of parity conservation in beta decay"
C.S. Wu *et al.* Phys. Rev., **105**, p.1413 (1957)
- [24] "Observations of the failure of conservation of parity and charge conjugation in meson decays: the magnetic moment of the free muon"
R.L. Garwin *et al.* Phys. Rev., **105**, p.1415 (1957)
- [25] "Nuclear emulsion evidence for parity nonconservation in the decay chain $\pi^+ - \mu^+ - e^+$ "
J.I. Friedman and V.L. Telegdi Phys. Rev., **105**, p.1681 (1957)
- [26] "Evidence for the 2π decay of the K_2^0 meson"
J. H. Christenson *et al.* Phys. Rev. Lett., **13**, p.138 (1964)
- [27] "Unitary symmetry and leptonic decays"
N. Cabibbo Phys. Rev. Lett., **10**, p.531 (1963)
- [28] "CP-violation in the renormalizable theory of weak interaction"
M. Kobayashi and T. Maskawa Prog. of Theor. Phys., **49-2**, p.652 (1973)
- [29] "Experimental observation of heavy particle J "
J.J. Aubert *et al.* Phys. Rev. Lett., **33-23**, p.1404 (1974)
- [30] "Discovery of a narrow resonance in e^+e^- annihilation"
J.-E. Augustin *et al.* Phys. Rev. Lett., **33-23**, p.1406 (1974)
- [31] "Elementary particles and SU(4)"
B.J. Björken and S.L. Glashow Phys. Lett., **11**, p.255 (1964)
- [32] "Weak interactions with lepton-hadron symmetry"
S.L. Glashow *et al.* Phys. Rev., **D2**, p.1285 (1970)
- [33] "A possible decay in flight of a new type particle"
K. Niu *et al.* Prog. of Theor. Phys., **46-5**, p.1644 (1971)
- [34] "Neutrino cross-sections"
J. Panman private communications
- [35] "Mesonium and antimesonium"
B. Pontecorvo Soviet Phys. JETP, **6**, p.429 (1958)
- [36] "Inverse beta processes and nonconservation of lepton charge"
B. Pontecorvo Soviet Phys. JETP, **7**, p.172 (1958)
- [37] "High precision measurement of the tritium β spectrum near its endpoint and upper limit on the neutrino mass" ; Ch. Weinheimer *et al.* Phys. Lett., **B460**, p.219 (1999)
- [38] "Direct search for mass of neutrino and anomaly in the tritium beta-spectrum"
V.M. Lobashev *et al.* Phys. Lett., **B460**, p.227 (1999)
- [39] "Upper limit of the muon-neutrino mass and charged-pion mass from momentum analysis of a surface muon beam" ; K.Assamagan *et al.* Phys. Rev., **D53-11**, p.6065 (1996)
- [40] "Determination of an upper limit for the mass of the τ -neutrino at LEP"
OPAL collaboration; R. Akers *et al.* Zeit. für Phys., **C65**, p.183 (1995)

- [41] "An upper limit on the τ neutrino mass from three- and five-prong tau decays"
ALEPH collaboration; R. Barate *et al.* Euro. Phys. J., **C2**, p.395 (1998)
- [42] "Remarks on the unified model of elementary particles"
Z. Maki *et al.* Prog. of Theor. Phys., **28**, p.870 (1962)
- [43] "Neutrino oscillations in matter"
L. Wolfenstein Phys. Rev., **D17-9**, p.2369 (1978)
- [44] "Resonance enhancement of oscillations in matter and solar neutrino spectroscopy"
S.P. Mikheev and A.Yu. Smirnov Soviet J. of Nucl. Phys., **42**, p.913 (1985)
- [45] "Incorporating systematic uncertainties into an upper limit"
R.D. Cousins and V.L. Highland Nucl. Inst. & Meth., **A320**, p.331 (1992)
- [46] "Gravitational microlensing by the galactic halo"
B. Paczynski Astrophys. J., **304**, p.1 (1986)
- [47] "Galactic microlensing as a method of detecting massive compact halo objects"
K. Griest Astrophys. J., **366**, p.412 (1991)
- [48] "The MACHO project: microlensing results from 5.7 years of Large Magellanic Cloud observations" ; MACHO coll.; C. Alcock *et al.* Astrophys. J., **542**, p.281 (2000)
- [49] "Limits on galactic dark matter with 5 years of EROS SMC data"
C. Afonso *et al.* Astronomy & Astrophys., **400**, p.951 (2003)
- [50] OGLE homepage: <http://www.astrouw.edu.pl/~ogle/>
- [51] "Wilkinson Microwave Anisotropy Probe (WMAP) three year results: implications for cosmology" ; D.N. Spergel and R. Bean, arXiv preprint astro-ph/0603449 (2006)
- [52] "Neutrinos in cosmology"
S. Hannestad New J. of Phys., **6**, p.108 (2004)
- [53] "Cosmological neutrino bounds for non-cosmologist"
M. Tegmark Phys. Scripta, **T121**, p.153 (2005)
- [54] "The formation of deuterons by proton combination"
H.A. Bethe and C.L. Critchfield Phys. Rev., **54**, p.248 (1938)
- [55] "Energy production in stars"
H.A. Bethe Phys. Rev., **55**, p.434 (1939)
- [56] "Astrophysical neutrinos: 20th century and beyond"
J.N. Bahcall Nucl. Phys. B(proc. suppl.), **91**, p.9 (2001)
Proceedings of the 19th international conference on neutrino physics and astrophysics, Sudbury, Canada, June 2000
- [57] "Helioseismological implications of recent solar abundance determinations"
J.N. Bahcall *et al.* Astrophys. J., **618**, p.1049 (2005)
- [58] "On the photon diffusion time scale for the sun"
R. Mitalas and K.R. Sills Astrophys. J., **401**, p.759 (1992)
- [59] "On the time scale of energy transport in the sun"
M. Stix Solar Phys., **212**, p.3 (2003)
- [60] "Solar neutrino detection by the ^{37}Cl - ^{37}Ar method", R. Davis Jr. and D.S. Harmer
proceedings of the informal conference on experimental neutrino physics, Geneva, Switzerland, January 1965, p.201 CERN (Geneva), 1965
- [61] "Solar neutrinos II: experimental"
R. Davis Jr. Phys. Rev. Lett., **12**, p.303 (1964)

