

INFORMATION TO USERS

This manuscript has been reproduced from the microfilm master. UMI films the text directly from the original or copy submitted. Thus, some thesis and dissertation copies are in typewriter face, while others may be from any type of computer printer.

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleedthrough, substandard margins, and improper alignment can adversely affect reproduction.

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

Oversize materials (e.g., maps, drawings, charts) are reproduced by sectioning the original, beginning at the upper left-hand corner and continuing from left to right in equal sections with small overlaps. Each original is also photographed in one exposure and is included in reduced form at the back of the book.

Photographs included in the original manuscript have been reproduced xerographically in this copy. Higher quality 6" x 9" black and white photographic prints are available for any photographs or illustrations appearing in this copy for an additional charge. Contact UMI directly to order.

U·M·I

University Microfilms International
A Bell & Howell Information Company
300 North Zeeb Road, Ann Arbor, MI 48106-1346 USA
313/761-4700 800/521-0600

Order Number 9215003

Study of upward-going muons with the IMB detector

Becker-Szendy, Ralph Attila, Ph.D.

University of Hawaii, 1991

U·M·I
300 N. Zeeb Rd.
Ann Arbor, MI 48106



Study of Upward-going Muons with the IMB detector

A dissertation submitted to the Graduate Division of the
University of Hawaii in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

in

Physics

December 1991

By

Ralph A. Becker-Szendy

Dissertation Committee:

John C. Learned, Chairperson

Pui K. Lam

Sandip Pakvasa

David E. Yount

R. Brent Tully

© 1991 by Ralph A. Becker-Szendy
All Rights Reserved

Acknowledgements

A particle physics experiment is not a large and complicated piece of equipment, it is a group of people. They have the ideas, build the hardware, use its data, and do physics with it. Behind every cable, water pump, circuit board, unistrut and subroutine there is a person, to which the tangible installation is secondary. The fact that the IMB laboratory exists and has been and will be successfully used to do particle physics and astrophysics is due a dedicated group.

The IMB collaboration unifies physicists, engineers and technicians from very varied backgrounds, with very different styles of doing physics, having many visions of what they shall accomplish and how it can be done. In spite of this, there is a common spirit of making this undertaking work, doing so in a manner pleasant to all involved, and producing good physics. We have worked together for many years (it was started nearly twelve years ago), and not only do we still talk to each other, we usually consider our collaborators to be friends. So foremost I have to be grateful to the members of the IMB group who made it all possible:

John Learned has been with the IMB experiment from its beginning, and his contributions are visible nearly everywhere in the experiment. He is a wonderful friend, an advisor in the literal sense of the word, and a physicist full of insights and intuition. He moved me into the direction of particle astrophysics when I knew hardly anything about the field, taught me most of what I know about it now (plus a lot about other things), and in the process remained an example of what a real hands-on physicist should be like. Bob Svoboda is a product of John's education, and his thesis is the direct predecessor of this work. He qualifies for being honorary co-advisor. Steve Dye taught me a lot of common sense (inside and outside physics), got me started on IMB data analysis work, and is always willing and able to help. Dave Casper is the most amazing graduate student I have met. He is as productive as two post-docs, and very careful in his work. He is the ultimate authority on the IMB detector, its operation and the analysis of its data. His thesis is an example of how well physics can be done, and this work fades

in comparison. Shige Matsuno labors in the quiet, resourceful, efficient and reliable fashion which is unique to him; he not only performed the ungrateful task of keeping the detector operating, he deserves high praise for making it run so well, and pushing both detector and collaboration towards new physics opportunities. Woiciech Gajewski combines the good features of Dave and Shige: he works away in the shadow, producing reliable results, and in his spare time answers any question authoritatively. Plus, he has the correct attitude about life, the universe, and all the rest: It sucks.

Fred Reines is the perfect head for this unruly collaboration: he is wise, gentle, respected by all, always there to listen and give good advice, and speaks up to settle problems and disputes without ever using brute force. If only all physics leaders were like him. Bill Kropp not only administrates the experiment, he also cares for it, both personally and by helping to keep it going. Bud Bratton shares Bill's duties and involvement; plus by being the physicist located closest to the experiment he has to deal with many real-life problems. Hank Sobel gets an unusual laudation: he is a master of operating the jackhammer. Which shows the unusual amount of energy he puts into operating this experiment. Maurice Goldhaber is a source of good judgement, as shown by the insights he gives when attending collaboration meetings. Coleen Cory provides that bit of down-to-earth-ness which prevents the collaboration from becoming airborne. She manages the data analysis in Irvine in spite of the hectic life of her own, and in spite of everybody else trying to make her job harder. Without her there would be no law-and-order in Irvine. Clark McGrew is always a source of good ideas and interesting discussions. Cheryl Clem, Nicole Beach and in particular Gary McGrath scanned large fraction of the data used in this work; Gary is not only a cool dude, but also the fastest scanner in the west. Joe Reese is the person who is actually most crucial to keeping the detector running and in one piece. He is rarely thanked, thought of, or mentioned by his collaborators. But as usual he got just about the last place in this list, which shows that the world is just not fair. Robert Render is the other half of the technical team in the mine (yes, we have forgiven you for going off on other ventures for a while, after all you came back). Last but not least, there are generations of graduate students and post-docs who actually built the IMB experiment; some of them I have never even met, but heard everything

about them. I can not do their labor justice here, so I have to refer to the works of my predecessors for that.

The DUMAND collaboration has been my second “home” (and sometimes my first one) during the last five years. They are the most crazy and unorthodox, but also smartest and most ingenious collaboration I have ever seen or heard of. They are the kind of people who think they can pull off any feat, and usually they are right. Many of them are good friends and models for me, and have helped me during various stages of my career in Hawaii, in particular John Clem, Dave Harris, Shige Matsuno, Bob Mitiguy, Dan O’Connor, Art and Janice Roberts, and Mark Rosen. Above all Hugh Bradner has to be mentioned: he should have been on my dissertation committee, but that was prevented by bureaucratic requirements; therefore I would like to make him a honorary member.

All the positive qualities listed above are even more obvious in comparison to another physics administrator, who tried to prevent this work from being done; but since he failed at that, he at least tried to make life miserable for all involved. But lets forget the bleak and get back to the positive side.

Chuck Hayes not only made it possible for me to come to the University of Hawaii, he also helped me stay here and do my work, in spite of problems which sometimes seemed unsurmountable to me. But to Chuck there are no unsurmountable problems, and he can make all the friction within the department vanish magically. His only fault is not being the chairman any longer. Pui Lam, Sandip Pakvasa, Brent Tully and Dave Yount all agreed to be on my dissertation committee, and gave their precious time and effort to read, correct and improve this work. I am grateful for their advice. Diane Ibaraki reliably manages the computer systems on which this work has been done; I have tried to compensate her by assisting in her duties, but have been able to return only a small fraction of the effort she spends. Josephine Asselin, Caroline Chong, Claudette Jaurigue, Cheryl Kamiyama, Kathleen Kaya, Margie Kidani, Teri Mimori and Glenna Sumiye run the offices of the high-energy physics group and the physics department; they helped on innumerable problems of daily life.

Morton International's Fairport mine is the hospitable place where our experiment is located. The miners (above and below ground, high or low on the totem pole, management or union) could not be more helpful nor forgiving when we hinder their work, need their help, or just provide amusement by being so clumsy and weird.

Somehow the most important contributions always get mentioned last; but I already pointed out above that life isn't always fair. My wife Linda puts up with me (which is hard enough), answers stupid questions, and solves intelligent problems. My parents, sisters and friends helped in any way they could, in spite of a distance of over tenthousand miles.

Abstract

The IMB detector is the worlds largest water Čerenkov particle detector, situated in the Fairport mine in Grand River, Ohio. It consists of about 8000 tons of ultrapure water instrumented with 2048 eight-inch photomultipliers. It has an effective area of 360 m^2 , and registers upward-going muons at a rate of about 0.5 events per day. These muons are interpreted to be the result of neutrino interactions in the rock below the detector. The flux and angular distribution of these muons is compared to the expectation from the production and interaction of atmospheric neutrinos; good agreement is found.

Pattern-sensitive event fitting techniques are used to determine the direction of the muons. An estimate of the muon energy is made using the measured energy deposition of the muons. That muon energy estimate is used to estimate the scattering angle between neutrino and muon, which determines the angular resolution of the muons when used for neutrino astronomy.

The celestial arrival direction of the muons is studied for correlations, using a twopoint correlation technique. A correlation of 2.5σ statistical significance is found at small angles. This is compatible with about 17 of the neutrino events originating from a few point sources.

The ensemble of the muons is studied by an ideogram technique in search of point sources. No statistically significant neutrino point source candidates are found; the most unlikely point source candidate has a chance probability of 12%. The flux of neutrinos from the direction of eight astronomical objects which are expected to emit neutrinos is estimated; two sources (Cen X-3 and Her X-1) show a flux excess of neutrinos with individual chance probabilities less than 2%.

Contents

Acknowledgements	iv
Abstract	viii
List of Tables	xii
List of Figures	xiii
1 Neutrino Astronomy	1
1.1 History	1
1.2 Atmospheric Neutrinos	5
1.3 Point Source Neutrinos	11
1.4 Candidate Sources	15
2 The IMB Detector	20
2.1 Photodetection	21
2.2 Trigger and Readout	24
3 Statistical Interlude I	27
3.1 Hypothesis Testing	27
3.1.1 Testing Multiple Hypotheses	30
3.2 Circular Statistics	32
4 Analysis Methods	35
4.1 Conventions Used	35
4.2 Data Selection	37
4.2.1 Data Set, Livetime	42

4.3	Timing and Pulse Height Resolution	47
4.3.1	Timing Resolution	50
4.3.2	Pulseheight Resolution	53
4.4	Muon Direction Fitting	60
4.4.1	Initial Estimate	60
4.4.2	Precision Fitting	64
4.4.3	Direction Resolution	68
4.5	Muon Energy Estimation	72
4.6	Detector Effective Area	79
5	Statistical Interlude II	86
5.1	Two Point Correlation	90
5.2	Skymap Ideograms	97
6	Results	104
6.1	Atmospheric Neutrinos	104
6.2	Neutrino Direction Correlation	105
6.3	Flux Excess from the Galactic Plane	112
6.4	Point Source Search	113
6.5	Flux from Point Source Candidates	125
7	Conclusions	138
A	Flux Limits and Excesses	141
B	Geometry of Expected Timing	145
C	Event Format	147
C.1	Tape and Event Format	147
C.2	Layout of the INBUF Data	154

C.3 Layout of the IX, IA and RN Data	157
D Pauli's 1930 Letter	160
Bibliography	162

List of Tables

4.1	Summary of live time and event rate	46
4.2	Number of events used for resolution calibration	49
4.3	Correction factors used for the estimation of muon energy from observed light output	74
4.4	RMS muon to neutrino scattering angles	78
6.1	Total flux of muons from atmospheric neutrinos	105
6.2	Peaks found in the flux excess significance skymap	122
6.3	Peaks found in the flux excess significance skymap, excluding secondary peaks	125
6.4	Significance and muon flux of source candidates	135
6.5	Neutrino flux for sources with an observable flux estimate	136
6.6	Neutrino flux and luminosity limits from point sources	136
C.1	Event format layout	151
C.2	Blocked event format	152
C.3	Summary of storage formats	153

List of Figures

2.1	Fraction of photomultiplier hits due to scattered light	23
4.1	Layout of the coordinate system for the IMB detector	36
4.2	Display of a through-going muon event	40
4.3	Display of a stopping muon event	40
4.4	Display of a contained stopping event	41
4.5	Event rate versus data tape number for IMB-3	46
4.6	Timing residuals: mean value and early/late resolution	51
4.7	Mean expected pulseheight as a function of impact angle and distance .	57
4.8	Pulseheight resolution as a function of impact angle and distance	58
4.9	Hit probability as a function of impact angle and distance	59
4.10	Distribution of angle deviation between generated monte-carlo muon and reconstruction	71
4.11	Distribution of of relative light output ϵ	74
4.12	Change in distribution of muon energy and muon-to-neutrino angle with ϵ	76
4.13	Mean muon energy as a function of ϵ and pathlength	77
4.14	RMS muon-to-neutrino angle as a function of ϵ and pathlength	77
4.15	Energy distribution of simulated atmospheric neutrino events	80
4.16	Energy-dependent recovery efficiency	82
4.17	Effective area as a function of zenith angle	84
4.18	Effective area as a function of declination	85
5.1	Distribution of pair separation angle	94
5.2	Twopoint correlation as a function of pair separation angle	95
5.3	Cumulative twopoint correlation as a function of pair separation angle .	95

5.4	Standard deviation of the cumulative twopoint correlation	96
5.5	Sample of the grid used to calculate ideograms	101
6.1	Flux of muons from atmospheric neutrinos	108
6.2	Distribution of pair separation angle	109
6.3	Twopoint correlation as a function of pair separation angle	110
6.4	Cumulative twopoint correlation as a function of pair separation angle .	110
6.5	Standard deviation of the cumulative twopoint correlation	111
6.6	One-dimensional muon flux excess search in galactic coordinates	114
6.7	Event positions in equatorial coordinates	116
6.8	Smearred event flux	117
6.9	Mean expected event flux	118
6.10	Significance of flux excess	119
6.11	Significance of flux excess in galactic coordinates	120
6.12	Distribution of peak number in the background	124
6.13	Chance probability of aggregate peak flux excess significance as a function of rank	124
6.14	Flux distributions for background data with injected signal events . . .	128
6.15	Flux distributions for background data with injected Poisson-distributed muons	130
6.16	Conversion from μ to ν flux	132
6.17	Comparison of neutrino flux from the Crab pulsar to gamma-ray measure- ments	134
6.18	Comparison of neutrino flux from Hercules X-1 to gamma-ray measure- ments	137
A.1	Flux limits for sources 1A0620 through AE Aqr	141
A.2	Flux limits for sources Cen A through Cyg X-3	142
A.3	Flux limits for sources Gal. Cen. through LMC X-3	143

A.4 Flux limits for sources LMC X-4 through Vela X-1	144
B.1 Geometry of muon, photomultiplier and Čerenkov light	146

Chapter 1

Neutrino Astronomy

Among all elementary particles, the neutrino* probably has the dubious distinction of having the most misconceptions associated with it. It is often termed a tiny particle. But so far we have not found indications (like finite form-factors) that any of the known quarks, leptons and gauge bosons are not pointlike, so the notion of size has no meaning. Even worse, it is sometimes called the lightest particle. Physics has only indirect evidence for masses of any neutrinos; but the limits on other particles are much more stringent: from observation of the magnetic field of galaxies²¹ we know the mass of the photon to be $< 2 \cdot 10^{-27}$ eV,[†] whereas the best limits on the masses of neutrinos are on the order of 10^{-1} eV. So it cannot fairly be called the lightest. It may be the only one about which a famous author wrote a poem.⁹⁴ It has been dubbed “ghostlike”; I fail to find it particularly spooky. It has also been named “the elusive neutrino”. That, however, is true.

