

**A search for weakly-interacting neutral
particles produced in 450 GeV/c
proton-nucleus collisions**

Jeremy Robert Dodd
University College London

Submitted to the University of London
for the degree of
Doctor of Philosophy
July 1990

ProQuest Number: 10611087

All rights reserved

INFORMATION TO ALL USERS

The quality of this reproduction is dependent upon the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if material had to be removed, a note will indicate the deletion.



ProQuest 10611087

Published by ProQuest LLC (2017). Copyright of the Dissertation is held by the Author.

All rights reserved.

This work is protected against unauthorized copying under Title 17, United States Code
Microform Edition © ProQuest LLC.

ProQuest LLC.
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106 – 1346

Abstract

This thesis presents results of a search for anomalous production of long-lived weakly interacting neutral particles in 450 GeV/c proton-nucleus collisions. The search is based on calorimetric measurements of visible event energy, relying on the general signature of 'missing' energy-momentum which characterises the production of these particles. A comparison between experimental data and Monte Carlo predictions based on known sources shows no significant excess, allowing upper limits to be placed on new sources of such particles.

Acknowledgements

I would like to thank the many people who made this thesis possible and who helped provide some welcome diversions along the way.

Firstly, my thanks to Bill Willis, Chris Fabjan and Norman McCubbin for their guidance and support, and for welcoming me to the HELIOS collaboration. Thanks also to my supervisor, Mike Esten, for his help at CERN and, especially, at UCL.

I am particularly indebted to Aris Angelis for showing me the ropes and for remaining reasonably calm despite having to share an office with me. I am equally grateful to Eddy Mazzucato and Pablo Yepes who helped in the analysis and were always a pleasure to work with. Thanks to Sacha Kalinovsky, Sacha Nikitin and François Lamarche for much advice and practical assistance. It has been a privilege to work in HELIOS and I am grateful to everyone in the collaboration.

Of the many people I have been fortunate to meet at CERN, I would like to mention Maria Varela Diaz, Joe Kroll, Affie Khan, Christos Markou, James Gillies, Lucas Taylor, Phil Burrows, Tony Cass, Bob Black, Andy Belk, Loïc Vinet and Pedro Resende; not for their tireless support and encouragement, but for sharing some good times in Geneva.

Finally, thank you to my family, without whom this thesis would not have been possible in the first place.

This work was supported in part by the UK Science and Engineering Research Council, by CERN and by University College London.

Contents

1	Introduction	1
2	Weakly interacting neutral particles	4
2.1	Theoretical motivation for a weakly interacting neutral particle search . . .	4
2.1.1	Standard Model: weak sector	4
2.1.2	Standard Model: Higgs sector	6
2.1.3	Beyond the Standard Model: supersymmetry	9
2.1.4	Beyond the Standard Model: axions	17
2.2	Experimental status	19
2.2.1	Weak processes	19
2.2.2	The Higgs	19
2.2.3	The gluino	20
2.2.4	The axion	24
2.3	Experimental detection of weakly interacting neutral particles	26
3	The HELIOS detector	30
3.1	Survey of the experimental programme	30
3.1.1	The proton beam	32
3.1.2	Beam telescope and general trigger	33
3.1.3	The silicon pad array	37
3.1.4	Calorimetry	39
3.1.5	The muon spectrometer	52
3.2	Calorimeter performance	56
3.2.1	Calibration	56
3.2.2	Intrinsic calorimeter properties	61
3.2.3	ULAC peak ADC performance	65
3.2.4	Resolution of the calorimeter ensemble	66
3.3	Muon spectrometer performance	71
3.3.1	Muon detection	71
3.3.2	Muon reconstruction	71
3.3.3	Reconstruction efficiency	72

3.4	The data sample	76
3.4.1	The missing energy data sample	76
3.4.2	The excess energy data sample	76
3.4.3	Data structure	77
4	The missing energy spectrum	79
4.1	The raw spectrum	79
4.2	Upstream interactions	82
4.3	Shower leakage	84
4.3.1	Lateral energy containment	84
4.3.2	Longitudinal energy containment	86
4.3.3	Calorimeter cuts	87
4.3.4	Spectrometer cuts	89
4.4	Pileup	94
4.4.1	On-line protection	96
4.4.2	Evidence for residual contamination	97
4.4.3	Off-line protection	100
5	Monte Carlo prediction of the missing energy spectrum	111
5.1	Standard Model neutrino production	112
5.1.1	Semi-leptonic heavy flavour decay	112
5.1.2	π , K decay	113
5.1.3	Hyperon decay	113
5.1.4	τ lepton decay	113
5.1.5	Vector boson decay	114
5.2	The event generator	115
5.3	Detector simulation	118
5.4	The contribution from charm production	120
5.4.1	Total charmed particle production cross-section	121
5.4.2	Charm parametrisation used in the Monte Carlo	127
5.5	The single muon sample	129
5.5.1	Muon selection	130
5.5.2	Missing energy spectrum for the single muon sample	131
5.6	Comparison of the final missing energy spectrum with Monte Carlo	132
6	Limits on the production of weakly interacting neutral particles	135
6.1	Comments on the measured missing energy spectrum	135
6.2	Cross-section limits	137
6.2.1	Simple limits	137
6.2.2	Model for WINP production	138
6.2.3	Total cross-section limits	139
6.2.4	Comparison with previous experiments	142

7	Conclusions	144
A	Resolution minimisation	146
A.1	Equivalence of resolution minimisation technique and Etot-flatness	146
A.1.1	Etot-flatness	147
A.1.2	Resolution minimisation	147
A.2	Results of resolution minimisation for the HELIOS calorimetry	148
B	The muon chamber efficiency	150
C	The muon hodoscope efficiency	152
D	Monte Carlo comparisons with data	156
E	Algorithm used for deriving upper limits	162

List of Figures

2.1	The Higgs potential	7
2.2	Loop diagrams contributing quantum corrections to m_f	10
2.3	Quadratically divergent loop diagrams contributing to δm_H^2	11
2.4	Graphs for the production of gluinos in hadron-hadron collisions.	16
2.5	Cross-section for gluino production via $gg \rightarrow \tilde{g}\tilde{g}$	16
2.6	Excluded region of gluino and squark masses	21
2.7	Limits on mass and lifetime of an axion-like particle.	25
2.8	Experimental sensitivity to weakly interacting particles.	27
3.1	Layout of the HELIOS (NA34) detector.	31
3.2	The H8 beam line.	33
3.3	Beam-defining counters.	34
3.4	Block diagram of the principal trigger logic components.	36
3.5	The silicon pad array.	38
3.6	The HELIOS calorimeter layout.	40
3.7	Section through the ULAC.	41
3.8	Towers in the electromagnetic part of the ULAC.	42
3.9	Isometric view of a U/Cu/scintillator module.	45
3.10	Artist's view of the calorimetrised dipole magnet.	47
3.11	The splitting of the PM anode signal and input stage to the first level summing.	49
3.12	Schematic representation of the summing layout for the BEAM calorimeter.	49
3.13	The EFL digital summing logic.	50
3.14	Schematic diagram of the hodoscope planes H3 and H2.	54
3.15	The σ_{rms}/E of the signal distribution as a function of B/A.	58
3.16	The e/h signal ratio as a function of B/A.	59
3.17	UCAL module response to hadronic showers as a function of energy.	62
3.18	e/h as a function of energy for the ULAC.	63
3.19	e/h as a function of energy for the UCALs.	64
3.20	Resolution as a function of energy for the ULAC.	64

3.21	Resolution as a function of energy for the UCALs.	65
3.22	Punchthrough probability in an iron absorber as a function of incident hadron energy.	74
4.1	Calorimeter response (FADCs) to minimum-bias interactions.	80
4.2	The raw missing energy spectrum.	81
4.3	The missing energy integrated rate spectrum.	81
4.4	The missing energy spectrum after cuts on upstream interactions.	84
4.5	Shower profile in the ULAC calorimeter.	85
4.6	Longitudinal shower centroid as a function of missing energy.	86
4.7	Hit distribution in the spectrometer wire chambers.	87
4.8	Energy detected in the first $1.5 \lambda_i$ of the dump.	88
4.9	Integrated rate for events with $E_{mis} > 50$ GeV versus maximum BEAM calorimeter energy.	89
4.10	Number of H2 hodoscope slats hit versus missing energy (no track).	91
4.11	Hit distributions in H2 as a function of E_{mis} (no track).	92
4.12	Number of H2 hodoscope slats hit versus missing energy (1 track).	94
4.13	Hit distributions in H2 as a function of E_{mis} (1 track).	95
4.14	The missing energy spectrum after cuts on leakage.	96
4.15	The missing energy spectra for low and high beam rate data.	98
4.16	Schematic illustration of calorimeter pileup.	99
4.17	Beam counter (B6) pulse-height spectra, indicating cuts.	101
4.18	Pileup as seen in the ULAC history FADC.	101
4.19	PADC/FADC correlation in the ULAC.	102
4.20	PADC-FADC energy difference (minimum-bias events).	103
4.21	PADC-FADC energy difference versus time separation of pileup.	104
4.22	Approximate parametrisation of calorimeter pulse.	105
4.23	The excess energy spectrum.	106
4.24	Missing energy rate as a function of beam flux squared.	108
4.25	The final missing energy spectrum, after rate extrapolation.	109
5.1	Hadronic production and semi-leptonic decay of charmed D mesons.	112
5.2	Drell-Yan production of lepton pairs via $q\bar{q}$ annihilation.	114
5.3	J/Ψ production and decay to neutrinos.	115
5.4	Preliminary simulation of the missing energy spectrum.	120
5.5	The variation of α with x_F for production of light flavour hadrons.	123
5.6	The variation of α with x_F for charm production.	124
5.7	Charm cross-section in pp interactions as a function of centre-of-mass energy.	125

5.8	Simulation of the missing energy spectrum using different charm parametrisations.	127
5.9	Final simulation of the missing energy spectrum.	129
5.10	The missing energy spectrum and Monte Carlo comparison for the single muon sample.	132
5.11	Simulation of the single muon E_{mis} spectrum showing contributions from π , K and charm decays.	133
5.12	The final missing energy spectrum and Monte Carlo prediction.	133
6.1	The single muon E_{mis} spectrum, with Monte Carlo predictions based on charm parametrisations from EHS and CCFRS.	136
6.2	Bourquin-Gaillard invariant cross-section distributions for a particle of mass 2.5 GeV.	140
6.3	Limits on the cross-section for WINP production in 450 GeV/c pp collisions as a function of mass for different lifetimes.	141
6.4	Comparison of this experiment with previous calorimeter searches.	143
C.1	Schematic of the hodoscopes and iron wall.	153
D.1	Differential cross-sections $\frac{d\sigma}{dx_F}, \frac{d\sigma}{dp_T^2}$ for π^0 production in pp collisions.	157
D.2	Differential cross-sections $\frac{d\sigma}{dx_F}, \frac{d\sigma}{dp_T^2}$ for π^- production in pp collisions.	158
D.3	Inclusive K_S^0 production in pN collisions.	159
D.4	Differential cross-section $\frac{d\sigma}{dx_F}$ for ρ^0, ω^0 production in K^+p collisions.	160
D.5	Differential cross-sections $\frac{d\sigma}{dx_F}, \frac{d\sigma}{dp_T^2}$ for charm production in pp collisions.	161

List of Tables

2.1	The supersymmetry spectrum.	12
2.2	Discovery potential for supersymmetric particles.	15
2.3	Experimental limits on neutrino production without accompanying charged lepton.	19
3.1	Dimensions and positions of the beam counters.	34
3.2	Characteristics of the proportional wire chambers.	53
3.3	Optimised resolutions for target and beam dump modes.	69
4.1	Summary of the calorimeter readouts.	98
4.2	The missing energy differential rate $\frac{1}{N} \frac{dN}{dE_{miss}}$ in 450 GeV p- ²³⁸ U collisions.	109
5.1	Monte Carlo decays giving final state neutrinos.	116
5.2	Compilation of data on x_F and p_T distributions for D mesons.	126
6.1	Cross-section limits for production of WINPs in pp collisions with energies above a certain threshold.	138
A.1	Results of resolution minimisation on ‘minimum-bias’ target interactions.	148
B.1	Spectrometer chamber efficiencies.	151
C.1	H3 hodoscope plane efficiencies.	154
C.2	H2 hodoscope plane efficiencies.	155

Chapter 1

Introduction

Modern science shares with Greek philosophy the conviction that the observed Universe is founded on simple underlying principles which can be understood and elaborated through disciplined intellectual endeavour. In elementary particle physics, the hope is that we will eventually achieve a unified scheme which combines all particles and all their interactions into one consistent theory.

At present, there is a very successful model which describes nearly all available data pertaining to the strong, weak and electromagnetic phenomena. One piece of the puzzle is still missing however - the spin-zero elementary Higgs boson needed by the Standard Model for spontaneous symmetry breaking (giving masses to the W^\pm , Z^0 and fermions). The Higgs has well-defined couplings to other particles but its mass is not fixed. If it is light, then its coupling to stable matter will be very small, i.e. the Higgs will be a weakly interacting neutral particle. Although one may argue that it is only a matter of time before the Higgs boson is discovered, there are indications that deeper problems exist, suggesting it is necessary to look beyond the Standard Model to understand the Higgs sector of the theory.

If there is new physics beyond the Standard Model, one would like to know what is the nature of this new physics, at what energy scale it enters and how it will be identified. By far the most intensely studied class of theories as a possible candidate for physics beyond the Standard Model has been supersymmetry. In addition to having properties needed to solve various theoretical difficulties, supersymmetric theories contain a large number of experimental consequences. One of these is that there should exist a stable weakly interacting neutral particle, the neutralino.

Admitting the possibility of new sources of weakly interacting neutral particles (both within and beyond the Standard Model), it is natural to ask whether there is any experi-

mental evidence for such particles. Over the past decade, several independent results have consolidated the belief that most of the mass in the Universe exists in an as yet unknown form of matter that is invisible to electromagnetic radiation and is thus termed dark matter. This evidence relies mainly on observations of the motions of large-scale structures in the Universe (optical rotation curves for galaxies, galactic cluster dynamics,...) which are in disagreement with expectations from the visible matter alone. The nature of dark matter is uncertain. If the mean density of the Universe (in units of the critical density) $\Omega = 1$ as some recent observations indicate [1,2], then the bulk of the dark matter cannot be baryonic since nucleosynthesis constrains $\Omega_{baryonic} \lesssim 0.2$ [3]. The simplest assumption that does not conflict with experimental evidence is that all of the dark matter is of the same type and non-baryonic. This is interesting in view of the predictions mentioned above for the existence of stable weakly interacting neutral particles that might be left in sufficient number after the Big Bang to make $\Omega = 1$. Current results from attempts to detect dark matter indirectly (for example, the detection of very energetic neutrinos from dark matter annihilation in the solar core [4]) do not constrain particles such as the neutralino.

This thesis presents results from a direct search for the production of weakly interacting neutral particles in 450 GeV/c proton-nucleus collisions¹. The research was carried out as part of the experimental programme of the HELIOS collaboration. The experiment (NA34/1) was designed to make a detailed study of prompt lepton production, combining good electron and muon identification, accurate determination of ‘missing’ energy (due to undetected particles), and detailed information on event topology (charged multiplicity and energy flow). This analysis focuses on the measurement of missing energy, since it provides the most general signature of weakly interacting neutral particle production. Charged lepton identification will also be seen to be crucial since it allows one to reject events from known Standard Model sources.

Chapter 2 reviews the theoretical arguments which motivate a weakly interacting neutral particle search, firstly for sources within the Standard Model, and secondly from new physics beyond the Standard Model. The experimental status of each candidate source is then outlined, including present mass and production cross-section limits where relevant. Having shown that a new search may be profitable, the remainder of the chapter is devoted to a discussion of experimental requirements and a demonstration that we are able to improve on previous experiments.

In Chapter 3, the HELIOS detector is described in detail. The emphasis is on calorimetry and its optimisation to achieve the best possible energy information for each event.

¹Throughout this thesis, energies are given in GeV ($1 \text{ GeV} = 10^9 \text{ eV} = 1.602 \times 10^{-10} \text{ J}$) and cross-sections in barn ($1 \text{ b} = 10^{-28} \text{ m}^2$). A complete list of SI units may be found in the ‘Review of Particle Properties’, M. Aguilar-Benitez et al., Phys. Lett. B204 (1988).

The detector performance in two different modes, target and beam dump, is compared to determine which is more sensitive to possible new sources of missing energy. The chapter concludes with a brief description of the 1988 data sample used in this analysis.

The treatment of the beam dump missing energy data and the associated correction procedures are the subject of Chapter 4. Three sources of contamination are identified: interactions upstream of the target, energy leakage through incomplete shower containment in the calorimeters, and problems associated with the overlap of two or more particle interactions in the target.

The interpretation of the final missing energy spectrum in terms of detector properties and known Standard Model physics forms the basis of Chapter 5. This includes a simulation of expected 'backgrounds' from Standard Model sources and in particular, a detailed discussion of the expected contribution to the missing energy spectrum from semi-leptonic decays of charmed particles.

Chapter 6 compares the observed missing energy spectrum with that expected from Standard Model sources. Based on this comparison, some comments on charm production are made and the question addressed of whether the data shows any evidence for anomalous production of neutrino-like particles. Model-dependent upper limits on new sources of weakly interacting neutral particle production in proton-proton interactions are derived.

Conclusions are presented in Chapter 7.

Chapter 2

Weakly interacting neutral particles

2.1 Theoretical motivation for a weakly interacting neutral particle search

This section reviews theoretical arguments for possible hitherto unrecognised sources of weakly interacting neutral particles, firstly within the Standard Model and secondly, in extensions to the Standard Model. The relevance of these arguments for hadronic interactions at centre-of-mass energy ~ 30 GeV (corresponding to 450 GeV/c proton-nucleus collisions) will be outlined where appropriate.

2.1.1 Standard Model: weak sector

Over the last few years there has been a remarkable synthesis in the theories used to describe interactions of elementary particles. A Standard Model, consistent with all existing experimental data, has been constructed to describe interactions of the known quarks and leptons.

The strong interaction, responsible for binding quarks in hadrons, is described in terms of a non-Abelian gauge theory, Quantum Chromodynamics (QCD), based on the group SU(3). In this scenario, the strong force between quarks is carried by massless, spin-one gluons which transform according to the octet representation of SU(3)_C. There are six types of strong charge, or ‘colour’: a quark can carry one of three colours (say red, blue or green) and an antiquark the corresponding anticolour. The theory explains both the phenomenon of confinement, whereby free quarks are not observed, and of asymptotic

freedom, according to which the interactions of quarks become weak at large momentum transfers or small distances.

The electromagnetic and weak interactions are described in the Standard Model by a theory based on the group $SU(2) \times U(1)$, in which there are four spin-one gauge bosons. This symmetry is not exact however, resulting in non-zero masses for three of the bosons (W^\pm, Z^0) while the fourth, the photon, remains massless. The symmetry-breaking is due to the existence of a scalar field, the Higgs field, transforming as a doublet under $SU(2)_I \times U(1)_Y$ ¹. The state of minimum energy occurs for non-zero value of this field and the $SU(2)_I \times U(1)_Y$ symmetry ‘spontaneously’ breaks down to $U(1)_{EM}$, corresponding to electromagnetism. A relic of this breakdown is a neutral scalar, the Higgs.

The weak sector of this model is responsible for the decays of long-lived particles, which are inhibited from decaying via the strong or electromagnetic interactions by conservation rules or kinematic considerations. As an example, consider the decay of the pion, which is a member of the lightest strongly interacting multiplet. It is found that:-

$$\begin{aligned} \pi^0 &\rightarrow \gamma\gamma & \text{with } \tau &= 8.4 \times 10^{-17} \text{ sec,} \\ \pi^- &\rightarrow \mu^- \bar{\nu}_\mu & \text{with } \tau &= 2.6 \times 10^{-8} \text{ sec,} \end{aligned} \quad (2.1)$$

which should be compared to typical lifetimes for decays through the strong and electromagnetic interactions of around 10^{-23} sec and 10^{-16} sec respectively. Decays via the strong interaction are excluded since the π is the lightest hadron and whereas the neutral π can decay into photons (an electromagnetic process), the charged pions cannot. As a result, the weak decay of the π^- in Equation 2.1 is the dominant one.

Neutrinos in the Standard Model

The example of the weak decay given above involves a neutrino. Neutrinos are unique in that they can interact only through the weak interaction. They are colourless and electrically neutral and, within current experimental limits, also massless. The effective neutrino interaction at low energies is best illuminated by looking at the structure of the relevant currents in the Lagrangian. For charged current processes (that is, those interactions involving charged gauge boson (W^\pm) vertices), one has

$$j_\mu^{CC} = \bar{u}_{\nu_l} \gamma_\mu \frac{1}{2} (1 - \gamma_5) u_l + \text{Hermitean conjugate,}$$

where u denotes a Dirac spinor and l signifies one of the fermion generation doublets. For neutral current processes (those containing a Z^0 or photon vertex), the corresponding term is

$$j_\mu^{NC} = \frac{1}{2} \left(\bar{u}_{\nu_l} \gamma_\mu \frac{1}{2} (1 - \gamma_5) u_{\nu_l} - \bar{u}_l \gamma_\mu \frac{1}{2} (1 - \gamma_5) u_l \right) - \sin^2 \theta_W \left(-\bar{u}_l \gamma_\mu \frac{1}{2} (1 - \gamma_5) u_l \right),$$

¹Here I denotes weak isospin and Y hypercharge.

where θ_W is the weak mixing angle, related to the weak isospin and hypercharge coupling constants g, g' by

$$g \sin \theta_W = g' \cos \theta_W = e.$$

With the exception of Z^0 decay, all first-order weak processes involving neutrinos (the spinors u_{ν_l}) are therefore also accompanied by a charged lepton l^\pm (u_l). This property of the electroweak theory will be seen to provide a useful ‘tag’ on neutrino production.

The number of different neutrino species is not prescribed in the Standard Model. Three have so far been identified: ν_e, ν_μ and ν_τ , one for each of the known fermion generations. Recent measurements on the width of the Z^0 peak at LEP have allowed the number of light neutrino species to be determined for the first time. Taking the average for the four LEP detectors[5], the number of neutrinos is found to be $N_\nu = 2.95 \pm 0.11$, excluding the existence of a fourth generation at the 9 standard deviation level.

Neutrino production through Z^0 decay is not possible at SPS fixed target energies ($\sqrt{s} \sim 30$ GeV) since $m_Z \sim 90$ GeV/ c^2 . Virtual Z^0 production (and subsequent decay) is however possible through resonance production of the heavy flavour quarkonia (e.g. $J/\Psi, \Upsilon$). This mechanism, and an estimate of its cross-section, will be given in Section 5.1.

In conclusion, neutrinos are expected to be produced in conjunction with their corresponding charged lepton partners. The only first-order process for which this is not the case is Z^0 decay (to $\nu\bar{\nu}$), kinematically forbidden at the centre-of-mass energy of this experiment. Since the number of light neutrinos is now known to be three, neutrinos should be produced via (l, ν_l) pair production with $l = e, \mu, \tau$ only. In addition to these known sources however, it may be of interest to search for rare or forbidden processes such as $\pi^0 \rightarrow \nu\bar{\nu}$. Limits on such processes will be reviewed in Section 2.2.1.

2.1.2 Standard Model: Higgs sector

The rôle of the Higgs

Particles in the Standard Model can only acquire masses when the $SU(2)_I \times U(1)_Y$ electroweak gauge symmetry is spontaneously broken. This is achieved through the vacuum expectation value of a Higgs field, and there are arguments based on perturbative unitarity (see below) that indicate there must exist a spin-zero particle with $m_H = O(10^{0\pm 1})m_W$ and couplings analogous to those of the Higgs boson in the minimal Weinberg-Salam




Figure 2.1: Forms of the Higgs potential for $\mu^2 > 0$ and $\mu^2 < 0$.

model. There is as yet no experimental evidence for the Higgs mechanism of spontaneous symmetry breaking, still less for the existence of any Higgs boson. Both are necessary features of the Standard Model. In this section, an outline of the Higgs mechanism will be given, with reference to the theoretical constraints on its mass.

The mechanism for generation of mass for the gauge bosons can be illustrated using the simplified example of a single real scalar field ϕ . The potential for this field has the form

$$V(\phi) = \frac{1}{2}\mu^2\phi^2 + \frac{1}{4}\lambda\phi^4,$$

where the corresponding Lagrangian $\mathcal{L} (\equiv T - V)$ is required to be invariant under the symmetry operation $\phi \rightarrow -\phi$. Here λ and μ^2 are parameters, and λ is taken to be positive. Figure 2.1 shows the form of the Higgs potential for $\mu^2 > 0$ and $\mu^2 < 0$. The case of $\mu^2 > 0$ describes a scalar field with mass μ . The ϕ^4 term shows that the four-particle vertex exists with coupling λ (i.e. ϕ is a self-interacting field). Although the mass term apparently has the wrong sign for $\mu^2 < 0$ (the relative sign of the ϕ^2 term and the kinetic energy T being positive), it is this case that is of interest. For $\mu^2 < 0$, the minimum of the potential is not at $\phi = 0$ but at $\phi \equiv v = \sqrt{-\mu^2/\lambda}$. The extremum $\phi = 0$ does not correspond to the energy minimum. A perturbation expansion can only be done around a local minimum, so the field must be translated to $\pm v$. Writing

$$\phi(x) = v + \phi'(x),$$

where $\phi'(x)$ is the fluctuation around $+v$, the form of the potential is

$$V(\phi') = \lambda v^2 \phi'^2 + \lambda v \phi'^3 + \frac{1}{4} \lambda \phi'^4 + \text{constant}.$$

There is now a mass term for ϕ' with the correct sign

$$m_{\phi'} = \sqrt{2\lambda v^2} = \sqrt{-2\mu^2}.$$

Hence, by introducing a potential whose minimum is not at zero, the field develops a non-zero vacuum expectation value v and a conventional mass term in the Lagrangian. This is spontaneous symmetry breaking.

By extending this technique of field translation to a complex scalar field, given by $\phi = (\phi_1 + i\phi_2)/\sqrt{2}$, two fields are introduced, one of which acquires a mass as before, the other remaining massless. This massless mode corresponds to excitations around the two-dimensional potential minimum (the ‘bottom of the wine bottle’).

From here, the final step is to study spontaneous breaking of a local gauge symmetry. In the case of the U(1) gauge symmetry, the particle spectrum (with the correct gauge choice) is two interacting massive particles, a vector boson A_μ and a massive scalar H , the Higgs. (The massless Goldstone boson has been ‘eaten’ to give the longitudinal polarisation state of A_μ). For the SU(2) group, ϕ is a doublet of complex scalar fields

$$\phi = \sqrt{\frac{1}{2}} \begin{pmatrix} \phi_1 + i\phi_2 \\ \phi_3 + i\phi_4 \end{pmatrix}.$$

In this case, three massless Goldstone boson fields are eaten by the gauge fields and become massive ($m = \frac{1}{2}vg$). The remaining field is the massive Higgs scalar.

The same procedure can be applied to electroweak interactions, that is, to the group SU(2) \times U(1). In order to generate masses for the weak bosons but have a massless photon (as observed experimentally), the field which is allowed to acquire a vacuum expectation value is taken to be electrically neutral. Thus the gauge symmetry of the U(1)_{EM} subgroup is left intact, ensuring that the photon does indeed remain massless. In so doing, some of the terms in the Higgs sector Lagrangian become mass terms for the charged and neutral gauge bosons, where

$$m_{W^\pm} = \frac{1}{2}vg, \quad m_Z = \frac{1}{2}v\sqrt{g^2 + g'^2}, \quad m_A = 0.$$

Thus, we find massive W and Z bosons, with the relation $m_W/m_Z = \cos\theta_W$, and a massless photon A . The Higgs coupling to other particles depends only on their masses:-

$$\begin{aligned} g_{fH} &= 2^{\frac{1}{4}} \sqrt{G_F} m_f && \text{for fermions,} \\ g_{VH} &= 2^{\frac{1}{4}} \sqrt{G_F} m_V^2 && \text{for vector bosons.} \end{aligned}$$

Consequently, its coupling to stable matter is very small, and production and detection are difficult. Though the masses of the gauge bosons and the Higgs couplings are well-defined, the mass of the Higgs boson is almost completely unconstrained, and depends on the value of λ . If λ becomes too large, perturbation theory breaks down and the Higgs will have strong self-couplings. This gives a qualitative upper bound on the mass of $m_H \lesssim 1 \text{ TeV}/c^2$. Radiative corrections contribute several GeV to m_H , and this is generally taken as a qualitative lower bound on the mass. A detailed analysis of the thermal history of the Universe [6,7] also requires $m_H \gtrsim 10 \text{ GeV}/c^2$. Thus

$$\text{a few GeV}/c^2 \lesssim m_H \lesssim 1 \text{ TeV}/c^2.$$

In general, the lower mass limits weaken as the top quark mass increases, disappearing altogether when $m_t \gtrsim m_W$, or if there are more undiscovered fermions (such as in a fourth generation).

The above picture is modified if one goes beyond the Standard Model. All such schemes involve additional Higgses. In the case of minimal supersymmetry (see Section 2.1.3), there are two complex Higgs doublets for a total of eight real fields. After removing the three fields which are eaten by the W^\pm and Z^0 to give them masses, one is left with five physical Higgs bosons. Two of these physical Higgs bosons are neutral scalars, one is a neutral pseudoscalar, and two are charged particles. In this scheme, at least one, and possibly two, of the neutral Higgses must have masses less than the mass of the Z^0 i.e. $m_{H^0} \lesssim 90 \text{ GeV}/c^2$.

If the Higgs is light (mass \sim a few GeV), its production in hadronic collisions is expected to be predominantly through the decay of heavy vector and B, K mesons. Since it couples directly to the masses of other particles, it will decay into the heaviest available 'daughters'. As was seen above, the Higgs is predicted to have a mass between $O(10) \text{ GeV}/c^2$ and $1 \text{ TeV}/c^2$. The leading-order decays are therefore $H^0 \rightarrow l^+l^-, q\bar{q}$ (and $H^0 \rightarrow W^+W^-, Z^0Z^0$ if $m_{H^0} > 2m_{W,Z}$); $H^0 \rightarrow \gamma\gamma, gg$ are both higher-order decays.

In summary, a light Higgs may exist. If it does, one may be sensitive to its production in a search for weakly interacting neutral particles. The experimental limits on the Higgs mass will be reviewed in Section 2.2.2.

2.1.3 Beyond the Standard Model: supersymmetry

Motivation for supersymmetry

The Standard Model successfully describes fundamental interactions up to the presently explored level of about 100 GeV. From a theoretical standpoint there are, however, sev-



Figure 2.2: Loop diagrams contributing quantum corrections to m_f .

eral outstanding problems left unsolved. These include the problems of including gravity within a consistent framework, of understanding the proliferation of different flavours of quarks and leptons, and of the origin of particle masses. Superstring and supermembrane models offer potential solutions to the the first two problems while supersymmetry (SUSY) addresses the question of particle masses.

Many different motivations for supersymmetry have been put forward in recent years. It is the only remaining symmetry of the \mathcal{S} -matrix. It unifies particles and interactions. It reduces the divergences in quantum gravity. It provides an elegant resolution of the hierarchy problem, namely that of reconciling ‘small’ mass scales such as m_W with ‘large’ ones such as the grand unification scale $m_X \sim 10^{15}$ GeV/ c^2 or the Planck scale $m_P \sim 10^{19}$ GeV/ c^2 associated with gravitation. The hierarchy problem occurs at two levels: one is that of creating the hierarchy, i.e. the origin of m_W , and the other is that of maintaining it once it has been created, i.e. the natural stability of m_W .

The value of a physical parameter is said to be natural if quantum corrections to it are no larger than its physical value. For example, a fermion mass, m_f , gets corrections from the loop of Figure 2.2

$$\delta m_f = O\left(\frac{\alpha}{\pi}\right) m_f \ln\left(\frac{\Lambda}{m_f}\right),$$

where Λ is the cut-off in the loop integral. As long as $\Lambda \lesssim m_f \exp O(\frac{\pi}{\alpha})$, then $\delta m_f \leq m_f$ and the fermion mass parameter can be naturally small.

In the case of the W , the smallness of m_W is linked to the smallness of m_H , which is unstable in the Standard Model. Radiative corrections to its squared mass due to loops like those of Figure 2.3 are quadratically divergent

$$\delta m_H^2 = O\left(\frac{\alpha}{\pi}\right) \Lambda^2,$$




Figure 2.3: Quadratically divergent loop diagrams contributing to δm_H^2 .

where Λ is a cut-off representing the threshold for some new physics beyond the Standard Model and may be, a priori, as large as the Planck mass $m_P \sim 10^{19} \text{ GeV}/c^2$. However, in order for δm_H^2 to be less than m_H^2 , and hence technically natural, requires $\Lambda \lesssim \text{O}(1)\text{TeV}$.

In itself, the divergence of δm_H^2 is not a problem (it is renormalisable and can be absorbed in the definition of the bare parameters). It *is* a problem, however, when attempting to construct more unified models. Grand Unified Theories (GUTs), for example, predict corrections to the Higgs mass of

$$\delta m_H^2 = \text{O}(m_X^2) \gtrsim \text{O}\left(10^{15} \text{ GeV}\right)^2,$$

which is much larger than the range allowed in Section 2.1.2. It is also a problem if one considers the effects of quantum gravity, which are predicted to give corrections

$$\delta m_H^2 = \text{O}(m_P^2) = \text{O}\left(10^{19} \text{ GeV}\right)^2.$$

Two strategies have been proposed for removing the large corrections to m_H^2 . The first is to give the Higgs a composite structure (of massless fermions) on a distance scale

$$R_T = \text{O}\left(\frac{1}{\Lambda_T}\right) \text{ where } \Lambda_T \sim 1\text{TeV}.$$

This model, called technicolour, introduces a cut-off $\Lambda \sim \Lambda_T$ in the loops of Figure 2.3. The scale of compositeness is where new technicolour interactions become strong. Attempts to give non-zero masses to the elementary fermions fail, however, and extended technicolour models predict unobserved phenomena (flavour-changing neutral-current interactions and light, charged, composite scalars).

The second approach, supersymmetry, cancels the loops of Figure 2.3 among themselves. The boson and fermion loops have opposite signs, leading to

$$\delta m_H^2 = \text{O}\left(\frac{\alpha}{\pi}\right) \left[\left(\Lambda^2 + \text{O}(m_b^2) \right) - \left(\Lambda^2 + \text{O}(m_f^2) \right) \right].$$

Particle	Spin	Sparticle	Spin
quark $q_{L,R}$	$\frac{1}{2}$	squark $\tilde{q}_{L,R}$	0
lepton $l_{L,R}$	$\frac{1}{2}$	slepton $\tilde{l}_{L,R}$	0
photon γ	1	photino $\tilde{\gamma}$	$\frac{1}{2}$
gluon g	1	gluino \tilde{g}	$\frac{1}{2}$
W	1	wino \tilde{W}	$\frac{1}{2}$
Z	1	zino \tilde{Z}	$\frac{1}{2}$
Higgs H	0	higgsino \tilde{H}	$\frac{1}{2}$
graviton G	2	gravitino \tilde{G}	$\frac{3}{2}$

Table 2.1: The supersymmetry spectrum.

By postulating boson/fermion pairs with the same quantum numbers (and hence identical couplings) and similar masses, the diagrams of Figure 2.3 give a residual

$$\delta m_H^2 = O\left(\frac{\alpha}{\pi}\right) [m_b^2 - m_f^2].$$

This correction is naturally small if the effective cut-off

$$\Lambda^2 \sim [m_b^2 - m_f^2] \lesssim O(1 \text{ TeV})^2.$$

Identical couplings and similar masses also remove the problems associated with GUTs and quantum gravity. The naturalness condition means that supersymmetric partners of the known particles can weigh no more than $O(1) \text{ TeV}/c^2$.

Structure of supersymmetric models

This section describes the basic characteristics of supersymmetric theories, emphasising those aspects important for a discussion of its phenomenological implications. Only minimal models will be considered. For a detailed treatment of supersymmetry formalism and a summary of extended models, the reader is directed to [8,9].

Supersymmetry is a symmetry which relates fermionic and bosonic degrees of freedom. It therefore predicts new supersymmetric particles which are partners of all the known particles and differing in spin by half a unit. In addition, the total number of bosonic and fermionic degrees of freedom must be equal. The complete set of supersymmetric partners for the known particles is shown in Table 2.1. No known particle can be the supersymmetric partner of any other. Considering the case of the spin- $\frac{1}{2}$ quark for example, there is no known elementary spin-0 particle and the squark cannot have spin 1, since in a renormalisable field theory all vector bosons must be gauge bosons. In the case of the gluon, its supersymmetric partner would have spin $\frac{1}{2}$, like the quark, but be an octet of colour. The supersymmetric particles have quantum numbers (with the exception of

spin) identical to those of their ordinary particle partners. One interesting consequence of supersymmetry is the necessity for two Higgs doublets, along with their supersymmetric partners. They are needed to give mass to both charge states of the elementary fermions.