- [62] "Detection of solar neutrinos by means of the $\text{Ga}^{71}(\nu, e^-)\text{Ge}^{71}$ reaction"
V.A. Kuz'min Soviet Phys. JETP, **22**, p.1050 (1966)
- [63] "How uncertain are solar neutrino predictions?"
J.N. Bahcall *et al.* Phys. Lett., **B433**, p.1 (1998)
- [64] "Report on the chlorine solar neutrino experiment", R. Davis Jr. *et al.*
proceedings of the thirteenth international conference on neutrino physics and astrophysics, Boston Massachusetts, June 1988
J. Schneps *et al.*(eds); World Scientific (Singapore), 1989 ISBN: 9-971-50778-1
- [65] "Time dependence of the solar neutrino flux observed at Homestake", R. Davis Jr. *et al.*
proceedings of the 21st International Cosmic Ray Conference, Adelaide Australia, January 1990, p.143
Adelaide University (Adelaide), 1990 ISBN: 0-947-29807-X
- [66] "Search for neutrinos from the sun using the reaction $^{71}\text{Ga}(\nu_e, e^-)^{71}\text{Ge}$ "
A.I. Abazov *et al.* Phys. Rev. Lett., **67-24**, p.3332 (1991)
- [67] "Solar neutrinos observed by GALLEX at Gran Sasso"
GALLEX collaboration; P. Anselmann *et al.* Phys. Lett., **B285**, p.376 (1992)
- [68] "Results from one thousand days of real-time, directional solar-neutrino data"
K.S. Hirata *et al.* Phys. Rev. Lett., **65-11**, p.1297 (1990)
- [69] "Solar models, neutrino experiments, and helioseismology"
J.N. Bahcall and R. Ulrich Rev. Mod. Phys., **60**, p.297 (1988)
- [70] "Constraints on neutrino-oscillation parameters from the Kamiokande-II solar-neutrino data"
K.S. Hirata *et al.* Phys. Rev. Lett., **65-11**, p.1301 (1990)
- [71] "Implications of the GALLEX determination of the solar neutrino flux"
GALLEX collaboration; P. Anselmann *et al.* Phys. Lett., **B285**, p.390 (1992)
- [72] "Where do we stand with solar neutrino oscillations?"
J.N. Bahcall *et al.* Phys. Rev., **D58**, p.096016 (1998)
- [73] "Über Beobachtungen der durchdringen Strahlung bei sieben Freiballonfahrten"
V.F. Hess Phys. Zeit., **13**, p.1084 (1912)
- [74] "End to the cosmic-ray spectrum?"
K. Greisen Phys. Rev. Lett., **16**, p.748 (1966)
- [75] "Upper limit of the spectrum of cosmic rays"
G.T. Zatsepin and V.A. Kuz'min JETP Lett., **4**, p.78 (1966)
- [76] "Experimental study of the atmospheric neutrino flux"
K.S. Hirata *et al.* Phys. Lett., **B205**, p.416 (1988)
- [77] "Observation of a small atmospheric ν_μ/ν_e ratio in Kamiokande"
K.S. Hirata *et al.* Phys. Lett., **B280**, p.146 (1992)
- [78] "Calculation of atmospheric neutrino induced backgrounds in a nucleon decay search"
T.J. Haines *et al.* Phys. Rev. Lett., **57-16**, p.1986 (1986)
- [79] "Measurement of atmospheric neutrino composition with the IMB-3 detector"
D. Casper *et al.* Phys. Rev. Lett., **66-20**, p.2561 (1991)
- [80] "Electron- and muon-neutrino content of the atmospheric flux"
R. Becker-Szendy *et al.* Phys. Rev., **D46-9**, p.3720 (1992)
- [81] "Experimental study of atmospheric neutrino flux in the NUSEX experiment"
NUSEX collaboration; M. Aglietta *et al.*, Europhys. Lett., **8**, p.611 (1989)

- [82] "Experimental study of upward stopping muons in NUSEX"
M. Aglietta *et al.*, *Europhys. Lett.*, **15**, p.559 (1991)
- [83] "A study of atmospheric neutrino oscillations in the Fréjus experiment"
Fréjus collaboration; Ch. Berger *et al.* *Phys. Lett.*, **B245**, p.305 (1990)
- [84] "Determination of the atmospheric neutrino spectra with the Fréjus detector"
Fréjus collaboration; K. Daum *et al.* *Zeit. für Phys.*, **C66**, p.417 (1995)
- [85] "Search for muon neutrino oscillations with the Irvine-Michigan-Brookhaven detector"
R. Becker-Szendy *et al.* *Phys. Rev. Lett.*, **69-7**, p.1010 (1992)
- [86] "Evidence for oscillation of atmospheric neutrinos"
Super-Kamiokande coll.; Y. Fukuda *et al.* *Phys. Rev. Lett.*, **81-8**, p.1562 (1998)
- [87] "Light neutrinos as cosmological dark matter. A crucial experimental test"
H. Harari *Phys. Lett.*, **B216**, p.413 (1989)
- [88] NOMAD homepage: <http://nomad-info.web.cern.ch/nomad-info/>
- [89] "Final NOMAD results on $\nu_\mu \rightarrow \nu_\tau$ and $\nu_e \rightarrow \nu_\tau$ oscillations including a new search for ν_τ appearance using hadronic τ decays"
NOMAD collaboration; P. Astier *et al.* *Nucl. Phys.*, **B611**, p.3 (2001)
- [90] "New results from a search for $\nu_\mu \rightarrow \nu_\tau$ and $\nu_e \rightarrow \nu_\tau$ oscillation"
CHORUS collaboration; E. Eskut *et al.* *Phys. Lett.*, **B497**, p.8 (2001)
- [91] "Present status of KamLAND"
kamLAND collaboration; A. Suzuki *Nucl. Phys. B(proc. suppl.)*, **77**, p.171 (1999)
Proceedings of the 18th international conference on neutrino physics and astrophysics, Takayama, Japan, June 1998
- [92] "Experimental investigation of geologically produced antineutrinos with KamLAND"
T. Araki *et al.* *Nature*, **436**, p.499 (2005)
- [93] "The Sudbury Neutrino Observatory"
SNO collaboration; J. Boger *et al.* *Nucl. Inst. & Meth.*, **A449**, p.172 (2000)
- [94] "The Super-Kamiokande detector"
Super-Kamiokande coll.; S. Fukuda *et al.* *Nucl. Inst. & Meth.*, **A501**, p.418 (2003)
- [95] "What do we (not) know theoretically about solar neutrino fluxes?"
J.N. Bahcall and M.H. Pinsonneault *Phys. Rev. Lett.*, **92**, p.121301 (2004)
- [96] "The atmospheric neutrino flavor ratio from a 3.9 fiducial kiloton-year exposure of Soudan 2"
W.W.M. Allison *et al.* *Phys. Lett.*, **B449**, p.137 (1999)
- [97] "Measurement of the atmospheric neutrino flavour composition in Soudan 2"
W.W.M. Allison *et al.* *Phys. Lett.*, **B391**, p.491 (1997)
- [98] "Atmospheric muon neutrino fraction above 1 GeV"
R. Clark *et al.* *Phys. Rev. Lett.*, **79-3**, p.345 (1997)
- [99] "MACRO results on atmospheric neutrinos"
MACRO coll.; G. Giacomelli *et al.* *Nucl. Phys. B(proc. suppl.)*, **145**, p.116 (2005)
Proceedings of the Neutrino Oscillation Workshop, Otranto, Italy, September 2004
- [100] "First observations of separated atmospheric ν_μ and $\bar{\nu}_\mu$ events in the MINOS detector"
MINOS collaboration; P. Adamson *et al.* *Phys. Rev.*, **D73**, p.072002 (2006)
- [101] "Measurement of atmospheric neutrino oscillation parameters by Super-Kamiokande I"
Super-Kamiokande coll.; Y. Ashie *et al.* *Phys. Rev.*, **D71**, p.112005 (2005)
- [102] "Evidence for an oscillatory signature in atmospheric neutrino oscillation"
Super-Kamiokande coll.; Y. Ashie *et al.* *Phys. Rev. Lett.*, **93**, p.101801 (2004)