1.1 History

The neutrino is typical for the new paradigm of particle physics: theory is confirmed by experiment, not experiment explained by theory. The neutrino was predicted first and only observed 25 years later. The premature announcement of its supposed existence was based on symmetry and conservation principles, which particle physicists hold in quasi-religious esteem: one chooses to violate Occam’s razor, make the world more

*In this work, the term *neutrino* refers both to neutrinos and anti-neutrinos, unless noted otherwise. It is also used as a generic term referring to all three species of neutrinos, the *electron-*, *muon* and *tau-neutrino*.

[†]In the following I will always use the convention of natural units by setting $\hbar = c = 1$, and measure energies, momenta and masses in units of eV.

complicated and invoke a few more particles, rather than giving up the notions of beauty and simplicity which are tried and tested.

As bad as this sounds, this method has worked very well. The different species of neutrinos are the foremost example of particles discovered long after they were predicted; others are \bar{p} , Ω^- , W^\pm , Z^0 and so on. There are a few counter-examples (like the axion^{35,69,98}) of particles which were never found, and after a while the theories which predicted them are usually modified, purged or just forgotten, to hide the evidence for a mistake having been made. A few predicted particles are still outstanding: the tau neutrino (observed only indirectly), the top quark (several accelerators have been built expressly to look for it, so far in vain), and the Higgs boson (the search for which will cost the world many billions of dollars).

The problem which was solved by the invocation of the neutrino was one of the black clouds on the horizon of physics in the 1930s. In beta decay, a nucleus undergoes a transformation into another nucleus and emits a “beta-particle”, called an electron today. The masses of both nuclei (before and after) are known and fixed, and the mass of the electron is a constant too. Therefore by the laws fundamental to relativity (conservation of four-momentum) the energy and the momentum of the emitted electron are fixed; in modern language: a two body decay has a line spectrum. Unfortunately, Chadwick had already observed by 1914 that beta particles have a continuous spectrum, with energies below a sharp cutoff energy, which happens to be where one would expect the line from two-body decay. Furthermore, by the 1930s it was known that the electron has a spin of $1/2$, but that the nucleus changes its spin by an integer in beta decay. To solve these two problems without having to give up all these conservation laws, in 1930 Pauli⁶⁸ proposed that an additional invisible particle is emitted in beta decay; the letter in which he first made that suggestion is translated in appendix D. The neutrino went public in 1934, when Fermi’s³⁷ theory of beta decay incorporated the new particle; his choice of the name “neutrino” stuck to it. Many of its properties could already be predicted in those years: it had to be fairly light (much lighter than the electron), otherwise the upper limit of the beta spectrum would not match predictions; it had to be electrically neutral (hence the name). And it had to interact weakly, to explain that it hadn’t been observed yet. It is a solution sure to please everyone: all conservation laws are saved,

another particle is added, its properties are crafted to prevent it from being visible, so as to explain the failure to detect it.

And indeed the smallness of its interaction cross-section prevented its observation until the early 1950s, when Fred Reines and Clyde Cowan finally announced the discovery of the neutrino.^{25,80} This experiment remains a textbook example of how physics should be done. Ray Davis²⁸ had already shown that neutrino and antineutrino must be separate particles. Shortly thereafter, Maurice Goldhaber⁴³ measured the helicity of the neutrino to be lefthanded.

The next few years saw the first “neutrino identity” crisis: the discovery of the neutrino had been through inverse beta decay, where the neutrino is produced in association with an electron in a reactor, and detected through the reaction $\bar{\nu} + p \rightarrow e^+ + n$. Later another type of process was observed at Brookhaven²⁷: a neutrino is produced in association with a muon in $\pi^+ \rightarrow \mu^+ + \nu$ and then detected through the creation of a muon: $\nu + n \rightarrow \mu^- + p$. But are these two neutrinos the same? Sakata⁸⁷ had already suggested in 1943 that there may be a difference between the neutrinos associated with electron and muon, and in 1958 Feinberg³⁶ showed that the two types of neutrinos should be distinct. The question was hotly debated, until the absence of $e^- + p$ from the final state of the second reaction proved that the two neutrinos must have different identities. A set of new quantum numbers was introduced to describe this conservation of electron-ness and muon-ness. In 1975 the τ -lepton was discovered at SLAC,⁷⁰ which requires another quantum-number, tau-ness. To date, only the first two neutrino species have been observed directly; but results from LEP and SLAC³⁸ show conclusively that there are three species of light neutrinos, where the third and as yet unobserved neutrino is probably associated with the tau lepton.

Let us discuss the properties of neutrinos which are of interest here, based on the standard model of particle physics. It was already pointed out that they are electrically neutral leptons, and that there is one species of neutrino associated with each species of charged lepton (electron, muon and tau). In the language of multiplets, the lefthanded neutrino and the lefthanded part of the charged lepton form a weak isospin doublet, while the righthanded part of the charged lepton is a weak isospin singlet of its own.

The association within such a lepton family stems from the way the weak currents are defined: they only couple particles within a family.

Since neutrinos are neutral leptons, they only interact through the weak interaction, which explains the smallness of their cross-sections.

The big question is the mass of the neutrinos. In the simplest case one assumes their masses to be all zero. This makes life easy for practical work, but seems unnatural. Why should the mechanism which gives mass to the other quarks and leptons ignore the neutrino? On the other hand, that mechanism is not well understood anyhow, and the standard fashion of invoking a Higgs boson and magically adjusting its coupling to the fermions to make their masses come out right is very arbitrary. In the case of completely massless neutrinos life is indeed easy: neutrinos cannot interact with each other, therefore they can also not change identity, there are no neutrino oscillations. Neither can they flip their spin (only a massive particle can do that), and since only lefthanded neutrinos and righthanded antineutrinos are coupled to the weak current, these are the only ones which have to be described by the theory.

The suggestion to allow neutrinos to have a mass and to mix was already made very early, initially by Pontecorvo⁷¹ for the case of a sterile neutrino and independently by Sakata⁶⁰, and it does complicate matters.^{1,57} First of all, since now neutrinos are allowed to flip their spin, one has to add righthanded neutrinos to the lepton families. The simplest assumption is to move that righthanded neutrino into an isospin singlet similar to the righthanded lepton singlet, thereby preventing it from interacting with anything, since there is no weak current coupled to weak isospin singlet particles. In such a model there is no need to introduce a righthanded current, except maybe aesthetic reasons; that leaves the righthanded neutrinos (and lefthanded antineutrinos) unable to interact at all, or “sterile”.

The real complication comes from the fact that a neutrino can now change into a different neutrino species. To describe that phenomenon one writes down a mixing matrix, equivalent to the Kobayashi-Maskawa (K-M) matrix in the case of quark interactions. Similar to the textbook case of the K^0 system one now has to look at neutrinos in two different ways: when they interact, they do so in flavor eigenstates, for example ν_e interacts

with an e^\pm , whereas ν_μ interacts with an μ^\pm . On the other hand, when propagating the neutrinos are in mass eigenstates ν_1 through ν_3 . The transformation of a state between the two bases $(\nu_e, \nu_\mu, \nu_\tau)$ and (ν_1, ν_2, ν_3) is done with a unitary mixing matrix, analogous to the K-M matrix. The easiest way to describe this system is not to specify the whole mixing matrix, but rather to approximate it by mass differences and mixing angles between pairs of neutrinos.

If the neutrinos travel through matter, the mixing is modified because for electron-type neutrino neutral and charged currents create an effective index of refraction of the medium, whereas for the other neutrinos only neutral current interactions contribute to that; this introduces a different “effective mass” for the electron neutrino. The resonant form of these matter oscillations known as the MSW effect^{65,99} may explain the presumably observed deficit of solar neutrinos.

For the purpose of this work, neutrino mixing is mostly irrelevant and will be ignored, although one could set limits on the mixing parameters, using the fact that observed and expected neutrino flux agree.

1.2 Atmospheric Neutrinos

On earth one can in principle observe neutrinos from many sources. Most copious are the neutrinos left over from the big bang; they are the equivalent of the 2.7 K microwave background. The typical energy of these neutrinos is $1.9 \text{ K} = 160 \mu\text{eV}$, which makes detecting them unimaginable with the means available today. They are nevertheless of great interest, since with even a small mass *per neutrino*, of the order of eV, they could partially or entirely solve the dark matter problem.^{34,100}

The sun and all other stars emit neutrinos of low energy in the process of their energy generation; their energies are typical of their nuclear physics origin, below 10 MeV. Solar neutrinos have been observed, both radiochemically, for example by the Homestake mine experiment,²⁸⁻³⁰ and directly by the Kamiokande experiment.⁴⁸

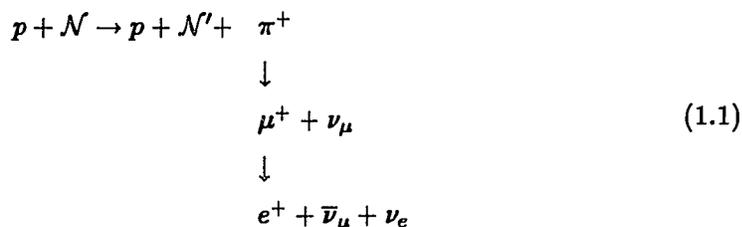
The next set of neutrinos are atmospheric neutrinos, which are of interest for this work, since they make up the background to the search for neutrinos from astronomical point sources. The earth is continuously bombarded by “cosmic rays”, with energies ranging up to 10^{21} eV. The term “cosmic rays” here refers to protons and nuclei (typical ones are for example alpha particles and ^{56}Fe nuclei). When these strongly interacting cosmic rays hit the atmosphere, they create showers of particles, many of which decay creating neutrinos. Thereby a smoothly distributed uniform flux of neutrinos is created, the angular and energy distribution of which can be predicted from the known properties of cosmic rays and of particle interactions. This will be discussed in greater detail below.

There are two more sources of neutrinos observable on earth: supernova explosions, or more accurately stellar collapse. The neutrinos from SN1987a detected by the experiment used in this work⁹ and by the Kamiokande group⁴⁷ were the first detection of neutrinos from outside our solar system, and opened the door to neutrino astronomy.

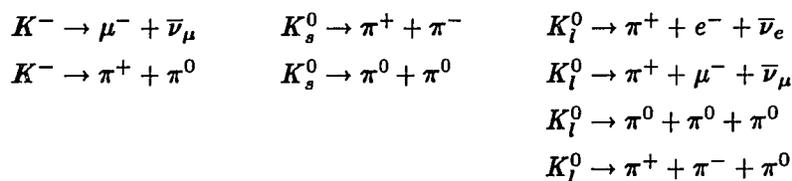
The last and to us most important set of observable neutrinos are high-energy neutrinos from astronomical sources. They constitute the desired signal of this work, and will be discussed in greater detail in section 1.4.

What exactly accelerates cosmic-ray proton's and nuclei to high energies is an open question. Most models are built around strongly time-dependent magnetic fields, the most popular being the shock wave of a supernova explosion.^{10,18,50,54}

When the cosmic ray protons and nuclei reach the earth's atmosphere, they interact inelastically with it, creating an atmospheric shower. In the shower kaons and pions are produced, their number being roughly proportional to the energy of the primary. The charged pions and kaons decay to muons, producing a muon neutrino as a byproduct. Some of the muons decay in flight into electrons, producing another muon neutrino and an electron neutrino; the fraction of muons which gets absorbed (usually by making it all the way to the earth's surface) depends on their energy. The dominant interaction (there are many others) is:



Kaons which are produced decay similarly, via



plus the charge-conjugate reactions. As the energy of the reaction increases, the fraction of kaon secondaries increases. As is easily seen, typically two muon neutrinos are created for each electron neutrino, but this ratio increases towards higher energies, since muons of high enough energy can reach the earth's surface and be absorbed without emitting a high-energy neutrino. This gives rise to a ratio of muon to electron neutrinos of about $2 \div 1$ at 1 GeV, rising to about $10 \div 1$ at 100 GeV.