Supersymmetry breaking

Supersymmetry cannot be an exact symmetry. If one has fermions $|F\rangle$ and bosons $|B\rangle$ related by the supersymmetry charge Q , $Q|B\rangle = |F\rangle$ and $Q|F\rangle = |B\rangle$, and if supersymmetry is a good symmetry, $[Q,H] = 0$, then $m_B = m_F$. However, no supersymmetric particle has been found experimentally, so $m_{\tilde{e}} \gg m_e$, $m_{\tilde{q}} \gg m_q$, etc. Symmetry breaking may be either explicit, by terms in the Lagrangian, or spontaneous, by having a non-invariant vacuum $Q|0\rangle \neq 0$. Explicit breaking is theoretically unattractive due to unitarity problems and the emphasis has therefore been on spontaneous breaking, in analogy with existing gauge theories. In addition to lifting the mass degeneracy, symmetry breaking allows mass eigenstates which are a linear combination of particles carrying the same conserved quantum numbers. The correct mass eigenstates would have to be determined by experiment (or by a theoretical model). An example of this is the mixing of photino, zino, higgsino,... in the physical 'neutralino' which is assumed to be the lightest supersymmetric particle (see below).

R-symmetry

Most supersymmetry models define a continuous global symmetry called an R-symmetry (essentially a generalisation of fermion number conservation). This implies that all particles carry a new quantum number called R-parity which is multiplicatively conserved. Its conservation can be related to that of other quantum numbers, namely spin S , baryon number B , and lepton number L

$$R = (-1)^{2S+3B+3L}$$

All ordinary particles are assigned an R-parity of +1, whereas the supersymmetric partners have an R-parity of -1. R-parity conservation is assumed in this review. In principle, it is possible to violate R-parity through a vacuum expectation value for some scalar sparticle such as a sneutrino but this possibility would also violate lepton number conservation which is highly constrained by experiments confirming L conservation. Conservation of R-parity has three important phenomenological consequences:-

1. Sparticles are always produced in pairs, e.g.

$$e^+e^- \rightarrow \tilde{e}^+\tilde{e}^- \text{ or } p\bar{p} \rightarrow \tilde{q}\tilde{q} + X.$$

2. Heavier sparticles decay into lighter sparticles, e.g.

$$\tilde{e} \rightarrow e\tilde{\gamma} \text{ or } \tilde{g} \rightarrow q\bar{q}\gamma.$$

3. The lightest supersymmetric particle (LSP) is stable.

The lightest supersymmetric particle

If supersymmetry *is* a symmetry of nature, then this last property means that the Universe should contain supersymmetric relics from the Big Bang, and provides a powerful cosmological constraint on the LSP. It must be electrically neutral and have no strong interactions. Otherwise it would have condensed in ordinary matter producing anomalously heavy isotopes. Experimental limits [10] on the abundance of such isotopes rule out this possibility.

Candidates for the LSP in the minimal supersymmetric Standard Model include the sneutrino $\tilde{\nu}$ of spin 0, the photino $\tilde{\gamma}$ and/or Higgsino \tilde{H} of spin $\frac{1}{2}$, and the gravitino \tilde{G} of spin $\frac{3}{2}$. In most models the LSP is a $\tilde{\gamma}/\tilde{H}$ combination. It will be referred to as a neutralino ($\tilde{\gamma}$) for this discussion. The $\tilde{\gamma}$ interaction cross-section is of weak-interaction size or less since it can only interact with matter via exchange of a heavy scalar electron or scalar quark. Hence, the $\tilde{\gamma}$ has the properties of a weakly interacting neutral particle. Any such relic would escape from some experimental apparatus in the same way as an unseen neutrino, producing a signature of missing energy-momentum. This observation is the basis of almost all experimental searches for supersymmetric phenomena.

Supersymmetric phenomenology: the light gluino scenario

The possibility of observing supersymmetry in particle physics experiments has been studied by several authors [8,9]. Predictions for production cross-sections and probable decay modes of the supersymmetric particles have been compared with Standard Model backgrounds to determine the most favourable domains for a supersymmetry search. Though these predictions are somewhat model-dependent, it is useful to note their conclusions concerning ‘discovery potential’ for the new particles. Table 2.2 gives a rough guideline for each of the supersymmetric particles, indicating the most favourable production

Sparticle	Favoured production mechanism	Cross-section	Signature
squark $\tilde{q}_{L,R}$	e^+e^- , pp , ep	\checkmark	?
slepton $\tilde{l}_{L,R}$	e^+e^-	\checkmark	\checkmark
photino $\tilde{\gamma}$	e^+e^- , pp	?	\checkmark
gluino \tilde{g}	pp	$\checkmark\checkmark$?
wino \tilde{W}	e^+e^-	\checkmark	?
zino \tilde{Z}	e^+e^- , pp	?	\checkmark
higgsino \tilde{H}	e^+e^- , pp	?	\checkmark

Table 2.2: Discovery potential for supersymmetric particles (\checkmark and ? indicate favourable and marginal conditions respectively).

mechanism taking into account expected cross-sections and experimental signatures (i.e. signal/background considerations).

The fact that no charged sparticle has been observed at LEP [11,12] indicates that $m_{\tilde{t}\pm}$, $m_{\tilde{W}\pm}$, $m_{\tilde{H}\pm} \gtrsim 40 \text{ GeV}/c^2$. The absence of any new strongly interacting sparticle in hadron-hadron collisions [13,14,15] indicates that $m_{\tilde{q}} \gtrsim 74 \text{ GeV}/c^2$. The case for the gluino is less clear. Results from the Sp \bar{p} S collider [13,14,15] exclude the region $4 \text{ GeV}/c^2 \lesssim m_{\tilde{g}} \lesssim 79 \text{ GeV}/c^2$, and several experiments have attempted to exclude gluino masses below this range, yielding $m_{\tilde{g}} \gtrsim 2.5 \text{ GeV}/c^2$. A detailed account of gluino mass limits is given in Section 2.2.3. Here, it is noted only that there remains a ‘window of opportunity’ for a light ($\sim 2 - 5 \text{ GeV}/c^2$) gluino to exist and have not (yet) been discovered.

This thesis considers the case of gluino production, and subsequent decay into hadrons plus a neutralino. As can be seen from Table 2.2, gluinos are expected to be copiously produced in hadron machines. Production cross-sections may be comparable to or greater than those for heavy quarks of the same mass [16]. R-parity conservation means that gluinos should be produced in pairs. The elementary processes involving two gluinos in the final state are shown in Figure 2.4. Computations of the expected cross-sections [17] show that the dominant contribution is from the gluon fusion diagrams; this is shown in Figure 2.5.

Once produced, gluinos (which are colour octets) form bound states in colour singlet R-hadrons. Depending on whether they bind with quarks, gluons, gluinos,... a complete spectrum of R-hadrons is expected. However, for sufficiently heavy gluinos (say $m_{\tilde{g}} \gtrsim 1 - 3 \text{ GeV}/c^2$), the mass and lifetime of such hadrons would be approximately equal to those calculated for the free gluino. In what follows, it will be assumed that this is the case. The dominant decay of the gluino is believed to be a three-body decay into $q\bar{q}\tilde{\gamma}$ with the lifetime

$$\tau_{\tilde{g}} \simeq \frac{48\pi m_{\tilde{q}}^4}{\alpha\alpha_s e_q^5 m_{\tilde{g}}^5} \simeq 1.3 \times 10^{-11} \text{sec} \left(\frac{0.15}{\alpha_s} \right) \left(\frac{m_{\tilde{q}}}{100 \text{GeV}} \right)^4 \left(\frac{1 \text{GeV}}{m_{\tilde{g}}} \right)^5$$




Figure 2.4: Graphs for the production of gluinos in hadron-hadron collisions. Two sub-processes are considered: (a) $gg \rightarrow \tilde{g}\tilde{g}$, (b) $q\bar{q} \rightarrow \tilde{g}\tilde{g}$.




Figure 2.5: Cross-section for gluino production via $gg \rightarrow \tilde{g}\tilde{g}$.

(summing over u, d and s quarks and assuming $m_q, m_{\tilde{\gamma}} \ll m_{\tilde{g}}$). For light gluinos (mass $\sim 1 - 5 \text{ GeV}/c^2$), the predicted lifetime then falls in the range $\sim 10^{-11} - 10^{-6} \text{ sec}$. It should be emphasised that the lifetime is a function of both squark and gluino mass. This explains why limits are usually presented (Figure 2.6 for example) in the squark/gluino mass plane. The neutralino is produced isotropically in the gluino rest-frame. Hadronisation of the $q\bar{q}$ system then results in a shower of hadrons in association with a neutralino.

Neutralino-matter interactions occur dominantly through $\tilde{\gamma}q \rightarrow \tilde{g}q$ with the cross-section

$$\sigma_{\tilde{\gamma}N} \simeq 8.9 \times 10^{-38} \text{ cm}^2 \left(\frac{\alpha_s}{0.15} \right) \left(\frac{100 \text{ GeV}}{m_{\tilde{q}}} \right)^4 \left(\frac{E_{\tilde{\gamma}}}{1 \text{ GeV}} \right) \tilde{F}$$

where \tilde{F} , the analogue of the standard structure function F_2 , ranges from 9.6×10^{-2} to 1.8×10^{-2} for $m_{\tilde{g}}$ between 2 and 5 $\text{ GeV}/c^2$. The cross-section thus has the same form as the neutrino interaction cross-section, and is also similar in magnitude for the range of $m_{\tilde{g}}$ and $m_{\tilde{q}}$ considered here.

In conclusion, a light gluino, if it exists, may be abundantly produced in hadron interactions. It is subsequently expected to decay into the lightest supersymmetric particle (which escapes detection) plus hadrons. In this scenario, gluino production and decay would be observed as a new source of weakly interacting neutral particles.

2.1.4 Beyond the Standard Model: axions

Attempts to extend the $SU(3) \times SU(2) \times U(1)$ model of lepton and quark interactions by including specific dynamical mechanisms for spontaneous symmetry breakdown lead to the appearance of nearly massless, spinless Goldstone bosons. Peccei and Quinn [18,19,20] have formulated an example of such a model for preventing CP violation in QCD and which predicts the existence of a light pseudoscalar boson, the axion.

The motivation for the axion comes from considering the consequences of a theoretical picture in which the weak and electromagnetic as well as the strong interactions are based on underlying non-Abelian gauge theories. The gauge theory of strong interactions is assumed to be quantum chromodynamics (QCD), based on an exact $SU(3)$ colour symmetry of quarks and gluons. The weak and electromagnetic interactions are assumed to stem from a spontaneously broken gauge theory based, in the simplest case, on the group $SU(2) \times U(1)$. The existence of instanton solutions for non-Abelian gauge theories raises a potential problem for this theoretical picture. In QCD, they allow the appearance of additional terms in the Lagrangian which produce strong CP symmetry violation, in contradiction with experimental observation. It is possible to eliminate these terms if either:-

1. at least one of the quarks in the theory is massless, or
2. the Lagrangian for the full theory has an overall global chiral U(1) symmetry.

The first option is inconsistent with present current-algebra estimates. A consequence of the second is that the axion should exist.

The axion is light for two reasons. In the absence of non-perturbative instanton effects, the axion would be a Goldstone boson associated with the spontaneously broken chiral U(1) symmetry of the theory. In reality the symmetry is broken and the axion is therefore expected to acquire a small mass, of the order of the pion mass. However, the mass scale characterising the broken weak symmetry is much larger than that associated with ordinary chiral symmetry-breaking and consequently the axion's mass is further reduced. Mass estimates, based on current-algebra techniques, have been made by several authors [21,22]. They find a mass in the range of 100 - 200 keV/c². If the axion is that light, it can only decay into two photons, and its lifetime should be around 10⁻¹ sec. For masses \gtrsim the MeV scale, the preferred decay mode is to an e^+e^- pair. The detailed properties of the axion depend, however, on which model of the weak interaction is adopted.

The production of axions in hadronic experiments is expected to be via the decays of heavy vector mesons ($J/\Psi, \Upsilon \rightarrow \gamma a$) or kaons ($K \rightarrow \pi a$), in analogy with the Higgs (see Section 2.1.2). Axion 'bremsstrahlung' from electrons may also contribute. The coupling of the axion to nucleons is very much smaller than the corresponding pion-nucleon coupling ($g_{a^0 NN}^2 \sim 10^{-7} g_{\pi NN}^2$) and one may expect such particles, if they exist, to escape from an experimental detector without interaction. Present limits on the existence of axions are reviewed in Section 2.2.4.

2.2 Experimental status

This section reviews current mass limits for the hypothetical new particles of Section 2.1. In the case of neutrinos, a brief summary of limits for rare (high-order) or non-Standard processes will be given.

2.2.1 Weak processes

In Section 2.1.1, it was noted that Standard Model neutrino production should also be accompanied by production of a charged lepton (at energies where Z^0 production is kinematically impossible). Present experimental limits [23,24,25,26,27] on rare processes involving neutrino production *without* an associated charged lepton are shown below:-

Process	Limit
$\pi^0 \rightarrow \nu\bar{\nu}$	$< 3.1 \times 10^{-6}$
$K^+ \rightarrow e^+ \nu\nu\bar{\nu}$	$< 6 \times 10^{-5}$
$K^+ \rightarrow \mu^+ \nu\nu\bar{\nu}$	$< 6 \times 10^{-6}$
$K^+ \rightarrow \pi^+ \nu\bar{\nu}$	$< 1.4 \times 10^{-7}$
$n \rightarrow p\nu\bar{\nu}$	$< 9 \times 10^{-24}$

Table 2.3: Experimental limits on neutrino production without accompanying charged lepton.

The fact that there are relatively weak limits on these processes is a direct consequence of the experimental difficulties involved in detecting neutrinos.

2.2.2 The Higgs

Until recently, the only firm experimental limit on the Higgs mass was $m_{H^0} > 14 \text{ MeV}/c^2$, arising from the non-observation of a long-range component to the nuclear force [28,29] and from the negative result of direct searches for Higgs bosons emitted in nuclear decay [30].

Higgs production has been searched for extensively in the decays of heavy mesons, for example $K \rightarrow \pi H^0, B \rightarrow K H^0$. In all cases, there are serious theoretical difficulties, particularly in the calculation of expected branching ratios. Of the more reliable results, the CUSB group [31] has looked for the decay $\Upsilon \rightarrow H^0\gamma$ and used this process to exclude

the range $600 \text{ MeV}/c^2 < m_{H^0} < 3.2 \text{ GeV}/c^2$. In conjunction, the bounds from B , Υ and K decays, probably exclude the range below $4 \text{ GeV}/c^2$, though it should be emphasised that this is not a firm limit. Further details can be found in [32]. Limits on a low-mass Higgs have been studied [33] by examining the possibility of detecting a Higgs from pion disintegration through its decay into e^+e^- in neutrino experiments. The limit thus obtained is $m_{H^0} > 40 \text{ MeV}/c^2$.

The complex question of limits on the Higgs mass has recently been clarified by results from LEP. By searching for decays of the Z^0 into a Higgs plus a fermion-antifermion pair ($Z^0 \rightarrow H^0 l+l-$, $Z^0 \rightarrow H^0 \nu\bar{\nu}$, $Z^0 \rightarrow H^0 q\bar{q}$), unambiguous constraints on the Higgs mass have been obtained [34,35]. At the 95% confidence level, a Higgs mass in the range $32 \text{ MeV}/c^2 \leq m_H \leq 19.3 \text{ GeV}/c^2$ is excluded. In conjunction with the limits described above, this gives a rather strong (albeit model-dependent) indication that a light Higgs is excluded. A definitive result from LEP (requiring higher statistics) is expected in the near future but at this stage, the possibility of a light Higgs cannot be conclusively eliminated.

2.2.3 The gluino

Existing limits on the gluino mass come from four types of experiment. They are stable particle searches [36,37,38], low-energy e^+e^- colliders [39,40], high-energy hadron colliders [13,14,15] and fixed-target hadron beam dumps [41,42,43]. Figure 2.6 shows the presently excluded regions in the gluino/squark mass plane. In Section 2.1.3 it was noted that light gluinos, if they exist, should be abundantly produced in hadron-hadron interactions. In this section, we will review gluino mass limits from hadron experiments, show that the ‘window of opportunity’ (regions I, II and III in Figure 2.6) is still open, and present the case for further experiment to close this window.

Mass limits from hadron colliders

Gluino production at hadron colliders should be seen as monojet, dijet and trijet events with missing energy-momentum (due to escape of neutralinos). This missing energy signature is valid only for gluino lifetimes less than $\sim 10^{-10}$ sec; for longer lifetimes, a significant fraction of produced gluinos will reach the calorimeter before decaying, resulting in a degradation of the missing energy signal. In practice, collider detectors are only sensitive to missing *transverse* momentum and this signature has been used to select candidate gluino events. Analysis of monojet and dijet events at the CERN Sp \bar{p} S collider [13,14,15] has shown that the observed rates are compatible with Standard Model




Figure 2.6: Excluded region of gluino and squark masses. Regions I and II are not yet excluded. Region III is excluded by experiment NA3 [43] for certain assumptions on the gluino interaction cross-section.

processes ($\tau, b\bar{b} + c\bar{c}, W, Z + g, t\bar{t}, \dots$). Since the gluino production cross-section falls with increasing gluino mass, this can be used to set a lower limit on its mass. Using a model for supersymmetry in which the photino is the lightest supersymmetric particle (assumed massless), the squarks are mass-degenerate, and ignoring processes involving the supersymmetric partners of the W, Z , Higgs and known leptons, the model thus contains only two free parameters, the squark and gluino masses $m_{\tilde{q}}$ and $m_{\tilde{g}}$. A limit of $m_{\tilde{g}} > 79 \text{ GeV}/c^2$ (independent of the squark mass) has been obtained. However, it has been argued [44,45] that uncertainties in the cross-section for light \tilde{g} production and decay, particularly in the way an energetic light \tilde{g} fragments to give an R-hadron, are such that a very light gluino (mass $\lesssim 4 \text{ GeV}/c^2$) cannot be excluded. The final mass limit corresponds to the area marked 'UA1' in Figure 2.6.

Mass limits from fixed-target hadron experiments

Limits on light gluinos are provided by hadron beam-dump experiments. These fall into two general classes. In the first case, the experimental technique is such as to be sensitive only to $\tilde{g} \rightarrow q\bar{q}\tilde{\gamma}$ decay occurring before the \tilde{g} (or more accurately, the R-hadron containing the confined \tilde{g}) has a chance to interact. Then the $\tilde{\gamma}$ decay product escapes from the dump giving the characteristic missing energy signature. Two variants of this scheme are commonly used. First, the dump is sometimes replaced by an active calorimeter to allow direct measurement of the energy lost via weakly interacting particles. Second, a 'neutrino' detector downstream of the dump may be employed to detect $\tilde{\gamma}N$ interactions. This option however, restricts the range of $\tilde{\gamma}$ interaction cross-sections to which the experiment will be sensitive. Finally, one can look for R-hadrons in the lifetime range such that they travel an observable distance before decay. We will review limits from each of these categories in turn.

The calorimeter search provides the most general approach in attempting to detect gluino production in hadron interactions (subject to the lifetime constraint noted above). The expected signature is undetected energy and momentum carried by neutralinos in events without final charged leptons. The best limit to date [46] comes from an experiment using a 400 GeV/c diffracted proton beam incident on a fine-grained iron target-calorimeter with energy resolution $\sigma/E \sim 70\%/\sqrt{E}$. The total data sample corresponds to approximately 2×10^5 proton interactions. The spectrum of the observed hadron energy in events without a muon is very nearly a Gaussian. A broadening of the low-energy tail is seen however and this is used to place an upper limit on R-hadron production of $33 \mu\text{b}$ ($8 \mu\text{b}$) for R-hadron masses of $1 \text{ GeV}/c^2$ ($3 \text{ GeV}/c^2$) respectively. This analysis is limited primarily by event statistics. Inability to trigger on interesting events with large

missing energy results in poor statistics in the low-energy tail. In addition, the background contribution from charm decays cannot be quantified by measuring the associated muon spectrum. This introduces large uncertainties in understanding the origin of any non-Gaussian behaviour in the low-energy tail.

Several groups [41,42,43] have used beam-dump experiments to search for evidence of gluino production and decay. In this mode, one searches for the interaction of a decay neutralino downstream of the dump in an appropriate detector (typically a bubble chamber or spectrometer). From Section 2.1.3, we recall that neutralino-matter interactions are expected to be via the channel $\tilde{\gamma}q \rightarrow \tilde{g}q$, a process which experimentally resembles neutral-current neutrino interactions. In all cases, the rate of observed neutral-current events is compared with predictions based on the measured rate for charged-current events. This allows one to distinguish between neutrino interactions and a possible excess signal due to interactions of some new particle. Over the past decade, beam-dump experiments have become increasingly sophisticated, looking not only for a simple excess of neutral-current events, but studying in detail the differences between neutrino and neutralino induced events (e.g. $\tilde{\gamma}$ induced events should have a higher average energy than neutral-current events), or identifying decay products (allowing a more precise determination of backgrounds). These experiments have used proton beams at 400 GeV/c and pion beams at 300 GeV/c. The most restrictive limit [42] rules out gluino masses below 4 GeV/c² for $m_{\tilde{q}} \simeq m_W$, but still permits $m_{\tilde{g}} < 4$ GeV/c² for larger squark masses. This is shown by the area marked 'BEBC' in Figure 2.6. All the beam-dump experiments have the disadvantage that they can only measure the product of gluino production cross-section and neutralino re-interaction cross-section. This imposes additional assumptions on an analysis which is already rather model-dependent.

The most recent light gluino search [47] uses a calorimeter in conjunction with a nuclear emulsion target, aiming to observe both the secondary vertex and the missing energy signature associated with gluino (or R-hadron) decay. The total data sample comprises approximately 5×10^5 pions of 350 GeV/c momentum, incident on the emulsion target. Missing energy is measured by a uranium-iron-scintillator calorimeter one metre downstream of the target, with angular acceptance of 250 mrad and energy resolution $\sigma/E \sim 90\%/\sqrt{E}$. A downstream muon spectrometer identifies and vetoes muon candidates, thus reducing background from semi-leptonic decays of heavy-flavour particles. The observed decay topologies are claimed to be compatible with those expected from $D\bar{D}$ background, allowing a lower limit $m_{\tilde{g}} > 3.5$ GeV/c² to be placed for $m_{\tilde{q}} \simeq m_W$, with a weaker limit as $m_{\tilde{q}}$ increases. The short-track search has the disadvantage that it is sensitive only to a window of lifetimes $\sim 10^{-14}$ sec. There are, additionally, problems due to secondary interactions and the limited acceptance of the calorimeter.

In summary, the ‘window of opportunity’ for a light ($m_{\tilde{g}} \sim 3 \text{ GeV}/c^2$) gluino is still open and is amenable to further experimental investigation.

2.2.4 The axion

In the original axion model [18], the axion mass and its coupling to stable particles are inversely proportional to the scale of the Peccei-Quinn symmetry breaking Λ_{PQ} , taken to be equal to the scale of the electroweak symmetry breaking Λ_{EW} . The mass and couplings are then completely fixed in terms of one parameter, the ratio of the vacuum expectation values of two Higgs particles (the minimal model having two Higgs doublets). Extensive searches have been made for such a very light, neutral pseudoscalar boson, covering a range of techniques from reactor neutrino experiments to beam dump experiments [48,49,50]. The results of these searches have been negative, although one beam dump experiment [51] claims to have an indication for axions with $m_a = 250 \pm 25 \text{ keV}/c^2$. A more recent search, prompted in part by the apparent observation of a $1.8 \text{ MeV}/c^2$ state decaying into an electron-positron pair in heavy ion collisions [52,53], has been made using an 800 GeV proton beam dump experiment. This experiment [54], in conjunction with earlier $g - 2$ measurements, rules out the existence of an axion lighter than $2.4 \text{ MeV}/c^2$ coupling only to e^+e^- . A summary of current limits on the axion mass and lifetime is presented in Figure 2.7.

This simple model has since been extended, by the introduction of either:-

1. a new strong interaction (making the axion heavy enough to avoid the experimental constraints), or
2. a new scale of the Peccei-Quinn symmetry breaking.

In the second case, the axion couplings are weak enough to avoid all experimental limits; this is the invisible axion model.

If the axion is found not to exist, the theoretical alternatives are threefold. Firstly, that $m_u = 0$; secondly, that a different weak interaction model is necessary, and thirdly, that weak symmetry breaking does not proceed via the Higgs mechanism but through some dynamical symmetry breaking.




Figure 2.7: Limits on mass and lifetime of an axion-like particle from a). E605 [54], b). KEK [55], c). E613 [56] and d). SLAC E56 [57]. Also shown is the limit from $g - 2$ experiments.

2.3 Experimental detection of weakly interacting neutral particles

Following the review of theoretical and experimental motivations for a weakly interacting neutral particle search, this section provides an outline of the experimental requirements for such a study. The general principles only will be discussed: Chapter 3 describes the extent to which these principles are realised in the HELIOS detector.

Given the theoretical and experimental constraints of Sections 2.1 and 2.2, the emphasis will be on a gluino search. The experimental approach adopted will, however, be shown to be sensitive to production of any particle with small interaction cross-section ($< 1\%$ of the typical hadronic cross-section) produced either directly or in secondary decays (provided only that the decay length is short compared to the interaction length).

The most general signature for production of a stable weakly interacting neutral particle will be missing energy-momentum in some idealised hermetic detector. To obtain the best possible measurement of the 'missing' energy associated with WINP production requires an accurate determination of the energies of *all other* particles produced in the interaction. For energy measurement of particles with electromagnetic and hadronic interactions (including neutrals), a calorimeter is the preferred device. In addition to being a fast, compact detector, it has the advantage of improved measurement accuracy with increasing particle energy. It should provide 4π -coverage (in the centre-of-mass frame) of the interaction point. This provides the motivation for calorimeter searches. With some loss of generality, a search might be based on the detection of interactions of weakly interacting particles in matter. Immediately, a constraint is placed on the range of possible interaction cross-sections. This is the idea behind beam dump searches and would be interpreted as excess neutrino induced neutral current events. In the case of secondary production (for example, $\tilde{g} \rightarrow q\bar{q}\tilde{\gamma}$), it may be possible to observe decay vertices displaced from the interaction vertex. This again entails some loss of generality, with constraints on the lifetime of the parent particle. Emulsion experiments have adopted this approach.

Figure 2.8 indicates the lifetime range accessible to a calorimeter search. For direct production of a weakly interacting neutral, the lifetime limit is imposed by the requirement that the particle escape from the detector before decay. For secondary production, the constraint is that the parent particle decay before interaction. Also shown is the sensitivity for other experimental configurations. This highlights a key advantage of the calorimeter search, namely its sensitivity over the widest possible range of lifetimes.

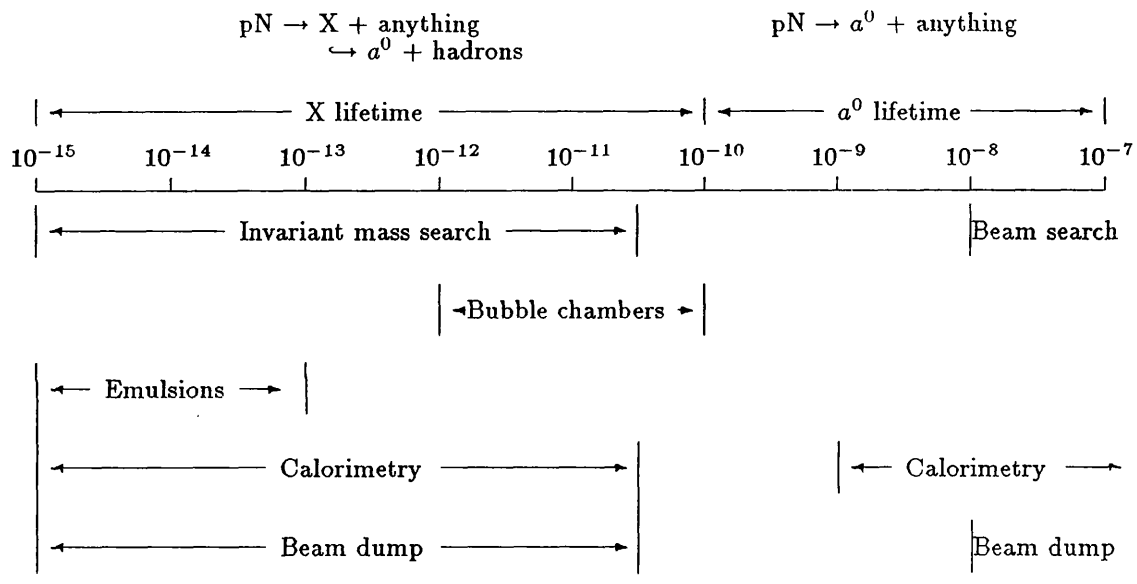


Figure 2.8: Experimental sensitivity to weakly interacting particles (a^0 is a weakly interacting neutral particle produced in the interaction, X is an unstable hadron which decays into an a^0).

With the expectation of Standard Model sources of weakly interacting neutral particles (neutrinos), it is important to be able to identify those events and, ideally, veto them. As was shown in Section 2.1, such events should be accompanied by a charged lepton. Realisation of a lepton veto therefore requires identification of electrons and muons over as much of the phase space as possible. This is in principle possible though the demands on the detector are somewhat conflicting: for electrons, one needs identification close to the interaction vertex (to suppress conversions and Dalitz decays); for muons, one needs a hadron absorber (dense matter) close to the interaction vertex to suppress pion and kaon decays. Given an ideal detector, with good efficiency and complete phase space coverage, this would be the most appropriate apparatus for this study. In practice, simultaneous electron and muon identification forces a compromise to be made between signal cleanliness in one channel and backgrounds in the other.

It is useful at this point to distinguish between prompt and non-prompt weak decays. Prompt will be taken to refer to decay times typical of the heavy quark flavours (decay length $c\tau \sim 10^{-4}$ m); non-prompt to decay times for π 's and K 's ($c\tau \sim 10^1$ m). For an electron-optimised detector (measurement close to the interaction vertex with minimal amount of material), the rate of 'background' in the muon channel is increased due to non-prompt decays of π 's and K 's. This problem is particularly acute if the muon coverage is not complete since some of these decay muons may then escape detection. For a muon-optimised detector (dump close to the interaction vertex), electron information is lost completely due to absorption.

This outlines the factors involved in the choice between target ('electron-optimised') and beam dump ('muon-optimised') modes in a calorimeter search. An additional consideration in beam dump mode is the question of a variable-density dump to determine experimentally the remaining background from pion and kaon decays. Since the flux of these non-prompt neutrinos is expected to scale with the inverse dump density, one can extrapolate to infinite density ($\lambda_i = 0$)², thereby subtracting the non-prompt contribution. The flux of prompt neutrinos from the semi-leptonic decay of heavy quarks remains unaffected because the decay path of the heavy quarks is negligible compared to λ_i (typically ~ 15 cm).

Within the framework of the HELIOS physics programme, we have chosen to study both possibilities, but have concentrated on the beam dump where backgrounds are expected to be smaller. As will be seen in Chapter 3, empirical results on the detector performance (calorimeter resolution) reinforce this decision and show the beam dump to be the preferred running mode for a weakly interacting neutral particle search.

² λ_i is the nuclear interaction length.

Having reviewed the requirements for a weakly interacting neutral particle search, we conclude with a brief introduction to the HELIOS detector, showing the extent to which it meets these requirements and the improvements it offers over previous experiments.

- The experiment uses a primary proton beam of 450 GeV/c momentum, a higher energy than has previously been used in searches for weakly interacting particles.
- The calorimeter provides the densest possible beam dump (uranium), thereby reducing background from non-prompt decays.
- The calorimeter energy resolution in beam dump mode is $\sigma/E \sim 45\%/\sqrt{E}$, considerably better than in any previous experiment.
- Muon identification is provided by a powerful charged-particle spectrometer.
- The detector is able to trigger on missing energy, yielding improved statistics on interesting events.
- The detector is capable of running in both target and beam dump modes, with the corresponding optimisations in electron and muon channels. This allows thorough cross-checking on the nature of any candidate signal and an improved understanding of Standard Model backgrounds.

A detailed description of the detector is given in the following chapter.

Chapter 3

The HELIOS detector

In this chapter, the HELIOS detector is described, emphasising the components relevant for the missing energy analysis. The chapter is divided into four sections: a review of the detector itself, a study of both the calorimeter and muon spectrometer performances, and a brief description of data-taking procedures.

3.1 Survey of the experimental programme

The HELIOS (NA34) experiment is situated in the H8 beam line of the CERN Super Proton Synchrotron (SPS). In normal fixed-target mode, the SPS accelerates protons to a maximum momentum of 450 GeV/c, at intensities up to 10^{12} protons per burst (2.4 s). Beams of primary protons and secondary hadrons (produced by means of a production target) are both available at high intensities. In 1986, the spectrum of SPS physics capabilities was further enhanced when light ions were successfully accelerated as part of the CERN relativistic heavy-ion programme. Fully-stripped oxygen and sulphur ions are now available at momenta of up to 200 GeV/c per nucleon.

The HELIOS detector was designed to measure observables in fixed-target collisions with both hadronic and nuclear incident beams. The lepton programme, using the proton beam, aims to settle open questions in the production of electrons, muons and neutrinos. Prominent among these are e/μ universality, anomalies in the production of single leptons, the contribution of charm decay to lepton pair (Drell-Yan) production, and the existence of 'anomalous' low mass pairs. The ion programme seeks to examine details of ultra-relativistic nucleus-nucleus collisions. It is expected that extended volumes of high-energy density hadronic matter should be produced in such collisions, and that this may provide




Figure 3.1: Layout of the HELIOS (NA34) detector.

one of the necessary conditions for the formation of a plasma of 'deconfined' quarks and gluons.

Figure 3.1 shows a plan-view of the detector.

Though the two physics programmes look rather independent at first sight, the HELIOS detector was designed as a multi-purpose device, able to operate in both environments and to take advantage of the areas of overlap between the two modes.

The set-up for the lepton programme has four major components:-

- the target and vertex-detector system, including silicon charged-multiplicity counters,
- the electron spectrometer, using high-performance drift chambers, a 'weak-field' calorimetrised magnet, a transition-radiation detector (TRD) and a high-resolution electromagnetic calorimeter,
- a large-acceptance muon spectrometer,
- 4π calorimetry coverage for measurement of missing energy.

This set-up gives rise to two alternative sub-configurations: ‘target’ mode, in which all of the above components are used, and ‘beam dump’ mode, in which the target/vertex system and electron spectrometer are unused.

For the ion programme, the 4π calorimetry allows detailed study of the energy flow and is a powerful tool for selecting events of interest. This is complemented by the muon spectrometer, an external few-particle mass spectrometer at large angles, and several soft-photon detectors. The electron spectrometer is unused.

The HELIOS co-ordinate system takes the beam direction to define the z -axis. The y -axis is vertically upwards and we use a right-handed co-ordinate system. The origin of the co-ordinate system is taken to be at the centre of the wire target (see Figure 3.6) in ‘target’ configuration running. This convention is maintained for the beam dump when the target is removed. The polar angle θ is measured with respect to the beam line and ϕ is the azimuthal angle measured from the x -axis in the x - y plane.

The following sections describe the beam and detector components which are relevant for the missing energy study. The emphasis is on operation in beam dump mode (the results of which form the basis of this thesis). Reference will be made, however, to data-taking in the ‘target’ configuration with a wire micro-target since it is not until results on the respective calorimeter performance are compared (Section 3.2.4) that a final choice between the two modes can be made.

3.1.1 The proton beam

This analysis relies heavily on the missing energy signature associated with the production of weakly interacting neutral particles in an interaction. To obtain the best possible measurement, it is important to have precise knowledge of both incoming *and* outgoing particle energies.

The HELIOS direct-lepton study requires high-intensity beams incident on a thin (50 μm diameter) wire target. To achieve this, the beam is focussed to a spot of similar size (transverse size of 50 μm full-width at base) with intensities up to 10^6 protons per second. The H8 beam line is shown in Figure 3.2.

The experiment uses a primary beam giving an excellent definition on the incident proton energy, corresponding to a momentum resolution $\frac{\delta p}{p}$ of $\pm 0.1\%$.




Figure 3.2: The H8 beam line.

3.1.2 Beam telescope and general trigger

This section reviews the definition of a valid beam particle, the general trigger and data-acquisition systems. The description concentrates on running in beam dump mode; modifications for the case of 'target' running are included where relevant.

In beam dump mode, beam particles are defined by a series of plastic scintillator counters placed between 50 and 281 cm upstream of the nominal target position (and ~ 5 m upstream of the calorimeter dump). The counters and the associated local shielding are shown schematically in Figure 3.3. The beam telescope not only gives the definition of a beam particle, but must also provide the time reference t_0 ($\equiv 0$ secs) for the experiment and recognise beam 'halo' and upstream interactions. As will be seen in Sections 3.1.4 and 4.2, the correct design of this telescope is essential for missing energy studies.