- [103] "Neutrino decay as an explanation of atmospheric neutrino observations"
V. Barger *et al.* Phys. Rev. Lett., **82**-13, p.2640 (1999)
- [104] "Probing possible decoherence effects in atmospheric neutrino oscillations"
E. Lisi *et al.* Phys. Rev. Lett., **85**-6, p.1166 (2000)
- [105] "Tau neutrinos favored over sterile neutrinos in atmospheric muon neutrino oscillations"
Super-Kamiokande coll.; S. Fukuda *et al.* Phys. Rev. Lett., **85**-19, p.3999 (2000)
- [106] "Measurement of the solar electron neutrino flux with the Homestake chlorine detector"
B.T. Cleveland *et al.* Astrophys. J., **496**, p.505 (1998)
- [107] "Measurement of the solar neutrino capture rate by the Russian-American Gallium solar neutrino experiment during one half of the 22-year cycle of solar activity"
SAGE collaboration; J.N. Abdurashitov *et al.* JETP, **95**, p.181 (2002)
- [108] "GALLEX solar neutrino observations: results for GALLEX IV"
GALLEX collaboration; W. Hampel *et al.* Phys. Lett., **B447**, p.127 (1999)
- [109] "Complete results for five years of GNO solar neutrino observations"
GNO collaboration; M. Altmann *et al.* Phys. Lett., **B616**, p.174 (2005)
- [110] "Final results of the ^{51}Cr neutrino source experiments in GALLEX"
GALLEX collaboration; W. Hampel *et al.* Phys. Lett., **B420**, p.114 (1998)
- [111] "Verification tests of the GALLEX solar neutrino detector, with ^{71}Ge produced in-situ from the beta-decay of ^{71}As "
GALLEX collaboration; W. Hampel *et al.* Phys. Lett., **B436**, p.158 (1998)
- [112] "Measurement of the response of a Gallium metal solar neutrino experiment to neutrinos from a ^{51}Cr source"
SAGE collaboration; J.N. Abdurashitov *et al.* Phys. Rev., **C59**-4, p.2246 (1999)
- [113] "Measurement of the response of a Ga solar neutrino experiment to neutrinos from an ^{37}Ar source" ; J.N. Abdurashitov *et al.* Phys. Rev., **C73**, p.045805 (2006)
- [114] "Measurement of the total active ^8B solar neutrino flux at the Sudbury neutrino observatory with enhanced neutral current sensitivity"
SNO collaboration; S.N. Ahmed *et al.* Phys. Rev. Lett., **92**, p.181301 (2004)
- [115] "Electron energy spectra, fluxes, and day-night asymmetries of ^8B solar neutrinos from measurements with NaCl dissolved in the heavy-water detector at the Sudbury neutrino observatory"
SNO collaboration; B. Aharmim *et al.* Phys. Rev., **C72**, p.055502 (2005)
- [116] "Solar neutrino measurements in Super-Kamiokande-I"
Super-Kamiokande coll.; J. Hosaka *et al.* Phys. Rev., **D73**, p.112001 (2006)
- [117] "Measurement of neutrino oscillation with KamLAND: Evidence of spectral distortion"
kamLAND collaboration; T. Araki *et al.* Phys. Rev. Lett., **94**, p.081801 (2005)
- [118] "First results from KamLAND: Evidence for reactor antineutrino disappearance"
kamLAND collaboration; K. Eguchi *et al.* Phys. Rev. Lett., **90**, p.021802 (2003)
- [119] "New solar opacities, abundances, helioseismology, and neutrino fluxes"
J.N. Bahcall *et al.* Astrophys. J., **621**, p.L85 (2005)
- [120] "Surprising sun: A new step towards a complete picture?"
S. Turck-Chièze *et al.* Phys. Rev. Lett., **93**, p.211102 (2004)
- [121] "Detection of accelerator-produced neutrinos at a distance of 250 km"
S.H. Ahn *et al.* Phys. Lett., **B511**, p.178 (2001)

- [122] "Evidence for muon neutrino oscillation in an accelerator-based experiment"
K2K collaboration; E. Aliu *et al.* Phys. Rev. Lett., **94**, p.081802 (2005)
- [123] "Observation of muon neutrino disappearance with the MINOS detectors in the NuMI neutrino beam" ; MINOS coll.; D.G. Michael Phys. Rev. Lett., **97**, p.191801 (2006)
- [124] OPERA homepage: <http://operaweb.web.cern.ch/>
- [125] T2K homepage: <http://neutrino.kek.jp/jhfnu/>
- [126] "Evidence for neutrino oscillations from muon decay at rest"
LSND collaboration; C. Athanassopoulos *et al.* Phys. Rev., **C54-5**, p.2685 (1996)
- [127] "Results on $\nu_\mu \rightarrow \nu_e$ neutrino oscillations from the LSND experiment"
LSND coll.; C. Athanassopoulos *et al.* Phys. Rev. Lett., **81-9**, p.1774 (1998)
- [128] "Evidence for neutrino oscillations from the observation of $\bar{\nu}_e$ appearance in a $\bar{\nu}_\mu$ beam"
LSND collaboration; A. Aguilar *et al.* Phys. Rev., **D64**, p.112007 (2001)
- [129] "Upper limits for neutrino oscillations $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ from muon decay at rest"
KARMEN collaboration; B. Armbruster *et al.* Phys. Rev., **D65**, p.112001 (2002)
- [130] "Statistical analysis of different $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ searches"
E.D. Church *et al.* Phys. Rev., **D66**, p.013001 (2002)
- [131] BooNE homepage: <http://www-boone.fnal.gov/>
- [132] "Observation of a neutrino burst from the supernova SN1987A"
K. Hirata *et al.* Phys. Rev. Lett., **58-14**, p.1490 (1987)
- [133] "Observation of a neutrino burst in coincidence with supernova 1987A in the Large Magellanic Cloud" ; R.M. Bionta *et al.* Phys. Rev. Lett., **58-14**, p.1494 (1987)
- [134] "Observation in the Kamiokande-II detector of the neutrino burst from supernova SN1987A"
K.S. Hirata *et al.* Phys. Rev., **D38-2**, p.448 (1988)
- [135] "Angular distribution of events from SN1987A"
IMB collaboration; C.B. Bratton *et al.* Phys. Rev., **D37-12**, p.3361 (1988)
- [136] Baikal homepage: <http://www.ifh.de/baikal/baikalhome.html>
- [137] ANTARES homepage: <http://antares.in2p3.fr/>
- [138] NESTOR homepage: <http://www.nestor.org.gr/>
- [139] NEMO homepage: http://nemoweb.lns.infn.it/index_1.htm
- [140] AMANDA homepage: <http://amanda.uci.edu/>
- [141] IceCube homepage: <http://icecube.wisc.edu/>
- [142] RICE homepage: <http://www.bartol.udel.edu/~spiczak/rice/rice.html>
- [143] ANITA homepage: <http://www.phys.hawaii.edu/~anita/>
- [144] GLUE homepage: <http://www.physics.ucla.edu/~moonemp/public>
- [145] "Status of global fits to neutrino oscillations"
M. Maltoni *et al.* New J. of Phys., **6**, p.122 (2004)
- [146] "Global analysis of neutrino data"
M.C. Gonzalez-Garcia Phys. Scripta, **T121**, p.72 (2005)
- [147] "Limits on neutrino oscillations from the CHOOZ experiment"
M. Apollonio *et al.* Phys. Lett., **B466**, p.415 (1999)
- [148] "Models of neutrino masses and mixings"
G. Altarelli and F. Feruglio New J. of Phys., **6**, p.106 (2004)