Obviously these reactions also produce muons copiously, and at higher energies they can reach the earth's surface and proceed deep underground. They are the dominant component of sea-level cosmic rays, where their flux is about $2 \cdot 10^6 / \text{cm}^2 \text{s}$. By the time they reach the depth of the IMB detector, their flux is down by over four orders of magnitude, but that still leaves an event rate of 2.7 events/s over the whole detector. The method used to isolate neutrino-produced muons from these millions of cosmic-ray muons is to use the earth as shielding, and study only event geometries where that shielding is sufficient.

There are two possible ways of distinguishing neutrino-induced muons from direct cosmic-ray muons. At low energies one just looks for events where the interaction takes place in the detector, that is events with no entering particles. At higher energies a much more efficient mechanism is available: the flux of cosmic-ray muons falls off very sharply as the rock overburden between the atmosphere and the detector increases. At a deep

underground detector, the overburden increases rapidly when going from zenith to the horizon. For upward-going muons to come from the surface of the earth, the distance they would have to travel through earth (90 km to 12000 km) is so large that the flux of direct muons is sufficiently reduced. On the other hand, the flux of muons generated by neutrino interactions in the rock surrounding the detector is homogeneous over zenith angle within a factor of two. The technique is now to restrict oneself to muon zenith angles where the direct muons have been removed. The advantage of this method over the study of contained events is that the target volume for neutrino interactions is much larger than just the detector volume. The exact angle below which the neutrino signal can be observed nearly free of the cosmic-ray muon background depends on the detector depth. In this work I will use only events originating below the horizon; at the depth of the IMB detector (600 m or 1570 m water equivalent) this virtually ensures that no cosmic-ray muons are accepted. Note that other experiments such as Case-Witwatersrand-Irvine,^{26,83} Kolar Gold Fields⁵³ and Fréjus⁷ were not able to reliably measure the sense of the direction a muon travels in; therefore they use an angle band around the horizon to study neutrino-induced events.

In principle one can observe both electron and muon neutrinos in this fashion. However, a charged current interaction of an electron neutrino will generate an electron, which will create an electromagnetic shower in the medium, which in turn is rapidly absorbed in the rock. A muon neutrino on the other hand can create a muon, which can travel large distances since it is a minimum ionizing particle, and has a much higher chance of reaching the detector. A tau neutrino (if it exists) would create a tau, which at the energies of interest here immediately decays either into an electron, a muon, or into hadrons, which create a hadronic shower; therefore tau neutrinos cannot easily be distinguished from electron and muon neutrinos. Neutral current interactions will have electrons or hadrons in the final state, so they also fall into the class of showers. The data studied here contains several events which are probably upward-moving particle showers. These showers can also be produced by the "debris" (hadrons, jets, nuclear waste) emitted from the interaction; in which case the observed final state will be the electron or muon plus the hadronic shower. In all cases, the bulk of the observed data will be single muons, with a small admixture of electromagnetic and hadronic showers.

In principle there should also be multiple muon events (for example muon pairs from Ψ decays); however, no upward-going multiple muon event has been observed so far.

The energy spectrum of atmospheric neutrinos can be estimated simply from the known energy spectrum of primary cosmic rays: $dF_{\mathcal{N}}/dE \propto E^{-2.7}$ and the energy scaling of hadronic interactions at high energies, so the neutrino spectrum should be $dF_{\nu}/dE \propto E^{-3.7}$. In practice it is a little steeper, mostly due to increase in absorption cross-sections for pions and kaons at higher energies. The angular distribution varies by a factor of two (at the energies typical of this experiment) between nadir and the horizon, due to the increase in decay length in the horizontal geometry.

The minimum energy for a muon from an atmospheric neutrino to be observed is set solely by the observational constraints on the muon detection. In the case of this work, the muon has to penetrate at least several meters of detector, which leads to an energy threshold of about 2 GeV. This is discussed in more detail in section 4.5. The maximum energy is set by absorption in the earth: at extremely high energies (PeV or more) the neutrino interaction cross-sections increase rapidly enough to prevent neutrinos from penetrating the earth. This effect is not of interest for atmospheric neutrinos due to their steeply falling energy spectrum.

A detailed calculation of the flux of upward-going muons from atmospheric neutrinos is required for comparison to the observed flux, which will be done in section 6.1. That requires a detailed understanding of interaction of pions and kaons is required; furthermore knowledge of the incoming cosmic ray flux (which has been measured very well) and the composition and structure of the atmosphere are needed. The measured muon flux and energy distribution at sea level is used as another normalization for such a calculation. The atmospheric neutrino flux and its angle and energy dependence has been calculated by several authors; I use the both the calculations by Volkova⁹⁵ and by Lee and Koh.⁵⁸ As will be shown in section 6.1, these calculations differ by about 10%; the systematic uncertainty of these calculations is about 20% or less in the absolute normalization of the muon flux.

To calculate the rate of detected muons expected from the flux of atmospheric neutrinos two more steps are required. One needs the cross-section for the interaction of neutrinos

with matter; the EHLQ³³ cross-section for conversion from neutrinos to muons are used; they are probably the best parametrization available today. One also requires knowledge of the muon energy loss (which determines muon range, the effective target volume and thereby the observed muon rate); I use the calculation of Bezrukov and Bugaev⁸; the muon energy loss calculation of Lohmann⁵⁹ has also been tried, the results are not substantially different. The resulting muons are then processed through a simulation of the detector and the resulting events manually scanned, just like real events; this yields an expectation of the observable muon rate. This simulation was carried out for a detector live time of 10,000 days each for IMB-1 and IMB-3.

The results are shown in section 6.1; the angular distribution is fairly flat, varying by less than a factor of two between nadir and the horizon. The total expected muon flux integrated over the lower hemisphere is about $2.5 \cdot 10^{-13} \text{ cm}^{-2} \text{ s}^{-1} \text{ sr}^{-1}$ for the energy threshold and efficiency of this detector.

Inversely, to calculate a neutrino flux from an observed number of muons also requires knowledge of cross-sections and muon range. Neutrino interaction cross-sections increase linearly with energy at low energies and logarithmically at high energies; the break between the two occurs at the W^\pm production threshold, which is at about 3 TeV in the laboratory system for $\nu + p$ interactions. Cross-sections and muon ranges are usually treated together in the literature. Detailed calculations can be found in references [39,40,76,84,91]. I use the resulting conversion factor from neutrino flux to observed muons in the detector from reference [40] which agrees well with the other references given above; the result is displayed in figure 6.16 in section 6.5.

Atmospheric neutrinos have been observed and their energy and angular distribution measured in many experiments, either as upcoming or horizontal muons, among them CWI,^{26,83} KGF,⁵³ Baksan Neutrino Telescope,²² Kamiokande,⁶⁷ and this experiment. All these results have threshold energies determined by the ability of the muon to penetrate the detector, typically a few GeV. Atmospheric neutrinos can also be observed as contained events in large detectors¹⁶; however with the much smaller target volume available to contained events and the absence of a geometric threshold, these events are at lower energies, typically 10 MeV–2 GeV.

1.3 Point Source Neutrinos

The previous section discussed the background of this study, this one is devoted to the signal. The idea of detecting neutrinos from astronomical point sources was developed early on. In the 1960 Berkeley instrumentation conference, Ken Greisen⁴⁴ outlined the program required to detect neutrinos from a “supernova” (which is to be taken as a supernova remnant in today’s language), including the suggestion of placing a detector in a salt mine. At the same time Fred Reines’ group⁸¹ was already building the first neutrino detector in the Fairport mine. In the same year Greisen,⁴⁵ Markov and Zheleznykh⁶² and Reines⁸² published fairly detailed accounts of how to detect neutrinos from outside the solar system and what fluxes to expect. Ray Davis²⁸ had previously tried to observe solar neutrinos in 1955; see reference [1] for more details about the history of solar neutrinos.

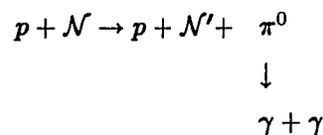
It was clear from the outset that the only possibility to detect neutrinos from objects farther than the earth’s atmosphere and the sun would require very large and very well shielded detectors. Large because the rate at which neutrinos interact is very small, and well shielded to have control over the cosmic-ray muon background. These two requirements naturally lead the experimenter to the two techniques which have characterized the field of neutrino astronomy: the use of water Čerenkov detectors, since water is the cheapest material from which a particle detector can be built; and the installation of experiments in locations deep underground or in the deep ocean, in order to use naturally available shielding.

A further requirement of the detection system is its ability to reconstruct the directions of muons and thereby neutrinos with sufficient precision. As we have seen in the previous section, there is a near-homogeneous flux of atmospheric neutrinos coming from all directions; they form the background to the search for neutrinos from astronomical sources. To find these point-source neutrinos in the background one employs the uniformity of the atmospheric neutrino background, and looks for a local enhancement in the flux of neutrinos in a particular direction. Clearly, that requires knowing the directions of the muons and neutrinos; in fact, the background rejection and therefore the signal-to-noise in the search for neutrino excess depends quadratically on the angular resolution with which the muon and neutrino directions can be reconstructed.

The desirability of doing astronomy using neutrinos is fairly easily explained too. Of all known particles, only neutral particles have the ability to point to their source, all others will be swept around by the magnetic field of the earth or the galaxy. Furthermore, only particles which can survive the lengthy trip from the source at astronomical distances to detection on earth are interesting here. Of the known particles, that already restricts us to photons and neutrinos. The history of using photons for the purposes of astronomy began when some caveman wondered what the light dots in the sky might be. Since then, the field has been exclusively owned by photons, ranging in wavelength from hundreds of meters at radio frequencies to $2 \cdot 10^{-31}$ m at the highest energy air showers observed. The reason for the dominance of photons is simple: they interact electromagnetically with matter, so their interaction cross-sections are large, and they can easily be observed, for example when they create charge in the retina of the human eye by ionization. The cross-section for neutrino interactions are much smaller, so they are not readily observed; that requires large amounts of matter for even a minute chance of stopping a neutrino. Furthermore, since unlike photons neutrinos cannot be focused or imaged, their interaction is naturally limited to some angular resolution which is typically on the order of degrees. The single advantage they hold from the detection point of view is that a neutrino detector can be made with an angular acceptance of 2π sr or more, which has to be compared with the tiny aperture of a telescope.

The physics disadvantage of photons stems from the same fact which makes their detection advantage, that they are easily absorbed. Not only does this apply to the photon detector, but also to matter around the location where they have been created. All the visible light which reaches us from normal stars is generated in the outermost layer; it has been said that with photons one can only probe the “weather” on the surface of an object. Compare that to neutrinos, which once they are produced, will in most cases leave their production site freely. They carry information about the conditions inside interesting sources, behind a column density of matter which makes escape of photons impossible. For example, the flux of solar neutrinos is extremely sensitive to the temperature in the core of the sun, and is in essence being used as a probe of the innermost part of the solar interior.

There are two ways to generate high-energy photons: by electromagnetically accelerating electrons, which then emit photons by radiating in a magnetic field, or from the decay of π^0 generated in hadronic reactions as a beam of accelerated protons and nuclei strikes some other matter, similar to a particle beam dump. If all photons are generated electromagnetically, there may be no observable neutrinos generated in the process. However, there is no model which can explain the energies and fluxes observed in high-energy gamma rays solely through electromagnetic creation of photons. Therefore one has to assume that they are generated in the hadronic process equivalent to reaction (1.1):



In those two reactions the number of π^0 and π^\pm generated is the same, and since either pion creates two observable neutral particles in the final state ($\pi^0 \rightarrow \gamma$ and $\pi^\pm \rightarrow \dots \rightarrow e^\pm + \nu_\mu + \bar{\nu}_\mu + \nu_e$), the total number of neutrinos produced is already about twice as high the number of photons. On top of that, if the production region (where the proton beam hits the target) is obscured from our view by matter with a column density of more than about 100 g/cm^2 (several radiation lengths), a large fraction of the photons will be absorbed. On the other hand, the production of neutrinos becomes more efficient with increasing column density, up to about 200 g/cm^2 . In summary, the production ratio λ of neutrinos to photons is a strong function of the details of the region in which the hadronic interaction takes place and column density obscuring that region from our view; λ is almost always greater than one, with typical expected values of 10 to 30 for binary pulsars.^{20,41,42,52}

Detailed models of certain categories of sources will be discussed in the next section. Here we will concentrate on the features common to all of them. The flux from all point sources can be very well approximated by a power law spectrum $dF/dE \propto E^{-\gamma}$, where γ is known as the spectral index. Atmospheric neutrinos have a spectral index of about 3.8. Neutrinos from astronomical sources are expected to have much harder spectra, with typical spectral indices of 2.0 to 3.0. Since the neutrino interaction cross-section increases essentially linearly with energy, and the muon energy is mostly proportional to the neutrino energy, and since the muon range increases again with energy, it

follows that the observed muon spectrum is much harder than the neutrino spectrum. A point source with a differential spectral index of 2.0 for neutrinos (which is about what is expected) gives rise to a muon spectrum which is close to being flat. Another important feature of point-source spectra is the upper cut-off energy. The mechanism which accelerates the neutrinos usually has some geometry-dependent maximum energy. At much higher energies (significant only above PeV) the absorption by the earth limits the flux of detectable neutrinos. The ultimate limit is the absorption of neutrinos by the uniform 2.7 K photon and 1.9 K neutrino background which remains from the big bang; this absorption is the equivalent of the Greisen-Zatsepin cutoff⁴⁶ for cosmic rays, but does not set in until energies of 10^{21} eV. The lower energy threshold is determined by the detection technique and will be discussed in more detail in section 4.5.