The dimensions of all the counters are given in Table 3.1. They are read out by fast photomultiplier tubes, the signals from which are discriminated for use in the beam counter logic and for off-line recording in TDCs. The counters B7, B6 and B3 are additionally instrumented with ADC and multi-hit TDC readouts. The latter provide time information on additional hits within a window of $\pm 1\mu\text{s}$ around the trigger particle and are important in removing pileup (see Sections 3.1.4 and 4.4).




Figure 3.3: Beam-defining counters (not to scale).

Counter	Position (z-axis)	Geometry	Active Area	Beam hole
B7	-281 cm	Square	40 x 40 cm ²	radius 0.5 cm
B6	-269 cm	Circular	radius 0.6 cm	-
B1	-140 cm	Rectangular	2.0 x 0.5 cm ²	-
B2	-140 cm	Rectangular	8.5 x 1.5 cm ²	-
B3	-50 cm	Rectangular	7.0 x 6.0 cm ²	radius 0.6 mm

Table 3.1: Dimensions and positions of the beam counters (z-axis positions are relative to nominal target centre).

The three counters, B6, B1 and B2, are aligned on the nominal beam axis. A beam particle is required to satisfy the beam coincidence $\equiv B6 \cdot B1 \cdot B2$. A 2-fold coincidence is necessary to remove fake triggers due to either non-beam particles (radioactivity of uranium calorimeter stacks, cosmic rays) or noise in the photomultiplier tubes and associated electronics. The 3-fold coincidence is used for technical reasons (a coincidence of the signals from counters B1 and B2 having been pre-defined in the trigger logic). Two additional counters, B7 and B3, act as veto counters. They are also centred on the nominal beam line, but have holes such that on-axis particles do not hit the active areas. The combined beam counter/shielding layout is designed to achieve the following aims:-

1. define genuine on-axis beam particles (using B6, B1 and B2)
2. protect the calorimeter from upstream interactions and beam halo (using beam-pipe shielding and B7)
3. prevent fake vetoes due to backsplash from the calorimeters (using B7 shielding)
4. reject interactions in beam-defining counters (using B3)

The overall trigger requirement from the beam counters is therefore

$$B = \overline{B7} \cdot B6 \cdot B1 \cdot B2 \cdot \overline{B3}.$$

The coincidence generates a valid beam (VB) strobe only if a beam coincidence has not occurred within a 'before protection' interval $\Delta_{1b} = 800$ ns. A late beam abort is issued if a further beam coincidence occurs within an 'after protection' window of $\Delta_{1a} = 300$ ns. This combined before and after protection is intended to prevent signal overlaps from different events in the various detectors. A detailed account of its importance, in particular for the calorimetry, is given in Section 3.1.4.

In beam dump mode, the beam particles are absorbed and interact in the calorimeter dump. For 'target' running, the beam is focused on a 2 cm long, 50 μm diameter wire of beryllium or tungsten. The components of the trigger logic relevant for the present discussion are shown in Figure 3.4. Trigger decisions are made at each of three stages: the pre-trigger, the first level, and the second level trigger or VFB, each of which is a computer-programmable logic array. The function of each is described briefly below.

In beam dump mode, the beam logic is the main component of the pre-trigger which causes the information from each of the detectors to be strobed in to ADCs, TDCs, etc.

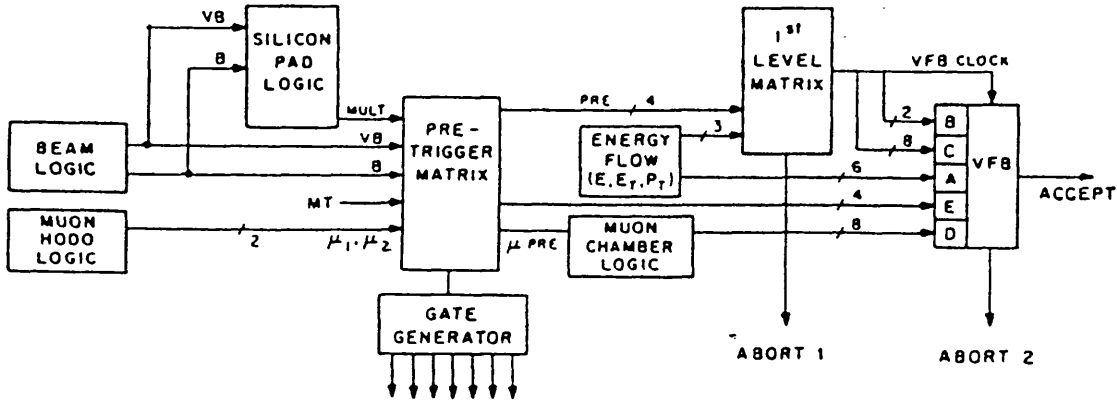


Figure 3.4: Block diagram of the principal trigger logic components.

This happens at approximately $t = t_0 + 340$ ns. Once issued, the pre-trigger sets a BUSY flip-flop which remains set until the event is either aborted or a reset is received from the data acquisition system (CLEAR). In target mode, the silicon pad (Section 3.1.3) is the other crucial element in the pre-trigger logic, defining an interaction in the wire target. Discriminator readout of each of the 400 pads is initiated by a beam coincidence and allows a fast measurement of the total charged-multiplicity to be made. Under normal run conditions, the interaction requirement is for ≥ 3 hits in the pad array. In both modes, a pre-trigger may also be an artificial ‘empty’ (MT) pre-trigger generated randomly and used for controlling pedestal stability etc. of the various detector components.

Issuance of a pre-trigger begins the first level trigger logic. For the analysis presented in this thesis, the only relevant components here are the Energy Flow Logic (EFL) and the muon logic. The Energy Flow Logic (Section 3.1.4) begins to process the calorimeter information which is digitised and combined to obtain values of E_{tot} (and additionally E_t and p_t). These values are compared to a series of comparator thresholds and, if a trigger requirement is satisfied, a trigger flag signal is generated and sent to the first level trigger, where it is logically-ANDed with the signals from the beam counters. The muon logic is also activated at this level, the spectrometer chambers being used to make a fast estimate of the momentum of any detected muon. This information may be included in the final EFL energy summation if desired. At the end of the first level trigger, the system either issues a first level abort or initiates the second level trigger. An abort causes a fast clear to be given to the data acquisition units which stored information at pre-trigger time.

Once all the trigger subsystems have completed their trigger evaluation ($\sim 1 \mu\text{s}$), a Very Fast Bus (VFB) is clocked to reach a final decision. This final element in the trigger logic is a Random Access Memory (RAM) Programmable Logic Array (PLA) which allows up to 51 different combinations of signals to be considered as interesting triggers. The relative populations of various types of triggered events are weighted with downscale factors n_i , such that only every $n_i + 1$ events of type i are recorded. In this way, the fraction of rare events written to tape can be optimised with respect to more frequent triggers which would otherwise flood the data-acquisition system. The division of the trigger logic into a first and second (VFB) level is a consequence of the fact that some detector signals are available after approximately 300 ns while others take more than 1 μs . This improves the live-time of the trigger system since uninteresting events may often be rejected before the complete event information is available.

The complete information for an event (typically ~ 7.5 kbytes of data) is transferred under CAMAC control to a recording buffer in a VAX data acquisition computer.

A complete description of the HELIOS trigger electronics may be found in [58].

3.1.3 The silicon pad array

Approximately 2 m downstream of the beam-defining counters and 16.1 cm downstream of the nominal target position ($z = 0$) is a fine-granularity silicon detector with segmented pad readout. The device is circular (diameter 30 mm), containing a hole of 0.5 mm diameter through which the beam passes. It is centred on the nominal beam axis, to a precision of a few microns, by means of an optical telescope. For the purposes of the beam dump, it is used only in the off-line analysis as a check on off-axis beam halo particles or secondary particles produced by interactions upstream in the beam line. For 'target' running, it serves as an interaction counter (measuring charged-multiplicity) and forms an important part of the electron trigger.

The detector is 300 μm thick and consists of an array of 400 pads varying in size from $0.02 \times 0.167 \text{ cm}^2$ to $0.167 \times 0.66 \text{ cm}^2$. The detector is operated at full depletion and the boundaries between individual segments are made of 40 μm of fully active oxide. The layout of the pad is shown in Figure 3.5.

The pad elements are connected to preamplifiers located close to the detector itself to reduce the input capacitance which is the major source of noise. The signal from each individual pad is then split, a fraction going to discriminators (used in the wire

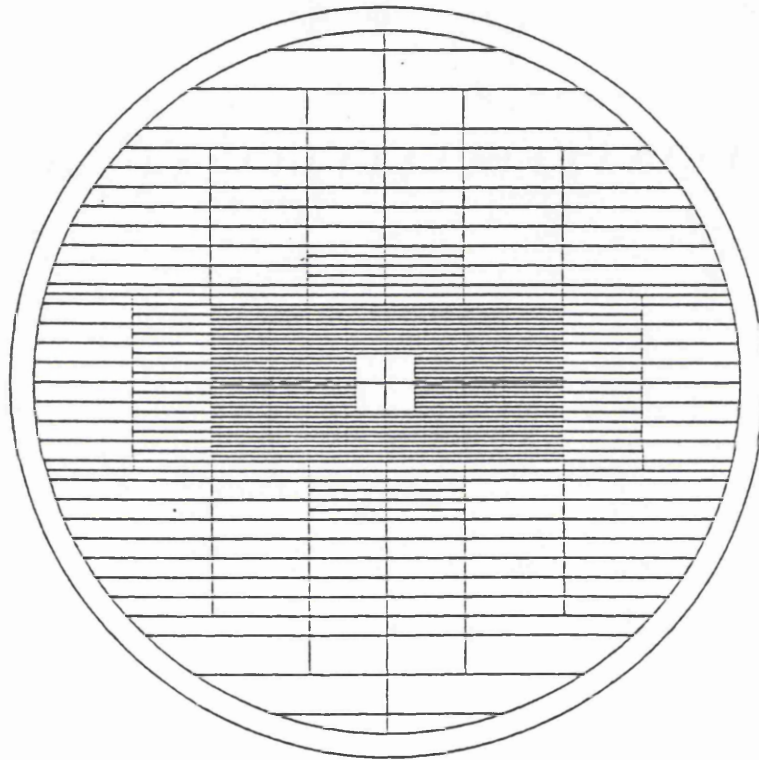


Figure 3.5: The silicon pad array.

target interaction and electron triggers), and the remainder to charge-integrating ADCs for precise measurement. The signal-to-noise ratio, that is, the ratio of the most probable energy loss to the r.m.s. width of the pedestal, is ≥ 10 for single particles. During the 1988 data-taking run, about 12% of the active area of the pad was not operational due to either excessive noise or dead ADCs.

In ‘target’ configuration running, the silicon pad is 16.1 cm downstream of the wire target and covers the pseudorapidity interval $3.1 < \eta_{lab}^1 < 7.2$.

More details on the silicon pad design and performance can be found in [59,60].

3.1.4 Calorimetry

For the purposes of this analysis, the HELIOS calorimetry is required to produce the best possible measurement of the *total* energy of an event. The emphasis is therefore on good energy containment (minimising leakage) and accuracy (optimising the calorimeter resolution).

The following section contains a detailed description of the calorimeters and their corresponding readouts. The forward calorimeters (situated in the beam-line) will be described first; they provide the dump for ‘beam dump’ geometry missing energy running. The backward or wide-angle calorimeters are of use only to provide a veto on upstream interactions in beam dump mode. When operating with the ‘target’ geometry, they provide part of the $\sim 4\pi$ calorimeter coverage (in the centre-of-mass frame) for interactions in the wire target, corresponding to opening angles between 6.3° ($\eta = 2.9$) and 95.7° ($\eta = -0.1$). For orientation, the reader is referred to Figure 3.6.

The uranium/liquid argon calorimeter

The first 4.5 nuclear interaction lengths (λ_i) of the forward calorimetry are provided by a uranium/liquid argon sampling calorimeter (ULAC). The choice of depleted uranium as the passive medium allows the construction of a compact calorimeter. At the time of its design, it was also known that shower-induced fission processes in uranium helped to equalise the response of the calorimeter to electromagnetic and hadronic components of the shower, thereby improving the energy resolution of the device. Further work [61] has confirmed this result but indicates that the path to achieving the compensation condition

¹The pseudorapidity η_{lab} is defined as $\eta_{lab} = -\ln(\tan(\theta_{lab}/2))$.




Figure 3.6: The HELIOS calorimeter layout. The target position indicated corresponds to 'target' configuration running; the target is removed in beam dump mode.

in a calorimeter requires careful tuning of both the passive *and* active components. The calorimeter operates in ionisation chamber mode (i.e. without any internal signal amplification), detecting ionisation charge of shower particles. For efficient operation, it is important to use a liquid in which the electrons released by the ionised atoms or molecules may drift freely over sufficiently long distances under the application of an electric field. Fulfilling these criteria, liquid argon was chosen as the active medium because it is dense and relatively inert. Absence of radiation damage also implies long-term operating stability.

The calorimeter is subdivided into an electromagnetic (EM) section, with tower readout, and an hadronic (HAD) section, with strip readout, as indicated in Figure 3.7. Viewed along the beam line, the instrumented area of the electromagnetic section is circular (with diameter 80 cm) whilst that of the hadronic section is square (with edge length 120 cm). At the centre of the electromagnetic section is a 5 cm-wide hole. This is important when HELIOS is running in 'target' mode where the interactions of interest take place in a target 4.2 m upstream of the ULAC. Non-interacting beam particles and/or projectile fragments do not then hit the electromagnetic section but shower in the hadronic section, allowing relative gain optimisation of the two sections. For the beam-dump experiment, the beam is steered through the hole to interact directly in the hadronic section.




Figure 3.7: Section through the ULAC, showing division into electromagnetic and hadronic sections.

The electromagnetic section has 1920 towers, mostly $2 \times 2 \text{ cm}^2$, and arranged in 2 floors. The first (most upstream) floor corresponds to 8 radiation lengths (L_{rad}), the second to $12 L_{rad}$. The longitudinal structure is based on a repeated 'unit' of

uranium	liquid argon	readout plane	liquid argon
1.7 mm	2.0 mm	1.7 mm	2.0 mm

giving a sampling fraction², f_{samp} , of about 8.5 %. The tower structure for one quadrant of the electromagnetic section is shown in Figure 3.8. It consists of four types of tower: triangular towers on the periphery of the octagonal central hole, normal towers ($2 \times 2 \text{ cm}^2$), large towers ($2 \times 4 \text{ cm}^2$) at the outer edge of the active area, and 'veto' towers, which are the pads with irregular shapes near the periphery.

The hadronic section is segmented laterally into x and y strips 2.5 cm wide and is divided in depth into 3 floors of $1.5 \lambda_i$ each. The repeated 'unit' in this case is

uranium	liquid argon	readout plane	liquid argon
3.4 mm	2.5 mm	1.7 mm	2.5 mm

²The fraction of the deposited energy detected in the active absorber layers of the calorimeter.




Figure 3.8: Towers in the electromagnetic part of the ULAC.

giving $f_{samp} \simeq 6.0\%$. The plates are grouped electrically into two interleaved, independent calorimeters, designated the primed and unprimed sets. This redundancy provides a backup if one set of planes is shorted out.

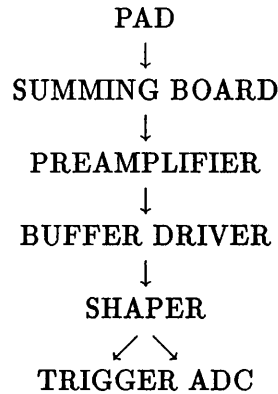
Liquid-argon calorimetry requires a precise control on the purity of the argon. Contamination by electronegative elements, oxygen in particular, may severely degrade the calorimeter response. The argon is enclosed in an hermetically-sealed vessel, cooled by liquid nitrogen and maintained at a slight overpressure. Purity levels with this system (oxygen contamination less than 0.3 parts per million) are such that the cryostat is normally filled only once per SPS running period (approximately 3 months). The ULAC maintains a voltage of 10 kV/cm under normal run conditions.

Analogue electronics

The following section describes the analogue electronics from the calorimeter electrodes to the input of:-

1. the digital trigger logic, and
2. the 'off-line' analogue-to-digital converters (ADCs).

The signal chain for the electromagnetic section is indicated schematically below



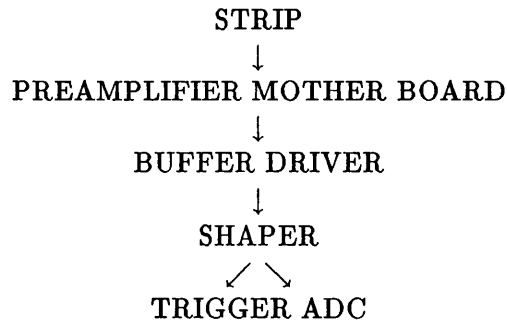
Under nominal operating conditions, a typical shower particle in the calorimeter produces an ionisation signal from the liquid argon with a rise-time of order 100 ns. The readout planes are multilayer printed circuit boards (PC) containing, in the case of the electromagnetic section, five layers of copper with G10 sandwiched between them. The planes are, in order, readout, ground, signal, ground, readout; the individual pads are connected to the signal plane via through-plated holes. High voltage is distributed through a resistive layer of epoxy silk-screened to the G10 boards. The pad signals are then fed via multilayer PC summing boards to charge-sensitive preamplifiers, where the negative input pulse is amplified and produces a positive-going signal for input to the shaping amplifiers.

The shaper module performs the following functions:-

1. produces a bi-polar signal (rise time around 200 ns)
2. allows both coarse and fine gain control under microprocessor control
3. sums signals for input to the trigger

The units themselves are modified FASTBUS modules. At this stage, the signal is split; one output going to the fast digital electronics of the Energy Flow Logic (see Section 3.1.4), the other to slow peak-sensing ADCs (PADCs) for off-line analysis. These ADCs are 12-bit peak-sensing LeCroy S2281 modules. The gate width for these units is ~ 540 ns under normal operating conditions.

The analogue signal chain for the hadronic section is shown below



The relevant differences between electromagnetic and hadronic sections are outlined below. Firstly, the readout planes are sandwiches of 3 (rather than 5) copper layers. The two outer planes are combined high voltage and readout planes while the middle plane supports the signal feedthrough. In this case, the high voltage is placed on the readout planes by means of a resistive layer silk-screened on to the printed circuit boards. Secondly, for the hadronic section, the preamplifiers sit next to the electrodes in the liquid argon. This is a response to constraints imposed by the HELIOS physics programme which demands calorimetry with a high rate capability in the very forward region. The shaping time of ionisation chambers is dictated by the length of cable linking the detector and the preamplifiers. To achieve the desired shaping times of the order of $1 \mu\text{s}$ or less, the preamplifiers must therefore sit as close as possible to the calorimeter, that is, in the liquid argon. They are maintained at optimal operating temperature of 140K (50K above the liquid argon temperature) by enclosure in an alloy box. Thirdly, the hadronic chain has an additional low-noise intermediate amplifier, the buffer driver, incorporating an additional computer-controlled gain switch. Finally, the shaping amplifier modules again produce signals for both the trigger and ADC readout chains. In this case, however, the respective shaping times are different: 100 ns rise-time for the trigger and 200 ns rise-time for the ADCs. Signals from the two interleaved hadronic sections are kept separate throughout the signal chain.

For more detailed information, the reader is directed to [62].

The uranium/scintillator calorimeters

The BEAM and VETO calorimeters

Immediately downstream of the ULAC are two calorimeters, the BEAM and VETO, built of U/scintillator and U/Cu/scintillator modules (UCAL), which add a further $6.0 \lambda_i$. Both types of module (see Figure 3.9) consist of metal plates interleaved with 0.25 cm




Figure 3.9: Isometric view of a U/Cu/scintillator module. The enlargement shows the arrangement of the optical readout and the sampling for the electromagnetic and hadronic sections.

thick scintillator plates. These modules are divided into an electromagnetic and hadronic section. The electromagnetic section is 6.4 radiation lengths deep and contains 2.0 mm thick depleted U plates. The hadronic section for the U/scintillator modules is $4.0 \lambda_i$ deep and has 3.0 mm U plates. For the U/Cu/scintillator modules, the hadronic section is $3.8 \lambda_i$ deep and has the U plates alternating with 5.0 mm Cu plates in the ratio 2:1. The pure uranium modules have the advantage of giving a more uniform response to electromagnetic and hadronic showers and were therefore placed near the nominal beam axis where the energy deposition is expected to be greatest. The physical positions of the two types of module are indicated in Figure 3.6.

The optical readout is designed to optimise the light yield to the photomultiplier tubes (PM). The scintillator plates consist of a PMMA³ matrix doped with 1% PBD-butyl scintillating agent and 10% naphthalene to increase the ultraviolet (UV) light yield. These plates ($\sim 120 \times 20 \text{ cm}^2$) are read out on both sides by six independent 2.0 mm Wave Length Shifter (WLS) plates, giving a $20 \times 20 \text{ cm}^2$ tower sub-structure to each module. In order to limit the light sharing between the towers of the most forward modules, the 6 towers are optically decoupled by slits of 0.5 mm width made in the scintillator plates with a 600 W CO₂ laser. The light containment in optically decoupled towers allows better energy localisation in the high energy and multiplicity environment of the forward calorimetry. The WLS is Plexiglass to which 80mg/l of BBQ⁴ shifting agent has been added. Doping in both the scintillator and WLS is such that the maximum absorption efficiency of the WLS coincides with the peak of the emission spectrum from the scintillator ($\sim 368 \text{ nm}$). The green light (peak $\sim 500 \text{ nm}$) from the WLS is brought to the PM by means of a light guide made of UV-absorbing acrylic glass. The optical coupling between light guide and PM is made by a 3 mm thick disc of elastic silicone rubber, again doped with a UV-absorbing chemical. This readout chain produces approximately 1.5 photoelectrons per minimum-ionising particle (mip) per scintillator plate.

The signal from the PM is divided between two readout chains; a third of the charge going to the Energy Flow Logic (see Section 3.1.4) for use in the trigger, the remainder to 12-bit charge integrating ADCs (QADCs) for off-line analysis. The gate width for these ADCs is 130 ns under normal operating conditions.

The backward calorimeters

The wide-angle calorimetry comprises mostly U/Cu/scintillator modules (described above) arranged in an almost hermetic box around the wire target. They form the WALL and

³Polymethyl Methacrylate

⁴Benzimidazo-benzisoquinoline-7-one




Figure 3.10: Artist's view of the calorimetrised dipole magnet.

BOX in Figure 3.6. The perpendicular distance from the target to the various walls of the box is approximately 120 cm in all directions.

The MAGnetised iron CALorimeter (see Figure 3.10) provides both the magnetic field for the electron spectrometer and the calorimeter coverage over the range $2.2 < \eta_{lab} < 2.9$ for 'target' interactions. It consists of 15 mm iron plates interleaved with 5 mm scintillator sheets, each divided into 24 'petals' covering 15 degrees in azimuth. The 24 towers are read out via wavelength shifter bars with optical fibres completing the path to the photomultiplier tubes. The MAGCAL has a maximum depth of $3.2 \lambda_i$.

Energy Flow Logic

The anode signals of the calorimeters are used for two purposes. A fraction of the signal is sent to analogue-to-digital converters which record either the peak voltage (for the ULAC) or total integrated charge (for the UCAL). The remaining fraction is diverted to the Energy Flow Logic (EFL) and is used on-line for triggering purposes.

The goal of the EFL is to provide a trigger on the total energy of an event and on its topology (that is, its transverse energy or momentum imbalance). For this experiment, the EFL logic can easily be configured to provide a trigger on *missing* energy i.e. events

containing particles which do not interact in the calorimeter. This is achieved by fast analogue and digital summing of the physical quantity E_{tot} obtained from the various calorimeters selected for inclusion in the trigger. In addition, there is the possibility to include in the trigger a fast measurement of the momentum of any penetrating charged particle seen in the muon spectrometer (the muon-update, see Section 3.1.5). The system has been designed to provide fast energy information, with the emphasis on achieving the best possible accuracy.

In the ULAC, the same cable provides signals for both the fast Energy Flow logic and for the slower PADC readouts (see Section 3.1.4). This is in contrast to the UCALs, where the signals are split at the base of the PM tubes (see below). The first stage in the ULAC EFL summing is an analogue passive sum at the board level. Each board contains resistors for the physical quantities E_{tot} , E_t , p_x and p_y (E_t for the electromagnetic section only). For each readout crate (20 boards), the EFL signals are merged on a common backplane connection and amplified. The final stage in the summation consists of linear summing units (merging signals from up to 4 crates) followed by programmable attenuators that provide up to 100 dB attenuation under computer control. The outputs from these attenuators (one per section of the ULAC) provide the signals for the flash ADCs (FADCs) in the Energy Flow Logic (see Figure 3.13).

For the UCALs, the actual analogue summing is carried out in three or four stages depending on the calorimeter component in question, and the corresponding electronic summing units are termed Σ_1 , Σ_2 , Σ_3 and Σ_4 . The splitting of the PM signal and the input stage to the first level (Σ_1) summing are shown schematically in Figure 3.11. The very first stage is a passive element, a resistor. Each resistor corresponds to a weight calculated according to the geometrical position of the particular calorimeter cell and the physical quantity of interest. The conductance of the resistor is proportional to that weight. In the case of E_{tot} , the weights are always unity. The additional outputs marked P_x , P_y and E_t are available for triggering on specific transverse-energy topologies but were not used for this experiment. A schematic representation of the cabling and summing layout for part of the calorimeter is shown in Figure 3.12. The summing hierarchy adds signals from progressively larger subsections of the calorimeter so that by the output of the Σ_4 modules, information is available for the complete calorimeter. All the units use fast transistors as current amplifiers with unit gain. The design ensures a linear response for currents up to 5 mA.

The final analogue sums are sent, via a differential driver unit, to 150 ns shaping amplifiers, and from there, via adapter units which allow fine adjustment of the summing chain gain, into flash ADCs (FADCs). The linearity of the whole chain is monitored by

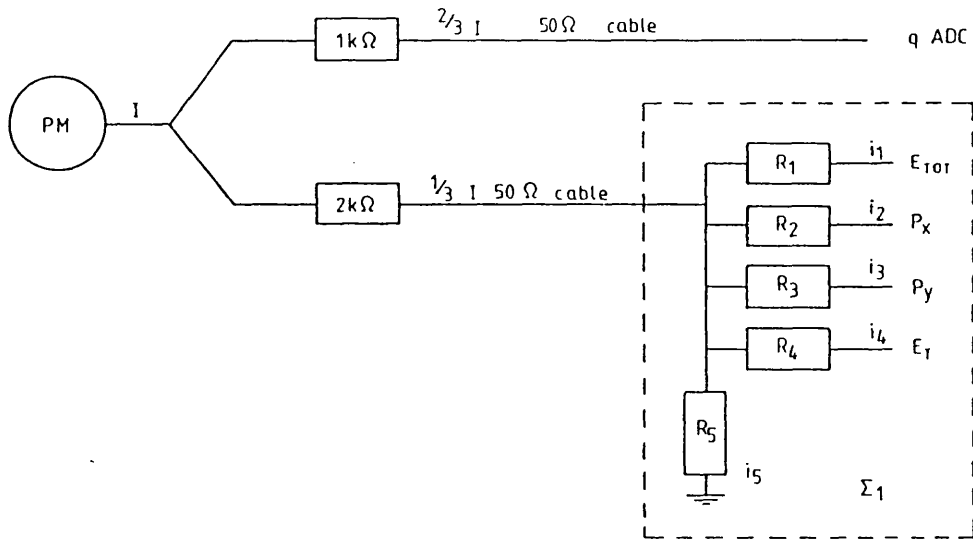


Figure 3.11: The splitting of the PM anode signal and input stage to the first level summing.

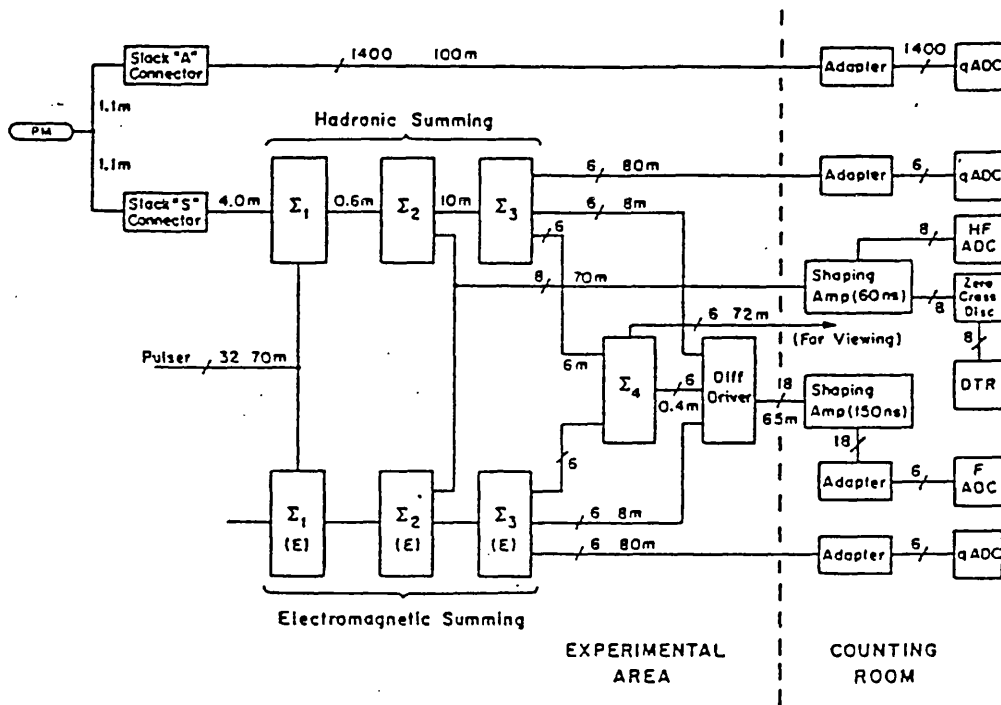


Figure 3.12: Schematic representation of the summing layout for the BEAM calorimeter.




Figure 3.13: The EFL digital summing logic.

injecting 50 ns pulses to test inputs at the Σ_1 summing level. The response at the FADC level is measured and found to be linear within 1%. The long-term stability is also of order 1%.

From the FADCs, signals are fed through a series of digital EFL modules to digital comparators for use in the parallel-trigger processor. Details of the digital logic are shown in Figure 3.13. These modules are under CAMAC control, which allows, via a selection of available comparator thresholds, a range of EFL triggers to be defined.

The FADC gates (obtained from the pretrigger with a suitable delay) occur precisely at the peak of the shaped signal. In this way, the calibration of the Energy Flow system is less susceptible to possible variations of timing in the trigger system. The timing of each FADC gate is studied during special runs using programmable delay generators. Once established, the optimum timing is fixed using cable delays.

The gains of the summing units corresponding to a given calorimeter component are optimised to keep electronic noise to a minimum while using the full FADC range for the signals expected from that component. The overall energy scale is fixed by the calibration of the charge ADCs and a knowledge of the beam energy (see Section 3.2.1).

Calorimeter before/after protection and the history FADC system

The performance of any calorimeter depends on the ratio of the time during which it is sensitive to shower particles from one event to the time between successive events. Clearly, if a particle interacts in the calorimeter before the signal collection from a previous particle is complete, the energy deposition measured by the calorimeter will be misleading. This problem, known as pileup, is particularly important for a missing energy analysis where contamination from overlapping events can create a fake missing energy signal.

Assuming it is possible to unambiguously identify such contamination off-line, this is not, a priori, a problem. However, in the case where one is aiming to trigger on rare events (those with large missing energy), an on-line veto on contamination is desirable to avoid flooding the data-acquisition system with corrupted data. In this experiment, the veto is provided by a protection window in the beam counter logic, which ensures event rejection if another particle is detected within the specified time window. For missing energy running, the most stringent constraint on the protection window comes from the uranium/liquid argon calorimeter. The PADC signals have a timebase of order $1 \mu\text{s}$. The before and after protection on the beam coincidence has been described in Section 3.1.2 ($\Delta_{1b} = 800 \text{ ns}$, $\Delta_{1a} = 300 \text{ ns}$). Its purpose is to remove pileup from genuine (on-axis) beam particles. Additional protection against pileup of off-axis particles ('halo' or upstream interactions) is provided by the veto counter B7; in this case only before protection $\Delta_{3b} = 50 \text{ ns}$ for technical reasons.

Off-line, the protection window can be extended to cover the full $\pm 1 \mu\text{s}$ using the multi-hit TDC readout of the beam counters (Section 3.1.2). However, it will be seen in Chapter 4 that this off-line protection is less-than-perfect. As an additional tool to study the effects of pileup, the HELIOS calorimeters are equipped with an array of sampling FADCs known as history FADCs. These units are 6-bit FADCs and read the time history of the signals at the level of the Σ_2 summing outputs in the EFL logic (the dashed lines in Figure 3.13). For the uranium/liquid argon calorimeter, the sampling is at 30 ns intervals over 960 ns (to be compared with the timebase of the single-particle signal response of $\sim 1 \mu\text{s}$). For the uranium/scintillator calorimeters, the sampling is at 10 ns intervals over 320 ns (the corresponding timebase is $\sim 200 \text{ ns}$ in this case). Pulse histories may then be studied off-line to look for evidence of pileup in the data sample. Details of the history FADCs may be found in [63].

3.1.5 The muon spectrometer

The muon spectrometer (see Figure 3.1) sits immediately downstream of the VETO calorimeter, providing position and momentum measurements of charged particles which are not absorbed by the calorimeters. For interactions in the wire target, the spectrometer is sensitive to particles produced with polar angles, θ , up to 100 mrad from the beam axis. For interactions at the front face of the ULAC in beam dump mode, the acceptance is improved, corresponding to $\theta \sim 130$ mrad.

The proportional wire chambers

The spectrometer is equipped with seven proportional chambers (PC0 to PC6), the principal features of which are described in Table 3.2.

The 32 planes each consist of a grid of anode wires, placed between two planes of graphited mylar which act as cathodes. The inter-wire spacing is 3 mm in all chambers giving a total of 28,544 wires in all. Anode-cathode gaps are 6 mm for all chambers except PC0 and PC1, where the separation is 7 mm. The gas mixture in the chambers is approximately 20.0% isobutane, 79.9% argon and 0.1% freon.

The nominal operating voltages, ie. on the plateau of the efficiency curves, were determined during beam tests. Where possible, these values (around 2700 to 2800 V) were maintained during running conditions. However, in some cases, these voltages drew unacceptably large currents from the chambers. In particular, the natural radioactivity from the uranium stacks of the VETO calorimeter induced large currents in PC0 and PC1. To reduce this effect, a 0.7 cm thick lead curtain was placed between the VETO and PC0.

Differences between the chambers are essentially due to the associated electronics and are described briefly below. PC0 and PC1 both comprise triplets of hexagonal chambers, recuperated from the NA4 experiment. For HELIOS, a new readout has been built, the choice of electronics being dictated by the need to obtain fast hit information for the second level muon trigger. Two readout systems (collectively termed FASTRO) process the chamber information; the FASTEN (FAST ENcoder) which calculates the hit cluster centre for use in the muon trigger, and the SIREN (SIRial ENcoder) which is a slow, serial readout recording information from all hit wires (see [64]). This circular chamber is positioned in the centre of the superconducting magnet. It plays an important rôle in track reconstruction. The slow readout is performed by JCF20 modules [65]. Of these

Chamber	Geometry	Sensitive Area	Planes		
			Type	Angle ϕ	Number of wires
PC0	Hexagonal	Inscribed circle of radius 76.8 cm	X	0°	512
			U	60°	512
			V	120°	512
PC1	Hexagonal	idem	X	0°	512
			U	60°	512
			V	120°	512
PC2	Circular	Circle of radius 67.2 cm	X	0°	448
			Y	90°	448
			U	45°	448
			V	135°	448
			Y'	90°	448
			X'	0°	448
PC3	Rectangular	307.2 x 259.2 cm ²	X	0°	896
			Y	90°	864
			U	45°	1024
			V	135°	1024
			Y'	90°	864
			X'	0°	1024
PC4	Rectangular	307.2 x 259.2 cm ²	X	0°	1024
			Y	90°	864
			U	14,25°	1024
			V	165,75°	1024
			Y'	90°	864
			X'	0°	1024
PC5	Rectangular	422.4 x 403.2 cm ²	X	0°	1408
			X'	0°	1152
			Y	90°	1344
			U	53,13°	1792
PC6	Rectangular	422.4 x 403.2 cm ²	X	0°	1408
			Y	90°	1344
			V	126,87°	1792
			X'	0°	1408

Table 3.2: Characteristics of the proportional wire chambers.




Figure 3.14: Schematic diagram of the hodoscope planes H3 and H2 (dimensions are in cm).

chambers, the first plane of PC3 participates in the muon trigger and is therefore equipped with the FASTRO readout described above. The remaining planes are equipped with the JCF20 system. The second plane of PC5 and the first of PC6 are read out by FASTRO modules for the trigger. Remaining planes use the LETI/FILAS system, developed by Saclay, which provide slow readout for each half-plane individually.