- [149] "Solar neutrinos and 1-3 leptonic mixing"
S. Goswami and A.Yu. Smirnov Phys. Rev., **D72**, p.053011 (2005)
- [150] "Three-flavour effects and CP- and T-violation in neutrino oscillations"
E.Kh. Akhmedov Phys. Scripta, **T121**, p.65 (2005)
- [151] "Future precision neutrino oscillation experiments and theoretical implications"
M. Lindner Phys. Scripta, **T121**, p.78 (2005)
- [152] "Neutrino beams from muon storage rings: Characteristics and physics potential"
S. Geer Phys. Rev., **D59**, p.6989 (1998)
Erratum: Phys. Rev., **D59**, p.039903 (1999)
- [153] "A novel concept for a $\bar{\nu}_e/\nu_e$ neutrino factory: the beta-beam"
P. Zucchelli Phys. Lett., **B532**, p.166 (2002)
- [154] KATRIN homepage: <http://www-ik.fzk.de/~katrin/>
- [155] "What can we learn from neutrinoless double beta decay experiments"
J.N. Bahcall *et al.* Phys. Rev., **D70**, p.033012 (2004)
- [156] "The CHORUS experiment to search for $\nu_\mu \rightarrow \nu_\tau$ oscillation"
CHORUS collaboration; E. Eskut *et al.* Nucl. Inst. & Meth., **A401**, p.7 (1997)
- [157] "Scintillating fiber trackers with optoelectronic readout for the CHORUS neutrino experiment"
S. Aoki *et al.* Nucl. Inst. & Meth., **A344**, p.143 (1994)
- [158] "Performance and calibration of the CHORUS scintillating fiber tracker and opto-electronics readout system" ; P. Annis *et al.* Nucl. Inst. & Meth., **A367**, p.367 (1995)
- [159] "The trigger system of the CHORUS experiment"
M.G. van Beuzekom *et al.* Nucl. Inst. & Meth., **A427**, p.587 (1999)
- [160] "The data acquisition system of the CHORUS experiment"
A. Artamonov *et al.* Nucl. Inst. & Meth., **A479**, p.412 (2002)
- [161] "The hexagonal toroidal air-core magnet of the CHORUS detector"
F. Bergsma *et al.* Nucl. Inst. & Meth., **A357**, p.243 (1995)
- [162] "The large area emulsion tracker of the CHORUS experiment"
S. Aoki *et al.* Nucl. Inst. & Meth., **A488**, p.144 (2002)
- [163] "Construction and test of calorimeter modules for the CHORUS experiment"
S. Buontempo *et al.* Nucl. Inst. & Meth., **A349**, p.70 (1994)
- [164] "Response to electrons and pions of the calorimeter for the CHORUS experiment"
E. Di Capua *et al.* Nucl. Inst. & Meth., **A378**, p.221 (1996)
- [165] "Electron, pion and multiparticle detection with a lead/scintillating-fiber calorimeter"
D. Acosta *et al.* Nucl. Inst. & Meth., **A308**, p.481 (1991)
- [166] "A new search for $\nu_\mu-\nu_\tau$ oscillation"
CHORUS collaboration; N. Armenise *et al.*, CERN report: SPSC/90-42
- [167] "A new search for $\nu_\mu-\nu_\tau$ oscillation"
CHORUS collaboration; M. de Jong *et al.*, CERN report: PPE/93-131
- [168] "Limits to $\nu_\mu, \nu_e \rightarrow \nu_\tau$ oscillations and $\nu_\mu, \nu_e \rightarrow \tau^-$ direct coupling"
E531 collaboration; N. Ushida *et al.* Phys. Rev. Lett., **57-23**, p.2897 (1986)
- [169] "Detector techniques for the application of topological and kinematical criteria for a $\nu_\mu \rightarrow \nu_\tau$ oscillation search"
I.M. Papadopoulos; University of Athens, Greece, 2001 Ph.D. Thesis

- [170] "The CHARON experiment: measurement of the near-elastic interaction mean free path of pions in nuclear emulsion"
A. Buelte; Humboldt University of Berlin, Germany, 2001 Ph.D. Thesis
- [171] "Prompt ν_τ background in wide band ν_μ beams"
B. Van de Vyver Nucl. Inst. & Meth., **A385**, p.91 (1997)
- [172] "The West area neutrino facility for CHORUS and NOMAD experiments (94-97 operation)"
G. Acquistapace *et al.*, CERN report: ECP/95-014
- [173] "Improved collection of secondaries from a long, narrow target by a horn doublet"
S. van der Meer, CERN report: PS/AA/80-12
- [174] "Measurement of the beam flux from chambercard neutrino interaction in the CHORUS calorimeter" ; R.G.C. Oldeman CHORUS internal note nr. 98001
<http://choruswww.cern.ch/Internals/Notes/archive/98001/beamnote.ps.gz>
- [175] "Measurement of differential cross-sections and structure functions in neutrino and anti-neutrino scattering on lead"
R.G.C. Oldeman; Universiteit van Amsterdam, Netherlands, 2000 Ph.D. Thesis
- [176] "Measurement of charged particle production from 450 GeV/c protons on beryllium"
NA56/SPY collaboration; G. Ambrosini *et al.* Euro. Phys. J., **C10**, p.605 (1999)
- [177] "Observations with electron-sensitive plates exposed to cosmic radiation"
R. Brown *et al.* Nature, **163**, p.47 (1949)
- [178] "Measurement of charged particle multiplicities in high energy neutrino- and anti-neutrino-nucleus interactions"
CHORUS collaboration; G. Öngüt *et al.* submitted to Euro. Phys. J. C (2007)
- [179] "Test results of the streamer-tube system of the CHARM II neutrino detector"
CHARMII collaboration; J.P. De Wulf *et al.* Nucl. Inst. & Meth., **A252**, p.443 (1986)
- [180] "Calorimetry: energy measurement in particle physics"
R. Wigmans; Clarendon press (Oxford), 2000 ISBN: 0-198-50296-6
- [181] "The response and resolution of an iron-scintillator calorimeter for hadronic and electromagnetic showers between 10 GeV and 140 GeV"
H. Abramowicz *et al.* Nucl. Inst. & Meth., **180**, p.429 (1981)
- [182] "Large planar drift chambers"
G. Marel *et al.* Nucl. Inst. & Meth., **141**, p.43 (1977)
- [183] "Muon momentum from range in the spectrometer"
G. Heyboer CHORUS internal note nr. 96011
<http://choruswww.cern.ch/Internals/Notes/archive/96011/cammor.ps.gz>
- [184] Tcl/Tk homepage: <http://tcl.sourceforge.net/>
- [185] TiX homepage: <http://tixlibrary.sourceforge.net/>
- [186] "CHORUS offline software" online documentation
<http://choruswww.cern.ch/Internals/Offline/software.html>
- [187] "Nuclear disintegration by meson capture"
D.H. Perkins Nature, **159**, p.126 (1947)
- [188] "Observations on the tracks of slow mesons in photographic emulsions"
C.G.M. Lattes *et al.* Nature, **160**, p.453 (1947)
- [189] "Evidence concerning the existence of the new unstable elementary neutral particle"
V.D. Hopper and S. Biswas Phys. Rev., **80**, p.1099 (1950)

- [190] "Anti-Lambda hyperon"
D.J. Prowse and M. Baldo-Ceolin Phys. Rev. Lett., **1**, p.179 (1958)
- [191] "Observation of the decay at rest of a heavy particle"
A. Bonetti *et al.* Nuovo Cimento, **10**, p.345 (1953)
- [192] "On the existence of unstable charged particles of hyperprotonic mass"
A. Bonetti *et al.* Nuovo Cimento, **10**, p.1736 (1953)
- [193] "Possible existence of a new hyperon"
Y. Eisenberg Phys. Rev., **96**, p.541 (1954)
- [194] "Observations with electron-sensitive plates exposed to cosmic radiation, part II"
R. Brown *et al.* Nature, **163**, p.82 (1949)
- [195] "Masses and modes of decay of heavy mesons. Part I. κ -particles"
C. O'Ceallaigh Phil. Mag., **42-II**, p.1032 (1951)
- [196] "Masses and modes of decay of heavy mesons. Part II. τ -particles"
P.H. Fowler *et al.* Phil. Mag., **42-II**, p.1040 (1951)
- [197] "The masses of positive κ -particles"
H.H. Heckman *et al.* Nuovo Cimento, **3**, p.85 (1956)
- [198] "Protons in the earth's magnetic field"
S.C. Freden and R.S. White Phys. Rev. Lett., **3**, p.9 (1959)
- [199] "The use of nuclear emulsion in association with spark chambers"
B.G. Duff *et al.* Nucl. Inst. & Meth., **29**, p.351 (1964)
- [200] "The location and analysis of neutrino interactions in photographic emulsion"
E.H.S. Burhop *et al.* Nuovo Cimento, **39**, p.1037 (1965)
- [201] "Investigation of the decay of charmed particles produced in neutrino interactions"
D. Allasia *et al.* Nucl. Phys., **B176**, p.13 (1980)
- [202] "Experimental details on lifetime measurements of neutrino-produced charmed particles
in a tagged emulsion spectrometer"
N. Ushida *et al.* Nucl. Inst. & Meth., **A224**, p.50 (1984)
- [203] "Fully automated emulsion analysis system"
S. Aoki *et al.* Nucl. Inst. & Meth., **B51**, p.466 (1990)
- [204] "Hybrid emulsion spectrometer for the detection of hadronically produced heavy flavor states"
K. Kodama *et al.* Nucl. Inst. & Meth., **A289**, p.146 (1990)
- [205] "Charm meson production in 800 GeV/c proton-emulsion interactions"
E653 collaboration; K. Kodama *et al.* Phys. Lett., **B263**, p.573 (1991)
- [206] "Measurement of beauty hadron pair production in 600 GeV/c π^- emulsion interactions"
E653 collaboration; K. Kodama *et al.* Phys. Lett., **B303**, p.359 (1993)
- [207] "Neutrino oscillations in emulsion experiment"
K. Niwa; in "Physics and astrophysics of neutrinos", p.520
M. Fukugita and A. Suzuki (eds); Springer (Berlin), 1994 ISBN: 0-387-70136-2
- [208] "Nuclear emulsions in a large, hybrid experiment (CHORUS) to search for $\nu_\mu \rightarrow \nu_\tau$ oscillations"
; S. Aoki *et al.* Nucl. Inst. & Meth., **A447**, p.361 (2000)
- [209] "Nuclear research emulsions. I. Techniques and theory."
W.H. Barkas; Academic Press (New York), 1963
- [210] "Nuclear research emulsions. II. Particle behaviour and emulsion applications."
W.H. Barkas; Academic Press (New York), 1973 ISBN: 0-12-078302-9