There are several other experiments which can detect neutrinos from astronomical sources, although none as large as IMB. The most important one is the Kamiokande detector,⁶⁷ which is a water Čerenkov detector very similar to IMB, located in a deep tin mine, about 300 km west of Tokyo. It is smaller than IMB by a factor of three, but with much better light collection (20" photomultipliers compared to our 8" photomultipliers). They have published results based on a sample of 146 events; the flux limits from that work will be compared to my results in chapter 6.5.

The Fréjus experiment,⁷ now decommissioned, was a 900 ton iron calorimeter with an effective area of about 100 m^2 , located in the Montblanc on the border of Italy and France. With its excellent angular resolution of 0.2° , it operated nearly background free, with a total exposure of $1.56\text{ kton}\cdot\text{year}$. It is unfortunately not able to distinguish downgoing from upcoming muons, so it has to cope with a much larger background. The flux limits from it are also shown in chapter 6.5.

Other detectors which have searched for neutrino emission from astronomical point sources have already been discussed in section 1.2 above; they are the CWI,^{26,83} KGF⁵³ and Baksan²² experiments.

So far no experiment has observed evidence for high-energy neutrinos to be anything other than atmospheric neutrinos; the experiments listed above only give upper limits on the possible flux of non-atmospheric neutrinos. The only two positive detections in

the field of neutrinos astronomy are at much lower energies than those of interest here: solar neutrinos and SN1987a.

1.4 Candidate Sources

Let us now turn to what kind of object one might observe in neutrino astronomy. In this section I will make up a short list of the most probable emitters of high-energy neutrinos potentially observable in the IMB detector, and a second list of objects which can be imagined to emit neutrinos. The purpose of the two lists, and the criteria for admission into them are quite different. The “short list” contains astronomical objects which have been observed to emit high-energy photons (high-energy is defined as ≥ 10 GeV here); as shown in section 1.3 any source which emits gamma-rays is usually expected to also emit neutrinos with a similar flux. For inclusion in the short list, the criteria for “observed” are to have been observed at least twice. Only objects from the short list will be seriously tested whether they show an excess of neutrinos. There are 8 sources in this short list; as we will see below, they all fall into the category of binary neutron stars.

The criteria for inclusion into the “long but loose” list are much less stringent. It includes sources which have been observed in gamma-rays, but only once. It also includes many sources which according to current prejudice might emit high-energy neutrinos with a flux observable in this experiment. This list could be made very long; I have arbitrarily chosen to include about 15 sources in this list. This list does not claim to be either complete or exhaustive.

Many of the possible astronomical sources of neutrinos fall into one category: rotationally accreting systems. They can generally be thought of as a binary, with a rapidly spinning object, for example a neutron star, usually with a large magnetic field or a shock-wave, and a second object, for example a red supergiant, which slowly loses mass onto the first object. Accretion onto the first partner takes place either in an accretion disk around it, or directly, for example through stellar wind. The interactions of accelerated protons (and from there the generation of pions, kaons, photons and neutrinos) takes place either

in the accretion disk, or on the limb of one of the partners. In other models, the neutron star is replaced with a black hole or with a whole galaxy nucleus.

The most promising category of sources are neutron-star binary systems. In the most popular model⁶ the accreted matter forms a flat accretion disk in the plane of the orbit, in which the accelerated protons interact, creating the particles of interest here. These systems have been observed to be copious producers of high-energy gamma-rays. The emission from most of these sources is periodical, sometimes with more than one period. These periods are usually explained from the rotation of the neutron star, the orbital movement of the companion, the precession of the orbit of either the companion or the accretion disc, or any combination thereof. All the sources in the “short” list are in this category; the selection criterion is that a source has been observed more than once in any form of high-energy radiation. These likely sources are:

4U0115 at position (1:15, +63.4°). This source has been observed by the Whipple,⁵⁵ Dugway¹⁹ and Gulmarg⁷⁷ air shower arrays and the Las Palmas¹⁴ gamma-ray telescopes. It is so far north that it is never below the horizon at the latitude of the IMB experiment, so it cannot be observed at all in this study.

Centaurus X-3 at position (11:18:55, -60.3°), at a distance of 5 to 10 kpc. This source has been observed by several experiments, including the BASJE⁹² and air shower array, and the Narrabri¹² and Potchefstroom⁶⁶ gamma-ray telescope.

Circinus X-1 at position (15:16, -56.9°), observed twice by the Buckland park array.²³

Crab pulsar also known as PSR0531+22, at a distance of 2 kpc, the remnant of the 1054 AD supernova, has been observed by many groups, including the spectacular work of the Whipple experiment.⁹⁷ This source is currently fashionable as the “standard candle” of VHE high-energy astronomy.

Cygnus X-3 at position (20:32, +40.96°) and a distance of about 11 kpc, everyone's favorite high-energy source, observed many times, but even more significantly often found to not emit radiation when it would have been expected to. This source shows large temporal variability (including radio outbursts during which it becomes the strongest radio source in the sky), which is often used to explain the unpredictability of its observations. It is the source which started the “muon-content” controversy: gamma showers observed from the direction of Cygnus X-3

seem to have as high muon content as if they were hadronic showers. It is also the only source which may have been observed in underground cosmic-ray muons (by this and other experiments^{3,63}), although that is still being debated. This source is so far north that at the latitude of the IMB experiment it is at or above the horizon most of the time; this study is unfortunately not very sensitive to it.

Hercules X-1 at position (16:57, +35.3°) and a distance of 5 kpc, after the Crab pulsar and Cygnus X-3 the most reliably observed source. It has the distinguishing feature of slightly but significantly different pulsar periods when observed in X-rays and in high-energy gamma-rays^{31,56,85}

LMC X-4 a binary system in the large magellanic cloud, at position (5:32, -66.4°) and a distance of 55 kpc. It has been observed in gamma-rays by the Buckland Park air shower array⁷⁵ and by the Narrabri experiment.¹³

Scorpio X-1 at position (16:19, -15.0°), relatively close at a distance between 0.3 and 1 kpc. It has been observed by the Mt. Chacaltaya,⁶⁴ Ooty⁸⁹ and Narrabri¹² experiments.

Vela X-1 at position (9:02, -40.5°) and a distance of 2 kpc, observed by the BASJE⁹² and Buckland Park⁷⁴ air shower arrays and by the Potchefstroom⁷⁸ and Narrabri¹¹ gamma-ray telescopes.

When deciding whether to include a source in this list sources in the southern hemisphere were treated more leniently, for two reasons. First, there are far less high-energy gamma-ray and air-shower experiments in the southern hemisphere, so fewer detections of a given source are possible. Second, the IMB experiment is in the northern hemisphere, and preferentially observes southern sources in neutrino-induced upward-going muons.

Other binary neutron star sources which have been observed in gamma-rays, but not repeatable, and which enter the long list, are

2A1822 at position (18:22, -37.1°), observed in air showers by Buckland park.²⁴

AE Aquarius a cataclysmic variable binary star, at position (20:38, -1.0°), right in our neighborhood at a distance of 53 pc (not *kiloparsec*), observed by the Potchefstroom gamma-ray telescope.¹⁵ This is a somewhat unusual binary source, since the accretion takes place on a white dwarf instead of on a neutron star.

SMC X-1 is a binary neutron star in the small Magellanic cloud at position (1:15, -73.3°) at a distance of 65 kpc, has been observed by the Narrabri gamma-ray telescope.¹³

As was already pointed out in the previous section, the production of high-energy neutrinos and of photons in such objects is intimately connected. The flux of neutrinos is expected to be at least as high as the flux of photons, and typically higher by a factor of 10 to 30.

The second most important class of possible neutrino sources can be summarized as being galactic nuclei; they follow the same pattern outlined above, but replacing the neutron star with a massive galaxy core. References [4,5,32,90] provide models for such sources. This class includes

3C273 is the quasar closest to us, at a distance of *only* $(6 \dots 9) \cdot 10^5$ kpc, at position (12:28, $+2.1^\circ$). Quasars are interpreted as very distant active galaxies. The neutrino flux expected from this source has been studied in reference [73].

Centaurus A one of the closest radio galaxies (a galaxy with extraordinarily strong radio emission), at position (13:25, -42.8°), at a distance of 4400 kpc.

The Galactic Center that is the center of the milky way, in equatorial coordinates is at (17:42, -30°), and in galactic coordinates obviously at $(0^\circ, 0^\circ)$, and at a distance of 10 kpc to us. It has been observed in low-energy gamma-rays (it is the most copious but variable emitter of the 511 keV positron line), but never at high energies. The reason to include it in the long list is obvious: if all galaxy cores emit neutrinos, this is the closest one.

That last class of sources catches all the remaining candidates. Black hole binary systems are similar to the first class of sources discussed. The last two sources were added for special reasons.

Cygnus X-1 a black hole candidate at position (19:58, $+35.0^\circ$) It is thought to be a binary system accreting onto a black hole of about $10 M_\odot$. It is a copious emitter of X-rays and low energy gamma-rays.⁸⁶

LMC X-3 a black hole candidate in the large Magellanic cloud, at position (5:38, -64.1°) and at a distance of 55 kpc.

Rho Ophi is the closest molecular cloud, at position (16:27, -24.0°), at a distance of 160 pc, with a mass of at least $2000 M_\odot$.⁴⁹ A molecular cloud generates neutrinos by the same mechanism which produces atmospheric neutrinos in the earth's atmosphere. In general the flux expected from such clouds is far too small to be observable, but the closest one may be worth a try.

SN 1987a the only observed source of neutrinos^{9,47} (unfortunately only of low energy gravitational collapse neutrinos so far), at position (5:35, -69.3°), at a distance of 55 kpc in the large Magellanic cloud. The neutron star which may have been formed in the gravitational collapse may be accreting matter and emitting high-energy particles.

A note of caution has to be added: as we will see in sections 4.4.3 and 4.5, the angular resolution of this study is on the order of several degrees. Sources which are spaced more closely than this distance cannot be individually resolved. That applies in particular to LMC X-1 through LMC X-4 and SN1987a, which are all within $\pm 2^\circ$ of each other, all at a distance of 55 ± 10 kpc. They can only be resolved in radio, visible and X-ray; gamma-ray astronomers manage to resolve LMC X-4 from the other three just because it is a pulsed source at a particular period.

Chapter 2

The IMB Detector

The IMB detector has been described so many times that I will keep this chapter quite short, concentrating on the features which are of particular importance to this work. References [16,93] contain more detail.

The Irvine-Michigan-Brookhaven laboratory is situated in Morton International's Fairport Mine, close to Mentor, a suburb of Cleveland. The collaboration which operates it consists of Boston University, Brookhaven National Laboratory, Cleveland State University, University of Hawaii, UC Irvine and Louisiana State University. The detector is at a depth of about 650 m, which is equivalent to 1570 m of water overburden, given the density of the rock above the detector. Its geographical position is ($41^{\circ}45'12''$ N, $81^{\circ}17'30''$ E). The detector itself is a rectangular pool of dimensions about $24 \times 17 \times 18$ m³, filled with about 8000 tons of extremely pure water.* Above the water level a framework of I-beams is attached to the walls; this frame defines the detector coordinate system, and has been surveyed with respect to an accurate magnetic compass in an empty part of the mine, far away from the laboratory and its magnetic disturbances. The layout of the coordinate system is shown in figure 4.1; it shall suffice to say here that the orientation of the coordinate system is well-defined to an accuracy below 1°.

Suspended from the frame are 2048 photomultiplier assemblies, each consisting of an 8" Hamamatsu photomultiplier attached to a 2' \times 2' wavelength shifter plate. The assemblies are on a 1 \times 1 m lattice on the six walls of the detector. The accuracy of the photomultiplier positions has been the source of some debate; one can safely state that the mean deviation from their nominal position is less than or on the order of 20 cm.

*Since the walls of the cavity are not perfectly flat, these dimensions are the minimum of a rectangular cube inscribed into the rounded corners of the pool, and the water volume is slightly larger than just the product of the three dimensions.

2.1 Photodetection

When a muon of high energy passes through the water in the detector, it emits Čerenkov light; the simple reason is that the muon (which travels essentially at the speed of light in vacuum at these energies) is moving faster than the speed of light in the medium; so in essence it creates an electromagnetic shock wave, which manifests itself as light emitted at an angle to the path of the muon. It is this light that the experiment is designed to detect, using the photomultipliers mentioned above.

These photomultiplier assemblies form the heart of the detector. The photomultiplier itself is an extremely sensitive light detector, with an efficiency of about 20% of detecting single photons. Its wavelength response (best in the blue and near ultraviolet) is well matched to the $1/\lambda^2$ spectrum of Čerenkov light. The wavelength shifter plate, which is optically coupled to the photomultiplier, absorbs light in the near ultraviolet around 340 nm and re-emits it in the blue-green around 430 nm. The wavelength shifter about doubles the sensitivity of the photomultiplier; at large distances between muon and photomultiplier assembly, the light has travelled a long distance through water, so the spectrum has been shifted towards the visible (water absorbs towards the ultraviolet), whereas at short distances the strong ultraviolet component of Čerenkov light is still present. Therefore the photomultiplier itself is the dominant part of the sensitivity at large distances, whereas at small distances the wavelength shifter plate contributes heavily.

The timing resolution of the photomultiplier is very good, on the order of 10 ns, and will be discussed in more detail in section 4.3. Unfortunately, the wavelength shifter adds uncertainty to the timing; the organic molecules which absorb ultraviolet and re-emit visible light need a finite amount of time to do so, plus the light can bounce around the wavelength-shifter plate. Since the readout electronics records the timing of the first edge of the photomultiplier signal, the timing resolution improves markedly at high pulseheights. There are two reasons for that:

- **Statistics:** the timing distribution of many photoelectrons is obviously tighter than the distribution of each photoelectron by itself.