The hodoscopes

Downstream of the chambers are two scintillator hodoscope planes, H3 and H2, separated by an 80 cm thick iron wall. The wall is equivalent to 4.8 nuclear interaction lengths and acts as a hadron filter. A large fraction of the hodoscope scintillator planes has been recycled from the old NA3 experiment.

The upstream hodoscope, H3, lies immediately in front of the iron wall. It is divided vertically in two halves, left and right, each composed of 23 horizontal slats 3 m long and 3 cm thick. The heights of each slat are indicated in Figure 3.14. The scintillators are read out at the outer edge by photomultiplier tubes. The two slats at the upper and

lower edges (numbers 1 and 23) of the hodoscope plane fall outside the acceptance of the spectrometer and were therefore not instrumented.

The downstream plane, H2, is of similar design to H3. For this plane, however, several of the old NA3 scintillators were replaced by new slats, made of ALTUSTIPE⁵, to improve efficiency. Here also, the two outermost slats were not instrumented, for the reason given above.

The muon magnet

The magnetic field is provided by a 2.5 m long superconducting dipole magnet. It has a cylindrical aperture, 1.6 m in diameter, opening out downstream to a conical section to improve acceptance. The field is essentially vertical (parallel to the y axis), with a value at the centre of 1.67 Tesla. The field integral is thus 4.1 Tm, providing a transverse momentum kick of 1.23 GeV/c.

The muon update to the EFL logic

In a purely calorimetric measurement of missing energy, there is in general a contribution to the 'signal' not only from neutrinos (and any other weakly interacting neutral particles) but also from muons. Muon energy loss in dense matter is primarily through ionisation (dE/dx) processes; in a calorimeter of approximately $10 \lambda_i$ this corresponds to an average energy deposition of only ~ 3.5 GeV. For triggering on events in which the missing energy is carried only by weakly interacting neutral particles, it is therefore desirable that the calorimeter missing energy be corrected for the muon energy. In the case of HELIOS, this is the purpose of the muon update to the EFL logic (see Figure 3.13). A micro-processor uses information from the proportional chambers before and after the magnet in the muon spectrometer to make a fast estimate of the momentum of any muon candidate. In the case where a muon is detected, the momentum information is added to the calorimeter energy sum in the EFL logic and input to the VFB (see Section 3.1.2) for the final trigger decision. It should be noted that since the spectrometer acceptance is ~ 100 mrad, muons at wide-angles may not be detected. These muons are, however, kinematically constrained to have low energies and the error in the total energy measurement is not significant at the trigger level.

⁵An acrylic material manufactured by Altulor, Paris

3.2 Calorimeter performance

3.2.1 Calibration

High-precision calorimetry requires a detailed understanding of the energy calibration of the system. Two aspects are important:-

1. tuning the individual channels such that the energy deposited by a given particle produces the same signal regardless of the particular calorimeter module in which it is detected
2. determination of the overall energy scale

The calibration procedure for each calorimeter is outlined below.

UCAL

For the UCALs, calibration problems include:-

1. instability of the PMs,
2. light attenuation in the scintillator and WLS plates which introduces a spatial dependence of the calorimeter response,
3. different energy-to-signal conversions for photons, electrons, muons and hadrons.

The use of uranium as an absorber plays an important rôle in resolving these problems. Its natural radioactivity provides a stable, evenly-distributed source of signals for relative calibrations of the complete optical chain. The relative gain of the two sections (electromagnetic and hadronic) is set using minimum-ionising muons. The calibration is described in detail below:-

1. The charge ADC channels are individually calibrated using a precision pulse generator to determine the charge equivalent per ADC count of each individual channel.

2. The current outputs of all PMs are balanced by measuring the signal resulting from the natural radioactivity of the uranium. To achieve sufficient accuracy with respect to the ADC pedestals, a long gate time (10 μ s instead of the usual 0.13 μ s for shower detection) is used. The high-voltage applied to each PM is adjusted to obtain a pre-determined mean integrated charge for the uranium noise signal distribution with an accuracy of better than 1%. This procedure ensures that the gains (ie. the ratios MeV/pC) for all towers of a given type are equal.
3. The absolute gain factors are determined by exposing the stacks to beams of different energies. Since the electromagnetic sections are thin, charged particles will, in practice, always produce a signal in both electromagnetic and hadronic sections of a module, allowing one to calibrate both sections simultaneously. Tagged electrons (at energies of 8, 17, 24, 32 and 45 GeV) and cosmic muons were used for this absolute calibration. The final choice between electron and muon calibrations is discussed in detail below.

The calibration constants A and B for the electromagnetic and hadronic sections were calculated by minimising the width of the total signal distribution, ie. by minimising the quantity

$$Q = \sum_{j=1}^N \left(E(\text{beam}) - A \sum_{i=1}^n S_{ij}^{em} - B \sum_{i=1}^n S_{ij}^{had} \right)^2,$$

where $\sum S^{em}$ and $\sum S^{had}$ are the sums of all the ADC counts in the towers i of the electromagnetic and hadronic sections that contribute to the measured signal for event j. A universal calibration is obtained via a global minimisation over different energies j of the reduced χ^2 , $\chi^2 = \Sigma(Q^j/\Delta E^j)^2$, where ΔE^j is an estimate of the sampling resolution ($\Delta E^j \sim 0.19\sqrt{E^j}$).

Comparing results obtained using electron beams and cosmic muons, the following points illustrate some of the problems involved in achieving high-resolution calorimetry:-

1. the values of A and B found from electrons were slightly energy dependent,
2. the calibration constants found from electrons were not consistent with those found from muons.

The first point, showing that different A and B factors, and also a different B/A ratio, need to be used for different electron energies, means signal linearity is violated. Figure 3.15 shows that the value of B/A which minimises the resolution is energy-dependent. A value for B/A of between 3 and 4 minimises σ_{rms} for the electron calibration. This implies that




Figure 3.15: The σ_{rms}/E of the signal distribution as a function of B/A from electron and pion measurements.

one cannot simultaneously optimise the energy resolution for electromagnetic showers while maintaining signal linearity.

Cosmic muons only allow a value for the ratio B/A to be derived, since the energy deposited in the different modules depends on the (unknown) energy of the muon. The ratio from the muon calibration was found to be $\sim 21\%$ larger than the average value from the electron run, giving $B/A = 4.0$ (the dashed line in Figure 3.15).

There are two contributions to these effects; one instrumental, the other arising from the physics of shower development. The instrumental contribution is due to light attenuation in the WLS plates of the hadronic section. Whereas light produced by both muons and uranium noise is uniform as a function of depth, that produced by electron showers originates in the part of the hadronic section furthest from the PM and is therefore attenuated more than the average for muons and uranium noise. The physics contribution is related to the fact that in a sampling calorimeter the fraction of the electromagnetic shower energy converted to a measurable signal changes with depth. In the case of uranium, a considerable decrease occurs. It is due to the soft γ component of the shower, for which the calorimeter response is much lower than for minimum-ionising particles. This means that a given amount of energy deposited by the fast part of the shower in




Figure 3.16: The e/h signal ratio as a function of B/A at 17 and 32 GeV.

the electromagnetic section will give a larger signal than the same energy deposited by the soft tail in the hadronic section. The result is that the relation between the energy deposited by a shower and the corresponding calorimeter signal is different for the two sections of the calorimeter and, moreover, energy-dependent. The combined instrumental and shower ageing contributions explain the observed discrepancies between electron and muon calibrations.

The requirement that the response to electrons and pions be equal provides a stronger constraint on the ratio B/A . For electrons at both 17 and 32 GeV, $e/h = 1$ is obtained for B/A of approximately 4 (see Figure 3.16).

The value of 4 was finally chosen. This value of B/A corresponds to that obtained from the muon calibration. Such a choice gives an almost ideal e/h ratio and preserves linearity at the expense of a less-than-perfect energy resolution. In this case, however, only the total shower energy is meaningful; the signals from the individual sections cannot be simply interpreted as energy information. The fact that the calorimeter response is shower-age dependent also means that weighting algorithms (see Section 3.2.4), used to improve the off-line energy resolution, should be energy dependent.

See [66] for further details of the uranium/scintillator calorimeter calibration.

ULAC

The overall energy scale of the ULAC was determined using a secondary beam of electrons and pions at 17, 45, 70 and 200 GeV/c momenta. Particle identification is achieved on-line through a coincidence in 6 ring-imaging Čerenkov detectors in the beam line and enhanced off-line using the TRD and longitudinal shower-profile. As will be seen in Section 3.2.2, the ULAC is not fully compensating ($e/h \neq 1$). Electron and pion calibrations will therefore set different energy scales. The choice was made to use the electron calibration. The beam calibration of the ULAC is performed once per running period. Within a running period, stability is maintained using a charge calibration with test pulses. An array of microprocessor-controlled analogue switches distributes a known test charge to each channel of the calorimeter. The corresponding peak ADCs are then read out and the information recorded on tape at the beginning of each data-taking run. This procedure is repeated with *no* signal input to the electronics chain and also written to tape for off-line analysis. The gains of each unit are under computer-control and may be tuned such that the test signal produces the same peak response in each channel. The peaks for different channels, however, do not occur at precisely the same time and it was decided instead to equalise the responses a fixed time after the injection of the test pulse. This has the advantage of optimising the Energy Flow Logic performance which also samples at a fixed time.

MAGCAL

For the magnetised calorimeter, the relative calibration is provided by a system of light diodes and a radioactive source. One fibre of each light guide is coupled to a light diode. Light passes down the fibre, along the WLS bar, where it is reflected at the end and returns to the photomultiplier tube. In order to control systematic effects, each diode feeds four channels and the temperature of the diodes is monitored. As a check on the scintillator response, a 740 MBq ^{60}Co source is pulled around the front face of the MAGCAL. It samples the first few layers of scintillator. The absolute scale is again set using electrons.

Intercalibration of all calorimeters

As will be seen in Sections 3.2.2 and 3.2.4, the optimal calibration for the *ensemble* of calorimeters is not yet defined and is, moreover, dependent on the topology of the events of interest. This is due in part to the non-compensated nature of components of the

calorimetry (the ULAC and MAGCAL) and in part to incomplete shower containment (the MAGCAL). These factors, and their relevance for a missing energy analysis, are discussed in detail below.

3.2.2 Intrinsic calorimeter properties

This section presents details of the intrinsic calorimeter properties (signal linearity, e/h ratio, resolution). For the uranium calorimeters, the information has been obtained from test beam data taken with beams of electrons and pions of momenta: 17, 45, 70 and 200 GeV/c (for the ULAC); 8, 17, 24, 32, 45, 70 and 200 GeV/c (UCAL). A low beam intensity (typically 10^2 to 10^4 particles per burst) was maintained to reduce the possibility of pileup and ensure that the results genuinely represent the single-particle response of the calorimeter. Electron/pion separation was achieved (up to 45 GeV/c) using tagged information from a Čerenkov differential counter with achromatic ring focussing (CEDAR). Additional particle identification was provided by the longitudinal and transverse profiles of the showers. For the MAGCAL, direct measurements of the electromagnetic energy resolution alone were possible (only 8 GeV/c electrons can be magnetically deflected sufficiently far from the beam line that they hit the MAGCAL petals). In this case, the e/h ratio has been deduced from other iron/scintillator calorimeters of similar design.

Linearity

Results from the ULAC beam tests show the FADC readout of the calorimeter to be linear within 1.0%. Larger non-linearities in the PADC readout chain are discussed in Section 3.2.3. The linearity of the the UCAL module response to electrons and pions was tested over the energy range 17 to 200 GeV/c. Results are shown in Figure 3.17. Within this energy range, linearity is maintained to better than 1.5%.

The e/h ratio

In order to optimise its performance (in terms of linearity and energy resolution) for the detection of hadronic showers, a calorimeter should have equal response to the electromagnetic (e.g. π^0) and hadronic components of the shower. This allows one to minimise the degradation in resolution and linearity caused by large event-by-event fluctuations in the fraction of the shower energy used in π^0 production. For most sampling calorimeters in which $Z^{absorber} \geq Z^{readout}$, one typically finds that $e/h \equiv (e/mip)/(h/mip) > 1$. The




Figure 3.17: The U/scintillator module response to hadronic showers as a function of energy in the range $17 \leq E_\pi \leq 200$ GeV/c.

suppression of the hadronic response is due to shower processes whose energy is either largely (soft neutrons and gammas, muons, neutrinos) or completely (binding energy from breakup of nuclei) undetected in the calorimeter. A compensation for this undetectable energy was first achieved in uranium/scintillator calorimeters and was attributed to the rôle played by detectable products (soft evaporation neutrons and gammas) of induced uranium fission. Detailed study of the mechanisms involved in achieving compensation in a calorimeter has shown that the dominant factor is not fission but a combination of densely-ionising protons (produced in spallation reactions) and neutrons (via energy transfer to recoil protons in the low-Z readout media). This work has indicated that fission (and therefore uranium absorber) is not a necessary requirement for compensation; the emphasis should be rather on the readout medium (in particular, its hydrogen content) and the tuning of the relative fractions of active *and* passive material.

The e/h ratio for the ULAC is shown as a function of incident particle energy in Figure 3.18. Clearly the ratio is somewhat greater than 1.0 and is, moreover, dependent on the beam energy. This result is a direct consequence of the use of argon as the active medium and is in quantitative agreement with the predictions of [61]. The e/h ratio is also dependent on the integration time of the readout ADCs. Since the time constants of the trigger FADC and off-line PADC readouts are different, a corresponding difference in




Figure 3.18: The e/h signal ratio as a function of energy for the U/liquid argon calorimeter.

the respective e/h ratios is seen. This introduces complications in determining the best calibration for missing energy studies (see Section 3.2.4). For the U/scintillator modules, the e/h ratio is shown in Figure 3.19. (The electron value at 200 GeV was obtained by assuming that it was possible to extrapolate the linear behaviour observed for electrons from 8 to 70 GeV). Over the energy range studied, the measured (e/h) ratio is constant within experimental error. An average value of 0.98 ± 0.01 (1.01 ± 0.02) is found for U/scintillator stacks with optically decoupled (coupled) towers. The U/Cu/scintillator modules have an e/h ratio = 1.11 for energies of 2 GeV or more (see [67]). The e/h ratio for the MAGCAL could not be measured directly. Based on measurements from similar calorimeters [68], $e/h \sim 1.4$.

Intrinsic resolution

Results on the resolution, σ/\sqrt{E} , are produced simultaneously with the linearity results above. The single-particle resolution of the ULAC is shown as a function of energy in Figure 3.20. For pions, $\sigma/E \sim 0.7/\sqrt{E}$ and for electrons $\sigma/E \sim 0.2/\sqrt{E}$. For the ULAC, multi-particle resolution is as good as the single-particle resolution. Results for the U/scintillator modules are shown in Figure 3.21. Within experimental uncertainties the value of σ/\sqrt{E} is independent of the particle energy. A least-squares fit to the data




Figure 3.19: The e/h signal ratio as a function of energy for U/scintillator modules (optically decoupled towers).




Figure 3.20: The resolution σ/\sqrt{E} as a function of energy for the U/liquid argon calorimeter. Results for pions are indicated by open circles, those for electrons by triangles.




Figure 3.21: The resolution σ/\sqrt{E} as a function of energy for the U/scintillator calorimeter. Results for pions are indicated by the solid circles, those for electrons by open squares.

gives $\sigma/E = (0.337 \pm 0.013)/\sqrt{E}$ for pions and $(0.215 \pm 0.007)/\sqrt{E}$ for electrons, with E in GeV. Corresponding values for the U/Cu/scintillator modules are $\sigma/E = 0.36/\sqrt{E}$ for pions and $0.16/\sqrt{E}$ for electrons. For the MAGCAL, only one experimental measurement was possible, corresponding to 8 GeV electrons. The measured resolution, σ/E is $0.25/\sqrt{E}$, which is as expected for our sampling thickness.

3.2.3 ULAC peak ADC performance

The slow PADC readout of the uranium/liquid argon calorimeter is designed to give a more precise energy measurement than the fast FADC readout used in the trigger. During operational studies of the calorimeter performance, two instrumental effects were found which showed the flash ADCs to be more reliable. They are discussed briefly below.

The first problem is an apparent dependence of the PADC response on the gate/fast clear frequency. This effect, simultaneously lowering all the individual channel responses by up to $\sim 10\%$, is most severe at high beam intensities. Not only does it create 'fake' missing energy in the PADC event reconstruction, but also degrades the energy resolution. An off-line correction algorithm has been developed, based on the response of channels

with no energy deposition, which introduces a systematic error of ~ 3 GeV in the total energy.

The second effect is non-linearity of the PADC response and is apparent at both high and low energies (the intermediate energy behaviour is linear). At high energies, the cause is probably saturation effects in the analogue summing chain. At low energies, the most likely explanation is noise though there could be a genuine non-linearity in the PADCs on top of this.

3.2.4 Resolution of the calorimeter ensemble

This section addresses the question of how to obtain the best possible energy resolution from a complex calorimeter having components which are not fully compensating. A brief description of shower development in a sampling calorimeter is given, outlining the characteristics which are important in attempting to optimise the resolution.

Particle showers in a sampling calorimeter

A calorimeter is basically a block of matter in which the particle that is to be measured interacts, and deposits all its energy in the form of a *shower* of successively lower-energy particles. The block is made such that a certain (usually small and hopefully constant) fraction of the initial particle signal is transformed into a measurable signal (light, electrical charge). Since calorimetry is based on statistical processes, the measurement accuracy increases with \sqrt{E} . The energy resolution for detecting the original particle is determined by fluctuations occurring in this process, which affect the size of the resulting signal. Important among these are:-

1. Sampling fluctuations, due to the fact that only a (small) fraction of the energy is deposited in the active layers. It has been shown [69] that if this fraction is larger than $\sim 1\%$, the effect is dominated by fluctuations in the total number of gap crossings by charged particles generated in the shower. For a given absorber material, the contribution from this source of fluctuations is therefore proportional (to first order) to the square root of the thickness of the absorber plates. In electromagnetic sampling calorimeters, sampling fluctuations are the major contribution to the energy resolution.

2. Fluctuations in the fraction of the energy deposited in the form of ionising particles. These effects usually dominate the energy resolution of hadron calorimeters. One component comes from incomplete shower containment, i.e. from secondary particles like μ, ν, n, K_L^0 , that (partly) escape detection. Another, and usually more important, source of fluctuations comes from energy losses at the nuclear level. The fraction of energy lost in this way is considerable, up to $\sim 40\%$ on average for the non-electromagnetic part of hadron showers in high- Z elements. The consequences of this for hadron detection are twofold. Firstly, the signal distribution for hadrons of energy E will be broader than that for electromagnetic showers of the same energy. Secondly, the average response will have a different value for electromagnetic and hadronic showers ($e/h \neq 1$).

In the development of a shower generated by a high-energy hadron, some fraction f_{em} of the energy is usually expended on production of π^0 's and η 's. The shower therefore has an electromagnetic and an hadronic component to which the calorimeter response is different. The fluctuations in f_{em} are large and non-Gaussian. Also, $\langle f_{em} \rangle$ increases (logarithmically) with energy. These fluctuations therefore give an additional contribution to the energy resolution, and introduce an energy dependence to the hadron response and e/h ratio. If $e/h = 1$, the contribution of fluctuations in f_{em} to the energy resolution are eliminated.

It has been suggested that there are other techniques to achieve this goal. The electromagnetic and hadronic shower components have very different characteristic dimensions, particularly in high- Z absorber materials such as uranium. If the calorimetry is sufficiently finely segmented it is possible to determine f_{em} on an event-by-event basis and hence eliminate (at least partly) the effects of fluctuations in this number. This approach is discussed below.

On-line resolution optimisation

In Sections 3.2.1 and 3.2.4, some of the difficulties in defining the optimum calibration for the HELIOS calorimetry were outlined. This section describes the weighting of signals from the individual calorimeters used on-line to minimise the energy resolution and simultaneously remove any dependence of the measured energy on the event topology. Using the calibration procedures of Section 3.2.1, the average measured energy of a 450 GeV proton-nucleus interaction is somewhat less than 450 GeV (typically $\sim 15\%$ less), and is moreover, dependent on the spatial distribution of the secondary particles. This is a direct

result of the 'invisible' component of hadronic showers in the non-compensating elements of the calorimeter. These elements (the ULAC and MAGCAL) cover restricted areas of phase-space and so the fraction of 'invisible' energy is (partly) determined by topological fluctuations.

For triggers on total (or missing) energy, it is clearly undesirable that the trigger decision should be affected by the event topology. It is useful to introduce the idea of Etot-flatness, that is, that the total measured energy should not be a function of the energy in one particular calorimeter. By weighting the signals from each component calorimeter, it is possible to eliminate any such dependence. This is shown in Appendix A. It is also shown that the same weights should be used to minimise the resolution of the calorimeter ensemble. Having determined the theoretical procedure for setting the overall calibration, there remain two practical questions: first, what is the most appropriate sub-division of the calorimetry, and second, what is the most appropriate event sample to be used in the optimisation. As described in the Appendix, the final sub-division was decided empirically (in essence, the calorimeters with small average energies being weighted in accordance with neighbouring calorimeters). 'Minimum-bias' events provided the data sample. In this way, one is certain that no bias is introduced from events with *genuine* missing energy (muons, neutrinos, etc.). However, since known sources of missing energy, such as heavy flavour production, may be associated with different event topologies (central or forward production for example), there may be a small offset in the measured energy for these events.

Application of this technique to the HELIOS calorimetry has been shown to remove topological correlations while simultaneously improving the total energy resolution. The weights obtained were used on-line (by suitable gain adjustments at the level of the FADC inputs) to improve the efficiency of the missing energy trigger.

Off-line compensation by π^0 -weighting

The previous section described procedures to improve the *average* energy resolution of a given calorimeter ensemble, thereby increasing the efficiency of calorimetric triggers. In this section, an event-by-event improvement will be described, known as π^0 -weighting, which uses longitudinal shower information to help compensate for fluctuations in the electromagnetic (π^0) fraction of the shower energy. The technique is, of course, useful only in the case of non-compensating calorimetry.

Local high energy concentrations in the shower development are due mainly to π^0 's and increase the total measured energy of an event since $e/h > 1$. For a calorimeter with

Configuration	Resolution		
	FADC	P/QADC	
		Before	After
Target	15.9 ± 0.7	18.3 ± 0.8	16.4 ± 0.7
Beam dump	9.7 ± 0.2	12.4 ± 0.2	9.6 ± 0.2

Table 3.3: Optimised resolutions for target and beam dump modes. Figures for the P/QADCs are given before and after π^0 weighting. (All measurements are with 450 GeV protons).

sufficiently fine segmentation, the local energy deposit may be assigned a ‘ π^0 -likelihood’ and weighted accordingly. The method chosen to correct for π^0 fluctuations is motivated by the approach of [68,70,71]. For each event, a ‘weighted energy’, E'_i , for each channel i is defined:-

$$E'_i = cE_i(1 - dE_i)$$

where the coefficient d is the correction for the π^0 contributions and c is the overall scale factor needed because the unweighted energy, E_i , is always smaller than the true value.

This algorithm was applied to the HELIOS calorimetry using the peak and charge ADC readouts since they provide finer segmentation than the flash ADCs. The improvement in terms of resolution can be seen in Table 3.3 where results are shown for target and dump modes before and after π^0 -weighting. The reduction in resolution width is of order 2 - 3 GeV in both cases.

Choice of running mode

This analysis aims to measure the missing energy spectrum corresponding to the production of weakly interacting neutral particles. At small missing-energies, this spectrum will be dominated by measurement fluctuations due to the finite calorimeter resolution. This should be Gaussian, reflecting the statistical nature of the shower processes in the calorimeter. For the best possible sensitivity to escaping particles, it is therefore important to look in the mode for which the energy smearing is smallest.

The energy resolution from both fast (flash ADC) and slow (peak and charge ADC) readouts was studied in ‘target’ configuration and also in beam dump mode. Table 3.3 shows the results corresponding to these four cases. In both cases, the weighting improves the PADC resolution by about 3 GeV. Two conclusions may be drawn: firstly, that the weighted signals from the peak and charge ADC readout do not offer an improvement over the flash ADC readout in terms of resolution, and secondly, that the resolution in beam

dump mode is significantly better than that in 'target' mode. This result, combined with the arguments outlined in Section 2.3, shows the beam dump to be the mode of choice in a search for weakly interacting neutral particles. Details of the beam dump data-taking run are given in Section 3.4.

3.3 Muon spectrometer performance

3.3.1 Muon detection

By way of an introduction, it is useful to consider the typical behaviour of a muon produced by a proton interaction in the beam dump. The first component of the detector seen by the muon is the beam dump itself. Here, the most likely outcome is that the muon will undergo multiple scattering, losing energy through ionisation (dE/dx losses). For the $\sim 10 \lambda_i$ of the dump, this amounts to about 3.5 GeV. Assuming the muon leaves the dump with polar angle $\theta_\mu \lesssim 130$ mrad (true in most cases), it will then encounter the first of the seven wire chambers in the spectrometer. Ionisation of gas molecules in the chamber will produce hits on the sense wires nearest to the muon trajectory which, by the use of wire planes oriented at different azimuthal angles, can be combined to give the muon position in the xy plane (perpendicular to the beam axis). Taking this information from each of the seven chambers, a 3-dimensional hit pattern can be constructed and used for tracking. Immediately downstream of the second chamber however, the muon passes through the centre of the spectrometer magnet where it receives a transverse momentum kick of 1.2 GeV/c. The resulting deflection, as determined by the chamber tracking, is used to measure the momentum of the particle. Finally, the muon hits the hodoscope arrays, causing ionisation in the scintillator material and in the $5 \lambda_i$ iron wall. dE/dx losses here are of order 1 GeV.

The combined effect of ionisation losses and the momentum kick of the magnet is to impose a lower threshold on the muon momentum if it is to be detected in all elements of the spectrometer. This lower threshold is approximately 5 GeV/c. Hence, the spectrometer acceptance is roughly $\theta_\mu \lesssim 130$ mrad for $p_\mu > 5$ GeV/c.

3.3.2 Muon reconstruction

The aim of muon reconstruction is to take the hit information from the 32 wire and 2 hodoscope planes and be able to identify charged particle tracks, giving also their associated momenta by measurement of the track bending in the spectrometer magnet. This is achieved using the standard HELIOS muon reconstruction software package MUREC[72].

Reconstruction takes place in several stages, outlined briefly below, using hit information from the SIREN chamber readout (Section 3.1.5). First, the hit clusters are used to search for straight lines behind the magnet (PC3-6) in the X and Y projections independently. Any found tracks are matched with the help of hit information from the remaining

U and V planes, inclined at various angles to the x and y axes (see Table 3.2). The reconstruction criteria are as follows: the track must have an associated hit in at least 5 (out of 8) X planes, 4 (out of 6) Y planes, 3 (out of 6) inclined planes, and at least one of the two X planes of PC3. MUREC then calculates the best fit for the track, giving corresponding χ^2 s in the X and Y projections. For each track found behind the magnet, a hit is tried in the first X plane (PC0). This gives a first value for the momentum p assuming that the magnet is a perfect dipole. Knowing p and the trajectory downstream of the magnet, MUREC then calculates the intersection of the particle track with each of the planes of PC2, PC1 and PC0, using the real field map. By looking at hits in some road around the trajectory, these hits may then be matched to give a reconstructed track before the magnet. If the match is unsuccessful, MUREC repeats the procedure using a different hit in the first X plane or, if necessary, a hit in the second X plane (PC1). The reconstruction criteria here are as follows: associated hits must be found in at least 2 (out of 3) planes in PC2, 2 (out of 3) X planes in PC0-2, and 4 (out of 6) planes in PC0-1. When a track has been found, a full fit is made and the corresponding χ^2 per degree of freedom calculated. Finally, the intersection of any reconstructed track with the two hodoscope planes H3 and H2 is calculated and the relevant scintillator slats checked for the corresponding hits. Further details may be found in [72].

The output from the reconstruction program includes complete information on any found tracks (momenta, fit quality, etc.) and the corresponding hit patterns in the downstream scintillator planes. In principle, this should be sufficient to identify muons in the spectrometer by their characteristic signature of a chamber track and one (and only one) hit in each of the two hodoscope planes. In practice however this is not the case, since the detector is not fully efficient and also because of contamination from charged hadrons in the spectrometer. The consequences of this are discussed below.

3.3.3 Reconstruction efficiency

Inefficiencies in track reconstruction can be divided into two categories: those arising from spectrometer hardware (poor scintillator performance because of radiation damage for example) and those due to the reconstruction software (typically through the use of imperfect algorithms). In order to properly determine the missing energy spectrum for events with associated muons, it is necessary to know these inefficiencies and to correct for them.

Looking at the hardware efficiency first, this includes the efficiencies of the wire chambers and the scintillator hodoscope planes. With the exception of the downstream hodoscope plane (H2), these efficiencies are calculated directly by MUREC, which returns

both a global value for the plane concerned and the position dependence of the efficiency within that plane. An important point to note here is that the data sample used to calculate the efficiencies is exactly that used for the missing energy analysis (Sections 3.4 and 4) and was consequently *not* triggered on muons but only on missing energy, avoiding the need to correct for trigger bias.

The technique used for the chambers is simple. A sample of tracks is defined according to whether the chamber is before or after the magnet. For planes upstream of the magnet, completely reconstructed tracks are used whereas for downstream planes, the sample comprises tracks reconstructed behind the magnet in either the X or Y projection. Then, for n_t tracks passing through the given plane, the number n_h having an associated hit is determined, giving the plane efficiency $\frac{n_h}{n_t}$. The values obtained are given in Appendix B.

In general, the plane efficiencies are quite reasonable (around 90% or more). Two of the chambers (PC2 and PC4) give rather poor performance however, with individual plane efficiencies as low as 25% in one case. Fortunately, the large number of planes in the spectrometer means there is a degree of redundancy which does not result in a serious deterioration in the overall track-finding efficiency.

Monte Carlo techniques are used to quantify the effect of chamber inefficiencies on the track reconstruction. A sample of single muon events is first generated using the LUND Monte Carlo package. Muons are then tracked and a full (GEANT-based) simulation of the spectrometer implemented to produce hits in the spectrometer chambers. At this stage the option exists for either fully efficient chambers or the measured values to be input into the Monte Carlo. Once the hit map has been produced, the standard MUREC reconstruction code is run, allowing the number of surviving muons to be counted. Taking as the initial event sample events having a muon which passes through *all* the 32 wire planes and the H3 hodoscope plane, this procedure is performed once with 100% efficient chambers and then a second time with the actual chamber efficiencies. In the case of fully efficient chambers, the fraction of surviving muons is in fact a direct measurement of the MUREC reconstruction algorithm efficiency. It is found to be 99%. Using the values of Appendix B, the corresponding fraction measures the combined effect of chamber and algorithm inefficiencies. In this case the efficiency is found to be 86%. Hence, a 16% correction must be made to the measured muon rate to take account of both the reconstruction algorithm and wire chamber efficiencies.

Turning now to the hodoscopes, the method used to determine their efficiencies is different for each of the two planes. This is due to the $5 \lambda_i$ iron wall between them. For the upstream plane H3, the technique is essentially that used for the chambers. A sample




Figure 3.22: Punchthrough probability in an iron absorber as a function of incident hadron energy. Punchthrough here is defined to be at least one charged particle. The dashed line corresponds to the HELIOS hodoscope wall.

of events with a reconstructed chamber track is used to point to the associated slat in H3. A straightforward count of the number of detected hits divided by the number of expected hits gives the efficiency as before. H3 slat efficiencies are given in Table C.1 (Appendix C) and indicate a mean value for the whole plane (integrated over all slats) of 70%.

For the downstream plane H2, the method of track pointing is not so obviously applicable because spectrometer tracks may be due either to muons *or* to charged hadrons and in the case of charged hadrons, the particle has a high probability of showering in the iron wall before H2. At first sight it might appear that hadrons interacting in the wall will be completely absorbed and consequently that no hits will be registered in the downstream hodoscope. This turns out not to be the case however, with a very high proportion of energetic hadron showers leaking out the back of the $\sim 5 \lambda_i$ wall. Figure 3.22 shows the results[73] of a range of punchthrough measurements for hadrons of momenta between 4.7 and 140 GeV/c incident on iron absorbers. The dashed line corresponds to the thickness of the iron wall in the HELIOS spectrometer and indicates that for our configuration, charged hadrons with momenta greater than ~ 50 GeV/c have essentially a 100% punchthrough probability. This will in fact be shown to be verified experimentally in Section 4.3 (Figure 4.12), where charged hadrons in the spectrometer can be recognised by the large number of (punchthrough) hits seen in H2. Muon/hadron separation can be

achieved by cutting on the number of hits in H2. Hence, by selecting only high momentum tracks (say $p_{tr} > 50$ GeV/c), the track pointing method may be used for H2 as well as H3. The H2 slat efficiencies obtained by this technique are given in Table C.2. The average efficiency for the plane is 76%.

The question of software efficiency has already been mentioned: inefficiencies arise primarily from the use of imperfect track-finding algorithms. Using a sample of single muon Monte Carlo events and the full detector simulation, the MUREC efficiency is found to be 99%.

3.4 The data sample

In this section, the run conditions and final data sample for the beam dump missing energy analysis are outlined.

3.4.1 The missing energy data sample

The HELIOS trigger system allows selection of as many as 51 parallel triggers in any one physics run. The missing energy data sample comprises 6 such triggers: one minimum-bias trigger, four missing energy triggers (with increasing missing energy thresholds at 20, 40, 60 and 100 GeV), and one empty ('MT') trigger (corresponding to a randomly-generated pretrigger and used to monitor detector and readout stability). Downscaling factors were applied to each of the six triggers to ensure roughly equal population of successive missing energy thresholds in the final data sample. All calorimeters were included in the on-line energy summation, with only the forward dump calorimetry used in the subsequent off-line analysis. The muon update to the Energy Flow Logic was not used. A requirement that the energy in the most downstream (VETO) calorimeter be less than 5 GeV was imposed on-line to reject events in which late shower development results in energy leakage. A low beam intensity (between $\sim 5 \times 10^4$ and 2×10^5 protons per burst) was chosen to reduce contamination by particle pileup. Under these conditions, the experimental 'live-time' is approximately 50%. The final data sample corresponds to an integrated flux of $\sim 1.2 \times 10^8$ protons.

3.4.2 The excess energy data sample

In parallel with the missing energy data-taking, *excess* energy triggers were configured and recorded. With no known (or predicted) physics mechanism for producing a net energy gain in an interaction, the purpose of this was to study detector effects; in particular, the behaviour of the calorimeters in the tail of the resolution spectrum. Where possible, the excess energy trigger requirements were identical to those for missing energy triggers. Important differences were: firstly, attenuation (2 dB) of the ULAC hadronic signal to the Energy Flow Logic (to avoid ADC overflows) and secondly, an attempted pulse-height cut on the beam counter signal (for on-line rejection of events with two particles in the same SPS radio-frequency bucket). Excess energy thresholds were set at 40 and 60 GeV.

3.4.3 Data structure

Under normal data-taking conditions, events satisfying the trigger requirements are stored in a buffer in the on-line VAX during beam spill to be written to tape at the end-of-burst. For the missing energy sample, each event should contain complete information from the trigger system, beam counters, silicon pad array, calorimeters and muon spectrometer, giving a total of about 7.5 kbytes of data per event. The on-line data structure is of course determined by the hardware modules. Data banks containing digitised information from each detector or sub-detector are given an identifying label and written to tape for subsequent use by the off-line software.

The burst structure is maintained on tape by additional labels defining the start-of-burst (SOB) and end-of-burst (EOB). An array of scalers (some 250 channels in all), counting various signals from the trigger logic and data acquisition system, is incremented throughout each beam spill to be read and reset at end-of-burst.

The final component of a raw-data tape is a series of labels written at the start of each run, containing details of the relevant trigger configurations, calibration files and general bookkeeping information.

Quality checks

Throughout any data-taking run, a subset of the events written to tape are analysed on-line to monitor detector performance and check the data quality. While this is usually sufficient to discover and correct major problems, it does not eliminate occasional 'glitches' in the data-acquisition system, where partial loss of event information may occur. For this reason, it is essential that the data be checked off-line. In this analysis, the first stage in data processing is therefore to verify that complete information *is* available for each event and that the necessary calibration files are present. Approximately 1.5 % of the total data sample is rejected in this way.

Formation of physical quantities

As with any particle physics experiment, the raw data is not easily interpreted in terms of physically meaningful quantities. Summarising event information in a concise form is the purpose of the data production process. The calorimeter production provides an

illustration of this stage in the analysis. For example, each of the 2432 individual PADC channels of the uranium-liquid argon calorimeter (ULAC) must be pedestal-subtracted (the pedestal values being obtained from the 'MT' events on tape) and subsequently multiplied by the appropriate gain factor (from the start-of-run calibration files) before yielding the actual energy deposition in that cell. Energy information from each cell may then be combined to give the total ULAC energy and perhaps combined with similar sums from other calorimeters to give the measured event energy. Condensed information from each event is written to data summary tapes (DSTs) for later analysis.