- [211] "Observation of neutrino induced diffractive D_s^{*+} production and subsequent decay
 $D_s^{*+} \rightarrow D_s^+ \rightarrow \tau^+ \rightarrow \mu^+$ "
 CHORUS collaboration; P. Annis *et al.* Phys. Lett., **B435**, p.458 (1998)
- [212] "The determination of linear distortion in nuclear emulsion"
 A.J. Apostolakis and J.V. Major British J. of Appl. Phys., **8**, p.9 (1957)
- [213] "Automatic analysis of nuclear emulsion" (in Japanese)
 T. Nakano; Nagoya University, Japan, 1997 Ph.D. Thesis
- [214] " D^0 production rate measurement in neutrino interactions and a limit on muon neutrino
 to tau neutrino oscillations"
 M. Güler; Middle East Technical University, Turkey, 2000 Ph.D. Thesis
- [215] "Determination of the semi-leptonic branching fraction of neutrino-induced charm
 hadrons using nuclear emulsion"
 B. Van de Vyver; Vrije Universiteit Brussel, Belgium, 2002 Ph.D. Thesis
- [216] "Development of the Netscan analysis technique for nuclear emulsion and study of
 neutrino interactions" (in Japanese)
 N. Nonaka; Nagoya University, Japan, 2002 Ph.D. Thesis
- [217] "The honeycomb strip chamber: the application in LHC/SSC experiments and the test results
 of a prototype" ; H. van der Graaf *et al.* Nucl. Inst. & Meth., **A307**, p.220 (1991)
- [218] "The honeycomb strip chamber: a two coordinate and high precision muon detector"
 H. Tolsma; Universiteit Twente, Netherlands, 1996 Ph.D. Thesis
- [219] "Multiwire and drift proportional chambers"
 G. Charpak Phys. Today, p.23, October (1978)
- [220] "Techniques for nuclear and particle physics experiments: a how-to approach (2nd edition)"
 W. Leo; Springer (Berlin), 1994 ISBN: 0-387-57280-5
- [221] "Particle detection with drift chambers"
 W. Blum and L. Rolandi; Springer (Berlin), 1994 ISBN: 3-540-58322-X
- [222] "Investigation of self-quenching streamer discharge in a wire chamber"
 G.D. Alekseev *et al.* Nucl. Inst. & Meth., **177**, p.385 (1980)
- [223] "Measurements on the 3 meter by 8 cells NIKHEF honeycomb prototype"
 J.W.E. Uiterwijk and R.G.C. Oldeman NIKHEF/CHORUS internal note
- [224] "The logic design of high precision time-to-pulse-height converters"
 M. Bertolaccini and S. Cova Nucl. Inst. & Meth., **121**, p.547 (1974)
- [225] "An integrated CMOS 0.15 ns digital timing generator for TDCs and clock distribution systems"
 ; J. Christiansen, CERN report: ECP/94-25
- [226] "Drift chamber electronics for time and pulse height measurements with multiple hit capacity"
 W. Farr and J. Heintze Nucl. Inst. & Meth., **156**, p.301 (1978)
- [227] "Readout of drift chambers with a 100 MHz flash ADC system"
 W. Farr *et al.* IEEE Trans. on Nucl. Sc., **NS-30**, p.95 (1983)
- [228] "Multiple time digitizers and a trigger system for drift chambers"
 K. Eggert *et al.* Nucl. Inst. & Meth., **176**, p.223 (1980)
- [229] "De voorversterker/discriminator voor Chorus"
 J.D. Schipper NIKHEF internal note
- [230] "Honeycomb tracking algorithm"
 R.G.C. Oldeman private communications

- [231] "Garfield - simulation of gaseous detectors" online documentation
<http://consult.cern.ch/writeup/garfield/>
- [232] "Report on momentum determination"
 B. Saitta CHORUS meeting notes September 2001
<http://choruswww.cern.ch/Internals/Collab-Meetings/200109-CM/>
- [233] "A C++ object-oriented toolkit for track finding with k -dimensional hits"
 J.W.E. Uiterwijk *et al.* Nucl. Inst. & Meth., **A560**, p.317 (2006)
- [234] "On microscope images, filtering and compression"
 P. Zucchelli CHORUS internal note nr. 97027
<http://choruswww.cern.ch/Internals/Notes/archive/97027/note.ps.gz>
- [235] JENOPTIK homepage: <http://www.jenoptik.com/>
- [236] "Optical microscope for nuclear emulsion readout - system design and results in application"
 K. Winkler *et al.*; SPIE 44th annual meeting, Denver, Colorado, July 1999;
 R.E. Fischer, W.J. Smith (eds); SPIE (Bellingham), 1999 ISBN: 0-819-43265-2
- [237] "Introduction to algorithms"
 Th. Cormen *et al.*; MIT press (Cambridge, MA), 1990 ISBN: 0-262-53091-0
- [238] "Proposal to study hadron production for the neutrino factory and for the atmospheric neutrino flux" ; HARP collaboration; M.G. Catanesi *et al.*, CERN report: SPCS/99-35
- [239] HARP homepage: <http://harp.web.cern.ch/>
- [240] "Spatial tessellations: concepts and applications of Voronoi diagrams"
 A. Okabe *et al.*; Wiley (New York), 1992 ISBN: 0-471-98635-6
- [241] "Centroidal Voronoi tessellations: applications and algorithms"
 Q. Du *et al.* SIAM Rev., **41** 4, p.637 (1999)
- [242] "Voronoi tessellations " online reference
<http://www.voronoi.com>
- [243] "Handbook of discrete and computational geometry", J.E. Goodman, J. O'Rourke (eds)
 CRC Press (Boca Raton), 1997 ISBN: 1-420-03531-2
- [244] "The art of computer programming; volume III, sorting and searching (second edition)"
 D.E. Knuth; Addison-Wesley (Reading), 1998 ISBN: 0-201-89685-0
- [245] "Handbook of computational geometry"
 J.R. Sack and J. Urrutia (eds); Elsevier (New York), 2000 ISBN: 0-444-82537-1
- [246] "C++ programming language: third edition"
 B. Stroustrup; Addison-Wesley (Reading), 1997 ISBN: 0-201-88954-4
- [247] "STL tutorial and reference guide: C++ programming with the Standard Template Library"
 D.R. Musser and A. Saini; Addison-Wesley (Reading), 1996 ISBN: 0-201-63398-1
- [248] "Active libraries: rethinking the roles of compilers and libraries"
 T.L. Veldhuizen and D. Gannon; proceedings of the SIAM Workshop on Object Oriented Methods for Inter-operable Scientific and Engineering Computing (OO'98), New York, October 1998
 SIAM Press (Philadelphia), 1999
- [249] "Techniques for scientific C++", T.L. Veldhuizen
 Indiana University Computer Science Technical Report: #542 (2000)
<http://www.cs.indiana.edu/cgi-bin/techreports/TRNNN.cgi?trnum=TR542>