- At high pulseheights the chance that at least one of the photons is observed directly from the photocathode of the photomultiplier (instead of from the wavelength shifter plate) is higher; these direct photons are not delayed and shifted by the organic dye.

This effect of the improvement of timing resolution is clearly visible in figure 4.6.

At very high pulseheights, the timing resolution deteriorates again; the reason being pre-pulsing. Usually the electrons emitted from the photocathode are accelerated towards the first dynode, where they knock out a few more electrons, which are accelerated to the second dynode, and so on. What happens in the case of very high light levels is that a photon directly strikes the first dynode and knocks out an electron, which then gets amplified in the dynode chain; or some of the electrons coming off the photocathode or one of the dynodes bypass one dynode. In either case the amplified pulse arrives too early compared to normal light levels, at which these effects are very unlikely.

In addition to the deterioration of the timing resolution at very small and at very large pulseheights the timing measurement is correlated with the pulseheight measurement. This effect is known as *slewing*, and is mostly due to the steepening of the leading edge of the photomultiplier pulse with increasing pulseheight. It is measured in the calibration procedure, and the event calibration removes it, using the measured pulseheight.

The negative effect of the wavelength shifter plates is that not all the blue-green light they generate gets coupled into the photomultiplier; some of it escapes back into the water. Furthermore, some of the direct Čerenkov light gets scattered, either by impurities in the water, or by parts of the detector infrastructure (surface of the photomultiplier, edge of the wavelength shifter plate and so on). Typically, for throughgoing muons* about 45% of the photomultiplier hits in the detector are due to scattered light; since the intensity of the scattered light is low that corresponds to about 15% of the observed light being from scatter. Figure 2.1 shows the total number of photomultiplier hits and the number

*These estimates were done for downgoing muons; the situation for upward-going muons should be similar.

of hits outside the Čerenkov light cone, which presumably are due to scattered light. The fraction of scattered light seems to be constant over time.[†]

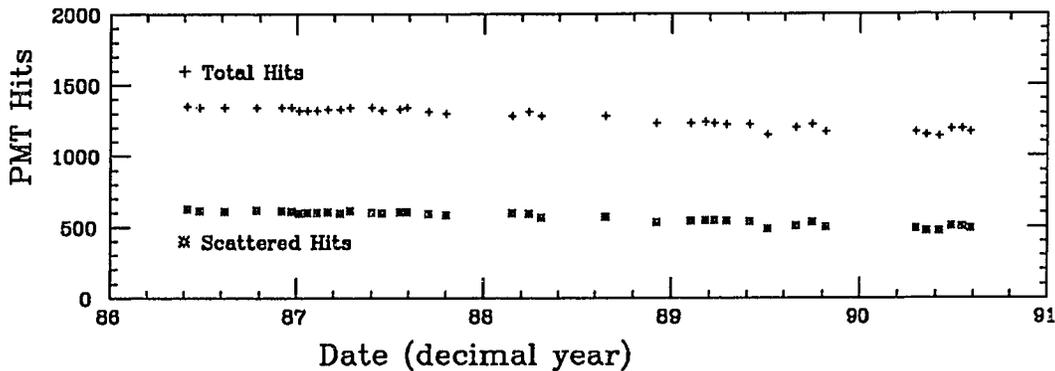


Figure 2.1: Number of photomultiplier hits as a function of time, for downgoing muons from special calibration runs. The top line is the mean total number of hits, the bottom line the mean number of hits outside the Čerenkov light cone, which are assumed to be due to scattered light.

As mentioned above, the wavelength spectrum of the Čerenkov light changes as the distance the light has to travel increases. Much more markedly, the light intensity changes. There are several factors which cause that. Most importantly, the light intensity from a line source drops off as $1/r$. Second, light is absorbed in water; even though the water in the IMB detector is extraordinarily clean, the absorption length is of the same order of magnitude as the diagonal size of the detector (absorption length is about 50 m). Lastly, as explained above, the relative weight of photomultiplier and wavelength shifter changes with distance. In principle one could try to describe the amount of light expected for a certain geometry of muon relative to the photomultiplier assembly analytically from first principles; after all, one *only* needs to know the number and spectrum of emitted Čerenkov photons (which is known), the absorption and scattering behaviour of water (which is partly measured, partly well modelled), the geometry of photomultiplier and wavelength shifter plate, the quantum efficiency of the photomultiplier, the absorption and emission spectra of the wavelength shifting dye, and the optical properties of the photomultiplier assembly. In practise, attempts to convolute all these more-or-less well

[†]The slight drop in both curves is due to the slowly increasing number of non-functional photomultipliers; due to the nature of DoE funding we were not able to replace them fast enough.

known numbers into a pulseheight estimate have failed to work completely, and still require correction factors¹⁶; instead in section 4.3 I will use real events to calibrate the details of the response of the detector.

2.2 Trigger and Readout

Each of the photomultiplier assemblies is connected with a single coaxial cable, which carries high-voltage in, with the output signal capacitively coupled into it. Outside the detector, the signal is decoupled from the HV, and fed into the readout electronics. The photomultiplier signal is first discriminated and gated to 50 ns, and that digital signal is then summed in two different fashions to form the trigger. First of all, all 2048 discriminated signals are summed up to form a global trigger; if that summed signal exceeds the value proportional to 20 photomultipliers firing, the so-called NPatch trigger is set off. The discriminated signals are also summed into “patches”, that is into square groups of 8×8 photomultipliers; each patch is again discriminated to a threshold of 9 hits, and the NPatch trigger is formed by at least two patch discriminators firing in a 150 ns window. All the discriminator levels are changeable under software control, but the above values have been used for most of the data under consideration here.

Anytime either of the two triggers fire, the event is digitized; for each photomultiplier both the pulseheight and the time of the leading edge of the pulse are recorded. A second hit on each PMT occurring at a later time (at least 500 ns later) can also be recorded; this capability is needed for finding muon decays, but is of no interest for this work. At the same time the event is digitized, the absolute time of the event is read from a radio- or satellite-controlled clock (initially a WWV radio clock, later replaced by a GEOS satellite clock). The event is then transferred to the CAMAC system and read by the online computer, in many contorted steps involving multiple buffering, which are described in gory detail in reference [16].

The online computer can now operate in several different fashions, which have to be discussed here. First of all, it can just grab all the incoming data and write it directly to tape; this is known as the uncut mode. However, the bulk of the data consists of

downgoing muons: the detector triggers on those at a rate of 2.7 Hz, but one expects only two interesting physics events (contained events and upward-going muons) *per day*. To get a head start on the data reduction, the online computer is equipped with some data analysis routines. First of all, it can run three different muon direction fitting algorithms on the data; they are known as **Cone**, **Waver** and **FastMu**, which will be described in more detail in section 4.4. It can also employ some simple algorithms, based on which it decides whether or not to save an event to tape. This is known as the cut mode. These algorithms save the following categories of event to tape:

- 1/8 or 1/16 of all events at random, so there is always an unbiased sample of events.
- Events with less than 900 photomultiplier hits. The bulk of the proton-decay candidate events will illuminate less than 900 photomultipliers, whereas most downgoing muons are at 1200 or more hits, so this cut efficiently saves contained events.
- Events with more than 1800 photomultiplier hits. These are very rare, and are typically caused by showers or multiple muons. The class of events one is interested in here is $N\bar{N}$ oscillations and muons of extremely high energy.
- Events where more than 20% of the photomultiplier hits are on the top face of the detector. This condition should be fulfilled for muons which exit through the top face of the detector or close to it; this condition is designed to save upward-going muons.
- Events for which the **FastMu** or **Waver** fit has a zenith angle of more than 70° . These are candidates for upward-going muons events.
- Events which occur within 10 ms of the previous event. This is of interest for monopole-catalyzed proton decay and stellar collapse event bursts.
- For a while, events for which the **FastMu** fit was within a band around the declinations of Cygnus X-3 were saved; this mode is known as the "Cygnus search". Since these are nearly all downgoing muons they are of no interest here.

These cuts typically reduce the amount of data by a factor of two to three, depending on operating conditions; but they are designed to save all interesting events; in particular they are extremely efficient at saving upward-going muon candidate events. Beginning with data tape number 2595, a compact summary including the information from the

three muon fitters is saved with the data even for events which are not saved; these are known as "muon summary records".*

The next step in sophistication was to add another computer to perform part of the data analysis immediately. Initially the online computer was a VAX 11/750 whereas this analysis computer was a microVAX II, so it became known as the microVAX, and the output from it are MV tapes; however in the meantime the actual processor has long been replaced by a VAXstation 3200 and the microVAX II has become the online computer. The online computer continues to write data to tape, either in cut or uncut mode. All events which go to tape are simultaneously written to the microVAX over a network link; the microVAX processes them through a part of the two data analysis chains used by both analysis groups (the west-coast and east-coast pass 0), and saves only the events which are accepted by either of those programs to tape. It also writes short summaries for all events, including the ones it rejected, for diagnostic purposes. This typically reduces the number of events by a factor of 10. Up to data tape number 3185, the upward-going muon pass 0 was not used on the microVAX; so for the purposes of this work the MV tapes before that time are not useful. Beginning with data tape number 3185 the offline analysis only has to process the MV tapes, which saves a lot of time. The other advantage of this form of processing is that a lot more data (in terms of detector live time) can be fit on one tape. When the online computer runs out of tape to write on (it typically ran with two tape drives, each with a 2400' reel of 6250 bpi tape) it can continue to gather data, except that only the MV tape is written. This extension of the cut mode is known as banzai mode.

*Unfortunately, these summaries were never saved on 8mm tapes.

Chapter 3

Statistical Interlude I

This chapter and chapter 5 will explain some of the concepts of statistics as they apply here, and outline the more esoteric techniques employed in this work.

3.1 Hypothesis Testing

Let me first clarify the language and concepts of hypothesis testing. This is frequently glossed over and poorly understood, but required for some applications in the next chapters.

Assume you have a certain idea you want to test; in our case that idea might be “there is an excess of neutrinos from a certain direction in the sky”. For the time being, let us assume the hypothesis to be tested is much simpler, say “it is darker at night than during the day”. Let us also assume that there is a measurable quantity associated with that hypothesis, called *brightness*. Obviously this brightness is not a binary variable, it takes on a whole range of values. The distribution of brightness has a certain shape, both for the day (higher at noon, lower on a rainy day) and for the night (compare a moonlit romantic evening to a cloudy night with no moon).

First we have to clarify what our hypothesis means in a more exact language. A first attempt might be “the mean brightness during the day is higher than the mean brightness during the night”. That is a very complicated hypothesis to test; it involves estimating the means of two distributions, comparing them, and assigning a significance to the comparison. An easier hypothesis to test is “the two distributions of brightness for the day and for the night are different”. Let us call this hypothesis \mathcal{H}_1 . In the case of astrophysical observations, most statistical testing is of this type. It involves a null

hypothesis \mathcal{H}_0 (the distribution of brightness is the same for day and night, the neutrino events are distributed evenly over the sky), which we hope to refute.

Therefore, a successful experiment will typically say “we have observed the distribution of brightness during day and night, and we find the two distributions to be incompatible at the 99.5% confidence level”, or equivalently “...to be different by 3 sigma”. Now what does that mean? Associated with the testing of a hypothesis such as \mathcal{H}_0 is a single number which determines whether the hypothesis seems to be right or wrong, known as the test statistic. The distribution of that test statistic can be predicted, given certain reasonable assumptions about the underlying distributions in the data. That really means: if we did many experiments of the same type we would find a certain distribution of this test statistic, and this distribution can be calculated from the known or assumed behaviour of the input data. Therefore a probability can be assigned to each value of the test statistic. Let us assume that the test statistic *increases* as the probability of \mathcal{H}_0 *decreases*. So in the case described above, the probability of observing such a high or higher value of the test statistic is only 0.5%, therefore we find that \mathcal{H}_0 has to be rejected with a probability of 99.5%, which is what the statement “rejected at the 99.5% confidence level” really means. The complement of this probability is often known as the significance or chance probability; that stems from the fact that if we repeat the experiment many times, the probability of observing such an improbable value of the test statistic *by chance even though the null hypothesis is actually correct* would be 0.5%.

To make that percentage number a little more palatable one can re-express it using the cumulative probability of the normal or Gaussian distribution as a number of standard deviations. The correct way to do that is to use the one-sided integral of the normal distribution, and express the probability p of observation of the null hypothesis (in our case $p = 0.5\%$) as a certain number of standard deviations σ , using the integral:

$$p = \int_{\sigma}^{\infty} e^{-t^2} dt.$$

Typical standard deviations and the corresponding confidence levels are $68\% \simeq 1\sigma$, $95\% \simeq 2\sigma$, $99.5\% \simeq 3\sigma$.

Note that we always accept or reject the null hypothesis: we find that the two distributions of brightness are not identical for day and night with a probability of 99.5%, or

that the chance probability of observing two distributions so different is only 0.5%. If and only if \mathcal{H}_1 is the exact logical inverse of \mathcal{H}_0 then we can conclude that \mathcal{H}_1 is correct from \mathcal{H}_0 having been refuted. As a counterexample, assume that \mathcal{H}_1 was the earlier hypothesis "it is lighter during the day than during the night". In fact, just because we have refuted \mathcal{H}_0 because it is only 0.5% likely, does not mean that \mathcal{H}_1 is true with 99.5% certainty. The problem is that there is an infinite universe of hypotheses which we have to test, and \mathcal{H}_0 and \mathcal{H}_1 are only two of those. For example, just because the distribution of brightness is different during day and night we cannot conclude that it is brighter during the day than during the night; it might have been the other way round. In most cases there are infinitely many alternate hypotheses, and if one assumes that all hypotheses are exclusive, the probability of \mathcal{H}_1 actually being true is

$$P(\mathcal{H}_1) = 1 - P(\mathcal{H}_0) - P(\mathcal{H}_2) - \dots - P(\mathcal{H}_n) - \dots - P(\mathcal{H}_\infty), \quad (3.1)$$

and the equal sign has to be replaced with a \geq if some of the hypotheses are not completely exclusive, that is they overlap. Clearly the result is unknown, and may be considerably smaller than 99.5% in our example case.