Chapter 4

The missing energy spectrum

In this chapter, the cuts applied to the data sample to reject events with ‘fake’ missing energy are described. The main experimental concerns here are threefold: off-momentum beam tracks, incomplete shower containment (punchthrough), and overlap of two or more particles within the sensitive readout window of the detector (pileup), all of which will be shown to contaminate the data sample. The aim throughout will be to identify and remove such sources of fake missing energy, while maintaining a high efficiency for events with genuine missing energy (corresponding to the production of non-interacting particles). The eventual goal of this chapter is to determine the cross-section for the production of such non-interacting particles as a function of their associated energies, that is, to produce a missing energy spectrum.

4.1 The raw spectrum

It is useful to begin by looking at the calorimeter response to minimum-bias (VB) interactions, that is, events not required to satisfy any missing energy threshold. Unless stated otherwise, the results in this chapter use the flash ADC (FADC) readout of the calorimeters, this having been shown in Section 3.2.4 to give superior energy resolution. A logarithmic plot of the measured energy is shown in Figure 4.1. This plot corresponds to a small fraction (roughly 15%) of the total data sample; only when the final missing energy spectrum is presented at the end of the chapter will results from the full data sample be shown.

In Section 3.2.4, the statistical basis of particle showers in dense matter was outlined. A series of calorimeter measurements on monoenergetic particles should therefore, in the

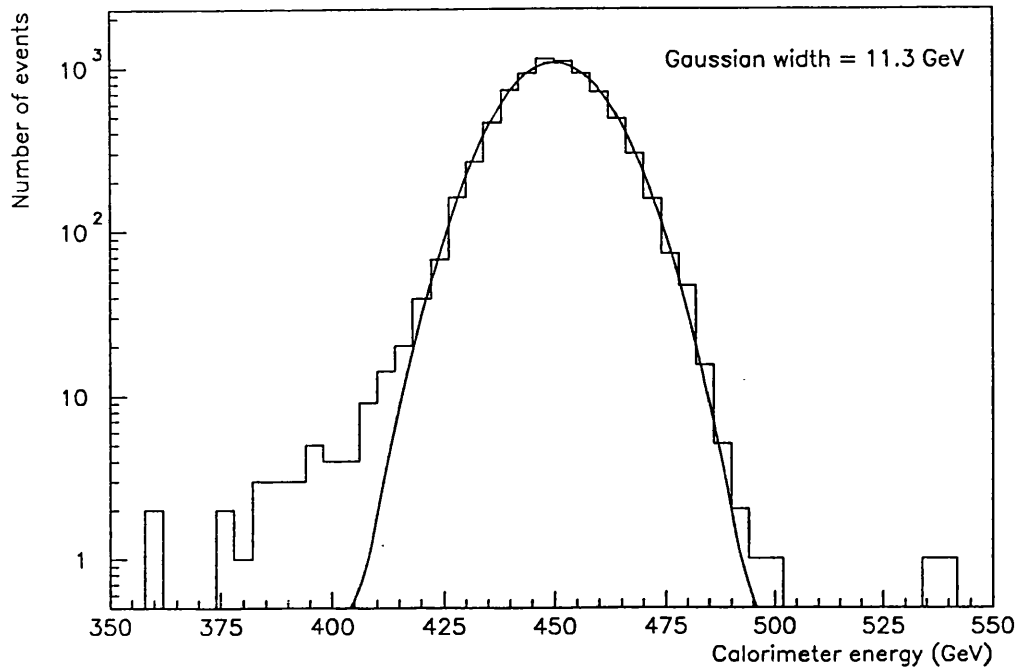


Figure 4.1: Calorimeter response (FADCs) to minimum-bias interactions.

absence of other effects, follow a Gaussian distribution. Superimposed on the plot in Figure 4.1 is the result of a fit to a Gaussian. The fit gives a mean of 450.1 GeV, with standard deviation (σ) of 11.3 GeV, i.e. an energy resolution of about 2.5%. There are, however, deviations from Gaussian behaviour at both the upper and lower edges of the distribution. On the high-energy side, there is evidence for discrete events having considerably more than nominal beam energy. On the low-energy side, the deviation from a Gaussian appears continuous: this is the ‘signal’ in which we hope to find evidence for production of non-interacting particles.

The complete missing energy spectrum is obtained by combining data from successive missing energy triggers (Section 3.4), taking into account the corresponding downscaling factors. This is shown in Figure 4.2. Since the on-line flash ADC readout is used also for the off-line calorimeter analysis, the individual E_{mis} thresholds are sharp step-functions. Within errors, the distributions from successive trigger thresholds overlap, excluding E_{mis} trigger inefficiencies of more than $\sim 20\%$.

Also shown (Figure 4.3) is the same data in terms of integrated rate (i.e. the relative cross-section for events with missing energy greater than a specific value). Superimposed on both plots is the Gaussian profile (fitted to minimum-bias interactions) expected from the calorimeter resolution alone (the double line indicates a $\pm 1\sigma$ error band on the fit).




Figure 4.2: The normalised event rate ($\frac{1}{N} \frac{dN}{dE_{mis}}$) as a function of missing energy (E_{mis}).




Figure 4.3: Integrated event rate ($\int_{E_{mis}}^{450} \frac{1}{N} \frac{dN}{dE_{mis}}$) as a function of missing energy (E_{mis}).

Several features are noteworthy. Firstly, the spectrum extends to around 450 GeV missing energy: apparently there are events where none of the incident energy is detected in the calorimeter! Secondly, deviation from resolution effects becomes evident for missing-energies greater than ~ 30 GeV, corresponding to an integrated rate of $\sim 10^{-2}$ of the total cross-section. Thirdly, the rate of zero-energy events ($E_{mis} \sim 450$ GeV) is an order of magnitude above the neighbouring ‘background’ level, indicating that these events may have a different origin.

It is useful here to make a rough comparison with cross-sections of Standard Model sources for producing missing energy. This is discussed in detail in Chapter 5, but is used here to provide a simple reference against which to assess the measured spectrum for evidence of additional sources of missing energy. For the energy domain accessible in this experiment ($\sqrt{s} \sim 30$ GeV) Standard Model neutrino production will be primarily through the semi-leptonic decay of charm quarks. Taking the charm production cross-section to be $\sim 10^{-3}$ of the total cross-section [74] and the semi-leptonic branching ratio to be $\sim 10^{-1}$ per (lepton) channel [75], neutrino production should comprise at most $\sim 4 \times 10^{-4}$ of the total inelastic cross-section. Thus, there is an apparent excess (more than two orders of magnitude) in the measured missing energy spectrum above crude expectation from known Standard Model sources.

The remainder of this chapter is devoted to studying contamination in the measured spectrum and its contribution to this excess. Sources of contamination are considered under three headings: upstream interactions, shower leakage and pileup.

4.2 Upstream interactions

This analysis is based on the determination of missing energy in an event, as defined by the difference between incoming and outgoing particle energies. While it is essential to have a reliable energy measurement of the produced secondaries, it is equally important to know the energy of the incident proton. With particle extraction from a primary accelerator beam line (Section 3.1.1), contamination arising from interactions upstream of the experimental area is expected to be small. The possibility of interactions during beam transport cannot be ignored however: this section studies such a possibility.

Particle interactions upstream of the calorimeter dump may create a missing energy signal if some or all of the produced secondaries fall outside the detector acceptance. The trigger requirements for this data sample already impose a rather strong veto on upstream

interactions: in addition to demanding that one of the produced secondaries hit the beam-defining counters ($B6 \cdot B1 \cdot B2$), a requirement is also made that no other secondary be detected in the veto scintillators ($\overline{B7} \cdot \overline{B3}$). Nevertheless, there may be event topologies which do fulfil these criteria (diffractive interactions far upstream or interactions in the beam counters themselves for example).

Three cuts have been made on the data in an attempt to eliminate upstream interactions. In each case, one component of the detector has been used to extend the 'veto' coverage of the scintillators B7 and B3.

The first cut is on the charged multiplicity in the silicon pad array (Section 3.1.3). This device sits 66 cm downstream of the last beam scintillator (B3), having a diameter of 30 mm. In addition to providing improved acceptance over the B3 scintillator counter, it has a useful rôle in identifying interactions in the beam counter array or the surrounding support material. The counter is aligned such that genuine beam particles pass through its central hole; for these events, no hits should be registered in the pad array (other than from stray particles or noise). A requirement of zero hits over the entire 400 pad elements is therefore made.

The two remaining cuts use different components of the backward calorimetry (Section 3.1.4). In the first case, the modules comprising the BOX, WALL and MAGCAL are combined to form a veto guard-ring covering the region from ~ 18 to 240 cm from the beam axis and about 1 to 3 m upstream of the beam dump. MT events (those corresponding to random triggers of the detector readout) are used to measure pedestal values (of finite width due to uranium and electronic noise) and define the energy cut. Fitting a Gaussian to the pedestal distribution, the cut is made at 3σ above the mean, corresponding to 7 GeV.

The second calorimeter cut is a restriction on the energy deposited in the electromagnetic section of the ULAC. Situated immediately upstream of the beam dump, it covers the region between 2.5 and 40 cm from the beam axis. In this case, minimum-bias (VB) events have been used to define the cut threshold rather than MT events. This distinction is necessary because the electromagnetic section is sufficiently close to the beam particle shower vertex to become activated (albedo). Fitting a Gaussian to MT and VB events, there is a systematic shift in the respective means of about 1 GeV due to this effect. A cut is again made at 3σ above the mean, in this case 6 GeV.

The cumulative effect of the three cuts described is to reject approximately 6% of the data, indicating that upstream interactions either occur at a low level or that they are




Figure 4.4: The missing energy spectrum after cuts on upstream interactions.

successfully vetoed by the on-line protection (Section 3.1.2). The missing energy spectrum after these cuts is shown in Figure 4.4 (the error bars are reduced compared to Figure 4.2 because of increased statistics in this plot).

4.3 Shower leakage

The second category of data cuts is concerned with incomplete containment of showering secondaries. Since the experiment uses a calorimeter of finite width and depth, energy leakage from the shower tails is certain to occur at some level. This section investigates techniques for identifying and rejecting punchthrough.

4.3.1 Lateral energy containment

We look first at the lateral energy distribution within the shower. Using the strip readout of the ULAC, an energy profile across the calorimeter may be built. An example of this for a minimum-bias event is shown in Figure 4.5. The width of the shower is clearly limited; indeed the fraction of ULAC energy detected in the four central strips is approximately 83%, and is independent of the total measured energy. By studying the lateral shower




Figure 4.5: Shower profile in the ULAC calorimeter as a function of strip and floor number (the area of each box is proportional to the energy in that strip).




Figure 4.6: Longitudinal shower centroid C_{long} as a function of missing energy (see text for centroid definition).

shape throughout the dump (including the BEAM and VETO), no significant energy deposition is seen near the calorimeter edges. That this is not a function of the measured event energy indicates that lateral shower leakage does not make a significant contribution to the experimental missing energy spectrum.

4.3.2 Longitudinal energy containment

The complete beam dump corresponds to approximately $10 \lambda_i$; hence energy leakage from the back of the calorimeter may be important. This section examines the evidence for both charged (including muon) and neutral particle punchthrough.

Leakage effects are clearly illustrated by examining the longitudinal shower centroid. Figure 4.6 shows the longitudinal centre-of-gravity in the three components of the dump as a function of missing energy. Here the centroid has been (arbitrarily) defined as

$$C_{long} = \frac{(1.E_{ULAC}) + (2.E_{BEAM}) + (3.E_{VETO})}{E_{ULAC} + E_{BEAM} + E_{VETO}}$$

The strong correlation between the shower centroid and measured event energy shows longitudinal leakage to be a major effect. This is confirmed in the charged-particle spectrometer immediately downstream of the beam dump, where events with large missing




Figure 4.7: Hit distribution in the spectrometer wire chambers (solid line: events with $E_{mis} > 100$ GeV, dashed line: MT events).

energy ‘light up’ the 32 wire chambers. Figure 4.7 shows the total number of hits in the chambers for events with more than 100 GeV missing energy and for MT events (solid/dashed lines respectively).

4.3.3 Calorimeter cuts

The first cut to improve shower containment is a restriction on the energy in the last $1.3 \lambda_i$ of the beam dump, that is, the VETO calorimeter. The cut was imposed on-line after preparatory studies had shown an unacceptably high missing energy trigger rate due to leakage. Fitting a Gaussian to the VB event sample, the cut is made 3σ above the mean, corresponding to a maximum energy deposition of 4 GeV. This rejects 19% of the minimum-bias data.

The second cut to reduce leakage is a requirement that the primary interaction should occur within the first $1.5 \lambda_i$ of the dump, that is, in the first floor of the ULAC. The peak ADC readout of the calorimeter is used to give energy information with the necessary fine segmentation for this cut. Figure 4.8 shows the energy distribution within the the first $1.5 \lambda_i$ and indicates the position of the cut at 10 GeV. 13% of the data is rejected by this




Figure 4.8: Energy detected in the first $1.5 \lambda_i$ of the dump, indicating cut to distinguish interacting and non-interacting particles.

cut (note that the quoted cut fractions are cumulative). For the unbiased data sample, in which there is no cut on the energy in the VETO calorimeter, the corresponding fraction is 18%, roughly as expected for an unbiased sample of interactions after $1.5 \lambda_i$ of material ($e^{-1.5} = 0.22$).

One additional calorimeter cut is applied, constraining the energy in the BEAM calorimeter (comprising interaction lengths 5 to 8 of the dump). It is observed that the missing energy event rate is a function of the energy in the BEAM calorimeter. This is illustrated in Figure 4.9, showing a direct correlation between late shower development and an increased number of events with missing energy. The cut is made so as to remove this correlation, in this case at 50 GeV. However, rather than imposing a fixed cut ($E_{BEAM} < 50$ GeV), the cut threshold for each event is weighted in direct proportion to the total measured energy E_{vis} of that event. This defines a ‘floating’ cut $E_{BEAM} < \frac{50 \cdot E_{vis}}{450}$ GeV. An energy-weighted threshold is necessary to ensure that events with genuine missing energy (due to production of non-interacting particles), and consequently reduced shower energy, are cut in equal proportion to events in which the full beam energy is available for the calorimeter shower. It may be noted that this cut should not discriminate in any way among events with different values of (genuine) E_{mis} and so does not affect the shape of the true E_{mis} spectrum. Cuts with this property have been used wherever possible. 63% of the data is rejected by the BEAM calorimeter cut.




Figure 4.9: Integrated rate for events with $E_{mis} > 50$ GeV as a function of maximum BEAM calorimeter energy.

4.3.4 Spectrometer cuts

The muon spectrometer downstream of the calorimeter dump provides an additional veto on particle leakage. Designed to give position and momentum information for energetic muons, it also provides powerful rejection of both charged and neutral hadron punchthrough. These three categories are discussed in turn below.

Muons

In Section 2.1.1, known sources of events with missing energy were outlined. Such events correspond to the production of one or more neutrinos in conjunction with an equal number of charged leptons. For this analysis, the models discussed in Chapter 2 predict sources of weakly interacting neutral particles *without* associated leptons and production of neutrino-charged lepton pairs may therefore be regarded as background. The choice of beam dump configuration, and the consequent optimisation of muon detection at the expense of electron detection (Section 2.3), allows rejection of background in the muon channel. In aiming to minimise contributions to the missing energy spectrum from known sources, a muon veto is therefore applied. No such rejection is possible in the electron

channel, with the result that the measured spectrum should contain a contribution from semi-leptonic decays to electron plus neutrino.

The muon veto is enforced by demanding quiet chambers and hodoscopes in the muon spectrometer. Based on a study of MT events, the total number of hits in the 32 wire chamber planes is required to be less than four (non-zero due to noise). For the hodoscope planes, the requirement is zero hits. 46% of the data is rejected by this cut. Taking the chamber efficiencies of Appendix B, the veto efficiency for charged particles within the spectrometer acceptance is essentially 100%.

Charged hadrons

In addition to providing a muon veto, the spectrometer can be used in a similar way to reject punchthrough from late-developing showers. With approximately $10 \lambda_i$ of material, such leakage from shower fluctuations may be important (note that the probability of an incident proton penetrating the entire calorimeter without interacting is around 0.005%). The charged component of any punchthrough (p's, π^\pm 's, K^\pm 's) will be rejected by the spectrometer chamber cut described above.

Neutral hadrons

Rejection of neutral particle punchthrough would not appear possible in a device based on the detection of ionisation produced by a moving charge. The only situation in which the spectrometer chambers might be useful is in the case of decays of short-lived neutral hadrons to one or more charged secondaries, for example $K_S^0 \rightarrow \pi^+ \pi^-$. In this example, ionisation caused by the secondary pions will be detected and vetoed as above. One component of the spectrometer, however, provides a much more powerful veto on neutral hadronic punchthrough: the $5 \lambda_i$ iron wall.

As described in Section 3.1.5, the 80 cm of iron sits downstream of the last wire chamber, sandwiched between the two scintillator planes H3 and H2. A hadron incident on the front face of the wall subsequently has a high probability (99.2%) of interacting and initiating a shower. Since the wall is not instrumented, the shower in the iron is not detected. It is seen, however, when shower particles escape from the back face of the wall and produce hits in the scintillator plane H2. Figure 4.10 shows the number of hits in H2 versus missing energy for events with quiet chambers and no hits in H3. Figure 4.11 shows the distribution projected on to the H2 axis for consecutive bins in E_{mis} . Also shown is




Figure 4.10: Number of H2 slats hit versus missing energy (no track found in spectrometer).

the hit distribution for MT events.

Focusing initially on events with large missing energy ($E_{mis} > 100$ GeV), the H2 hit population is clearly double-peaked. The lower ‘spike’ is compatible with noise in the hodoscope array (see MT event plot); however there is an additional, distinct class of events in which a large number of H2 slats register hits. Examining the spatial distribution of slat hits, these events are characterised by activation across most of the central region of H2 (see Figure 3.14, Section 3.1.5 for reference). Punchthrough of energetic neutral particles (n’s, K^0 ’s), and their subsequent showers in the iron wall, can therefore be identified and rejected. It is interesting to note that the lateral extent of these showers after $5 \lambda_i$ may be as large as ~ 1.8 m. This is in rough agreement with previous measurements on punchthrough [73] and is due to low energy, large angle particle leakage.

In order to determine if there is any residual punchthrough contamination after demanding zero H2 hits, the distributions of Figure 4.11 have been used. Estimates have been made for each of the five bins as follows:-

- i). The minimum-bias data ($E_{mis} < 20$ GeV) distribution is in agreement with the corresponding MT distribution. Contamination is estimated to be 0%.

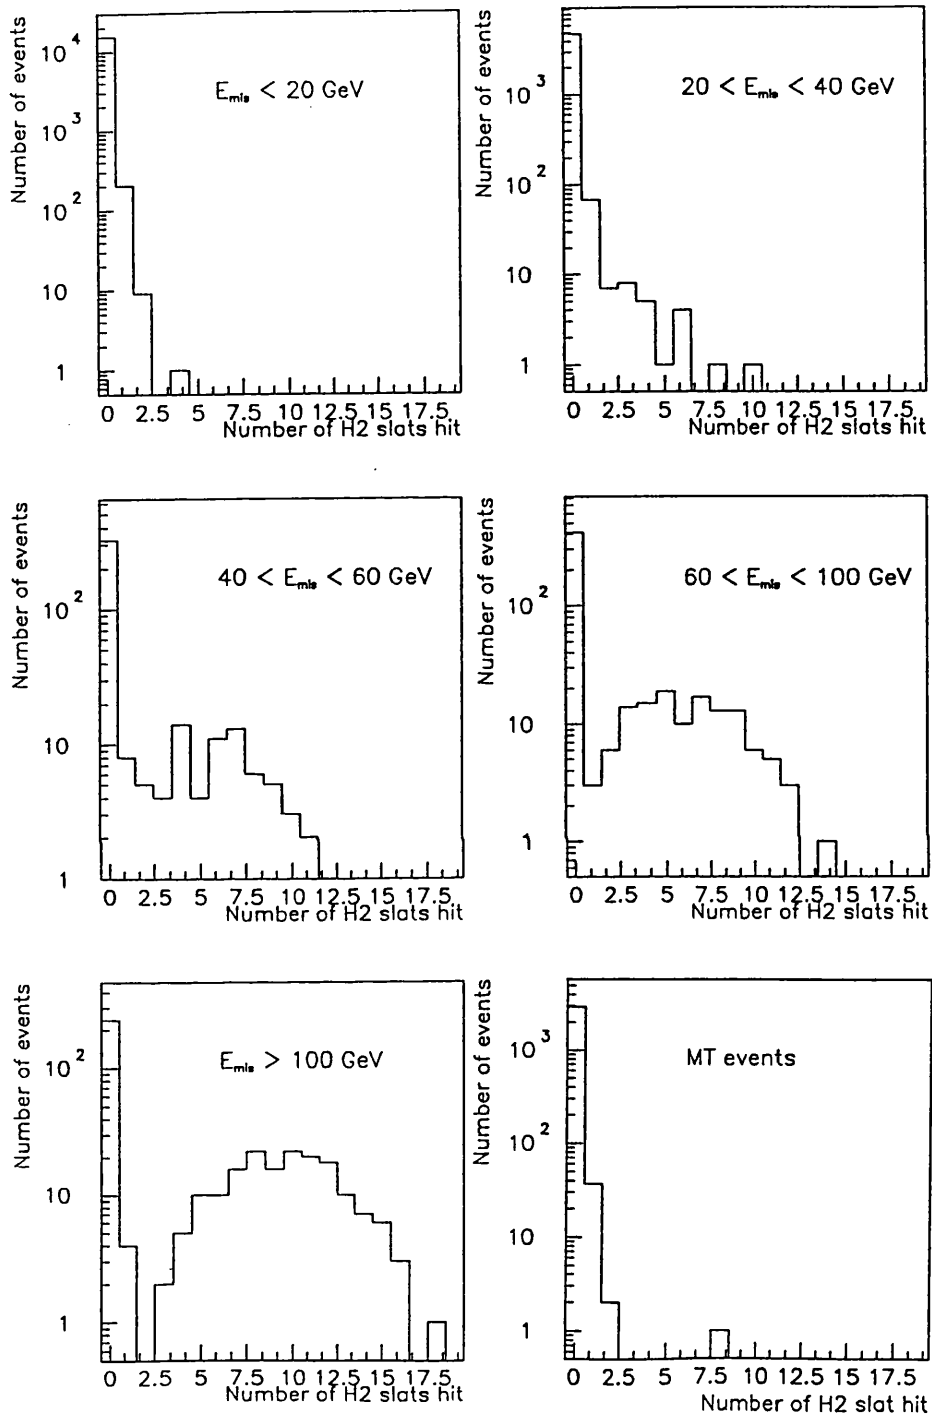


Figure 4.11: Hit distributions in H2 as a function of E_{mis} (no track found in spectrometer).

- ii). For $20 < E_{mis} < 40$ GeV, there is a small enhancement above 2 hits (but note the ratio of 0:1 hits is still compatible with the corresponding MT ratio). Contamination is estimated to be at most 1%.
- iii). For $40 < E_{mis} < 60$ GeV, the ratio of 0:1 hits has decreased. The upper limit on contamination is estimated to be 10%.
- iv). For $60 < E_{mis} < 80$ GeV, the ratio of 0:1 hits is again in agreement with the MT distribution and two distinct peaks (noise and punchthrough) are becoming evident. Contamination is estimated to be at most 1%.
- v). Finally, for $E_{mis} > 100$ GeV, it appears possible to identify unambiguously (on limited statistics) those events with neutral punchthrough. The upper limit on contamination is placed at 0.1%.

Clearly, using this technique it is not possible to determine whether there is residual punchthrough for individual events: rather, the contamination is estimated on a *pro rata* basis and errors assigned to the final spectrum. In summary, contamination levels are negligible at both large and small missing energy (the two lowest E_{mis} bins being less interesting in any case since they fall within the Gaussian of the calorimeter resolution). The level of remaining punchthrough is most significant at moderate missing energy ($40 < E_{mis} < 60$ GeV). For this bin a (pessimistic) 10% systematic error is assigned to the final data point.

It is useful here to show the analogous plot(s) for the case when a track is reconstructed in the spectrometer chambers. If the hodoscope distribution of Figure 4.10 is due to neutral hadron punchthrough, then the corresponding charged track plot should show a similar distribution due to charged hadron (p, π^\pm) punchthrough. Figure 4.12 shows the number of hits in H2 versus missing energy for events where one (and only one) track is found. The measured energy has not been updated with the track momentum from the spectrometer. Selection and quality cuts for good tracks are defined in Section 5.5.1; hits used in reconstructing the track are removed and cuts then made on the number of remaining hits in the chambers and upstream hodoscope plane (H3) as above. The corresponding projections are shown in Figure 4.13. The trend towards a double-peaked distribution with increasing E_{mis} is again seen for the single-track events. There is one significant difference however: the relative populations of the 0- and 1-hit events. The difference is due to muons, which are expected to form a significant fraction of the charged particles escaping from the calorimeter. Most muons will pass through the 80 cm iron wall losing energy by ionisation only. With 100% efficient and noiseless scintillator hodoscopes, these muons should produce one and only one hit in the downstream hodoscope plane H2.




Figure 4.12: Number of H2 slats hit versus missing energy (1 track found in spectrometer).

In reality, some of the hodoscope slats have efficiencies as low as 40% (see Appendix C) resulting in muon events with zero hits in H2.

The two distributions (Figures 4.10, 4.12) indicate that hadron punchthrough, both charged and neutral, is a significant effect but also show that it can be vetoed with high efficiency.

The missing energy spectrum with combined cuts on upstream interactions and leakage is shown in Figure 4.14.

4.4 Pileup

The final class of cuts is aimed at removing events in which two or more beam particles arrive within the time window during which the detector is sensitive. With a beam rate of up to 10^6 protons/burst and an array of detectors with readout pulse lengths of up to $1 \mu s$ (ULAC), it is clear that pileup may be a problem. In general, one might expect multi-particle pileup in a calorimetric device to result in events with measured energies significantly greater than those recorded for single-particle events. In this scenario, contamination due to pileup, though undesirable, would not be expected to affect a

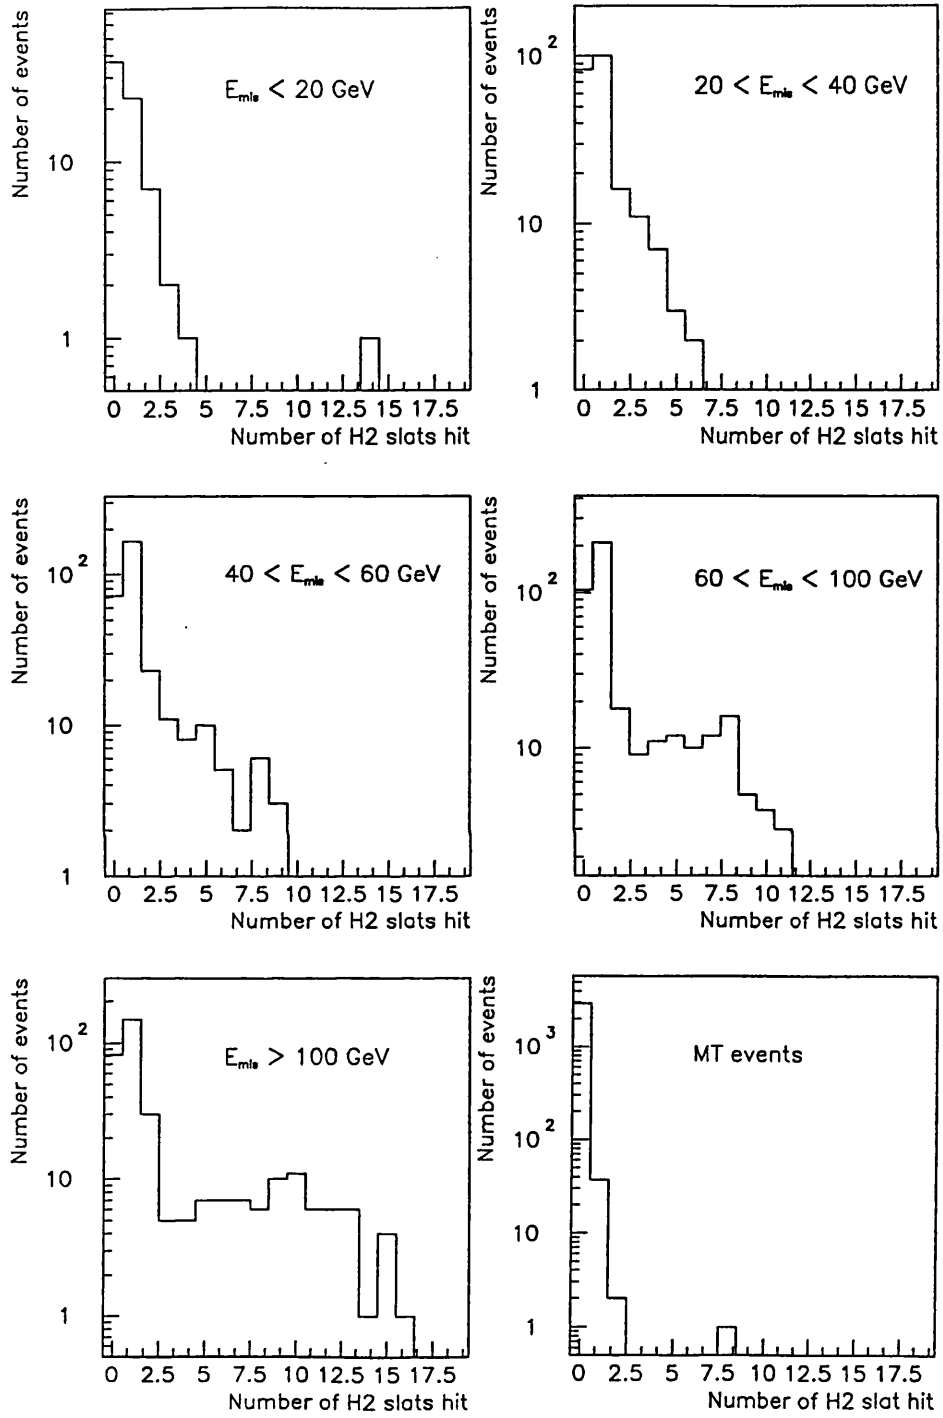


Figure 4.13: Hit distributions in H2 as a function of E_{mis} (1 track found in spectrometer).




Figure 4.14: The missing energy spectrum after cuts on leakage.

measurement of the missing energy spectrum. For the HELIOS calorimetry this is not the case. The design of the readout electronics (based on bipolar signals) is such that pileup can result in fake missing energy measurements. This section examines various techniques used in identifying and rejecting pileup, both on-line and off-line.

4.4.1 On-line protection

Preparatory studies of the detector performance in 'missing energy mode' had shown pileup to be a significant problem. The use of calorimeter triggers enhances the problem and demonstrates the need for on-line protection against pileup to avoid writing contaminated data to tape.

A first precaution taken for this data run was to try and reduce pileup at source. Instead of the usual beam rate of up to 10^6 protons/burst, the flux was reduced to $\sim 10^5$ protons/burst.

For the detector itself, an on-line veto is provided by the before/after protection window in the beam counter logic. This has been detailed in Sections 3.1.2 and 3.1.4. In summary, the logic should impose a veto on any event in which a second particle arrives within a window of:-

- i). 800 ns before and 300 ns after for on-axis particles, or
- ii). 50 ns before for off-axis particles.

The detector component most susceptible to pileup is the ULAC, having readout pulses of $\sim 1 \mu\text{s}$. In principle, the before/after protection should therefore veto all pileup within a window of $\pm 1 \mu\text{s}$. This was not implemented on-line due to rate considerations for the standard HELIOS physics programme.

A consequence of this incomplete on-line protection is the need to impose the full $\pm 1 \mu\text{s}$ window in the off-line analysis. This is done using the multi-hit TDC readout for each of the beam counters B7, B6 and B3, vetoing all events in which additional hits are recorded within the prohibited time window. As well as rejecting pileup of on-axis particles (using B6), the scheme also removes contamination from off-axis pileup (B7 and B3) in the same time interval.

4.4.2 Evidence for residual contamination

In order to check the efficiency of the pileup cuts described above, the remaining data was studied for evidence of beam rate-dependent effects. With burst-to-burst fluctuations of up to $\sim 30\%$ and long-term variations (over several hours) up to a factor of 3, it is possible to bin the data according to burst-averaged beam flux and examine the corresponding missing energy spectra. Figure 4.15 shows the data divided into only two rate bins: ‘low’ ($1 \times 10^4 < \text{flux} < 7 \times 10^4$ protons/s) and ‘high’ ($7 \times 10^4 < \text{flux} < 13 \times 10^4$ protons/s). There is a clear excess in the high rate data. The dependence of missing energy event rate on beam flux indicates that there is still contamination in the data due to pileup. The remainder of this section is devoted to studying ways of removing the contamination.

It is useful first to illustrate how the calorimeters respond to pileup and are therefore able to produce a fake missing energy signal. Table 4.1 gives a summary of the readout schemes for the two categories of HELIOS calorimeter. The effects of pileup are different in each of the three readout chains (FADC, PADC and QADC). However, since the on-line trigger is based on the FADC readout and the FADCs have so far been used for the off-line analysis, we concentrate on the FADCs here.

The on-line readout of both ULAC and UCAL calorimeters is based on flash ADC digitisations of bipolar pulse signals. As was seen in Section 3.1.4, the EFL logic is timed such that all FADCs are read out at the peak of the shaped signal, corresponding to a fixed interval after issuance of the pretrigger. If a second particle hits the calorimeter within its sensitive window, the two bipolar pulses will overlap. Depending on the time




Figure 4.15: The missing energy spectra for low and high beam rate data.

interval between these pulses, the amplitude of the original (trigger) signal may be either enhanced or diminished ('constructive' or 'destructive' interference). An illustration of this is given in Figure 4.16. The upper plot shows the calorimeter response to a single beam particle (the continuous curve being obtained by smoothing digitised values from an oscilloscope trace). Also indicated is the FADC gate (dashed line). In this particular example, a second particle has arrived before the trigger particle (centre plot). Its pulse shape is the same, and it has been chosen to have the same amplitude, corresponding to an equal energy deposition. The resulting pulse is shown in the lower plot. The negative lobe of the pileup particle overlaps with the positive lobe of the trigger particle to give a reduced amplitude at the FADC strobe, and hence a 'missing energy' signal. By including variations in particle time lag and energy deposition, almost any energy between zero and

Calorimeter	Readout pulses and digitisation					
	On-line			Off-line		
	Pulse shape	Rise time	Digitisation	Pulse shape	Rise time	Digitisation
ULAC	Bipolar	~ 100 ns	FADC	Bipolar	~ 200 ns	PADC
UCAL	Bipolar	~ 30 ns	FADC	Monopolar	-	QADC

Table 4.1: Summary of the calorimeter readouts.

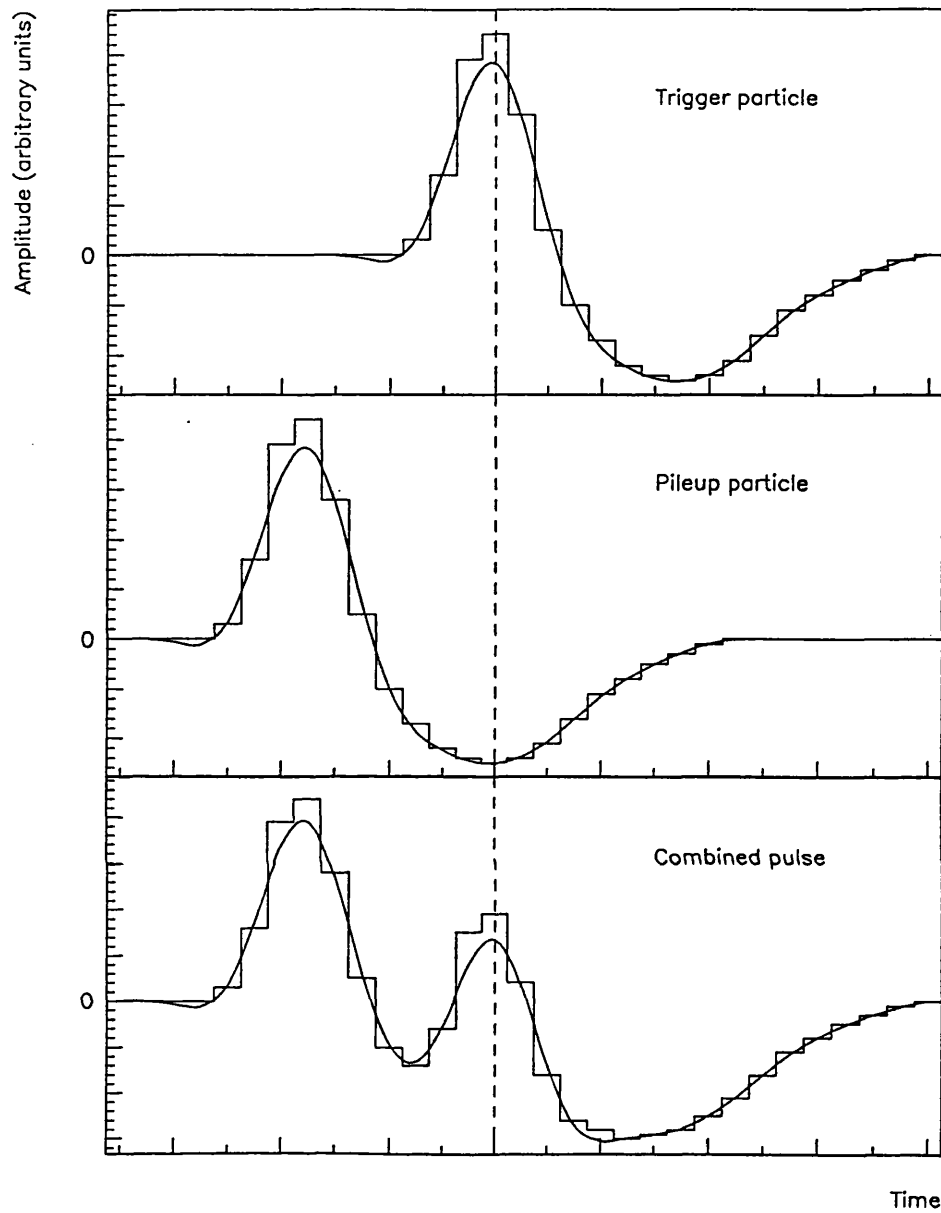


Figure 4.16: Schematic illustration of calorimeter pileup. The dashed line indicates the timing of the FADC gate. In this example, the two pulses interfere destructively, resulting in fake missing energy.

twice beam momentum can be recorded.