- [250] "BLITZ library" , T.L. Veldhuizen online reference
<http://osl.iu.edu/~tveldhui/papers/techniques/>
<http://www.oonumerics.org/blitz/>
- [251] "Integer sequences (Entries: A000073 and A000078)" , N.J.A. Sloane online reference
<http://www.research.att.com/~njas/sequences/>
- [252] "The Mathematica book, fourth edition"
S. Wolfram; Wolfram media (Cambridge), 1999 ISBN: 0-521-64314-7
- [253] "ControlHost: Package for distributed data handling", R. Gurin and A. Maslennikov
talk given at the HEPiX95 Meeting, Prague, June 1995
<http://www-hep2.fzu.cz/computing/HEPiX/HEPiX95/talks/conhost.ps>
- [254] "New results in linear filtering and prediction theory"
R.E. Kalman and R.S. Bucy J. of Basic Eng. (ASME), **83D**, p.95 (1961)
- [255] "Application of Kalman filtering to track and vertex fitting"
R. Frühwirth Nucl. Inst. & Meth., **A262**, p.444 (1987)
- [256] "Measurement of D^{*+} production in charged-current neutrino interactions"
CHORUS collaboration; G. Önençüt *et al.* Phys. Lett., **B614**, p.155 (2005)
- [257] "Study of D^{*+} and search for D^{*0} production by neutrinos in BEBC"
BEBC neutrino coll.; A.E. Asratyan *et al.* Zeit. für Phys., **C68**, p.43 (1995)
- [258] "Production of D^{*+} (2010) mesons by high energy neutrinos from the Tevatron"
E632 collaboration; A.E. Asratyan *et al.* Zeit. für Phys., **C76**, p.647 (1997)
- [259] "Study of D^{*+} production in ν_{μ} charged current interactions in the NOMAD experiment"
NOMAD collaboration; P. Astier *et al.* Phys. Lett., **B526**, p.278 (2002)
- [260] "Cross sections for neutrino production of charmed particles"
E531 collaboration; N. Ushida *et al.* Phys. Lett., **B206**, p.375 (1988)
- [261] "Production characteristics of charmed particles in neutrino interactions"
E531 collaboration; N. Ushida *et al.* Phys. Lett., **B206**, p.380 (1988)
- [262] "Determining the CKM parameter V_{cd} from νN charm production"
T. Bolton, CERN preprint hep-ex/9708014 (1997)
- [263] "Measurements of D^0 production and of decay branching fractions in neutrino-nucleon scattering" ; CHORUS collaboration; G. Önençüt *et al.* Phys. Lett., **B613**, p.105 (2005)
- [264] "GEANT 3.21: detector description and simulation tool" online documentation
<http://consult.cern.ch/writeup/geant/>
- [265] "Calorimetric techniques for the kinematical selection of events in CHORUS"
P. Zucchelli; University of Ferrara, Italy, 1995 Ph.D. Thesis
- [266] "LEPTO generator"
G. Ingelman, Uppsala University (Sweden) preprint TSL/ISV 92-0065 (1992)
- [267] "High energy physics event generation with PYTHIA 5.7 and JETSET 7.4"
T. Sjöstrand Comp. Phys. Comm., **82**, p.74 (1994)
- [268] "Background yield for charm searches from strange particles decays"
M. Sorrentino CHORUS internal note nr. 2000027
http://choruswww.cern.ch/Publications/Notes/charm_background.pdf
- [269] "Study of charm photoproduction mechanisms"
NA14/2 Collaboration; M.P. Alvarez *et al.* Zeit. für Phys., **C60**, p.53 (1993)
- [270] "Prediction of charm-production fractions in neutrino interactions"
G. De Lellis *et al.* Phys. Lett., **B550**, p.16 (2002)

Summary

In this dissertation, several different topics related to the CHORUS experiment are presented. The CHORUS experiment has been used to study neutrino oscillations using the neutrino beam at CERN. The neutrino oscillation hypothesis provided an explanation for the lower than expected fluxes of solar and atmospheric neutrinos. There are three neutrino species in nature corresponding to different weak eigenstates, namely, the electron neutrino (ν_e), the muon neutrino (ν_μ), and the tau neutrino (ν_τ). The lower fluxes could be interpreted as spontaneous oscillations between electron and muon neutrinos and between muon and tau neutrinos. The CHORUS experiment was designed to detect oscillation of muon neutrinos into tau neutrinos with small mixing probability down to $2 \cdot 10^{-4}$ and a mass difference square between ν_μ and ν_τ larger than 0.5 eV^2 .

In the last decade, several disappearance experiments have confirmed the neutrino oscillation hypothesis and showed that oscillations occur between mass eigenstates with mass differences of $8 \cdot 10^{-5} \text{ eV}^2$ and $2.5 \cdot 10^{-3} \text{ eV}^2$. The observed oscillations correspond to almost maximal mixing probability. Current and future efforts focus on detecting the remaining sub-dominant mixing.

In the CHORUS experiment, one looks for the typical decay signature of a τ particle in a large stack of nuclear emulsion plates. Emulsion can record with very high resolution the tracks and decays of short-lived particles, like τ leptons and charmed mesons. A hybrid setup is used consisting of electronic detectors downstream of a 770 kg emulsion target. The electronic detectors consist of an accurate tracking station followed by a hadron spectrometer, a calorimeter, and a muon spectrometer. These detectors are used to measure particle momenta and to select one track per event for scanning the emulsion. Hybrid detectors combined with automated microscopes have greatly increased the amount of emulsion that can be scanned and the information that can be extracted. The FKEN laboratory in Nagoya developed designated hardware that can find a track with known angle in a $100 \mu\text{m}$ slice of emulsion. Further upgraded new hardware can find all tracks with slopes within a given angular range. The hardware is used to locate the neutrino vertex by following an electronically measured track from one plate to the next. The new hardware is then used to reconstruct all tracks and vertices of the event. This is done by combining track segments found in the $100 \mu\text{m}$ upstream layers of the eight plates around the interaction point.

The suppression of background in the CHORUS experiment relies on the measurement of hadron momenta in the hadron spectrometer. A honeycomb drift-tube tracker was added to the detector to improve this measurement. This made it possible to reconstruct independently 3-D track segments in the hadron spectrometer. This improved the uncertainty on hadron momenta from $\Delta p/p = \sqrt{0.35^2 + 0.38p/\text{GeV}}$ to $\sqrt{0.35^2 + 0.25p/\text{GeV}}$, where the constant term is due to multiple scattering. The read-out of this detector is based on a continuous sampling of the over-threshold signal of all wires. Because all transitions of a wire's signal are recorded, both leading edge and pulse length can be recovered. The bit-stream of samples is stored in four phases in separate memory banks. Failure of some memory banks does not destroy the hit information but only lowers its timing accuracy.

A new automatic emulsion scanning system was developed at CERN. For this, new microscope optics with a large field of view were deployed and emulsion scanning was implemented in software. Using software for track finding in emulsion images requires an algorithm capable of handling of the order of 10^5 3-D hits in a few seconds. An algorithm was developed that searches a network of connected hits. The connection graph links hits in close proximity with a predefined restriction. In this case, the link should possibly be part of a track. This graph is created by sorting the hits in k -dimensional space using a specially designed ordered container with logarithmic search time. The track-finding algorithm considers any connection between two hits as a possible track and builds longer segments using a limited depth-first search of the connection graph. Although originally written for track finding in emulsion, the track-finding algorithm can work in any dimension and is decoupled from the actual hit and track model. The track model and the acceptance of hits in a track are handled by a specific implementation of an abstract acceptance class. The track-finding implementation has been used to find bent tracks in the HARP time-projection chamber. In the CHORUS emulsion application, it has enabled improvements in the vertex location efficiency. It can also be used to reduce human confirmation of secondary vertices by applying it to data taken over the full emulsion thickness.