In the same fashion, observing that the neutrino directions are not compatible with being "randomly" distributed (where the term "randomly" has to be carefully defined) does not mean that there is a point source of neutrinos. There are many other ways to make a non-random distribution. This is the fundamental predicament of hypothesis testing: it can never *prove* anything interesting, it can just refute the null hypothesis. The only way out of that would be to test every single alternative hypothesis \mathcal{H}_n , and after having done that infinitely many times employ equation (3.1) to calculate $P(\mathcal{H}_1)$, obviously not a practical approach. So typically all one can do is to make a statement similar to "we have observed the distribution of brightness during day and night, and we find the two distributions to be incompatible at the 99.5% confidence level. This result is compatible with daytime being lighter than nighttime".

Put more simply, the way the truth is found through statistics is the following. One wants to prove a certain hypothesis \mathcal{H}_1 ; but since statistics can never prove anything, one has to be content with refuting the corresponding null hypothesis \mathcal{H}_0 . The naive reader may ask "but what good does it do me to refute \mathcal{H}_0 , I want to prove \mathcal{H}_1 ?" Well,

tough, this is not Burger King, you can't have it your way. The only recourse lies in the method of scientific dialog. One refutes \mathcal{H}_0 , and can then go and claim to everyone that it seems that \mathcal{H}_1 must be correct. Some critic will come and say "hold it, you may be wrong, what about \mathcal{H}_2 ?" If \mathcal{H}_1 was indeed correct, one just goes and refutes \mathcal{H}_2 . More critics will spring up, and the successful scientist proceeds to refute \mathcal{H}_3 , \mathcal{H}_4 and so on. After a while the critics will get tired (or die of old age), and slowly but surely \mathcal{H}_1 will be accepted as the "truth". Naturally, if in the initial test \mathcal{H}_0 was refuted with an overwhelming statistical significance, refuting all the other critics can be done much more swiftly than in the unfortunately common case of weaker signals. As ridiculous as this process may sound, that's the way it goes.

A point at which practical hypothesis testing often deviates a little from the textbook approach is the significance level at which one chooses to refute \mathcal{H}_0 , which is the complement of the chance probability of observing a test statistic large enough to refute \mathcal{H}_0 . In principle one should choose a significance α (for example 0.5%) *a priori without regard to the data*, and then try to refute \mathcal{H}_0 at the $1 - \alpha$ level. In practice however, one usually inverts the process: measure the chance probability of wrongly refuting \mathcal{H}_0 , and then use that for the significance. So one usually finds experimentally that the test statistic yields a probability of \mathcal{H}_0 of only 0.5%, and then proceeds to say that \mathcal{H}_0 has been refuted at the 99.5% confidence level. This is not wrong, but just a shortcut.

3.1.1 Testing Multiple Hypotheses

Assume we test two hypotheses \mathcal{H}_1 and \mathcal{H}_2 against the corresponding null hypotheses $\mathcal{H}_{0,1}$ and $\mathcal{H}_{0,2}$, and find the two null hypotheses being rejected at confidence levels $1 - P_1$ and $1 - P_2$, which means the chance probability of wrongly accepting \mathcal{H}_1 and \mathcal{H}_2 are P_1 and P_2 . Furthermore, assume these two tests are independent of each other (whether they have been performed on the same or on different data, their results shall not be correlated, which means that knowledge of one test does not lead to an advantage in predicting the outcome of the other). What is the chance probability of wrongly rejecting both of the two null hypotheses?

At first one would assume it is just $P_1 \cdot P_2$. That is one possible valid answer; it is the answer to the question “what is the chance probability of the significance for test 1 being $\leq P_1$ and the significance of test 2 being $\leq P_2$ ”. However, it is usually not the answer one is interested in. The often more meaningful question is “what is the chance probability of observing two results, which when taken together have a significance which is $\leq P_1 \cdot P_2$ ”. If one is only interested in the compound statistical test, then the quantity of interest is only the product, not the sharing of chance probability between the two tests.

If we think of the chance probabilities p_1 and p_2 in the two-dimensional plane bounded by $0 \leq p_1 \leq 1$ and $0 \leq p_2 \leq 1$, the part of the plane where the product is smaller than $P = P_1 \cdot P_2$ is the area to the left of the hyperbola defined by $p_1 \cdot p_2 \leq P$. The distribution of confidence levels is by definition uniform over this plane: since both p_1 and p_2 are proper probabilities, their distributions are uniform over the interval $(0, 1)$. This area can be found by integrating below and to the left of the hyperbola:

$$\begin{aligned} P_{12} &= \iint_{p_1 \cdot p_2 \leq P} dp_1 dp_2 = \int_0^P dp_1 + \int_P^1 \left(\int_0^{P/p_1} dp_2 \right) dp_1 \\ &= P \cdot (1 - \ln P) \end{aligned} \quad (3.2)$$

Note that since $P = P_1 \cdot P_2 < 1$ this can be considerably larger than just P . Ignoring to use this formula and using P instead of P_{12} will therefore *underestimate* the chance probability, and may lead to refuting the null hypothesis erroneously.

This scheme, due to Wallis,⁹⁶ can be trivially extended to the case of three or more tests:

$$\begin{aligned} P_{123} &= P \cdot \left(1 - \ln P + \frac{\ln^2 P}{2} \right), \quad P = P_1 \cdot P_2 \cdot P_3 \\ P_{1234} &= P \cdot \left(1 - \ln P + \frac{\ln^2 P}{2} - \frac{\ln^3 P}{6} \right), \quad P = P_1 \cdot P_2 \cdot P_3 \cdot P_4 \\ P_N &= P \cdot \sum_{i=0}^{N-1} \frac{(-\ln P)^i}{i!}, \quad P = \prod_{i=1}^N P_i \end{aligned} \quad (3.3)$$

3.2 Circular Statistics

There is a wide field of research devoted to the statistics of quantities which are repetitive, periodic, or behave like angles; see references [2,61] for introductions. I will only outline the fundamentals and one result which will be employed in section 4.4.

Circular statistics is designed for dealing with situations where measurements are indeterminate modulo an integer multiple of some value. For example, angles are indeterminate modulo 360° , phases are indeterminate modulo 2π or modulo 1 depending on the convention used, the timing of periodic events is indeterminate modulo the period. The term indeterminate here means for example that an angle of 370° is identical to an angle of 10° or a pulsar orbital phase of 1.6 is the same position as phase 0.6. For simplicity I will express everything in terms of angles between 0° and 360° in this section; also the results can trivially be rewritten in terms of quantities of period 2π , period 1, or any other.

Obviously the usual methods of statistics have to be modified when dealing with circular data; the most trivial example to show that is: the mean value of the three angles 359° , 0° and 1° is not 120° but instead 0° , although it seems that $120^\circ = (359^\circ + 0^\circ + 1^\circ)/3$. I will ignore all the simpler questions of circular statistics, and proceed to the following problem.

How can one find out whether a circular distribution of measured values is symmetrical or not? Assume one has measured N values x_i , and each measurement has been done at a certain angle ϕ_i . As an example, say we measure blood pressure x_i at many times ϕ_i (take ϕ as being the time-of-day, expressed in degrees instead of hours); similar examples involve birds flying in directions ϕ_i for a distance x_i . The question we are interested in is: does blood pressure vary throughout the day, or putting it more precisely: is blood pressure correlated to the time of day. Note that the measurements may be taken at any time (for example whenever a patient walks into our medical office), so the distribution of ϕ may be very non-uniform.

A simple but powerful method of testing that is the Rayleigh test.⁷⁹ Think of each measurement as a two-dimensional vector \vec{v}_i of length x_i and pointing in direction ϕ_i ,

which means its components are $(x_i \cos \phi_i, x_i \sin \phi_i)$. Form the sum of all those vectors $\vec{V} = \sum_i \vec{v}_i$. The test statistic to use is the Rayleigh-power z , defined as

$$z = \frac{1}{N} |\vec{V}|^2 = \frac{1}{N} \left\{ \left(\sum_i x_i \cos \phi_i \right)^2 + \left(\sum_i x_i \sin \phi_i \right)^2 \right\}. \quad (3.4)$$

It can then be argued⁸⁸ that the expected probability distribution of z should be just

$$\frac{dn_z}{dz} = e^{-z}, \quad (3.5)$$

assuming the x_i are normally distributed (zero mean, unit variance). The corresponding confidence levels (complementary cumulative distribution function) is also just e^{-z} . This simple result already shows the use of this Rayleigh test: the chance probability of observing a value of the test statistic larger than z is e^{-z} ; so if one observes a large value of z , the confidence level at which one can refute the null hypothesis is $1 - e^{-z}$. The null hypothesis in this case is that the x_i are not biased towards one side of the circle.

Scargle⁸⁸ has improved this test. He introduces an auxiliary angle τ defined as

$$\tan 2\tau = \frac{\sum_i \sin 2\phi_i}{\sum_i \cos 2\phi_i}$$

and replaces equation (3.4) with

$$z = \frac{1}{2} \left\{ \frac{(\sum_i x_i \cos(\phi_i - \tau))^2}{\sum_i \cos^2(\phi_i - \tau)} + \frac{(\sum_i x_i \sin(\phi_i - \tau))^2}{\sum_i \sin^2(\phi_i - \tau)} \right\}$$

Although this expression looks rather more complicated, the actual values of z calculated from it usually don't differ significantly from those calculated using (3.4), except for pathological cases. Calculating it takes only about twice as much more computer effort: the additional pass through the data to calculate τ , and then having to keep four instead of two sums in memory; the number of trigonometric functions to be calculate stays the same.

The advantage of this method is that in this case, equation (3.5) for the expected distribution of the Rayleigh-power can be analytically proven. One has to assume that the x_i are distributed normally. The quantities to which we will apply this test in section 4.4, δ_i and κ_i , have been crafted to fulfill this requirement under the assumption that the muon

is correctly reconstructed; κ_i by definition is normally distributed, and δ_i is constructed to have zero mean and unit variance. Note that the result is completely independent of the distribution of the angles ϕ_i . However, the power of the test depends on the distribution of ϕ and decreases if that distribution is very irregular; for example if all the ϕ_i are identical, z will be the variance of the the x_i , or by definition of order one.

Summarizing, this algorithm enables one to do a quick test to determine whether the distribution of some quantity x , which is measured at angles ϕ_i is symmetrical over the angle or not, and it immediately yields the chance probability of the observed distribution being random.

Chapter 4

Analysis Methods

4.1 Conventions Used

The full geometric description of a single non-showering particle traversing the whole detector at the speed of light requires five parameters: the time t_0 at which it enters the detector, the position \vec{v} where it enters the detector, called the vertex here, and the direction of its travel \hat{d} . The direction is a unit vector, so it counts as two parameters. I will alternate between expressing it as the unit vector \hat{d} and as two angles (ϑ, φ) . The vertex position also counts as two independent parameters, since it is restricted to be on the surface of the detector, a rectangular box. These five parameters completely describe the movement of a throughgoing particle; I will in the following often refer to such a description as a “fit” (the term comes from the iterative fitting procedure usually required to arrive at such a description).

This geometrical description of events is complete only for throughgoing events; for stopping entering muons one additional parameter is required, for example the pathlength they travel from the entry point. For the complete description of contained events another additional parameter is required; for example one just removes the restriction that \vec{v} has to be on a detector wall, so it requires three independent parameters.

The coordinate system is shown in figure 4.1. The nominal center of the detector defines the origin. The six surfaces of the detector are defined by the planes in which the photomultipliers lie, and are at coordinates $x = \pm 1148$ cm, $y = \pm 842$ cm and $z = \pm 874$ cm. The coordinate system is aligned with the geographic directions as shown in figure 4.1; the procedure has been described in section 2, and is accurate to at least 1° . The absolute

an integer; but in practice the pulseheight resolution of the photomultipliers used in the IMB experiments does not allow distinguishing these integer levels. The calibration procedure, described for example in reference [16], determines the conversion factor from observed pulseheight back to the number of PEs; it employs the properties of the poisson-statistic of PEs; therefore after calibrating the data, light levels are measured in units of photo-electrons. In principle one could calculate an absolute illumination (in units of photons) from that, using the (more or less known) absolute quantum efficiency of the photomultiplier assembly; in practice, such a conversion is fortunately not necessary.

4.2 Data Selection

As shown above, the IMB detector gathers about $2.3 \cdot 10^5$ events per day, of which at least 10% are written to tape. The number of upward-going muons events expected is on the order of one every two days; so a data reduction by a factor of about one-half million is necessary. I will describe the procedure used for IMB-3 data; for IMB-1 and IMB-2, see reference [93].

The first procedure which has to be done to events after they are read is to *calibrate* them. The digitizers which transform the time and pulseheight of a photomultiplier hit into digital information are neither completely linear, nor do they all behave identically. They also introduce correlations between pulseheight and timing, known as “slewing”, see section 2.1. All these problems have long been solved by the collaboration and are explained in great detail in reference [16]. In this work I use the calibration procedures and constants of the west-coast analysis, analogous to the ones described in reference [93]. The resolution of timing and pulseheight information after calibration will be discussed in section 4.3.