The remaining contamination from pileup indicates that the beam counter cuts (above) are not sufficient. There are two possible reasons for this, either:-

- i). the counter protection is not fully efficient, or
- ii). the pileup particle falls outside the beam counter acceptance.

Since the largest beam counter (B7) covers an area up to ~ 20 cm from the beam axis, it is possible to use the highly-segmented readout of the ULAC calorimeter to examine the second possibility and ascertain whether energy is indeed deposited in two distinct regions $\gtrsim 20$ cm apart. No significant activity is seen outside the central strips and case ii). above can therefore be rejected. The residual pileup appears to be due to beam counter inefficiency.

4.4.3 Off-line protection

This section describes techniques used off-line to try and identify events in which the missing energy signal is produced by pileup.

Beam counter pulseheight

The beam counters provide one further piece of information used to remove multi-particle events: the pulse-height measurement from B6 ADCs. A cut has been made to reject events with pulse-heights corresponding to two or more singly-ionising particles. (A low pulse-height cut has also been made to veto any events triggered by noise in the beam-counter readout). This is shown in Figure 4.17. 7% of the data is removed by this cut.

History FADCs

The second tool used to look for pileup is the array of history FADCs which sample the on-line calorimeter pulses as a function of time (see Section 3.1.4). Pileup in any component of the calorimeter is seen as a deviation from the single-particle pulse shape of Figure 4.16. A clear example is shown in Figure 4.18. A simple slope-counting algorithm can then be employed to reject such events with good ($> 95\%$) efficiency. Technical difficulties during data-taking unfortunately resulted in poor history FADC performance; consequently, this cut has very low efficiency ($\sim 10\%$) and is used only after all other pileup cuts have been applied.




Figure 4.17: Beam counter (B6) pulse-height spectra for valid beam (solid) and MT (dashed line) events, indicating cuts.




Figure 4.18: Pileup as seen in the ULAC history FADC (apparent missing energy for this event = 244 GeV).




Figure 4.19: PADC/FADC correlation in the ULAC (PADC energies have been re-scaled by the mean FADC/PADC ratio from minimum-bias events).

PADC/FADC correlation

In Section 4.4.2, it was noted that the different readout chains respond differently to pileup. Taking the example of Figure 4.16 on Page 99, the FADC readout of the ULAC will record a reduced energy signal (at the dashed line) while the PADC readout (which measures the maximum pulse amplitude) records the larger energy corresponding to the first peak. Hence, we may expect the PADC/FADC correlation to show anomalous behaviour for multi-particle events. Figure 4.19 shows the scatter plot of PADC versus FADC measurements of energy in the ULAC. PADC values have been re-scaled by the mean FADC/PADC ratio from minimum-bias events to take account of the on-line weighting of the FADC energies in the EFL logic (see Section 3.1.4). While most of the data falls on the diagonal (PADC = FADC) in this plot, there are events lying both above and below the central band. Using the history FADCs, these ‘off-diagonal’ events are shown to be due to pileup. A cut on the PADC/FADC correlation is therefore made, rejecting all events with PADC-FADC values more than 3σ away from the mean as defined by minimum-bias interactions (see Figure 4.20).

The correlation cut is useful in eliminating pileup; however, it will not remove pileup events (if they exist) in which the PADC and FADC readouts are corrupted equally. In




Figure 4.20: PADCFADC energy difference (minimum-bias events), indicating cut.

this case, the correlation between PADCFADC will be maintained and the event will survive the cut. Figure 4.21 shows the PADCFADC readout difference versus time separation of the two particles for events in which the history FADCs were working and detected pileup. Time separation has been estimated directly from a visual scan of pulse shapes in the history FADCs. The plots shows a systematic dependence of PADCFADC difference on the time separation (as would be expected from Figure 4.16). More importantly, it shows that there is a time window (~ 400 to 600 ns separation) for which the PADCFADC readouts remain correlated. Hence, it is probable that the data still has contamination due to pileup. Three techniques have been used to try and estimate this contamination: they are outlined below.

The excess energy spectrum

It was noted above that pileup may fake not only missing energy (destructive interference) but also excess energy (constructive interference) in the calorimeters. There are no known physics processes allowing the sum of outgoing particle energies to exceed the sum of incoming particle energies in an interaction. In contrast to events with missing energy therefore, where the missing energy signature may be due to a combination of WINP production and detector effects, events with excess energy will arise from detector effects alone.




Figure 4.21: PADC-FADC energy difference versus time separation of pileup pulses.

Under the assumption that beam particles are distributed ‘randomly’ in time (the 5 ns separation of SPS radio-frequency buckets being negligible on the 1 μ s calorimeter pulse timescale), pileup giving rise to constructive interference should occur. A measurement of the corresponding excess energy spectrum, taken under the same experimental conditions as the missing energy data, can then be used to estimate the remaining pileup contribution in the missing energy spectrum.

The asymmetric shape of the calorimeter pulse does not give a ‘one-to-one’ correspondence between the pileup contribution to the excess energy spectrum at a certain energy offset (say, $E_{exc} = 100$ GeV) and the pileup contribution to the missing energy spectrum at the same offset ($E_{mis} = 100$ GeV). An approximate parametrisation is shown in Figure 4.22. The two lobes of the bipolar pulse have both a different peak amplitude and different length. This dictates the need for both an ‘offset’ and ‘rate’ factor (respectively) in the extrapolation. The offset factor reflects the relative amplitude shifts for constructive and destructive pileup. Roughly:-

$$\text{offset}(E_{mis}) \sim y - by$$

$$\text{offset}(E_{exc}) \sim y + y$$

and hence:-

$$\text{offset}(E_{mis}) \sim \frac{1-b}{2} \text{offset}(E_{exc}).$$




Figure 4.22: Approximate parametrisation of calorimeter pulse.

The rate factor reflects the relative frequency of occurrence of constructive and destructive pileup. Here;-

$$\text{rate}(E_{mis}) \sim a \text{rate}(E_{exc}).$$

Combining these two effects and taking realistic values for the factors ($a \sim 2, b \sim \frac{1}{2}$), the predicted pileup rate at, say, $E_{mis} = 100$ GeV is roughly twice the measured pileup rate at $E_{exc} = 400$ GeV. It should be emphasised that this provides only an approximate prescription for the extrapolation.

Section 3.4 gives details of the excess energy measurement. This data was taken simultaneously with a part of the missing energy data, to minimise systematic errors in the comparison. Trigger conditions in each data set were identical with two exceptions: firstly, attenuation of the ULAC hadronic FADC signal and secondly, a pulse-height cut on the beam counter signal. After off-line re-scaling of the ULAC hadronic energy (reversing the effect of on-line attenuation), the data is passed through the complete series of cuts described in this chapter to produce the excess energy spectrum. This is shown in Figure 4.23. The data is clearly consistent with calorimeter resolution effects alone up to 50 GeV excess energy. Beyond 50 GeV, there are no data points. This is a consequence of limited statistics in excess energy mode (around 5% of the missing energy data sample). An estimate of pileup in the missing energy data by extrapolation from the excess energy spectrum is therefore not possible.

Computer simulation

An attempt was subsequently made to try to estimate the remaining pileup by modelling pulse overlaps and simulating the readouts of peak and flash ADCs. By tuning the simulation so that it reproduces the observed off-diagonal distribution in the PADC/FADC plane, one might then have confidence in its prediction of the level of pileup falling on the diagonal, allowing a subtraction to be made.




Figure 4.23: The excess energy spectrum. Data cuts have been applied exactly as for the missing energy spectrum.

The pulse shape is taken from digitisation of the ULAC signal as seen by an oscilloscope (see Figure 4.16). Two such pulses are then superimposed, the time interval between them being generated randomly within the relevant range ($\pm 900\text{ns}$), and the peak and flash ADC readouts simulated by a peak-finding algorithm and fixed 'gate' respectively. Variation of the pulse amplitude is also introduced to simulate event-by-event fluctuations in energy deposition. The size and frequency of these fluctuations are taken from the measured ULAC energy distribution for events with no pileup. Finally, the overall energy scale is determined by setting the average pulse-height equal to the mean ULAC energy deposit in minimum-bias events.

The resulting PADC versus FADC distribution reproduces the gross characteristics of the corresponding data sample. However, the relative population of the PADC/FADC plane is rather strongly dependent on assumptions in the model. This level of uncertainty makes a correction based on simulation unreliable; the method is not used.

Beam rate extrapolation

The final method used to estimate pileup was to measure the missing energy spectrum as a function of burst-averaged beam rate, and remove the contribution from rate-dependent

effects (i.e. pileup) by extrapolation to zero flux.

As the first step in this process, it is necessary to hypothesise the form of the dependence of pileup frequency on beam rate. Assuming that the probability of a particle arriving at a given instant is not dependent on the time distribution of the preceding particles, the frequency of two-particle overlap in a certain time interval should have a quadratic dependence on the beam rate. The data can subsequently be used to test this hypothesis. The event sample is binned burst-by-burst according to the average beam flux within the burst (obtained from a scaler count of the total number of particles). Plotting the missing energy spectrum as a function of flux-squared, the quality (chi-squared) of a simple linear fit for each E_{mis} bin gives a direct test of the quadratic dependence. Note that this method is only valid under the assumption of an essentially uniform beam rate within the burst. This assumption was verified using a storage oscilloscope connected to the beam counters. To improve both the lever-arm and the statistical precision of the extrapolation, an additional data sample is included with the existing data. This data was taken during preparatory studies under identical experimental conditions but with an average beam flux about a factor of 3 higher than in the original sample. This 'high-rate' data is subsequently processed through the identical series of cuts as described above for the low-rate data.

Figure 4.24 shows the missing energy spectrum in 12 E_{mis} bins as a function of the squared beam rate. Three rate bins are shown in this instance. Also shown are the best straight-line fits for each of the twelve bins and the corresponding chi-squareds per degree of freedom. For E_{mis} greater than 240 GeV there is insufficient data to give measurements in each of the three rate bins. Analysis of the final spectrum is therefore based on data points in the range $0 < E_{mis} < 240$ GeV.

Looking at Figure 4.24, the data is in reasonable agreement with the hypothesis of quadratic rate dependence (as shown by the fit quality). The errors are large however, giving rise to (unphysical) negative slopes for some bins. The slope is not constrained to be non-negative in these cases. The zero-rate value of E_{mis} and its associated error are then taken directly from the straight-line fit. In view of the uncertainty associated with the extrapolation procedure and in order to give a more realistic estimate of the error on the zero-rate values, the statistical error on the lowest-rate bin is added in quadrature to the error defined by the fit. The extrapolation method gives at most a $\sim 30\%$ correction and in all cases the magnitude of the correction is less than its corresponding error.

Figure 4.25 shows the final missing energy spectrum after all cuts and corrections have been applied. The shape is by now familiar: the data points follow a Gaussian (with an

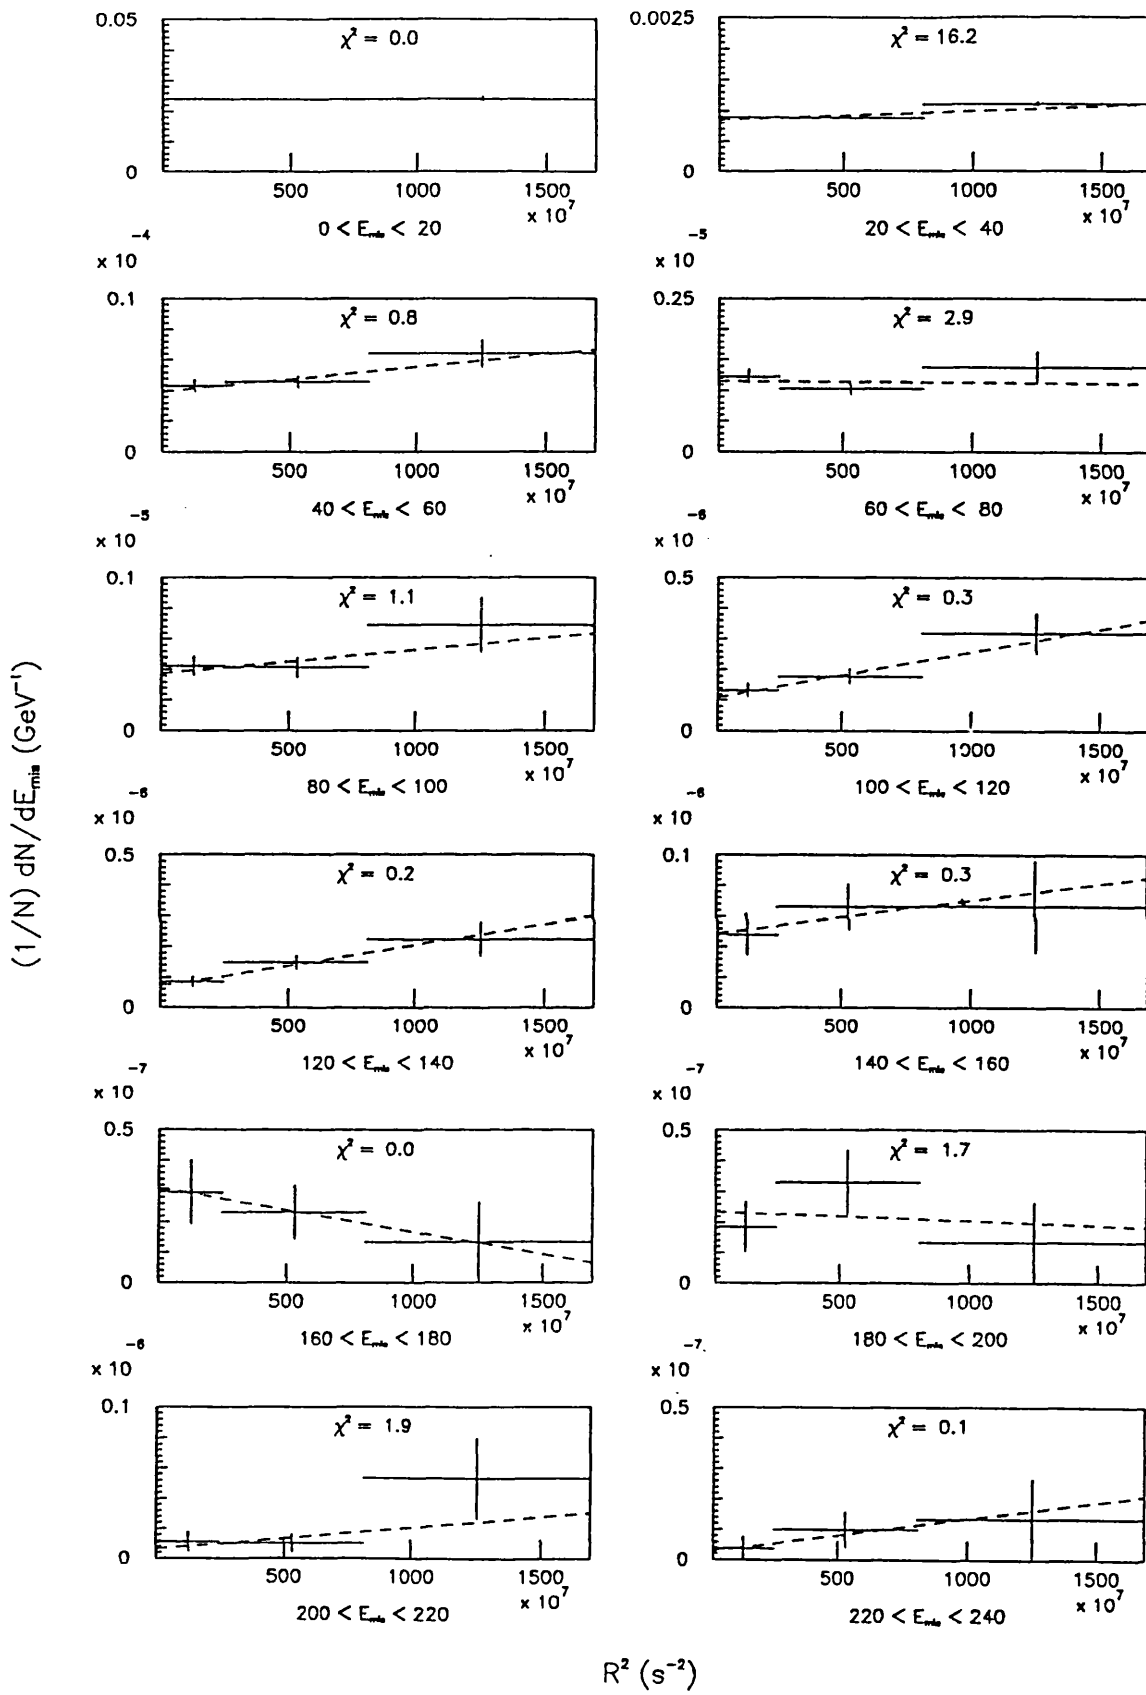


Figure 4.24: Missing energy rate as a function of beam flux squared.




Figure 4.25: The final missing energy spectrum, after rate extrapolation.

improved σ of ~ 10.6 GeV) up to about 4σ deviation from average beam energy, followed by a tail of events extending well beyond the expected calorimeter resolution to missing energies of 200 GeV or more. The actual values of the missing energy differential rate $\frac{1}{N} \frac{dN}{dE_{mis}}$ are given in Table 4.2.

This spectrum represents a measurement of the cross-section for production of WINPs as a function of their associated energies. It should be emphasised that this is a measurement of the *relative* cross-section for WINP production with respect to minimum-bias

E_{mis} (GeV)	$\frac{1}{N} \frac{dN}{dE_{mis}}$ (GeV^{-1})	Error (GeV^{-1})
0 - 20	2.41×10^{-2}	6.07×10^{-4}
20 - 40	8.40×10^{-4}	2.61×10^{-5}
40 - 60	3.96×10^{-6}	5.64×10^{-7}
60 - 80	1.15×10^{-6}	1.59×10^{-7}
80 - 100	3.77×10^{-7}	9.41×10^{-8}
100 - 120	1.09×10^{-7}	3.38×10^{-8}
120 - 140	7.02×10^{-8}	2.72×10^{-8}
140 - 160	4.86×10^{-8}	1.98×10^{-8}
160 - 180	3.12×10^{-8}	1.46×10^{-8}
180 - 200	2.34×10^{-8}	1.21×10^{-8}
200 - 220	6.21×10^{-9}	1.01×10^{-8}
220 - 240	2.70×10^{-9}	5.75×10^{-9}

Table 4.2: The missing energy differential rate $\frac{1}{N} \frac{dN}{dE_{mis}}$ in 450 GeV p - ^{238}U collisions.

interactions: no attempt has been made to determine an absolute ('millibarn') cross-section. In the following chapter, the expected contributions to this measured spectrum from known physics and detector processes will be determined. A subsequent comparison between observed and predicted spectra allows a search for anomalous WINP production to be made and is described in Chapter 6.

Chapter 5

Monte Carlo prediction of the missing energy spectrum

Figure 4.25 shows the final missing energy distribution after all data cuts and corrections have been made. The aim of this chapter will be to understand the contributions to that distribution from instrumental effects and from known physics processes giving rise to neutrino production. The comparison between measured and predicted rates of missing energy events can then be used to set limits on the production of weakly interacting neutral particles.

The expected contribution of finite calorimeter resolution to any energy measurement has already been described (Section 3.2.4 and beginning of Chapter 4). Purely statistical considerations imply that a calorimeter measurement of monoenergetic particles should follow a Gaussian distribution. This has been shown to be true in both the missing energy and excess energy data to more than 4σ (~ 45 GeV) from the mean. However, the low-energy side of the measured distribution shows a clear deviation from the calorimeter Gaussian, with event missing energies of up to ~ 300 GeV. Standard Model sources of neutrinos are expected to account for some fraction of this low-energy tail.

This chapter describes a Monte Carlo simulation of neutrino production. The first section gives quantitative estimates for Standard Model neutrino production at this energy. Then, having ascertained which mechanisms are important, the Monte Carlo event generator is optimised accordingly and interfaced with a simple detector simulation. The remaining sections describe tuning of the event generator to reproduce existing experimental data, an important cross-check of the Monte Carlo using single muon data from the beam dump, and finally, the predicted missing energy spectrum for comparison with data.




Figure 5.1: Hadronic production and semi-leptonic decay of charmed D mesons.

5.1 Standard Model neutrino production

5.1.1 Semi-leptonic heavy flavour decay

Hadronic production of heavy flavour quarks was first seen 1974, with the observation of the J/Ψ (bound $c\bar{c}$) state in p-Be collisions. Since then, open charm production, charmed baryons and charm decay have been observed and the existence of a fifth (beauty, or bottom) quark is experimentally well-established. A companion top quark is expected from the same theoretical arguments which postulated the charm quark. From the charmed quark onward they are collectively referred to as heavy quark flavours.

Since the strong interaction conserves flavour, the lightest hadrons of each heavy flavour species cannot decay strongly and decays are consequently via the weak interaction. With lifetimes of the order of 10^{-13} to 10^{-12} s, any heavy flavour hadrons produced in this experiment will decay before interacting in the beam dump. Decays are predominantly through hadronic channels due to the large masses of the heavy flavour quarks but semi-leptonic branching ratios are still significant, typically of order 10% per lepton channel. Hence, if the production cross-sections for heavy flavour quarks in proton-nucleus collisions are large, processes such as that shown in Figure 5.1 will lead to substantial neutrino production.

Heavy flavour production in proton-nucleus collisions has recently been reviewed by Tavernier [74]. Charm production has been studied extensively for many years now and although problems remain in comparing different experiments (factors such as A-dependence, detector acceptance,...), recent results indicate that $\sigma(c\bar{c}) \sim 20 \mu\text{b}$ in p-p interactions at 400 GeV/c. This is roughly 10^{-3} of the total p-p cross-section and charm is therefore expected to be an important source of neutrinos in the beam dump data. The hadronic production of beauty flavoured particles is several orders of magnitude smaller than the production of charmed particles. Neutrino production by semi-leptonic B decay is therefore negligible in comparison to charm and may be ignored in the Monte Carlo simulation. Mass limits on the top quark exclude its production at $\sqrt{s} \sim 30 \text{ GeV}$ and it has also been ignored in the Monte Carlo.

5.1.2 π, K decay

The choice of beam dump configuration for this search was motivated by the desire to minimise neutrino production from charged pion and kaon decays. By using a uranium target, the suppression factor is as large as possible and background from these ‘trivial’ sources should be reduced by several orders of magnitude. However, with abundant production of pions and kaons in proton-nucleus collisions, these decays are not eliminated completely and since there is no variable-density capability for the HELIOS beam dump, the remaining π, K contribution cannot be measured by extrapolation to infinite dump density. Pion and kaon decays have therefore been included in the Monte Carlo.

5.1.3 Hyperon decay

The strange baryons (Λ, Σ, Ξ) have semi-leptonic decay modes giving neutrinos in the final state. With production cross-sections at the millibarn level and typical branching ratios for semi-leptonic decays in the range 10^{-3} to 10^{-5} , hyperon decay may be significant as a source of missing energy. It has therefore been included in the Monte Carlo.

5.1.4 τ lepton decay

Hadronic production of lepton pairs is possible via the Drell-Yan mechanism, as shown in Figure 5.2. The final state leptons may be either electrons, muons or taus. Unlike the other two, the τ is very short-lived ($\tau_\tau = 3.3 \times 10^{-13} \text{ s}$) and decays into a τ neutrino plus




Figure 5.2: Drell-Yan production of lepton pairs via $q\bar{q}$ annihilation.

other particles, which may include additional neutrinos. The cross-section for tau production in proton-nucleus interactions has been calculated [76] using data from direct $\mu^+\mu^-$ production in hadronic interactions. The total cross-section $\sigma_{total}(pp \rightarrow \tau^+\tau^-)$ is estimated to be ~ 180 pb, roughly 5×10^{-9} of the inelastic p-p cross-section. Hence, production of tau leptons is insignificant and may be ignored in the Monte Carlo simulation.

5.1.5 Vector boson decay

With a mass of approximately 91 GeV, production of real Z^0 particles is not possible at this energy. However, virtual Z^0 s may be produced through resonance production of low mass vector mesons such as the J/Ψ , and their subsequent decay into a virtual Z^0 . This is shown in Figure 5.3. Estimates of the $\nu\bar{\nu}$ decay mode of the J/Ψ [77] give a ratio

$$\frac{\Gamma(J/\Psi \rightarrow \nu\bar{\nu})}{\Gamma(J/\Psi \rightarrow \mu^+\mu^-)} \simeq \frac{3\left(\frac{G_F}{\sqrt{2}}\right)^2}{\frac{4}{9}\left(\frac{\alpha}{M_{J/\Psi}^2}\right)^2} \simeq 10^{-3},$$

and taking a cross-section for J/Ψ production and decay to muons of 14 nb [78], then $\sigma_{total}(pp \rightarrow J/\Psi \rightarrow \nu\bar{\nu}) \sim 14$ pb. This is roughly 5×10^{-10} of the inelastic cross-section and has therefore also been ignored in the Monte Carlo.

In summary, the primary sources of neutrino production at this energy are expected to be semi-leptonic decays of charmed and strange particles, and decays of pions and kaons.

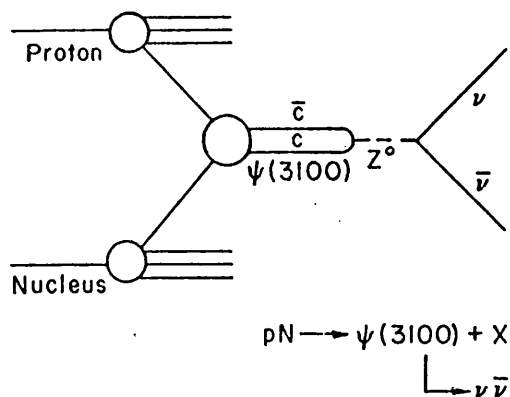


Figure 5.3: Hadronic production of J/Ψ resonance and subsequent decay to neutrinos via virtual Z^0 .

A complete list of the relevant decays (and their associated branching ratios) included in the Monte Carlo is given in Table 5.1.5. Other processes are expected to occur at such low levels that they would not be observable in this analysis.

5.2 The event generator

The aim of the missing energy Monte Carlo is to reproduce as accurately as possible the detector response to Standard Model processes in 450 GeV proton-nucleus collisions. The requirements are twofold: firstly, the primary interaction and any subsequent particle decays should be in accordance with experimental measurements and secondly, the interactions of these particles in the various detectors should be as realistic as possible. This section deals with the first component, the event generator, leaving the detector simulation to Section 5.3.

There are at present several Monte Carlo programs available which simulate hadron-hadron collisions at high-energy. They have in common the picture of an interaction as comprising a hard scatter (using QCD matrix elements and structure functions) followed by parton fragmentation into hadrons. In general, initial and final state radiation and

Decays	Branching ratio
$\pi^+ \rightarrow \mu^+ \nu$	1.00
$K^+ \rightarrow \pi^0 e^+ \nu$	4.82×10^{-2}
$K^+ \rightarrow \pi^0 \mu^+ \nu$	3.18×10^{-2}
$K_L^0 \rightarrow \pi^+ e^- \nu$	3.86×10^{-1}
$K_L^0 \rightarrow \pi^+ \mu^- \nu$	2.70×10^{-1}
$\Lambda \rightarrow p e^- \nu$	8.30×10^{-4}
$\Lambda \rightarrow p \mu^- \nu$	1.60×10^{-4}
$\Sigma^+ \rightarrow \Lambda e^+ \nu$	2.00×10^{-5}
$\Sigma^- \rightarrow \Lambda e^- \nu$	6.00×10^{-5}
$\Sigma^- \rightarrow n e^- \nu$	1.02×10^{-3}
$\Sigma^- \rightarrow n \mu^- \nu$	4.50×10^{-4}
$D^+ \rightarrow K^0 e^+ \nu$	9.90×10^{-2}
$D^+ \rightarrow K^{*0} e^+ \nu$	7.20×10^{-2}
$D^+ \rightarrow \pi^0 e^+ \nu$	9.00×10^{-3}
$D^+ \rightarrow K^0 \mu^+ \nu$	9.90×10^{-2}
$D^+ \rightarrow K^{*0} \mu^+ \nu$	7.20×10^{-2}
$D^+ \rightarrow \pi^0 \mu^+ \nu$	9.00×10^{-2}
$D^+ \rightarrow \mu^+ \nu$	6.00×10^{-4}
$D^0 \rightarrow K^- e^+ \nu$	3.85×10^{-2}
$D^0 \rightarrow K^{*-} e^+ \nu$	2.80×10^{-2}
$D^0 \rightarrow \pi^- e^+ \nu$	3.50×10^{-3}
$D^0 \rightarrow K^- \mu^+ \nu$	3.85×10^{-2}
$D^0 \rightarrow K^{*-} \mu^+ \nu$	2.80×10^{-2}
$D^0 \rightarrow \pi^- \mu^+ \nu$	3.50×10^{-3}

Table 5.1: Monte Carlo decays giving final state neutrinos (plus the appropriate charge conjugate decays where relevant).

the ‘underlying’ event are also simulated and the program may also include decays of secondary particles.

For this analysis, the LUND package has been used for event generation. The basic LUND program PYTHIA [79] is a Monte Carlo program intended for the study of high- p_T physics in hadronic interactions, but covers also the domain of low- p_T interactions as an integral part of the total cross-section. It includes matrix elements for the basic $2 \rightarrow 1$ and $2 \rightarrow 2$ Standard Model (strong and electroweak) subprocesses, elastic and diffractive scattering, structure functions and initial and final state parton showers. Whilst PYTHIA has been shown to be in good agreement with a range of hard-scattering data, it does not address the area of low- p_T interactions in any detail. The simulation of soft interactions is provided by FRITIOF [80]. Rather than adopting a QCD-inspired approach to this low- p_T region (a poorly understood domain of QCD), FRITIOF uses a model in which hadrons are treated as vortex lines in a superconducting vacuum. In this scheme, most of the produced particles are moving in a narrow cone along the beam direction, reproducing the dominant feature of minimum-bias hadron interactions. A more recent addition to the LUND package is the TWISTER program [81], again designed to simulate high- p_T hadron-hadron scattering but improved and extended for low cross-section processes such as heavy flavour and photon production.

The processes of interest for the missing energy Monte Carlo have been discussed in the previous section; those expected to be significant are semi-leptonic charm decays, pion and kaon decays and also minimum-bias events which do not in general contain neutrinos but contribute to the Gaussian resolution function of the calorimeter. In principle, all Standard Model processes may be studied by generating a sufficiently large number of minimum-bias events. This is rather inefficient for low cross-section processes however, and it is more convenient to be able to select particular modes which are then generated with the appropriate weight. Adopting this philosophy, the FRITIOF Monte Carlo has been used to generate minimum-bias events, including pion and kaon decays, and the optimised TWISTER program used to generate charm events. For the π s, K s and charmed particles, the program has been additionally optimised so that all decays include neutrinos in the final state, with weights allocated according to the actual branching ratio for that particular decay mode.

All events are generated with 450 GeV/c protons incident on a proton target (effects arising from the use of a nuclear target are discussed in Section 5.4). Secondary particles are subsequently tracked, taking account of the interaction length for the particle in the calorimeter and including the relevant decays where appropriate. Information from the complete event (particle masses, momenta, decay vertices etc.) is then stored prior to the final detector simulation.

At this stage it is interesting to make some comparisons between Monte Carlo predictions and experimental data. Appendix D shows several such comparisons (using both FRITIOF and TWISTER) for inclusive particle production in a range of hadron-hadron interactions. The FRITIOF comparisons (Figures D.1- D.4) show rather good agreement between data and Monte Carlo, allowing one to have confidence in the minimum-bias event generation. The more important comparison, that for charm production, is shown in Figure D.5. In this case, data points from the EHS collaboration [82] are superimposed on charm predictions from TWISTER. Experimental data on charm production is usually given (Section 5.4) in terms of the convenient parametrisation

$$\frac{d^2\sigma}{dp_T^2 dx_F} \sim (1 - |x_F|)^n e^{-ap_T^2},$$

where x_F denotes the Feynman x value¹ for the particle and p_T its transverse momentum. Figure D.5 also shows the appropriate curves using this parametrisation, where $n = 5$ and $a = 1$. The agreement here also looks rather good. However, the experimental situation on charm production is somewhat confusing [74] and the eventual parametrisation used for the simulation has different values for the variables n and a . This is discussed in detail in Section 5.4.

5.3 Detector simulation

Having chosen an appropriate event generator, the second requirement of the Monte Carlo is to give an accurate simulation of the experimental detector. For this analysis, the most important element of the detector is clearly the calorimeter. However, since we are also interested in identifying processes with ‘neutrino-charged lepton’ pairs in the final state, the simulation must additionally include the muon spectrometer.

Before a particle can be detected, it must undergo some sort of interaction in the material of a detector. Processes resulting from the electromagnetic interaction are the most important for many detectors although there are cases where the nuclear interaction may represent the dominant mechanism. Whilst it is in principle possible to simulate particle interactions in a detector at the atomic (or nuclear) level, in practice it is rarely necessary or even feasible. The case of particle showers in dense matter (calorimetry) is a good example here. EGS4[83] uses Monte Carlo techniques to give a full simulation of electromagnetic showers, including fluctuations, and may be used to determine the electromagnetic energy resolution of a particular calorimeter configuration. However, rather

¹Feynman’s x variable is given by $x_F = \frac{p_L}{p_{L\max}}$.

than use EGS4 for any electrons or photons generated in an interaction, a simple model or parametrisation (based on an initial EGS run) is normally adequate.

For the missing energy analysis there are two possible approaches to the question of calorimeter simulation. The key consideration is leakage. In Section 4.3, leakage was seen to be a significant problem, giving rise to a fake missing energy signal. One way of ensuring a valid comparison between data and Monte Carlo would be to include leakage in the Monte Carlo by means of a complete ('EGS-type') shower simulation. Aside from being prohibitively expensive in terms of computer time, this method has no guarantee of correctly reproducing rare shower fluctuations which may be an important background. The alternative is to use a simple Monte Carlo which simulates a 'perfect' detector and to eliminate leakage in the data with the appropriate cuts. The second approach has been chosen in this case.

Taking the idealised calorimeter described above, all secondary particles are completely contained *except* muons and neutrinos. The calorimeter energy is then simply the sum of secondary particle energies (excluding μ s and ν s) convoluted with the calorimeter resolution. The energy smearing due to calorimeter resolution is taken directly from minimum-bias data by generating points according to a Gaussian distribution with the relevant width ($\sigma = 10.6$ GeV).