An emulsion-specific algorithm was developed to look for tracks with a known direction. The algorithm is based on moving hits along the track direction and summing the number of hits in its angular acceptance region. Searching for hits in these acceptance regions can be done in logarithmic time using the aforementioned containers. The track-finding algorithm is then used to verify the existence of a track in the small volumes defined by a track-trigger sum over threshold. The k -D space ordering containers have been used throughout the work described in this dissertation, for example in several alignment procedures.

The emulsion volume, scanned around the neutrino vertex, defines a decay space for short-lived particles. The maximum decay length is about 4.5 mm along the beam direction. In the CHORUS experiment, 1048 charged-current ν_μ interactions, identified by a μ^- from the primary vertex, with a secondary decay vertex from a D^0 meson have been reconstructed. These events were used to study D^{*+} resonance production. The decay $D^{*+} \rightarrow D^0\pi^+$ can be identified using the π^+ momentum measured in the hadron spectrometer and the direction of both the D^0 and π^+ measured in the emulsion. From the observed 22.1 ± 5.5 events, the cross-section ratio for hadronization of a charm quark into D^{*+} and D^0 is derived: $\sigma(D^{*+})/\sigma(D^0) = 0.38 \pm 0.09(\text{stat}) \pm 0.05(\text{syst})$. The production cross-section of D^{*+} with respect to the total charged-current cross-section is $\sigma(D^{*+})/\sigma(\text{CC}) = [1.02 \pm 0.25(\text{stat}) \pm 0.15(\text{syst})] \%$.

Samenvatting

Dit proefschrift behandelt verschillende onderwerpen met als gemeenschappelijk thema het CHORUS experiment. CHORUS heeft gezocht naar $\nu_\mu \rightarrow \nu_\tau$ oscillatie in een neutrino-bundel op CERN. Neutrino oscillatie — mogelijk als neutrino's massa hebben — was een mogelijke verklaring voor de gemeten stromen van zonne- en atmosferische-neutrino's die kleiner waren dan de theoretische verwachtingen. Er bestaan drie neutrino's die overeenkomen met verschillende zwakke eigentoestanden, namelijk: het elektron neutrino (ν_e), het muon neutrino (ν_μ) en het tau neutrino (ν_τ). De gemeten lage stromen konden mogelijk verklaard worden door oscillatie tussen ν_e en ν_μ , en tussen ν_μ en ν_τ . CHORUS was specifiek ontworpen om oscillatie tussen ν_μ en ν_τ te onderzoeken met een menghoek groter dan $2 \cdot 10^{-4}$ en een massaverschil groter dan 0.5 eV^2 .

In het laatste decennium zijn de te kleine stromen van zonne- en atmosferische-neutrino's verklaard door verschillende experimenten die neutrino's zien verdwijnen. Deze experimenten hebben laten zien hoe neutrino's oscilleren met bijna maximale menghoeken en massaverschillen van $8 \cdot 10^{-5} \text{ eV}^2$ en $2.5 \cdot 10^{-3} \text{ eV}^2$. Huidige en nieuwe experimenten proberen nu de subdominante menghoek θ_{13} te meten.

CHORUS zoekt naar het typische signaal van een τ verval in een stapel emulsieplaten. Met emulsie is het mogelijk om de sporen en het verval van kortlevende deeltjes, zoals het τ lepton en charm mesonen, met hoge resolutie zichtbaar te maken. CHORUS heeft een hybride opzet met elektronische detectoren en emulsie. De elektronische detectoren zijn achtereenvolgens een zeer nauwkeurige spoorzoeker, een hadron spectrometer, een calorimeter en een muon spectrometer. Deze detectoren staan stroomafwaarts van een doel met 770 kg emulsie. Deze detectoren meten de impuls van de deeltjes en selecteren één spoor per evenement dat gevolgd zal worden in de emulsie. Hybride detectoren en de ontwikkeling van volledig geautomatiseerde microscopen hebben ertoe geleid dat veel meer gebeurtenissen onderzocht kunnen worden en dat veel meer informatie uit de emulsie gehaald kan worden. Het FKEN laboratorium in Nagoya heeft elektronische apparatuur ontwikkeld die heel snel een spoor met bekende hoek in een $100 \mu\text{m}$ dikke laag in de emulsie kan vinden. De apparatuur is later verder ontwikkeld om snel alle sporen met een gegeven maximale hoek te vinden. In CHORUS wordt deze apparatuur gebruikt om de oorsprong van de neutrino-interactie in de emulsie te vinden. Eén elektronisch gevonden spoor wordt stroomopwaarts van plaat naar plaat gevolgd tot de interactie positie. De nieuwe apparatuur wordt vervolgens gebruikt om alle sporen en vertices behorende bij het evenement te vinden door op acht platen rondom de neutrino-interactie $100 \mu\text{m}$ van het stroomopwaartse oppervlak gemeten. De gevonden spoorgedeeltes worden dan aan elkaar geknoopt en gecombineerd tot vertices.

De onderdrukking van achtergrond in CHORUS vereist een goede meting van de impulsen van hadronen in de hadron spectrometer. Een honinggraat driftbuizen detector is aan het experiment toegevoegd om deze meting te verbeteren. Deze detector was in staat om zelfstandig 3-D spoorelementen in de hadron spectrometer te reconstrueren. Hiermee is de impulsresolutie verbeterd van $\Delta p/p = \sqrt{0.35^2 + 0.38p/\text{GeV}}$ tot $\sqrt{0.35^2 + 0.25p/\text{GeV}}$. In de uitlezing worden bitreeksen voor alle 3888 draden in de detector opgeslagen in een geheugen. Elk bit in een reeks geeft aan of het signaal voor die draad beneden of boven

een ingestelde limiet is. Omdat elke overgang in het signaal van een draad gereconstrueerd kan worden, kan zowel de tijd als de lengte van elke puls gemeten worden. De bitreeks wordt opgeslagen in vier fases in verschillende geheugenbanken. Het uitvallen van één of meerdere geheugenbanken leidt daarom niet tot verlies van het signaal maar heeft alleen een slechtere tijdsresolutie tot gevolg.

In de ontwikkeling van scanning microscopen op CERN is nieuwe verbeterde optica met een groot beeldveld toegepast. Tevens is de apparatuur voor het zoeken van sporen vervangen door software. Spoorzoek software voor emulsie vereist een algoritme dat binnen enkele seconden in de orde van 10^5 3-D punten kan verwerken. Het ontwikkelde algoritme zoekt voor sporen in een graaf van gekoppelde punten. Punten die dicht bij elkaar in de buurt liggen en waarvan de lijn van de een naar de ander mogelijk tot een spoor behoort, zijn in dit connectienetwerk aan elkaar gekoppeld. Het netwerk wordt opgebouwd door de punten eerst te sorteren in een k -D ruimte met één van de speciaal ontwikkelde k -D containers. Deze containers zijn ontworpen om een logaritmische zoektijd voor punten in een opgevraagd volume te geven. Het spoorzoek-algoritme beschouwt elke koppeling tussen twee punten als een mogelijke start van een spoor en voegt dan meer punten toe door de verbindingen in het netwerk te volgen met een gelimiteerd diepte-eerst graafdoorzoek algoritme. Hoewel het algoritme ontworpen was voor het zoeken van sporen in emulsie, was het dusdanig algemeen dat het ontkoppeld kon worden van de precieze details van de punten en de sporen en met het aantal dimensies van de punten als een vrije parameter. Een model voor de verwachte sporen en toevoeging van punten aan een spoorsegment worden gedefinieerd door de implementatie van een abstracte klasse met alle acceptatie criteria. Dit spoorzoek-algoritme wordt gebruikt om gebogen sporen te vinden in de HARP tijdprojectie-detector. De toepassing in de CHORUS emulsie heeft geleid tot een hogere efficiëntie voor het lokaliseren van de interacties. Het aantal benodigde manuele operaties, voor het bevestigen van verval kandidaten, kan verminderd worden, indien het algoritme toegepast wordt op de volledige emulsielaag.