For efficiency reasons the data reduction is done in two steps.

The first step, also known as “pass 0” is designed for speed. It first rejects events which are “obviously” not upcoming muons: events with more than 1900 photomultiplier hits cannot be caused by single muons, and events below 400 photomultiplier hits are rejected as being too small to be accurately reconstructed. This cut is set at 250 photomultiplier

hits for IMB-1, and its effect will be discussed in section 4.6. It also rejects events with more than 27% of all hits in the bottom of the detector; these are typically downgoing muons, where the Čerenkov light cone lights up large fractions of the detector bottom. It then performs a simple entry-exit point fit (see section 4.4 for an explanation of fitting techniques), and rejects events with a zenith angle of less than 84° , which are downgoing muons. To increase the speed of this step, only a partial calibration is used, and nearly all processing is done on integer (instead of floating point) numbers; it typically reduces the data by a factor of 500 to 1000, using typically 3.5 hours of 11/780 equivalent CPU time for each day of detector live time.

The second step, known as “pass 1” executes a more elaborate fitting procedure, and reduces the data by another factor of 10. Computer usage is not a problem at this step any longer.

The philosophy in both automated processing steps is to save any event which is (even marginally) consistent with being a muon, and which is not clearly downgoing; both programs will save events a few degrees above the horizon (downgoing), to prevent losing too many upward-going events due to the fit resolution.

At this stage the events have been reduced to about one event per hour of live time. Further reduction by computer program would be risky; the danger of losing many good events in the process is too high. Instead, human “scanners” manually inspect each event, and decide whether it is an upcoming muon. The details of how a human will usually fit an event is explained below in section 4.4. Most events which have made it to this stage by confusing the fitting procedures can be rejected very rapidly with some experience; the lion’s share of the events are just a handful topologies which the scanner will immediately recognize. The other events are carefully hand-fitted; whether the fit is upward- or downward-going decides the fate of the event.

For further safety reasons, the manual scanning is also separated into two steps. In principle, the first pass should be a rough scan, performed for example by undergraduate student help; they are instructed to save any events which are even slightly suspect of being upward-going muons. In the second pass a “physicist” re-checks the events accepted by the scanners, and reaches the final verdict. Furthermore, the “physicist”

(in this case your truly) should recheck a significant fraction of the events the scanners rejected. In the case of the analysis at UH, the distinction between “scanner” and “physicist” has become blurred. A large part (nearly half) of the data processed here has been scanned by Gary McGrath, who is a graduate student working on this experiment himself; a small fraction was done by Cheryl Clem, who truly was an undergraduate student helper (and the wife of a physicist); the balance (about half the data) I scanned myself, plus all the physicist re-scanning and check-scanning. To prevent that process from becoming too incestuous, I usually allowed at least a few days to pass between scanning a file, re-scanning the accepted events, and re-checking the rejected events.

There are two more categories of events which require special care; both are problems in deciding whether either end of the muon track is inside the detector. *Stoppers* are low-energy muons which enter the detector and then stop. If the muon stops sufficiently far before it would exit it is easily distinguished from through-going events: what would usually be a Čerenkov light cone has deteriorated into a wide ring. Figure 4.3 is an example of a stopping event. The detector is shown in an isometric view, the top of the detector is at the top of the figure. Each little line or star symbolizes a photomultiplier hit; the number of cross-hatches in each little star indicates the pulseheight of the hit, or more accurately the number of photoelectrons in the hit scaled down by a factor of 3; so a single line is a hit below 3 photoelectrons, a cross is a hit between 3 and 6 photoelectrons and so on. The color of each little star would show the timing, except in this printed version the color is not visible. The outline of the Čerenkov light cone is drawn as a continuous line. The muon enters on the left, and would have left through the back wall if it hadn't run out of momentum first. Obvious are the large pulseheight hits close to the entry point, and also the lack of large hits in the center of the Čerenkov light cone. For comparison, figure 4.2 shows a throughgoing event. The photomultipliers at the entry point are in this view nearly exactly in front of the large pulseheight hits at the exit point, and are hard to see. The difference between the exit points in the top and bottom display is clearly visible.

More difficult are events which do not enter the detector, but the muon is created in the detector. If the vertex is clearly inside the detector (several meters inwards from the walls) the difference in event topology is unmistakable; the problem arises when the

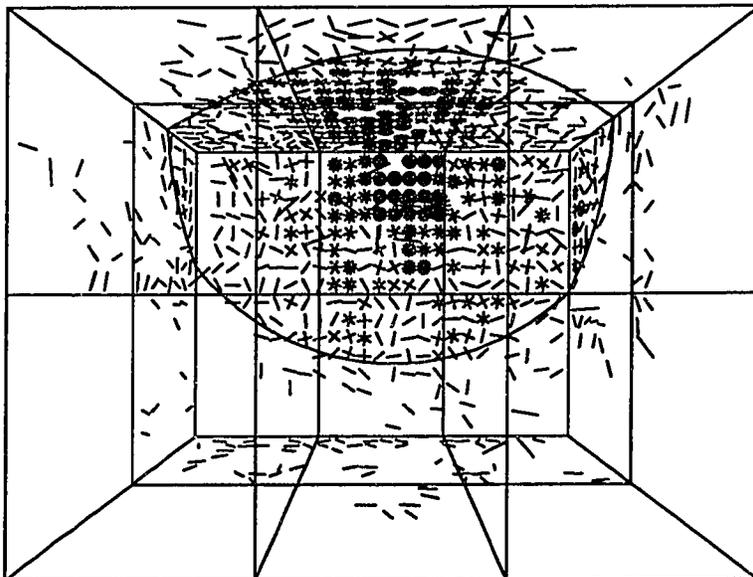


Figure 4.2: Display of a through-going muon event, number 2842/3716.

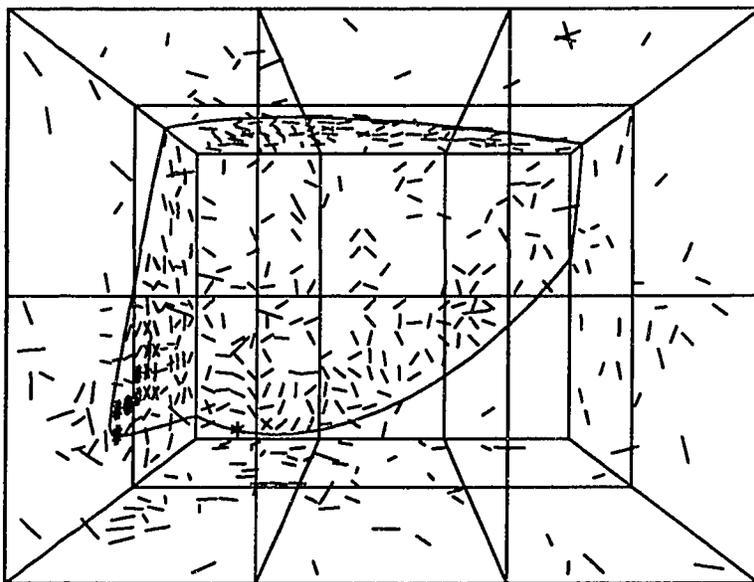


Figure 4.3: Display of a stopping muon event, number 2844/29558.

vertex is close to the wall. Figure 4.4 shows such an event. The Čerenkov light cone is drawn for the correct fit assuming the event entered the detector. There is clearly no entry point with high pulseheights; still, the largest pulses occur where the Čerenkov light cone comes closest to the wall on the right. The muon fit was done assuming that the muon enters the detector by minimizing the multipole-moments of the timing residuals. What gives away that this muon did not enter the detector is that the Čerenkov light does not extend all the way to the edge of the cone, instead is always far inside it. In the case of this event, there would be no fit such that the outline of the observed light matches the Čerenkov light cone expected for a muon.

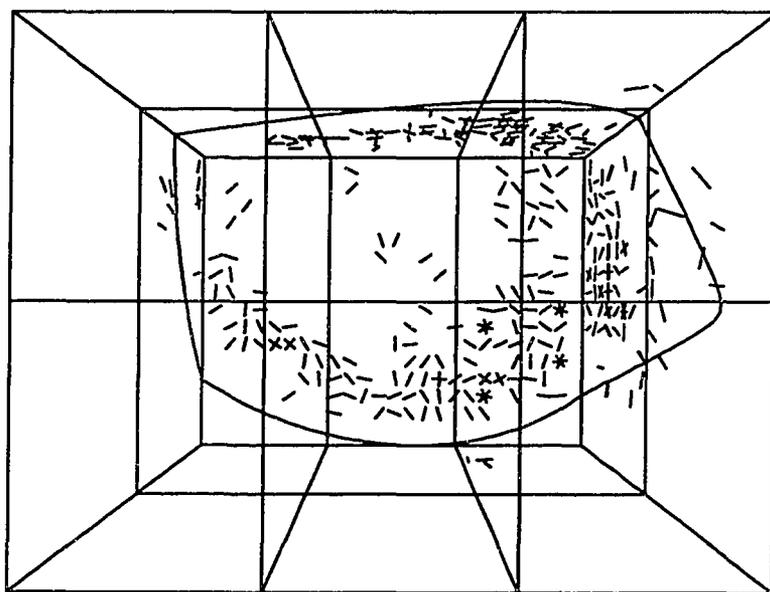


Figure 4.4: Display of a contained stopping event, number 3962/51375.

For data tapes which have been processed by more than one group, the resulting events are compared. The live time for each tape is also determined by more than one group; usually the west-coast analysis uses the livetime clock scaler of the detector and the east-coast analysis the detector logbooks. The at least two resulting live times are also compared. Disagreement about the events found or about livetime is ironed out; it is often caused to data which has to be discarded due to problems which are found in the comparison.

4.2.1 Data Set, Livetime

The raw data which has been processed as described above consists of several distinct periods, with the breaks between these periods defined by changes in the mode the detector was operated or the data was analyzed:

IMB-1 was the first stage of the detector, using 5" EMI photomultipliers, and a PDP 11/34 as the online computer. It was operated from October 1982 through May 1984, and corresponds to data tape numbers 1 through 1537.

IMB-2 added wavelength-shifter plates to the existing 5" EMI photomultipliers, and operated through February 1985, when the increasing failure rate of the photomultipliers required them to be replaced.

The upward-going muon data from IMB-1 and IMB-2 was analyzed by Bob Svoboda, and the results reported in reference [93]. The total live time for the IMB-1 and IMB-2 data used here is 455 days, and there are 200 events from this period.

IMB-3 uses 8" Hamamatsu photomultipliers, and initially a VAX 11/750 for the online computer. It began operating in April 1986 with tape number 2023.

Its data has to be divided into finer parts; the data-taking modes are explained above in section 2.2.

- Before tape number 2400, that is before November 1986, the upward-going muon data was analyzed by Bob Svoboda at UC Irvine. This data comes from raw data tapes, usually taken in cut mode.
- From tape number 2400 on, all the data has been analyzed at UH (but a large fraction of it has also been processed and/or scanned at Boston University, Louisiana State or UC Irvine).
- Around data tape number 2701 in April 1987 the pulseheight digitizers were modified. Before then, the gain was adjusted such that the maximum pulseheight they could record was about 20 photoelectrons. To better employ the larger dynamic range of the 8" photomultipliers the digitizer gain was reduced and the detector re-calibrated; note that the photomultiplier setup and high-voltage supply itself were not modified, only the conversion constant of the digitizer. This process took

several weeks (one resistor had to be replaced on each and every digitizer board) but was completed before tape number 2701. As will be seen in section 4.3 this makes a significant difference in pulseheight response and resolution.

- Before tape number 3185 upward-going muon event candidates were not saved by the data analysis computer at the experiment. Until that period the raw data tapes have been processed, most of them taken in cut mode.
- Since tape number 3185, that is since January 1987, the data analysis computer saves upward-going muon candidate events on its output tape. Since that time mostly microVAX tapes have been processed; only where those are damaged the raw tapes have been used.
- At the end of 1988, or equivalently after tape number 3634, our spineless collaborators from the University of Michigan left the group, and Shige Matsuno became the resident physicist at the experiment. The quality of the data and the live time fraction increased markedly; other than that the data-taking mode was not changed.
- In May 1989 the 6250 bpi reel-to-reel tapes were replaced with 8mm cartridge tapes. From now on, the data is written only by the data analysis computer; there is no distinction between raw and microVAX tapes any more. At the same time, some of the programs on the data analysis computer were modified. Since the microVAX II was not fast enough to handle an event rate of 2.7 Hz, data was still being taken in cut mode. As will be shown below, the data from this period has a problem, most likely in the live time estimate.
- In February 1990, beginning at tape number 4014, the microVAX II was replaced by a VAXstation 3200. From now on most of the data is written in uncut mode.
- At tape number 4196 in November 1990 the ailing VAX 11/750 was replaced by the microVAX II. Other than that no change in data taking operations took place.
- The data used in this work extends to the end of 1990, that is up to and including tape number 4216.
- In March 1991, the IMB-3 detector suffered a catastrophic leak; the last data tape is number 4277. The data from 1991 is not included in this work.