For the muon spectrometer, the simulation is similarly straightforward. A geometrical acceptance cut is made, requiring $\theta_\mu < 130$ mrad, and a lower momentum threshold applied, requiring $p_\mu > 5$ GeV/c (see Section 3.1.5). Muons outside these bounds are not detected. For the missing energy analysis, the spectrometer is used as a veto on muons. The veto efficiency in the Monte Carlo is:-

$$\varepsilon_{\bar{\mu}} = \begin{array}{ll} 100\% & \text{for } \theta_\mu < 130 \text{ mrad, } p_\mu > 2.5 \text{ GeV} \\ 0\% & \text{otherwise,} \end{array}$$

as outlined in Section 4.3. However, although its primary function is to veto muons, the spectrometer may also be used to measure those same muons and provide a useful cross-check on expected background processes (such as charm production) in the Monte Carlo. This is discussed in detail in Section 5.5 but the relevant Monte Carlo efficiencies are included here for completeness. The muon detection efficiency is:-

$$\varepsilon_\mu = \begin{array}{ll} 100\% & \text{for } \theta_\mu < 130 \text{ mrad, } p_\mu > 5 \text{ GeV} \\ 0\% & \text{otherwise.} \end{array}$$

In this case, the value of 100% is certainly too high since no account is taken of the various inefficiencies (chamber, tracking, hodoscope,...) in muon reconstruction. Here also, the approach has been to simulate a 'perfect' detector in the Monte Carlo while correcting for detector inefficiencies in the data.




Figure 5.4: Preliminary simulation of the missing energy spectrum, showing contributions from calorimeter fluctuations, π , K and charm decays.

5.4 The contribution from charm production

It is now possible to use the Monte Carlo to produce a preliminary missing energy spectrum. Figure 5.4 shows the results of the simulation for 450 GeV/c protons interacting in the uranium calorimeter, and indicates the relative contribution of different processes to the overall spectrum. The detector simulation is as described in Section 5.3, including the veto on muons falling within the spectrometer acceptance. Several points are noteworthy. Firstly, for small E_{mis} the dominant contribution to the missing energy spectrum comes from fluctuations in the calorimeter response. Secondly, there is a tail of events extending to more than 200 GeV E_{mis} due to charm decays. Since a muon veto is imposed, this tail should be predominantly due to semi-leptonic decays in the electron channel where the electron interacts and is 'lost' in the calorimeter. Finally, the contribution from pion and kaon decays is negligible for all E_{mis} being at least an order of magnitude less than the contribution from calorimeter fluctuations alone. This illustrates clearly the reasons for choosing the densest possible absorber material for the beam dump.

Figure 5.4 represents the expected Standard Model background for a WINP search using the missing energy technique. Outside the calorimeter resolution ($E_{mis} > 60$ GeV), this background is due exclusively to charm decays. Since this is the level above which

the search will be made for a candidate signal, it is important to understand the charm background as accurately as possible. The charm parametrisation used in Figure 5.4 is that given by the EHS collaboration[82]. This is not, however, the parametrisation used for the eventual comparison with data. The arguments governing this decision are discussed below.

5.4.1 Total charmed particle production cross-section

Since the first evidence for charm production in hadron-hadron interactions in the mid-1970s, more than fifty papers have been published on charmed particle production in proton-nucleon interactions. The results cover a wide range in centre-of-mass energy (11.5 to 63 GeV), target nucleus (p to W), phase space acceptance and detector design. However, no data has yet been published on charm production in 450 GeV/c proton-uranium collisions (the configuration used in this experiment). For comparison purposes, it is useful to establish a common reference point between the different data sets, in this case the *total charmed pair production cross-section in p-p interactions*.

As a first step in determining the most relevant charm results, only data taken at roughly similar centre-of-mass energies will be discussed below. This limits our scope to experiments using proton beams with momenta between 350 and 400 GeV/c. Two general comments are in order here.

First, many results have been obtained using heavy targets. To derive the cross-section on protons an assumption has to be made about the atomic number (A) dependence of the cross-section. Hadronic cross-sections are well described by the parametrisation $\sigma(pA) \sim A^\alpha$ for $A > 4$. Hydrogen data tends to fall somewhat below such a fit. For the total inelastic cross-section at high energy $\alpha \simeq 0.72$ [84], however for charm production the situation is much less clear. Theoretically, the production of heavy quarks can be described using either perturbative or nonperturbative techniques: in general, perturbative QCD models require A^1 while nonperturbative or diffractive models favour an $A^{\frac{2}{3}}$ dependence. Experimentally, the situation is similarly unclear, where a roughly linear A dependence is required to achieve compatibility between different experiments, in contrast with direct measurements of α which tend to favour $\alpha \simeq 0.75 - 0.80$. This is discussed in detail in Section 5.4.1 but is an important factor in the comparison between data sets.

The second comment concerns detector acceptance. Many experiments measure the charm cross-section in a limited part of the phase space (often at large x_F). The total cross-section is then obtained by extrapolation and is rather model dependent. Experimental data on x_F and p_T dependence is reviewed in Section 5.4.1.

Three types of experiment have been performed to measure charm production in the energy range of interest. They are:-

- i). neutrino beam dump experiments,
- ii). muon beam dump experiments, and
- iii). a bubble chamber experiment.

Neutrino beam dumps form by far the largest class of experiment, but only detect neutrinos with $E > 20$ GeV and in a small forward cone. This corresponds to the high x_F region, typically $\langle x_F \rangle \simeq 0.8$. Taking the two experiments with the smallest errors and assuming a linear A dependence, then the corresponding cross-sections are:-

E613[85,86]	400 GeV/c	W target	$\sigma(c\bar{c}) = 15.5 \pm 0.8 \pm 2.3 \mu\text{b}$
CHARM[87]	400 GeV/c	Cu target	$\sigma(c\bar{c}) = 15.5 \pm 2.6 \pm 1.2 \mu\text{b}$.

In the case of the muon beam dumps, the angular acceptance is larger than that in the neutrino experiments, and as a result the muon data should be less sensitive to the parametrisation of the charm cross-section. Taking again the experiment with the smallest errors and using an $A^{1.0}$ dependence:-

CCFRS[88]	350 GeV/c	Fe target	$\sigma(c\bar{c}) = 11.3 \pm 1.1 \pm 1.8 \mu\text{b}$.
-----------	-----------	-----------	---

Only one experiment has been performed in category iii)., using a hydrogen bubble chamber as both target and vertex detector, followed by a downstream spectrometer. For a measurement of the charm cross-section in proton-proton interactions this experiment has two advantages: firstly, it uses an H_2 target thereby avoiding an assumption on A dependence, and secondly, it has a good acceptance in x_F reducing uncertainties in the cross-section parametrisation. The charm cross-section from this experiment is:-

EHS[82]	400 GeV/c	H_2 target	$\sigma(c\bar{c}) = 15.1 \pm 1.7 \mu\text{b}$.
---------	-----------	---------------------	---

There is good agreement between the four experiments, indicating a charm pair cross-section of around $15 \mu\text{b}$ or around 5×10^{-4} of the inelastic p-p cross-section.

The corresponding charm cross-sections assuming an $A^{0.72}$ dependence (following the inelastic cross-section) are:-

E613	W target	$\sigma(c\bar{c}) = 66.7 \pm 3.4 \pm 9.9 \mu\text{b}$
CHARM	Cu target	$\sigma(c\bar{c}) = 49.6 \pm 8.3 \pm 3.8 \mu\text{b}$
CCFRS	Fe target	$\sigma(c\bar{c}) = 35.0 \pm 3.4 \pm 5.6 \mu\text{b}$
EHS	H_2 target	$\sigma(c\bar{c}) = 15.1 \pm 1.7 \mu\text{b}$.

Clearly the experiments with heavy targets derive a larger charm cross-section in this case. The discrepancy between experiments appears to suggest an A dependence greater than $A^{0.72}$. However the tendency for hydrogen data to fall below an A^α fit (noted above) should be recalled here.




Figure 5.5: The variation of α with x_F for production of light flavour hadrons. (Data points are for a range of incident proton energies between 24 and 400 GeV). The curve is a fit to the data.

Atomic number dependence

One experiment[89] has attempted to measure the A dependence of charm directly. Using 400 GeV/ c protons on targets of beryllium, copper and tungsten, the α parameter is determined to be $\alpha = 0.75 \pm 0.05$. A plausible explanation for the difference between this value of α and that suggested by the cross-section measurements above is that α is a function of x_F . Figure 5.5 shows results[90] on the variation of α with x_F for production of light flavour hadrons. Also shown (Figure 5.6) are the existing data points on the variation of α_{charm} with x_F . The value of 0.75 ± 0.05 given above is measured in a neutrino beam dump experiment in which the average x_F range $\langle x_F \rangle \simeq 0.45$. Conversely, several bubble chamber experiments (including EHS and others at lower energy) which have good acceptance for all $x_F > 0$, are mutually compatible only if linear A is assumed. Thus the data is very suggestive of a dependence of α on x_F such that $\alpha \simeq 1.0$ for $x_F = 0$, decreasing to $\alpha \simeq 0.7$ for $x_F \rightarrow 1$.




Figure 5.6: The variation of α with x_F for charm production. The curve is the curve from Figure 5.5 displaced so that $\alpha = 1$ at $x_F = 0$.

Energy dependence

Figure 5.7 shows a compilation of data on the charm pair cross-section for a range of centre-of-mass energies between $\sqrt{s} \sim 12$ and 63 GeV. All entries for $\sqrt{s} < 40$ GeV are from fixed-target experiments, the remainder are results from the CERN p-p colliding beam machine ISR. Clearly the cross-section is increasing with increasing \sqrt{s} and, with the exception of the ISR data, is in reasonable agreement with QCD predictions.

HELIOS, with a beam momentum of 450 GeV/c, sits at $\sqrt{s} = 29.1$ GeV. Its nearest neighbours are, on one side, the cluster of experiments using 400 GeV/c proton beams ($\sqrt{s} = 27.4$ GeV), and on the other side, a preliminary result from Fermilab experiment E743[91] at $\sqrt{s} = 38.8$ GeV. The charm cross-section appears to increase by a factor of about 2 in this range. Assuming a linear rise, the increase in cross-section from $\sqrt{s} = 27.4$ to 29.1 GeV is estimated to be at most $\sim 20\%$.




Figure 5.7: Charm pair production cross-section as a function of centre-of-mass energy \sqrt{s} . Superimposed are the results of a QCD parton model prediction[74] using the structure functions of Duke and Owens.

Experiment	Momentum (GeV/c)	Target	Acceptance	n	a (GeV/c) ⁻²	b (GeV/c) ⁻¹
NA16[93]	360	p	$x_F > 0.0$	1.8 ± 0.8	1.1 ± 0.3	
EHS[82]	400	p	$x_F > 0.0$	4.9 ± 0.5	1.21 ± 0.14	
CCFRS[88]	350	Fe	$x_F > 0.2$	6.0 ± 0.8	0.75 ± 0.2	2.0 ± 0.4
E613[85,86]	400	W	$\langle x_F \rangle \simeq 0.45$	3.2 ± 0.2		1.5 ± 0.2

Table 5.2: Compilation of data on x_F and p_T distributions for D mesons. In all cases, n is defined by $dN/dx_F \sim (1 - |x_F|)^n$ while the parameters a and b are defined by $dN/d|p_T|^2 \sim e^{-a|p_T|^2}$ and $dN/d^2p_T \sim e^{-b|p_T|}$ respectively.

Differential cross-sections

As was seen in Section 5.2, the longitudinal and transverse momentum dependence of charm production is often described in terms of the variables x_F and p_T . Typically, experimental data is parametrised in the form

$$\frac{d^2\sigma}{dp_T^2 dx_F} \sim (1 - |x_F|)^n e^{-ap_T^2}.$$

Models of charm production contain predictions of n and a , for example quark counting rules[92] suggest $n = 5$ for central, $n = 3$ for diffractive meson production, and $n = 1$ for diffractive baryon production. Furthermore, the value of n may also depend on the nature of the charmed particle in question.

Experimentally, the data on differential cross-sections is not easy to interpret. Different experiments use different parametrisations and are sensitive to different areas of the phase space. Also, there is the additional complexity of a heavy target in many cases, which may influence the x_F and p_T distributions. Four experiments have published results using proton beams in the energy range of interest. The data is summarised in Table 5.2.

The results on the transverse momentum distributions are broadly compatible and indicate that charmed particles are produced with $\langle p_T \rangle \sim 900$ MeV/c, significantly larger than the corresponding value $\langle p_T \rangle \sim 300$ MeV/c for non-charmed mesons. There is much poorer agreement concerning the x_F dependence however. This cannot easily be interpreted as a phase space effect since experiments covering similar x_F regions[93,82] report very different values for the parameter n .

In summary, there remains considerable uncertainty concerning certain aspects of charm production, in particular the x_F dependence and the effect of a heavy target. The final choice of charm parametrisation used in the Monte Carlo is discussed in Section 5.4.2 below.




Figure 5.8: Monte Carlo simulation of the missing energy spectrum using different charm parametrisations from published data. The hatched area indicates the range of variation at large E_{mis} .

5.4.2 Charm parametrisation used in the Monte Carlo

The need for a reliable estimate of the charm contribution to the missing energy spectrum was outlined in Section 5.4. As a quantitative illustration of the level of uncertainty in charm production, it is interesting to take a range of published charm parametrisations and calculate the corresponding missing energy spectra for the HELIOS beam dump configuration. Figure 5.8 shows the results of such a calculation using data from three experiments: EHS[82], CCFRS[88] and E613[85,86]. In each case, the total charmed pair cross-section on uranium is derived assuming an $A^{0.75 \pm 0.05}$ dependence (Section 5.4.1) and the differential cross-section parametrisations taken directly from published data. A $\pm 1\sigma$ error band, obtained by allowing all measured parameters to vary by one standard deviation, is subsequently calculated for each experiment. The hatched area of Figure 5.8 represents the expected variation in the missing energy spectrum due to these different parametrisations.

The choice of charm parametrisation clearly gives rise to very significant differences in the missing energy spectrum, with variations of almost two orders of magnitude for $E_{mis} > 200$ GeV. Whilst the total charm cross-section (and hence A dependence) is relevant

here, it is in fact the x_F dependence which is crucial. This can be seen from the diverging slopes of the upper and lower bounds in Figure 5.8, corresponding to x_F parametrisations of $\sim (1 - |x_F|)^3$ and $\sim (1 - |x_F|)^5$ respectively. The discrepancy between different parametrisations becomes increasingly significant for large x_F (and hence large E_{mis}). This highlights the importance of a reliable x_F parametrisation since it is in the large E_{mis} tail that the WINP search will be made.

It was noted above that the most important factors in the choice of charm parametrisation for the Monte Carlo are x_F dependence and atomic number (A) dependence. Looking at the A dependence first, we can minimise the uncertainty in any A -dependent extrapolation by using experimental data taken with a heavy target. Tungsten is the closest element for which results have been published[85,86], having an atomic number $A = 183.85$ as compared to $A = 238.03$ for uranium. Even taking the extremal values of 0.66 and 1.00 for α , the resulting uncertainty in the extrapolated cross-section for uranium is only of order 10%. The corresponding uncertainty for extrapolation from an hydrogen target is of order six! Details of the E613 tungsten measurement were given in Section 5.4.1 above. The same experiment also provided the direct measurement of charm A dependence discussed in Section 5.4.1 and their quoted value of $\alpha = 0.75 \pm 0.05$ is used for the extrapolation to uranium.

Turning to the question of x_F dependence, E613 measures charm production predominantly at large x_F (recall it is a neutrino beam dump experiment). They find an x_F dependence of $(1 - x_F)^{3.2 \pm 0.2}$. Although most of the cross-section is at small x_F it is the large x_F region which is of particular interest in this analysis, corresponding to the missing energy tail of Figure 5.4.

Hence, the E613 results[85,86,89] appear to be a very satisfactory choice for the Monte Carlo simulation, providing data with a heavy target, giving a direct measurement of the A dependence, and covering the x_F region of interest. To summarise, the details of the chosen charm parametrisation are as follows. The total charmed pair production cross-section is $57.2 \pm 2.9 \pm 8.5 \mu\text{b/nucleon}$ which, taking an $A^{0.75}$ atomic number dependence, gives a charm cross-section of $3466.3 \pm 175.7 \pm 515.1 \mu\text{b}$ on uranium. This is roughly 1.8×10^{-3} of the total inelastic cross-section. The x_F and p_T dependence is given by the invariant parametrisation

$$E \frac{d^3\sigma}{dp^3} \sim (1 - x_F)^n e^{-bp_T},$$

where $n = 3.2 \pm 0.2$ and $b = 1.5 \pm 0.2 \text{ (GeV/c)}^{-1}$. No attempt has been made to correct for the \sqrt{s} dependence of the charm cross-section. It was noted above that this is expected to be at most a 20% correction for an increase in beam energy of 400 to 450 GeV.




Figure 5.9: Final simulation of the missing energy spectrum using E613 data for parametrisation of charm production. The hatched area represents a $\pm 1\sigma$ error band.

The final Monte Carlo prediction of the missing energy spectrum is shown in Figure 5.9. The hatched area at large E_{mis} is again a $\pm 1\sigma$ error band around the mean value. This Monte Carlo spectrum can now be compared with the *measured* missing energy spectrum of Figure 4.25, and any discrepancy interpreted as possible evidence for WINP production.

There is, however, one important cross-check which can now be made on the validity of the Monte Carlo prediction. It is a direct measurement of charm production in 450 GeV/c proton-uranium collisions using the missing energy spectrum of the muon sub-sample from the HELIOS beam dump data.

5.5 The single muon sample

The missing energy analysis of Chapter 4 gives a measurement of weakly interacting neutral particle production in proton-uranium interactions. As was noted above (Section 5.1 onwards), charm production, and its subsequent semi-leptonic decay to neutrinos, is expected to make a significant contribution to that measurement and is in fact the dominant ‘background’ at large missing energy.

The experimental configuration chosen for this WINP search is that of a calorimeter beam dump and downstream muon spectrometer. For charm production, this means that semi-leptonic decays in the electron channel cannot be tagged by observing the electron directly (since the electron showers and is absorbed in the calorimeter) whereas decays in the muon channel *may* be tagged using the spectrometer. In Chapter 4, events with detected muons were rejected in order to reduce the level of Standard Model background. This includes semi-leptonic decays of charmed particles. Rather than veto such events however, they can be used to provide a direct measurement of the background and hence to check the validity of the Monte Carlo simulation upon which the identification of any WINP signal depends. This section describes the single muon measurement and its significance for the background simulation.

5.5.1 Muon selection

The goal of this section is to produce a missing energy spectrum (analogous to Figure 4.25) for events with one or more associated muons. Assuming there are no new physics processes giving rise to weakly interacting neutral particles in conjunction with a muon, the spectrum can then be used to measure Standard Model $\nu_\mu\mu$ production, in particular $\nu_\mu\mu$ production from semi-leptonic charm decay.

The techniques required to measure missing energy, and the rejection of the corresponding backgrounds, are of course similar to those described in Chapter 4. In addition however, any produced muons must now be identified and measured. This means that any spectrometer cuts used to reject leakage must be modified to take account of the additional muon. Details of the muon spectrometer have been given in Section 3.1.5: it consists of a series of 7 wire chambers (allowing charged particle track measurement), a superconducting magnet (particle momentum information) and 2 scintillator hodoscope planes separated by a $4.8 \lambda_i$ iron wall (for muon/charged hadron separation). The geometrical acceptance for muons produced at the front face of the beam dump is roughly $\theta_\mu < 130$ mrad.

With the conclusions of the reconstruction procedure (Section 3.3) in mind, criteria for muon selection can now be defined.

The first requirement is that a good track be found in the spectrometer chambers, satisfying the conditions defined by MUREC (Section 3.3). In addition, MUREC demands that any track candidate should not pass through the magnet iron. A cut is then made on track quality, with a requirement that the chi-squared per degree of freedom (χ^2/NDF)

for the fit be less than 2.0. (Note that this cut is also imposed in the Monte Carlo so that any correction factor required is already incorporated in the reconstruction efficiency). For tracks surviving these cuts, the measured momentum is required to be greater than 5 GeV/c, corresponding to the combined effect of dE/dx losses and magnet momentum kick.

The second requirement for muon selection is identification in the hodoscope planes. The upstream plane H3 is somewhat redundant here since it does not distinguish muons from charged hadrons, acting essentially as an additional plane in the spectrometer wire chambers. Also, the average plane efficiency is only 70%. Demanding a hit in the relevant H3 slat would therefore require a large, and unnecessary, correction to the measured muon rate. Hence, no hit requirement is made on H3. The downstream plane H2 on the other hand, is crucial. It allows separation of muons and charged hadrons using the number of associated hits seen in the scintillator slats. Figure 4.13 (Chapter 4) shows the number of hits in H2 as a function of missing energy for events with one (and only one) reconstructed track. The tendency towards a double-peaked hit distribution at large E_{mis} was noted in Chapter 4 and believed to be due to a combination of muons and charged hadrons in the single-track sample. A cut is therefore made on the number of hits in H2, demanding a hit in the slat referred to by the track pointing, but zero hits elsewhere.

As in the case of neutral hadronic punchthrough (Section 4.3), it is important to determine if there is any residual punchthrough contamination once the H2 cut has been made. Using the distributions of Figure 4.13, contamination is estimated to be at most 10% for large E_{mis} ($E_{mis} > 60$ GeV). Possible contamination is most significant at moderate E_{mis} ($40 < E_{mis} < 60$ GeV), for which a (pessimistic) 50% systematic error is assigned to the data. This is not crucial however, since the purpose of measuring the single muon missing energy spectrum is to examine the charm contribution which only becomes the dominant source at large E_{mis} .

In conjunction with the selection criteria defined above, the spectrometer cuts described in Section 4.3 to minimise leakage have to be modified to take account of the muon track. The same cuts are used (number of chamber hits < 4 , number of hodoscope hits = 0) but are applied *after* the hits used in the reconstruction have been discounted.

5.5.2 Missing energy spectrum for the single muon sample

Figure 5.10 shows the missing energy spectrum for the single muon data sample. Corrections for reconstruction efficiency (16%) and hodoscope efficiency (32%) have been made




Figure 5.10: The missing energy spectrum $\frac{1}{N} \frac{dN}{dE_{mis}}$ for the single muon sample. Also shown is the Monte Carlo simulation where the hatched area represents the $\pm 1\sigma$ error band.

and systematic errors added for possible punchthrough contamination. Also shown is the Monte Carlo prediction using the charm parametrisation of the E613 experiment. The hatched area at large E_{mis} is a $\pm 1\sigma$ error band around the central value.

There is good agreement between data and Monte Carlo predictions except at small missing energies ($20 < E_{mis} < 60$ GeV), where the measured rate is roughly twice that expected from the simulation. In this region, the dominant contribution to the spectrum is from pion and kaon decays (see Figure 5.11) and the excess may be due to an incomplete understanding of π, K suppression in the beam dump. For the large E_{mis} tail however, the agreement is excellent and shows that the E613 charm parametrisation is a good approximation to the charm signal as measured by this experiment.

5.6 Comparison of the final missing energy spectrum with Monte Carlo

Figure 5.12 shows the final data (from Chapter 4) and Monte Carlo prediction for the missing energy spectrum in 450 GeV/c proton-uranium collisions. The agreement is excellent, spanning more than six orders of magnitude in cross-section and extending to missing




Figure 5.11: Simulation of the single muon missing energy spectrum showing contributions from π , K and charm decays.




Figure 5.12: The final missing energy spectrum $\frac{1}{N} \frac{dN}{dE_{mis}}$ and Monte Carlo prediction. The hatched area indicates the $\pm 1\sigma$ error band on the Monte Carlo as before.

energies of more than 200 GeV. The slight excess (a 1.9σ effect) of data over Monte Carlo at moderate E_{mis} may be due to non-Gaussian tails in the calorimeter response function or to imperfect simulation of the π, K suppression in the beam dump.

Conventional physics processes therefore appear able to account for the missing energy spectrum as measured by this experiment. It is now possible to use the results of Figure 5.12 to place upper limits on new sources of weakly interacting neutral particles. This is discussed in Chapter 6.

Chapter 6

Limits on the production of weakly interacting neutral particles

This chapter addresses the question of how to use the distribution of Figure 5.12 to place limits on new sources of neutrinos or neutrino-like particles. Some comments on the measured spectrum are given first, followed by a description of a general model for WINP particle production in hadronic collisions. Given this model, the expected contribution to the missing energy distribution is calculated and, by comparing with the observed spectrum, upper limits on WINP production obtained.

6.1 Comments on the measured missing energy spectrum

The good agreement between the measured spectrum and Monte Carlo prediction indicates that Standard Model processes are sufficient to account for the observed missing energy distribution. In fact, the dominant contributions to Figure 5.12 are expected to come from only two sources: minimum-bias events, convoluted with calorimeter fluctuations, at small E_{mis} , and events in which charm particles are produced and subsequently decay, at large E_{mis} .

The minimum-bias region of the spectrum is perhaps not of great interest, since it demonstrates only that the calorimeter response follows a Gaussian distribution (albeit to more than four standard deviations from the mean) as expected. The large- E_{mis} tail is more interesting however, particularly in view of the experimental uncertainties concerning




Figure 6.1: The single muon missing energy spectrum, with Monte Carlo predictions based on charm parametrisations from the EHS and CCFRS experiments. The upper and lower curves indicate a $\pm 1\sigma$ error band in each case.

charm hadroproduction (Section 5.4), since it may shed light on certain aspects of charm production in proton-nucleus collisions.

The charm parametrisation used in the Monte Carlo simulation of Chapter 5 was that given by the E613 collaboration[85,86] as a best fit to their experimental data. The experiment used a 400 GeV/c proton beam, incident on a heavy target, tungsten. Looking at the single muon E_{mis} spectrum of Figure 5.10, for which only charm production is assumed to contribute at large missing energy, the HELIOS measurement is in good agreement with the charm parametrisation as determined by E613. A similar comparison can be made of the HELIOS single muon E_{mis} spectrum with Monte Carlo predictions using the charm parametrisations determined by experiments with lighter targets. Predictions from two such experiments discussed in Section 5.4 are given in Figure 6.1, the first (dashed line) using a hydrogen target[82], the second (dotted line) an iron target[88]. The agreement here is considerably worse than for the E613 parametrisation.

The HELIOS measurement of the single muon missing energy spectrum is consequently seen to favour the parametrisation of charm production determined by the E613 collaboration. In particular (Section 5.4), this result provides information on the x_F dependence of charm production, indicating that diffractive meson production ($n \sim 3$) rather than

central production ($n \sim 5$) is the dominant mechanism in 450 GeV/c proton-uranium collisions. The good agreement between the two experiments also reinforces the E613 measurement on the A dependence of charm production, namely $\alpha \sim 0.75$, similar to the A dependence of the inelastic cross-section.

Discrepancies between this measurement and those of the EHS and CCFRS collaborations (amongst others) may well be due to a combination of A dependent effects and the different x_F regions covered by the respective experiments. Hopefully it is only a matter of time before new data (from the E769 collaboration for example) resolves the still open question of charm production in hadronic collisions.

6.2 Cross-section limits

Although Standard Model processes can account for the observed missing energy spectrum, the data may still be used to set limits on the production of weakly interacting neutral particles. The approach here is to rely on the fact that the number of events in the large- E_{mis} tail from WINP sources is *at most* equal to the observed data. Limits will be derived for two cases:-

- i). no charm background in the large- E_{mis} tail, and
- ii). charm background estimated according to E613 results.

Case i)., although apparently unrealistic, corresponds to taking the lowest charm level implied by Figure 6.1. Since the data points lie well above this level, it can in practice be set to zero. Note that this is in fact the true upper limit.

The technique used to derive the limits consequently has to take account of the presence of background (with errors) in the observed signal. Appendix E gives details of the algorithm used in this case.

6.2.1 Simple limits

Using the results of Figure 5.12, it is rather straightforward to obtain upper limits on the cross-section for weakly interacting particles with energies above a certain threshold. In order to avoid possible backgrounds from such sources as non-Gaussian calorimeter tails or residual π, K contributions, only E_{mis} bins well outside the 'resolution' region of the spectrum are used in the limit evaluation. A lower threshold of 80 GeV is used in this case. Limits are quoted for pp collisions, assuming a target atomic number dependence of $A^{\frac{2}{3}}$ for the extrapolation from pU to pp .

Charm background assumed	Energy threshold (GeV)		
	≥ 80	≥ 100	≥ 120
a). zero	4.5×10^{-5}	2.0×10^{-5}	1.3×10^{-5}
b). E613	1.4×10^{-5}	5.4×10^{-6}	4.2×10^{-6}

Table 6.1: Cross-section limits (in mb) for production of WINPs in pp collisions with energies above a certain threshold (GeV). Limits are given a). with no charm background, and b). with the charm level as measured by E613.

The 95% confidence level upper limits on the cross-section for WINPs having energies greater than a certain threshold are given in Table 6.1. Three energy thresholds (at 80, 100 and 120 GeV) have been used, and limits derived assuming either no charm background or charm levels as measured by the E613 experiment.

6.2.2 Model for WINP production

Prior to calculating limits on the *total* WINP cross-section, a model of weakly interacting neutral particle production is needed. Rather than construct a model for each of the proposed WINP sources discussed in Section 2.1, a single, rather general prescription is used, based on a phenomenological description of hadron production in pp collisions by M. Bourquin and J.-M. Gaillard[94]. In this scheme, a single expression for the invariant cross-section as a function of E_T, p_T and y is given, which is in good agreement with inclusive production distributions of non-leading particles over a wide range of energies $6 < \sqrt{s} < 63$ GeV. Cross-sections for WINP production can then be calculated for a range of interesting WINP masses.

The Bourquin-Gaillard parametrisation takes the basic form:-

$$E \frac{d^3\sigma}{d^3p} = \frac{A}{(E_T + B)^C} f(y) \begin{cases} e^{-p_T}, & \text{for } p_T < 1 \text{ GeV}/c, \\ e^{-D(p_T-1)/\sqrt{s}} e^{-1}, & \text{for } p_T > 1 \text{ GeV}/c, \end{cases}$$

where the transverse energy E_T is defined by $E_T = (m^2 + p_T^2)^{\frac{1}{2}}$ and y is the particle rapidity¹. The function $f(y)$ is expressed in terms of a variable $Y = y_{max} - y$:

$$Y = \ln \left(\frac{E^{max} + p_L^{max}}{E + p_L} \right)$$

and takes the explicit form:

$$f(Y) = e^{-\alpha/Y^\beta}.$$

¹The rapidity variable y is given by $y = \frac{1}{2} \ln \left(\frac{E+p_L}{E-p_L} \right) = \ln \left(\frac{E+p_L}{E_T} \right)$

Results of fits to experimental data give $B = 2$ GeV, $C = 12.3$, $D = 23$ GeV, $\alpha = 5.13$ and $\beta = 0.38$. Substituting these values, the cross-section formula can be rewritten as:-

$$E \frac{d^3\sigma}{d^3p} = A \left(\frac{2}{E_T + 2} \right)^{12.3} e^{-5.13/Y^{0.38}} \begin{cases} e^{-p_T}, & \text{for } p_T < 1 \text{ GeV}/c, \\ e^{-23(p_T-1)/\sqrt{s}} e^{-1}, & \text{for } p_T > 1 \text{ GeV}/c, \end{cases} \quad (6.1)$$

where p_T is in GeV/c, and E_T in GeV. A is a normalisation factor which, in the case of pion data, takes the value $3.78 \times 10^{-24} \text{ cm}^2 \text{ GeV}^{-2}c$.

Using this parametrisation, the cross-section for production of a particle of given mass can be evaluated as a function of y and p_T . One such example is shown in Figure 6.2, where the Bourquin-Gaillard formula has been used to generate a production distribution for (hypothetical) particles of mass 2.5 GeV. The corresponding projections, giving the invariant cross-section as a function of y and p_T , are also given. Note that the overall normalisation is not needed for the purposes of setting cross-section limits: it is the relative shape of the distribution which is important.

6.2.3 Total cross-section limits

Having determined the simple cross-section limits (Section 6.2.1) and chosen a suitable model for WINP production, upper limits on the *total* WINP cross-section may now be calculated. Assuming that the production of WINPs of a certain mass follows the Bourquin-Gaillard parametrisation, the fraction of WINPs having energies greater than a certain threshold can be evaluated. Once this fraction is known, it can be used to extrapolate from the cross-section limits above a certain energy threshold (those in Table 6.1) to the total cross-section limit. The procedure is outlined below.

First, the Bourquin-Gaillard formula is used to generate invariant cross-section distributions for inclusive WINP production over a range of interesting WINP masses. For this analysis, the mass range chosen is $0 < m_{\text{WINP}} < 5 \text{ GeV}/c^2$, reflecting the theoretical and experimental considerations of Sections 2.1, 2.2. For each mass bin, Equation 6.1 then gives the cross-section as a function of y and p_T in the range $0 < y < 8$, $0 < p_T < 4 \text{ GeV}/c$, yielding distributions of the type shown in Figure 6.2.

A large number, n , of WINP ‘events’ are then generated for each mass point, with each particle’s rapidity and transverse momentum being chosen at random within the limits defined above, and with the frequency of generation at any (y, p_T) coordinate being weighted by the corresponding Bourquin-Gaillard cross-section at that point. Since the particle’s mass, rapidity and transverse momentum are now all known, its energy is easily determined.

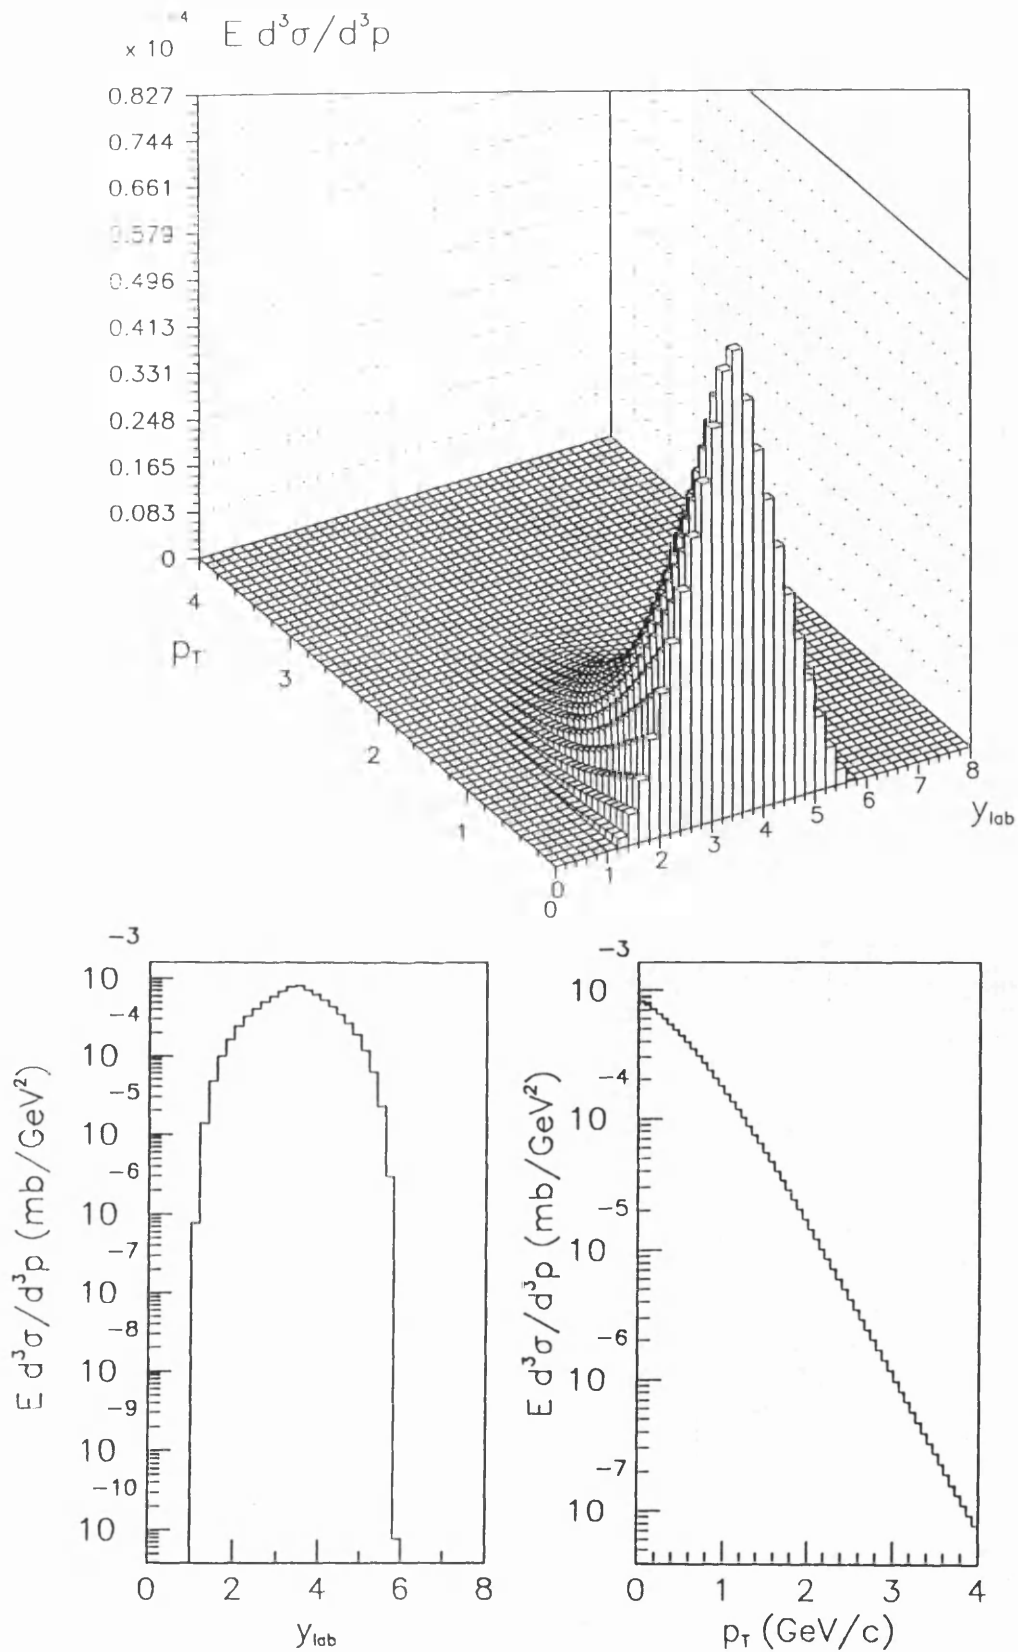


Figure 6.2: The invariant cross-section $E \frac{d^3\sigma}{d^3p}$ (mb/GeV²) as a function of y and p_T for a particle of mass 2.5 GeV. The corresponding projections are also shown.