Een ander algoritme, specifiek voor emulsie, genaamd de 'track-trigger', is ontworpen om een track met een bekende hoek te vinden. In dit algoritme worden de punten verschoven langs de richting van het te zoeken spoor en wordt het aantal punten geteld dat binnen de hoekacceptatie ligt. De punten in de acceptatievolumes worden opgezocht met één van de k -D containers. Punten in de track-trigger met een voldoende hoge som definiëren kleine volumes. Het globale spoorzoek-algoritme wordt gebruikt om te verifiëren of het gezochte spoor in een dergelijk volume aanwezig is. Verschillende andere algoritmes, bijvoorbeeld voor het uitlijnen van emulsieplaten, maken ook gebruik van de k -D containers.

Een groot deel van de kortlevende deeltjes vervallen in het emulsievolume van ongeveer 4.5 mm lengte dat door net-scan onderzocht wordt. In totaal heeft CHORUS 1048 ν_μ geladen zwakke interacties gevonden (d.w.z. met een μ^- komende van de neutrino-interactie) met een tweede verval van een D^0 meson. Deze evenementen zijn gebruikt om de productie van D^{*+} resonanties in ν_μ interacties te bestuderen. Het verval $D^{*+} \rightarrow D^0 \pi^+$ wordt herkend doordat de π^+ een lage transversale impuls heeft ten opzichte van de richting van de D^0 . De verhouding van de werkzame doorsneden voor de hadronisatie van een c quark tot een D^{*+} of een D^0 kan berekend worden uit de 22.1 ± 5.5 gevonden evenementen en geeft $\sigma(D^{*+})/\sigma(D^0) = 0.38 \pm 0.09(\text{stat}) \pm 0.05(\text{syst})$. De totale werkzame doorsnede voor de productie van D^{*+} als fractie van de totale werkzame doorsnede voor geladen zwakke processen is $\sigma(D^{*+})/\sigma(\text{CC}) = [1.02 \pm 0.25(\text{stat}) \pm 0.15(\text{syst})] \%$.

Acknowledgments

Finally, it is done. From here on, physics mode is off, no more checking every sentence for its correctness, no more doubts if every statement made is defensible. From now on, just free-flow writing to thank the people that have been involved at any one stage in finishing this dissertation. Then again, it seems to be a habit that people try to read in between the lines, trying to find out what was really meant. Well, don't bother, there isn't anything there. Even though this whole project has taken much much longer than expected or necessary, the only person responsible for that is me.

Now, where to start? As said, it has taken forever to finish this thesis and many people must have been desperate over the years but still have been supportive when I still didn't want to give up and continued doing things my way. Special thanks therefore has to go Giovanna, Jaap, Joop, Rene, Maarten, Jan, Teus and my parents and parents in law. Giovanna, my wife, had to stand all my mood swings, desperations and endless delays. As she rightly remarked, I had to be really desperate to prefer cleaning the house over working on my thesis. Jaap always had good ideas on any subject and was the solid rock during the experiment. Maarten was so kind to become my thesis adviser and read the whole manuscript very carefully. Jan found me (more than one) thesis adviser and helped me, way earlier, getting started with the first chapter with many useful suggestions. Teus has been most helpful in getting me past, ever more vigilant, CERN security guards (among other things). Till this day though, I don't quite understand what they are guarding us from. As I had to explore other options, I found many ways of entering CERN. My parents-in-law kept me watered and fed during the stressful times that we were both very busy. My parents, have never given up hope that I would finish my PhD and finally get to see the result of all the support they have given me and Giovanna.

There is of course more to a PhD than writing the final dissertation. The most enjoyable part of this PhD project has been all the actual work before. I want to thank all those people I have worked with in CHORUS, both at CERN and at NIKHEF. Special mention should go though to several people; in more or less chronological order. Joop, who hired me — I hope he doesn't regret it — and who always brought a lighthearted spirit to work (plus one orange and one appel). Jan and Kees, who introduced me to CERN and its surroundings. Jan and Bernardine later took me sailing to places I had never seen in the Netherlands and made my time in Amsterdam more pleasant. However, I still refuse to get on a horse as long as they don't come installed with an emergency break. If a horse would be only half as stubborn as me, who know where we end up.

During the honeycomb project I have worked with much pleasure with both Martin and Peter to get the electronics working in record time. Not only are they very meticulous

electronic engineers, they also are very pleasant people to work with (and not only for the constant supply of liquorice). We also got to see the total solar eclipse through the one gap in the clouds.

In the microscope project, I have enjoyed immensely working with Ioannis, Bart and Piero. Piero is a never ending source of new and good ideas. Ioannis is always happy to join software design bashing. Bart has dug up lots of interesting material and could always be found for helpful discussions. I've learned from Bart that there is usually a mathematical correct way of doing any analysis and had many discussions on whatever subject one can think off.

With Maarten de Jong, Maarten Litmaath, Rolf and Oliver, I had overlap in several projects during our time at CERN; it was always a pleasure to work together. With Rolf I had many philosophical discussions, late at night at NIKHEF, which made the time there more bearable. Rolf often had a a very different viewpoint from me, which makes you reexamine your own beliefs and assumptions. Maarten Litmaath is a very useful computer expert to have around at any time.

There is of course also a life outside work. I have enjoyed the many barbecues with Bart & Gaelle and Ioannis & Katarina. I remember very fondly the ridiculous hikes Aimo and I have undertaken. We must have been nuts, but lets do it again soon. Then there are all my opponents at the various squash clubs, who have kept me sane by letting me destroy a squashball instead of something else and with whom I many interesting discussions: Nal, Jane, Rene, Mark B., Ivan, Job, Wim, Jim, and many others. Last, but not least, there are my close friends, Maxime, Peter, Mark B., and M², who are always willing to help out.

And finally, very big thanks again to Gio, without her, this thesis would never have existed within my lifetime. Thank you very much for always being there.

Curriculum Vitae

Op 7 februari 1969 ben ik geboren te Maastricht. Al op de lagere school waren de exacte vakken een aantrekkingspunt en dat is nooit veranderd. In de derde klas van het Gymnasium op het Henric van Veldeke college te Maastricht, kwam ik via de computerclub van school in aanraking met één van de eerste thuiscomputers, de TRS80 met 1 kilobyte geheugen. Een “upgrade” naar 2 kilobyte kostte maar 500 euro, maar dan kon je wel space-invaders (firebird) spelen. Vanaf toen waren computer hardware en software een zeer intensieve hobby.

Na vrolijk het Gymnasium te hebben doorlopen, ben ik in 1987 technische natuurkunde gaan studeren in Eindhoven. Tijdens mijn studie heb ik me geamuseerd door elektronen boven vloeibaar helium te laten zweven, squash te leren, en ben ik afgestuurd door te laten zien hoe lokale spingolven toch kunnen bestaan in een Knudsengas waar de atomen bijna niet meer met elkaar botsen.

In 1995 ben ik, via NIKHEF, begonnen met mijn promotiewerk in het CHORUS experiment op CERN. Met mijn achtergrond van computer hardware en software, was het niet verrassend dat ik vooral bezig ben geweest met de ontwikkeling van elektronica en data-acquisitie software. Toen CERN zijn eigen emulsie scanning microscopen ging opzetten, ben ik me daar mee bezig gaan houden. Na lange perioden op CERN, onderbroken door verblijven van enkele maanden in Amsterdam, heb ik een postdoc positie op CERN aangenomen om verder te werken aan de microscopen en emulsie analyse. Dit proefschrift bevat onderwerpen van zowel mijn promotie als mijn postdoc periode.

Ik ben nu op zoek naar een nieuwe baan. Aangezien mijn interesse enigzins verschoven is richting moleculaire biologie, zoek ik op dit moment naar iets in de bio-informatica. Maar een heel ander veld van werk sluit ik niet uit.