Figure 6.3: The 95% confidence level limits on the cross-section (mb) for WIMP production in pp collisions as a function of mass (GeV/c^2) and lifetime (s).

At this point, the question of particle lifetime becomes important. So far no mention has been made of lifetime in the discussion on limits, but it is clearly relevant since only WIMPs which escape from the calorimeter before decaying will contribute to the measured missing energy spectrum. Following the conclusions of Section 2.3, the range of WIMP lifetimes for which this technique is expected to be sensitive is roughly $\tau_{\text{WIMP}} > 10^{-10}$ s. This limit is imposed by the requirement that the particle escape the calorimeter before decay. For simulation purposes, an effective upper limit of $\tau_{\text{WIMP}} < 10^{-7}$ s is also imposed: for lifetimes longer than this, the particle is effectively stable and the cross-section limit does not change.

Knowing the particle energy, and now its lifetime, the probability that it escapes from the calorimeter before decaying can be calculated. The total number of generated particles escaping from the calorimeter and having energies greater than the chosen E_{mis} threshold (see Section 6.2.1) is denoted $n_{E_{\text{mis}}}$. The total cross-section limit is finally obtained by taking the simple limit for the chosen E_{mis} threshold and scaling by the factor $\frac{n}{n_{E_{\text{mis}}}}$ for the given WIMP mass and lifetime.

Figure 6.3 shows the total WIMP cross-section limits as a function of mass and lifetime. Limits have been calculated at the 95% confidence level with, in this case, the charm

background as determined by the E613 charm parametrisation. The limits obtained lie roughly in the range 0.1 - 1 μb , depending of course on the WINP lifetime. Note that for $\tau_{\text{WINP}} < 5 \times 10^{-10}$ s, the limit tends to infinity as expected (Figure 2.8).

6.2.4 Comparison with previous experiments

This experiment was designed to search for evidence of weakly interacting neutral particle production in proton-nucleus interactions. With the aim of being sensitive to the widest possible range of WINP lifetime and interaction cross-section, the method of a calorimeter search was adopted (Section 2.3).

Previous experiments (see Section 2.2) have used the same technique, the best limit to date [46] coming from missing energy measurements in 400 GeV/c proton-iron interactions. In order to set general limits on WINP production, a cross-section dependence of the form

$$\frac{d^2\sigma}{dx_F dp_T^2} \sim (1 - |x_F|)^\alpha e^{-\beta p_T}$$

was assumed. Under the assumption of a linear A dependence for WINP production, cross-section limits in the range $\sim 1 - 10 \mu\text{b}$ were obtained for a range of WINP masses (1 MeV/c² to 2.5 GeV/c²) and variation in the parameters α (0 to 10) and β (1 to 5). The main constraint on this analysis was limited event statistics; the absence of a missing energy trigger allowing measurement of the missing energy spectrum over about 4 orders of magnitude only.

By contrast, the present experiment, with the ability to trigger on rare missing energy processes, provides a measurement over nearly 7 orders of magnitude. General limits on WINP production are therefore considerably improved, yielding values of order 0.1 - 1 μb for a similar range of WINP masses. The limitation in this case comes not from event statistics but from Standard Model backgrounds, in particular, from semi-leptonic decays of charmed particles.

To allow a *direct* comparison between results of the present experiment and previous searches, the missing energy spectrum taken from [46] has been used to set WINP cross-section limits assuming production according to the Bourquin-Gaillard parametrisation (see above). Figure 6.4 shows the results of this comparison. Limits are shown for 'stable' ($\tau > 10^{-7}$ s) WINPs only.

Hence this experiment allows limits to be set on weakly interacting neutral particle production which are at least an order of magnitude more stringent than existing limits obtained using calorimetric techniques.




Figure 6.4: Comparison of this experiment with previous calorimeter searches. The plot shows cross-section limits (mb) on stable WINP production in pp collisions from this experiment and those adapted from [46].

Chapter 7

Conclusions

This thesis has described a search for possible sources of weakly interacting neutral particle production in proton-nucleus collisions. At present, the only known particle having such properties is the neutrino, but there are theoretical arguments which admit the possibility of additional neutrino-like particles, both within the Standard Model (a light Higgs) and in extensions to the Model (a supersymmetric ‘neutralino’, and the axion). Despite the absence of direct evidence for any of these particles, there are indications that WINPs may play the rôle of the so-called dark matter in the Universe.

One manifestation of WINP production in an interaction would be the apparent non-conservation of energy in an idealised hermetic calorimeter. This method has the advantage of being sensitive to a wide range of possible WINP lifetimes and also imposes relatively weak constraints on the magnitude of the WINP coupling to ordinary matter. Adopting this approach, a calorimetry search has been performed using the HELIOS (NA34) detector to examine 450 GeV/c proton-uranium interactions. In order to maximise the experimental sensitivity to WINP production, a detailed study of the HELIOS calorimetry was undertaken. Excellent energy resolution (2.4 %) has been achieved by means of careful calibration of the individual calorimeter components and subsequent optimisation of the whole calorimeter ensemble.

The capability of the HELIOS calorimetry to trigger on rare events with large missing energies permitted a measurement of the proton-uranium missing energy spectrum over almost seven orders of magnitude. To achieve this result, a series of data cuts was applied, removing sources of ‘fake’ missing energy arising from instrumental effects in the calorimetry.

The expected contribution to the observed spectrum from known sources of weakly interacting neutral particles has been estimated using a detailed Monte Carlo simulation.

The dominant Standard Model ‘background’ sources are determined to be neutrinos from semi-leptonic decays of charged pions and kaons and, in particular, of charmed particles. The crucial question of a reliable estimate for the charm background is resolved with the help of a direct measurement of the single muon missing energy spectrum, taken under identical experimental conditions.

Comparing the observed and predicted missing energy spectra, no significant excess of WINPs is observed: the HELIOS measurement is in good agreement with Standard Model expectations.

Based on this result, it is possible to set upper limits on WINP production cross-sections in proton-nucleus interactions. Using a general model for WINP production, upper limits of between ~ 0.1 and $1 \mu\text{b}$ (at the 95% c.l.) are derived for production of WINPs with masses between $5 \text{ GeV}/c^2$ and $100 \text{ MeV}/c^2$ (respectively) in proton-proton collisions. These limits are at least one order of magnitude more stringent than previous limits obtained using calorimeter searches.

Appendix A

Resolution minimisation

This appendix demonstrates the equivalence of the principles of resolution minimisation and Etot-flatness for a given calorimeter ensemble. A brief mathematical proof is shown followed by empirical results from the HELIOS calorimetry.

A.1 Equivalence of resolution minimisation technique and Etot-flatness

Consider an ensemble of n calorimeters. The total energy E_{tot} measured is formed by a linear sum of the signals from the individual calorimeters:-

$$E_{tot} = \sum_i \alpha_i E_i$$

where α_i are weight factors applied to the $i = 1, 2, \dots, n$ calorimeters. In general, one may assume that E_{tot} is dependent on the energy deposited in one or more of the component calorimeters (i.e. not all calorimeters are Etot-flat). To first order, this dependence is linear, and one may write for *each* calorimeter i :-

$$y_j^{(i)} = a_{(i)} x_j^{(i)} + b_{(i)}$$

where $x_j^{(i)} \equiv \alpha_i E_i^{(j)}$, $y_j^{(i)} \equiv E_{tot}^{(j)}$ and $j = 1, 2, \dots, N$ events. The parameters a and b are the slope and intercept respectively. Finding the best fit for this line gives a and b :-

$$a = \frac{N \sum_j x_j y_j - \sum_j x_j \sum_j y_j}{N \sum_j x_j^2 - \left(\sum_j x_j\right)^2}$$

and

$$b = \frac{\sum_j x_j^2 \sum_j y_j - \sum_j x_j \sum_j x_j y_j}{N \sum_j x_j^2 - \left(\sum_j x_j\right)^2}$$

A.1.1 Etot-flatness

Etot-flatness is achieved when the coefficients $a \rightarrow 0$. The coefficients b are not constrained a priori; the obvious choice (on physical rather than mathematical grounds) is $b \rightarrow E_b$ where E_b is the beam energy. Hence,

$$\frac{1}{N} \sum_j y_j = E_b$$

and

$$\sum_j x_j y_j = E_b \sum_j x_j$$

Substituting x_j and y_j for calorimeter i :-

$$\sum_j \alpha_i E_i^{(j)} \left(\sum_k \alpha_k E_k^{(j)} \right) = E_b \sum_j \alpha_i E_i^{(j)}$$

Hence,

$$\sum_k \alpha_k \sum_j E_i^{(j)} E_k^{(j)} = E_b \sum_j E_i^{(j)} \quad (\text{A.1})$$

A.1.2 Resolution minimisation

Minimising the total energy resolution requires:-

$$\frac{\partial \chi^2}{\partial \alpha_i} = 0$$

where

$$\chi^2 = \sum_j \left(\sum_k \alpha_k E_k^{(j)} - E_b \right)^2$$

$$\frac{\partial \chi^2}{\partial \alpha_i} = 2 \sum_j \left(\sum_k \alpha_k E_k^{(j)} - E_b \right) E_i^{(j)}$$

Hence,

$$\sum_k \alpha_k \sum_j E_i^{(j)} E_k^{(j)} = E_b \sum_j E_i^{(j)} \quad (\text{A.2})$$

As can be seen from Equations A.1 and A.2, the weights α_i which define Etot-flatness also determine the optimal energy resolution.

Calorimeter grouping	BEFORE WEIGHTING		AFTER WEIGHTING		
	Slope a	Intercept b	Weight	Slope a	Intercept b
ULACEM	-0.096 ± 0.037	458.25 ± 0.39	0.95	0.024 ± 0.037	447.05 ± 0.39
ULACHAD	0.182 ± 0.015	394.29 ± 1.76	0.99	0.045 ± 0.015	436.44 ± 1.76
BEAM	0.239 ± 0.058	444.71 ± 0.30	1.11	0.027 ± 0.059	448.88 ± 0.30
VETO	1.603 ± 0.184	448.08 ± 0.27	-0.11	-0.514 ± 0.440	449.32 ± 0.27
BOX	-0.623 ± 0.142	450.87 ± 0.30	0.08	3.161 ± 0.638	448.91 ± 0.30
W+M+E	-0.450 ± 0.055	462.68 ± 0.34	1.96	0.025 ± 0.046	448.25 ± 0.34
Energy resolution	Mean = 449.04 ± 1.44 GeV Sigma = 20.73 GeV		Mean = 449.37 ± 1.17 GeV Sigma = 16.80 GeV		

Calorimeter grouping	BEFORE WEIGHTING		AFTER WEIGHTING		
	Slope a	Intercept b	Weight	Slope a	Intercept b
ULACEM	-0.095 ± 0.037	458.20 ± 0.39	0.95	0.025 ± 0.037	446.91 ± 0.39
ULACHAD	0.181 ± 0.015	394.50 ± 1.76	0.99	0.047 ± 0.015	435.77 ± 1.76
B+V	0.235 ± 0.057	444.67 ± 0.30	1.07	0.029 ± 0.059	448.83 ± 0.30
B+W+M+E	-0.401 ± 0.053	462.40 ± 0.34	1.82	0.027 ± 0.046	448.13 ± 0.34
Energy resolution	Mean = 449.06 ± 1.43 GeV Sigma = 20.59 GeV		Mean = 449.34 ± 1.19 GeV Sigma = 17.22 GeV		

Table A.1: Results of resolution minimisation on ‘minimum-bias’ target interactions (tungsten target). Flash ADC results are shown for two different sub-divisions of the calorimetry. (Figures for the energy resolution before weighting have been multiplied by an overall scale factor for comparison).

A.2 Results of resolution minimisation for the HELIOS calorimetry

The minimisation technique described above was used to obtain weights for the HELIOS calorimetry. ‘Minimum bias’ interactions in the target configuration were studied first. This sample deposits, on average, energy in *all* the component calorimeters and is therefore appropriate for determining weights for the ensemble. Table A.1 shows results from the FADC readouts corresponding to two different sub-divisions of the calorimetry (results for the energy resolution before weighting have been re-scaled to bring the total energy to 450 GeV). Three conclusions can be drawn. Firstly, that too fine a sub-division of the calorimetry gives unreliable results for calorimeters with very little energy. This can be seen, for example, with the VETO and BOX calorimeters in the upper table. Having studied results from different groupings, the eventual choice was to weight according to the lower table in Table A.1. Secondly, the technique does produce Etot-flatness. The slope parameters (a) after weighting are approximately zero. Thirdly, the same weights do improve the energy resolution.

Having ascertained the correct weights for interactions in the wire target, the same method was applied to interactions in beam dump mode. This essentially samples only the hadronic section of the ULAC and the BEAM calorimeters. The weights were compatible with those found above.

Appendix B

The muon chamber efficiency

The muon detection efficiency is dependent on the performance of the spectrometer wire chambers. Hence it is important to know the individual plane efficiencies and the resulting deterioration in track-finding capability.

The method used for the chambers is simple. A track sample is defined as follows: for planes upstream of the magnet, completely reconstructed tracks are used while for downstream planes, the sample comprises tracks reconstructed behind the magnet in either the X or Y projection. Then, for n_t tracks passing through the given plane, the number n_h having an associated hit is determined, giving the plane efficiency $\frac{n_h}{n_t}$. The results are shown in Table B.1.

Plane number	Chamber number	Plane within chamber	Efficiency
1	0	0	0.958 ± 0.004
2	0	1	0.750 ± 0.008
3	0	2	0.970 ± 0.003
4	1	0	0.936 ± 0.005
5	1	1	0.672 ± 0.009
6	1	2	0.907 ± 0.005
7	2	0	0.931 ± 0.005
8	2	1	0.960 ± 0.004
9	2	2	0.931 ± 0.005
10	2	3	0.876 ± 0.006
11	2	4	0.585 ± 0.009
12	2	5	0.721 ± 0.008
13	3	0	0.938 ± 0.004
14	3	1	0.893 ± 0.005
15	3	2	0.806 ± 0.007
16	3	3	0.949 ± 0.004
17	3	4	0.888 ± 0.005
18	3	5	0.969 ± 0.003
19	4	0	0.444 ± 0.008
20	4	1	0.846 ± 0.006
21	4	2	0.946 ± 0.004
22	4	3	0.725 ± 0.008
23	4	4	0.281 ± 0.008
24	4	5	0.488 ± 0.008
25	5	0	0.934 ± 0.004
26	5	1	0.919 ± 0.005
27	5	2	0.914 ± 0.005
28	5	3	0.883 ± 0.005
29	6	0	0.848 ± 0.006
30	6	1	0.911 ± 0.005
31	6	2	0.800 ± 0.007
32	6	3	0.898 ± 0.005

Table B.1: Spectrometer chamber efficiencies.

Appendix C

The muon hodoscope efficiency

For an accurate measurement of muon rates, it is clearly essential to have a good knowledge of the efficiency for muon detection, in particular the efficiency of the hodoscopes. These scintillator planes, recovered from the NA3 experiment, have had substantial radiation exposure and their performance may consequently have deteriorated. The missing energy data used in this analysis provides a subset of events containing energetic muons, obtained without triggering on the muon spectrometer. This unbiased subsample was used to quantify the hodoscope efficiencies; the technique is described briefly below.

Tracks in the spectrometer chambers are first reconstructed using the standard MUREC code (see Section 3.3.2). The efficiency of a particular hodoscope slat is then calculated by simply looking at the number of times that slat fires when a chamber track points to it. The efficiency ϵ for any H3 or H2 slat is then:-

$$\epsilon(\text{H3}) = \frac{\text{number}(\text{track} \cdot \text{H3})}{\text{number}(\text{track} \cdot \text{H3}) + \text{number}(\text{track} \cdot \overline{\text{H3}})}$$

and

$$\epsilon(\text{H2}) = \frac{\text{number}(\text{track} \cdot \text{H2})}{\text{number}(\text{track} \cdot \text{H2}) + \text{number}(\text{track} \cdot \overline{\text{H2}})}$$

respectively.

For the plane upstream of the iron wall (H3) this technique gives the slat efficiencies directly: they are shown in Table C.1. The overall plane efficiency is 70%. For some slats, there is a noticeable dependence of efficiency on the x co-ordinate within the slat. This is due to attenuation as light travels (up to 3 m) along the scintillator to the photomultiplier readout on the outer edge of the hodoscope plane. Unfortunately, the slats are most inefficient near the centre of the hodoscope plane where the hit frequency is greatest.




Figure C.1: Schematic of the hodoscopes and iron wall.

For the downstream scintillator plane (H2) the calculation of slat efficiency is complicated by the presence of the $\sim 5 \lambda_i$ iron wall. The track pointing method is still valid so long as high-momentum tracks ($p_{tr} > 50 \text{ GeV}/c$) are used, allowing separation of muons and charged hadrons. Further details are given in Section 3.3.3. The slat efficiencies obtained for H2 are shown in Table C.2. The average plane efficiency in this case is 76%.

Slat number	Efficiency	Slat number	Efficiency
1	-	1	-
2	no hits	2	no hits
3	"	3	"
4	"	4	"
5	"	5	"
6	"	6	"
7	"	7	"
8	"	8	"
9	"	9	"
10	0.400 ± 0.219	10	1.000 ± 0.000
11	0.875 ± 0.068	11	1.000 ± 0.000
12	0.513 ± 0.036	12	0.761 ± 0.028
13	0.788 ± 0.071	13	0.805 ± 0.062
14	1.000 ± 0.000	14	0.556 ± 0.166
15	no hits	15	no hits
16	"	16	"
17	"	17	"
18	"	18	"
19	"	19	"
20	"	20	"
21	"	21	"
22	"	22	"
23	-	23	-

Table C.1: H3 hodoscope plane efficiencies.

Slat number	Efficiency	Slat number	Efficiency
1	-	1	-
2	no hits	2	no hits
3	"	3	"
4	"	4	"
5	"	5	"
6	"	6	"
7	"	7	"
8	"	8	"
9	"	9	"
10	1.000 ± 0.000	10	0.500 ± 0.354
11	0.742 ± 0.079	11	0.667 ± 0.073
12	0.868 ± 0.025	12	0.679 ± 0.031
13	0.829 ± 0.064	13	0.854 ± 0.051
14	no hits	14	0.600 ± 0.219
15	"	15	no hits
16	"	16	"
17	"	17	"
18	"	18	"
19	"	19	"
20	"	20	"
21	"	21	"
22	"	22	"
23	-	23	-

Table C.2: H2 hodoscope plane efficiencies.

Appendix D

Monte Carlo comparisons with data

The following plots (Figures D.1 - D.5) show several comparisons made between the LUND event generator and published experimental data. Inclusive cross-section results are given for a range of hadron-hadron interactions and compared with predictions based on both FRITIOF and TWISTER. Further details of the Monte Carlo simulation can be found in Section 5.2.




Figure D.1: Differential cross-sections $\frac{d\sigma}{dx_F}$, $\frac{d\sigma}{dp_T^2}$ for π^0 production in pp collisions. Data points are taken from [95].




Figure D.2: Differential cross-sections $\frac{d\sigma}{dx_F}$, $\frac{d\sigma}{dp_T^2}$ for π^- production in pp collisions. Data points are taken from [96].




Figure D.3: Inclusive K_S^0 production in pN collisions. Data points are from [97].




Figure D.4: Differential cross-section $\frac{d\sigma}{dx_F}$ for ρ^0, ω^0 production in K^+p collisions. Data points are taken from [98].

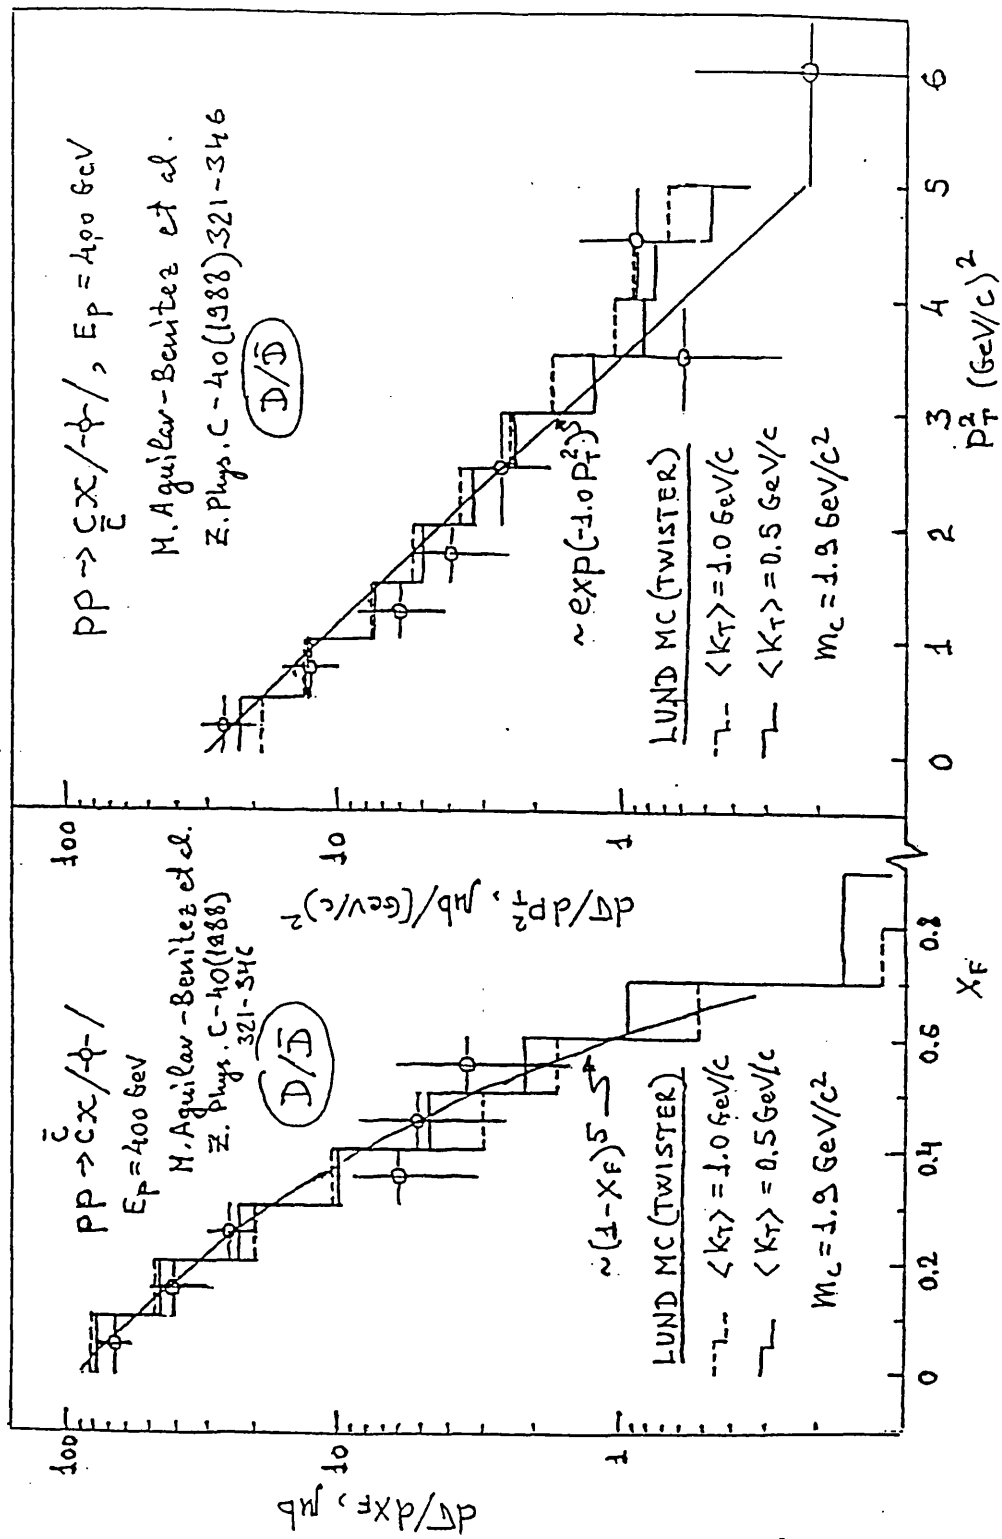


Figure D.5: Differential cross-sections $\frac{d\sigma}{dx_F}$, $\frac{d\sigma}{dp_T^2}$ for charm production in pp collisions. Data points are from the EHS collaboration[82].

Appendix E

Algorithm used for deriving upper limits

The algorithm used to evaluate upper limits for WINP production must take account of the presence of background contributions (charm) to the observed missing energy spectrum. In the case where the background contribution is of comparable size to the signal, there is the additional problem of statistical fluctuations associated with background subtraction. This can lead to unphysical solutions: most obviously, negative upper limits on new sources of weakly interacting neutral particles!

Let the number of observed events be n_0 , where n_0 is the sum of the contributions from background(b) plus any signal(s). The number of each type is not known but it is certain that the number of background events is at most n_0 . The probability that the total (signal plus background) exceeds n_0 with the added requirement that the background be less than or equal to n_0 is then:-

$$P((s + b > n_0) + (b \leq n_0)) = p - q \quad (\text{E.1})$$

where $p = P(b \leq n_0)$ and $q = P(s + b \leq n_0)$. Since $q \geq 0$, the probability given by Equation E.1 cannot exceed p . 'Rescaling' it to lie between 0 and 1 gives:-

$$\text{Rescaled prob} = \frac{p - q}{p} = 1 - \frac{q}{p} \quad (\text{E.2})$$

where it is this 'rescaled' probability which is the appropriate quantity to use when considering confidence levels between 0 and 1.

If signal and background are Poisson processes, then p and q can be replaced by the appropriate Poisson distributions in Equation E.2, yielding:-

$$\text{CL} = 1 - \frac{e^{-(\mu_B + N)} \sum_{n=0}^{n_0} (\mu_B + N)^n / n!}{e^{-\mu_B} \sum_{n=0}^{n_0} \mu_B^n / n!} \quad (\text{E.3})$$

where n_0 is the number of observed events, μ_B the mean background level, CL is the desired confidence level ($0 \leq \text{CL} \leq 1$) and N is the required upper limit on the candidate signal.

This corresponds to the case where the mean background contribution μ_B is known exactly. If μ_B is not known exactly, but has a probability density function p.d.f. = $f(B)$, then Equation E.3 generalises to:-

$$\text{CL} = 1 - \frac{\int f(B) e^{-(\mu_B + N)} \sum_{n=0}^{n_0} (\mu_B + N)^n / n! dB}{\int f(B) e^{-\mu_B} \sum_{n=0}^{n_0} \mu_B^n / n! dB}. \quad (\text{E.4})$$

Typically, as in the case of the charm background, $f(B)$ is a Gaussian of mean μ_B and variance σ_B^2 . In the limit $\mu_B + \mu_S \gg 1$, it is often more convenient to approximate *all* Poisson probability density functions as Gaussians. This approximation is valid for the limits calculated in Section 6.2.1 where there are typically several tens of events above the chosen missing energy thresholds.

Bibliography

- [1] A. Yahil, D. Walker and M. Rowan-Robinson, *Ap. J.* L1 (1986) 301.
- [2] E. D. Loh and E. J. Spillar, *Ap. J.* L1 (1986) 307.
- [3] A. M. Boesgard and G. Steigman, *Ann. Rev. Astr. Astrophys.* 23 (1985) 319.
- [4] J. Silk, K. A. Olive and M. Srednicki, *Phys. Rev. Lett.* 55 (1985) 257.
- [5] E. Fernandez, Talk presented at Neutrino '90, CERN, 1990.
- [6] A. H. Guth and E. J. Weinberg, *Phys. Rev. Lett.* 45 (1980) 1131.
- [7] M. Sher, *Phys. Rev. D*22 (1980) 2989.
- [8] H. E. Haber and G. L. Kane, *Phys. Rep.* 117, Nos.2-4 (1985) 75-263.
- [9] I. Hinchliffe, *Ann. Rev. Nucl. Part. Sci.* 36, 1986.
- [10] D. Nitz et al., Univ. Michigan preprint UMHE 86-16 (1986).
- [11] M. Z. Akrawy et al., CERN EP 89-176 (1989).
- [12] B. Adeva et al., CERN L3 002 (1989).
- [13] C. Albajar et al., *Phys. Lett.* B198 (1987) 261.
- [14] R. Ansari et al., *Phys. Lett.* B195 (1987) 613.
- [15] J. Alitti et al., CERN EP 89-151 (1989).
- [16] G. R. Farrar and P. Fayet, *Phys. Lett.* 76B (1978) 575.
- [17] P. R. Harrison and C. H. Llewellyn-Smith, *Nucl. Phys.* B213 (1983) 223.
- [18] R. D. Peccei and H. R. Quinn, *Phys. Rev. Lett.* 38 (1977) 1440.
- [19] S. Weinberg, *Phys. Rev. Lett.* 40 (1978) 223.
- [20] F. Wilczek, *Phys. Rev. Lett.* 40 (1978) 279.

- [21] S. Weinberg, Phys. Rev. Lett. 40 (1978) 223.
- [22] W. A. Bardeen and S.-H. Tye, Fermilab report no. 77/110-THY (unpublished).
- [23] C. M. Hoffman, Phys. Lett. 208B (1988) 149.
- [24] J. Heintze et al., Nucl. Phys. B149 (1979) 365.
- [25] C. Y. Pang et al., Phys. Rev. D8 (1973) 1989.
- [26] Y. Asano et al., Phys. Lett. 107B (1981) 159.
- [27] Barabanov et al., JETPL (Sov. Phys. ZETF Lett.) 32 (1980) 359.
- [28] R. Barbieri and T. E. O. Erikson, Phys. Lett. 57B (1975) 270.
- [29] J. Ellis, M. K. Gaillard and D. V. Nanopoulos, Nucl. Phys. B106 (1976) 292.
- [30] S. J. Freedman et al., Phys. Rev. Lett. 52 (1984) 240.
- [31] CUSB Collaboration, reported by J. Lee-Franzini, Proceedings of the 1987 Intl. Symposium on Lepton and Photon Interactions, Hamburg, 1987.
- [32] S. Raby, G. B. West and C. M. Hoffman, LA-UR 88-2667, Los Alamos National Laboratory, 1988.
- [33] P. Yepes, CERN EP 89-68 (1989).
- [34] M. Z. Akrawy et al., CERN EP 89-174.
- [35] D. Decamp et al., CERN EP 89-157.
- [36] B. Alper et al., Phys. Lett. 46B (1973) 265.
- [37] D. Cutts et al., Phys. Rev. Lett., 41 (1978) 363.
- [38] H. R. Gustafson et al., Phys. Rev. Lett. 37 (1976) 474.
- [39] H. Albrecht et al., Phys. Lett. 167B (1986) 360.
- [40] P. M. Tuts et al., Phys. Lett. 186B (1987) 233.
- [41] F. Bergsma et al., Phys. Lett. 121B (1983) 429.
- [42] A. Cooper-Sarkar et al., Phys. Lett. 160B (1985) 212.
- [43] J. Badier et al., Z. Phys. C31 (1986) 21.
- [44] G. Altarelli, B. Mele and S. Petrarca, Phys. Lett. 160B (1985) 317.
- [45] A. de Rujula and R. Petronzio, Nucl. Phys. B261 (1985) 587.

- [46] B. C. Barish et al., Phys. Lett. 85B (1979) 142.
- [47] R. Arnold et al., Phys. Lett. 186B (1987) 435.
- [48] T. W. Donnelly et al., Phys. Rev. D18 (1978) 1607.
- [49] S. Barshay et al., Phys. Rev. Lett., 46 (1981) 1361.
- [50] A. Barroso et al., Phys. Lett. 106B (1981) 91.
- [51] H. Faissner et al., Phys. Lett. 103B (1981) 234.
- [52] T. Cowan et al., Phys. Rev. Lett. 54 (1985) 1761.
- [53] T. Cowan et al., Phys. Rev. Lett. 56 (1986) 444.
- [54] M. R. Adams et al., FERMILAB-Pub-86/103-E.
- [55] A. Konaka et al., KEK preprint 86-9 (1986).
- [56] J. Bjorken, Proceedings of 1974 Rencontres de Moriond.
- [57] A. Rothenburg, SLAC report 147 (1972).
- [58] W. Cleland, HELIOS trigger electronics, Univ. of Pittsburgh (1987).
- [59] R. H. Beuttenmuller et al., Nucl. Instr. Meth. A252 (1986) 471.
- [60] R. H. Beuttenmuller et al., Nucl. Instr. Meth. A253 (1987) 500.
- [61] R. Wigmans, Nucl. Instr. Meth. A259 (1987) 389.
- [62] M. Murray, HELIOS note 174 (1987).
- [63] G. Delavallade, 'Type 347 ADC', CERN EP.
- [64] J. P. Bard, V. Cromelynk, G. Francinet, HELIOS note 96 (1986).
- [65] A. Lounis, HELIOS note 187 (1987).
- [66] T. Åkesson et al., Nucl. Instr. Meth. A262 (1987) 243.
- [67] T. Åkesson et al., Nucl. Instr. Meth. A241 (1985) 17.
- [68] H. Abramowicz et al., Nucl. Instr. Meth. 180 (1981) 429.
- [69] U. Amaldi, Physica Scripta 23 (1981) 409.
- [70] M. de Vicenzi et al., Nucl. Instr. Meth. A243 (1986) 348.
- [71] W. Braunschweig et al., DESY 87-098.

- [72] F. Gibrat and J. P. Pansart, HELIOS note 95 (1986).
- [73] A. Bodek, Proceedings of the Workshop on Muon Identification and Momentum Measurement at SSC and LHC Energies, University of Wisconsin, (1985).
- [74] S. P. K. Tavernier, Rep. Prog. Phys. 50 (1987) 1439-1489.
- [75] M. Aguilar-Benitez et al., 'Review of Particle Properties', Phys. Lett. B B204 (1988) 1.
- [76] M. J. Shochet, SLAC report 204 (1977).
- [77] J. Rich and D. R. Winn, Phys. Rev., D14, (1976) 1283.
- [78] E. J. Siskind et al., Phys. Rev., D21, (1980) 628.
- [79] H. Bengtsson and T. Sjostrand, LU TP 87-3, UCLA-87-001.
- [80] B. Nilsson-Almqvist and E. Stenlund, LU TP 87-6.
- [81] G. Ingelman, DESY 86-131.
- [82] M. Aguilar-Benitez et al., Z. Phys. C40 (1988) 321-346.
- [83] W. R. Nelson, H. Hirayama and D. W. O. Rogers, SLAC report 265 (1985).
- [84] A. S. Carroll et al., Phys. Lett. 80B (1979) 319-22.
- [85] R. C. Ball et al., Phys. Rev. Lett. 51 (1983) 743-6.
- [86] M. E. Duffy et al., Phys. Rev. Lett. 57 (1986) 1522-5.
- [87] J. Dorenbosch et al., Proc. 12th Intl. Conf. on Neutrino Phys. and Astrophys., (1986), pp. 402-12.
- [88] J. L. Ritchie et al., Phys. Lett. 126B (1983) 499-505.
- [89] M. E. Duffy et al., Phys. Rev. Lett. 55 (1985) 1816-9.
- [90] D. S. Barton et al., Phys. Rev. D27 (1983) 2580-99.
- [91] A. T. Goshaw, Proceedings of the 1987 NATO Advanced research workshop on QCD hard hadronic processes.
- [92] J. F. Gunion, Phys. Lett. 88B (1979) 150.
- [93] M. Aguilar-Benitez et al., Phys. Lett. 123B (1983) 98-102.
- [94] M. Bourquin and J.-M. Gaillard, Nucl. Phys. B114 (1976) 334-364.
- [95] I. V. Ajinenko et al., Z. Phys. C35 (1987) 7-14.

[96] M. Adamus et al., *Z. Phys.* C39 (1988) 311.

[97] P. Skubic et al., *Phys. Rev.* D18 (1978) 3115.

[98] M. Adamus et al., *Phys. Lett.* B198 (1987) 292.