Naturally while analyzing the data and before assembling the events for this work, all parts of the data have been checked. One of the most sensitive methods of finding problems with the raw data is just care when processing: bad data will usually show up as a lack or an excess of events found at the various steps in the processing chain. The amount of CPU time used, and the number of input-output operations performed by the processing programs is also checked; often bad data is easily found by extraneous values of these simple parameters. Even more sensitive in rejecting bad data is the human scanner; often when the data seems suspicious when scanning (too many or too few events to scan for this live time, or for example too many downgoing muons in a certain configuration), and in nearly all cases the intuitive judgement of the scanner is found to be based on some hardware or software problem. Furthermore, for data tapes which have been processed by more than one institution in the collaboration, a "silver bullet" philosophy is used: to declare a tape as bad it is sufficient if one group finds a problem with it (except obvious technical failures, for example the east-coast copy of the tape is physically damaged). Of the 2204 IMB-3 tapes (including calibration and test tapes) analyzed for this work, the most frequent causes for rejecting a tape are

- Patch timing problem, that is the relative timing of a group of 8×8 photomultipliers is shifted with respect to the master readout pulse, invalidated 127 tapes. Unfortunately a flaky connection of the master readout pulse to the backplane of one of the 16 digitizer crate (2 patches) was not repaired for about two years, and the intermittent problem was found in quite a few tapes.
- Electronics problems in the readout, 68 tapes.
- High-voltage supply failure, 27 tapes.
- Flashers, that is a photomultiplier which develops a small leak and turns from a light detector into a light emitter, 10 tapes. Note that the search for throughgoing muons is much less sensitive to low-level flashers than the search for contained events, since events with little light are rejected immediately. Only when a flasher overwhelms the data recording system does the tape become useless for the purposes of this work.
- Tape is unreadable (media damaged), 9 tapes.

Of the remaining 1505 data tapes without problems which were processed, 1129 were done at UH, the rest at other institutions. About 40% of the data processed at UH was also processed by another group (UC Irvine, Boston University and Louisiana State). For all the tapes which were processed by more than one group, the results were compared; the overlap between the events is usually very good. For example, of the 130 events from data tape numbers 3907 through 4100, all of which were processed at UH and at LSU, 15 events were missed at UH. Since not all data has been rigorously processed by two independent analysis groups, one cannot estimate the processing and scanning efficiency from comparison of two groups; instead the Monte-Carlo will be used to estimate the efficiencies in section 4.6.

The live times, numbers of events, and event rate per live day are summarized in table 4.1. It is clearly visible that the upward-going muon event rate increased from IMB-1 and -2 to IMB-3; this is due to the increase in effective area shown in section 4.6, which leaves the flux of muons essentially unchanged, see section 6.1. The last three lines of the table show the IMB-3 data, split into the era before 8mm tape drives, and then into cut and uncut data on 8mm tapes. If one assumes the data rate before 8mm tapes to be correct, the two data sets on 8mm tape show only 77% and 90% of the rate expected; the deficit on cut 8mm tapes is statistically significant at the 2.1σ level, whereas the deficit on the later uncut data is not statistically significant with only 0.9σ . Figure 4.5 shows the rate of upward-going muon events as a function of data tape number; each data point corresponds to 500 hr live time. Note that the data points are spaced 250 hr apart, so they overlap. The dotted line is the rate from reel tapes only. The bottom figure shows the latest data enlarged. The beginning of the cut and uncut 8mm tape periods are marked by arrows. A drop in event rate after tape number 3900 is visible; whether it returns to normal after tape number 4014 is unclear.

The events from these three periods have been compared extensively, without finding any significant difference between them or their distributions. All imaginable causes* for such a deficit have been checked; no explanation has been found. The most likely cause is either a problem in the live time determination, and/or an inefficiency in the human scanning. Due to this as yet unsolved problem, the data from the two 8mm tape periods

*At least the ones the collaborators could imagine.

Period	Tape number	Live hours	No. Evt.	Event rate per day	Comments
All	0-4216	30341	640	0.5062 ± 0.0200	
IMB-3	2023-4216	19421	440	0.5437 ± 0.0259	
Reel tape	2023-3907	11003	268	0.5846 ± 0.0357	Standard rate
8mm cut	3907-4013	4048	76	0.4506 ± 0.0517	77%, 2.1σ low
8mm uncut	4014-4216	4370	96	0.5272 ± 0.0538	90%, 0.9σ low

Table 4.1: Summary of live time and event rate

has not been used to calculate the total flux of muons from atmospheric neutrinos in section 6.1; however it is used for all the muon astronomy work of sections 6.2, 6.3, 6.4 and 6.5, since none of those are sensitive to the absolute normalization.

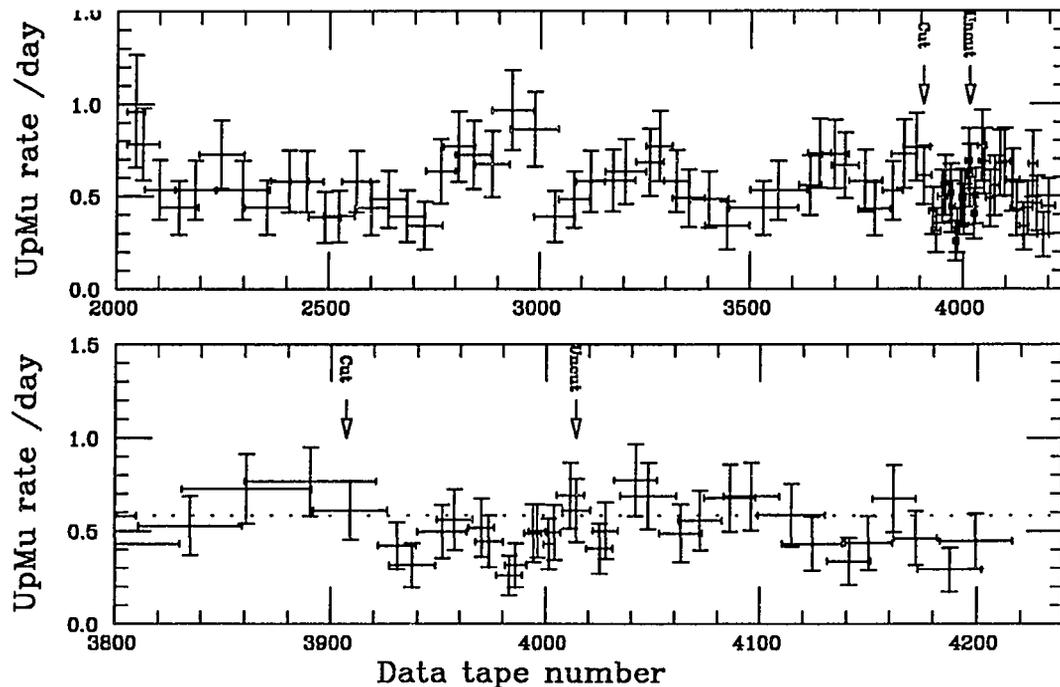


Figure 4.5: Event rate versus data tape number for IMB-3. Top figure shows all IMB-3 data, bottom figure is enlarged to show the 8mm tape data in more detail.

With this caveat, the data set used in this work can be summarized as follows: From the beginning of the IMB experiment through the end of 1990, there are a total of 640 events upward-going muon events, of which 496 are throughgoing and 144 stop in the detector.

The total live time is $1264.2 \text{ days} = 3.46 \text{ years}$. If one excludes the data from 8mm tapes (which is done for the calculation of total atmospheric neutrino flux) that reduces to 468 events in a live time of $913.5 \text{ days} = 2.50 \text{ years}$.

4.3 Timing and Pulse Height Resolution

For the fitting procedure outlined in section 4.4, knowledge of the resolution of the timing (known as t) and pulseheight (known as q) information is required. The usual procedure for fitting an event is the following:

1. Assume an initial fit.
2. For this fit, calculate the expected timing and pulseheight of all photomultiplier hits.
3. Compare the measured hits to the expected ones. Such a comparison can for example be done by subtracting the expected time from the measured time, resulting in the *timing residuals*, or by estimating the probability of measuring a certain pulseheight, given a certain expected pulseheight. Eventually all the differences should be expressed in some statistically easily manageable form, in this case in standard deviations.
4. Merge all the deviations of timing and pulseheight between measured and expected values into one estimator of fit goodness; let me call that number the *fit quality*.
5. Vary the fit position and direction in some fashion, maximizing the fit quality, until some sort of satisfaction is achieved.

In this section, I will summarize the information required for steps 2 and 3: whether the detector timing still depends on pulseheight, what the timing resolution is, how the expected pulseheight depends on the relative geometry of muon and photomultiplier, and what the pulseheight resolution is.

To measure these resolutions a set of standard events is required. In principle, if the detector simulation monte-carlo were of sufficient quality, it would require knowledge of all these parameters, and the resolutions could just be taken from the monte-carlo programs. However, it isn't of sufficient quality. Therefore I measure these resolutions

using a set of well-reconstructed downgoing muons, which I will call “calibration muons”. Clearly a large enough sample of calibration muons is required to achieve reasonable statistics in the measurement of resolutions; also, the calibration events encompass the same variety of event geometries and types as the data under study in this work, so as to ensure that for each possible data event there is a calibration. Furthermore, the events used here have to be selected to only include single throughgoing muons (since that the class of events these calibrations are going to be used on), and have to have good fits (since both t - and q -information depend on the relative geometry of muon and photomultiplier).

Such a set is required for each of the six periods of substantially different detector operation, which are IMB-1, IMB-2, IMB-3 separately before and after the Q-scale was extended (see section 4.2.1 for a discussion of that change), and the detector simulation for IMB-1 and for IMB-3.

The set of calibration events used here consists of:

- For IMB-1, 106 downgoing muon events from tape number 785, which were hand-scanned and hand-fitted at UC Irvine.
- For IMB-2, 110 similar events from tape number 1749.
- For IMB-3, 810 downgoing muon events, selected and hand-fit by me, from several data tapes in the 2400 through 2900 range. Of those, only 70 are from the data with the extended Q-scale.
- For IMB-3, downgoing muons gathered at the end of each calibration run, and automatically selected and fitted at Boston University. These muons are selected to not enter or leave the detector within 2m of any edge, so they do not include certain geometries (for example corner clippers, see section 4.4). There are 33 such calibration tapes used here, each with 400 to 500 muons.

Since the statistics of this data is fairly low for IMB-1 and IMB-2, the upward-going muon data set was also employed for the study of Q-resolution. Together, these data sets have sufficient statistics to establish resolution calibration for all muon geometries of interest here.

Since the monte-carlo simulated data has to be processed through the same data analysis as the real data, a similar calibration was done for it. For IMB-1 and IMB-2 (they are treated together as far as the monte-carlo is concerned) I use the monte-carlo detector simulation program written by Woiciech Gajewski of UC Irvine; for IMB-3 I use Dave Casper’s excellent program written at Boston University. Since the detector-simulation takes a lot of time, this calibrations were done using the same sample of simulated upward-going muon events used in section 4.6; this means that the data is calibrated using mostly downward-going muons, whereas the monte-carlo is calibrated using upward-going muons. This difference will be discussed and compensated for in section 4.5. For the calibration of the monte-carlo, the reconstructed muon fit is used as a reference, not the “true” generated muon position and direction (which is known for monte-carlo, obviously not for the data). The reason is that when using the true muon fit instead of the (slightly uncertain) reconstructed muon fit, the geometry of the muon is too well known, and the resolution of timing and pulseheight will be slightly underestimated, which leads to problems when comparing data and monte-carlo.

The final number of calibration events for each of the six separate data sets is shown in table 4.2. The abbreviation “IMB-3 loQ” is used here (and throughout this chapter) for the data before the extended Q-scale, and analogous “IMB-3 hiQ” for the data after the extended Q-scale.

Calibration Period	No. calibration events	
	t-calibration	q-calibration
IMB-1	106	268
IMB-2	110	129
IMB-3 loQ	5292	5519
IMB-3 hiQ	12907	11938
IMB-1 MC	2724	2595
IMB-3 MC	4126	4117

Table 4.2: Number of events used for timing and pulseheight resolution calibration

4.3.1 Timing Resolution

As will be shown in section 4.4, the most important quantity used in fitting an event are the *timing residuals*, which are the difference between measured and expected time of each hit. The fitting procedure requires knowledge of the resolution of these residuals, so the statistical significance of any deviation between measured and expected time can be quantified.

As has already been pointed out in section 2.1 the timing resolution should improve for higher pulseheights, until at very high pulseheights pre-pulsing sets in. Furthermore, the measured (uncalibrated) timing tends to be early for high pulseheights (slewing). The calibration should take care of that, but it may not have done so perfectly.

The timing resolution is also asymmetrical. Due to causality light cannot arrive at the photomultiplier before it should, but it can arrive substantially delayed (either as scattered light, or delayed in the wavelength-shifter plate). To account for this, I will measure the timing resolution separately for early light (observed before it was expected, equivalent to negative timing residuals) and late light (positive timing residuals).

To measure the timing shift (slewing) and the negative (early) and positive (late) resolution, all as a function of pulseheight, the following method is used: for each event, calculate all the timing residuals, using the best-fit track. Bin the photomultiplier hits as a function of their observed pulseheights, on a pseudo-logarithmic scale (the bin limits are 0, 1, 2, 3, 4, 5, 6, 8, 10, 12, 15, 18, 21, 25, 30, 35, 40, 50, and above, all measured in photo-electrons). For each bin in pulseheight, the mean timing residual is calculated in a first pass. In a second pass, the mean timing residuals is subtracted from the measured residuals, and the the RMS timing residuals are calculated separately for hits with positive and negative corrected timing residuals. Figure 4.6 shows the results, separately for IMB-1, IMB-2, IMB-3 (both before and after extended Q-scale), and the monte-carlo simulation for IMB-1 and IMB-2 (using the Irvine simulation program, labelled "IMB-1 MC") and for IMB-3 (using the Boston simulation program, labelled "IMB-3 MC"). Note that the statistics are low at high pulseheights for the first three.

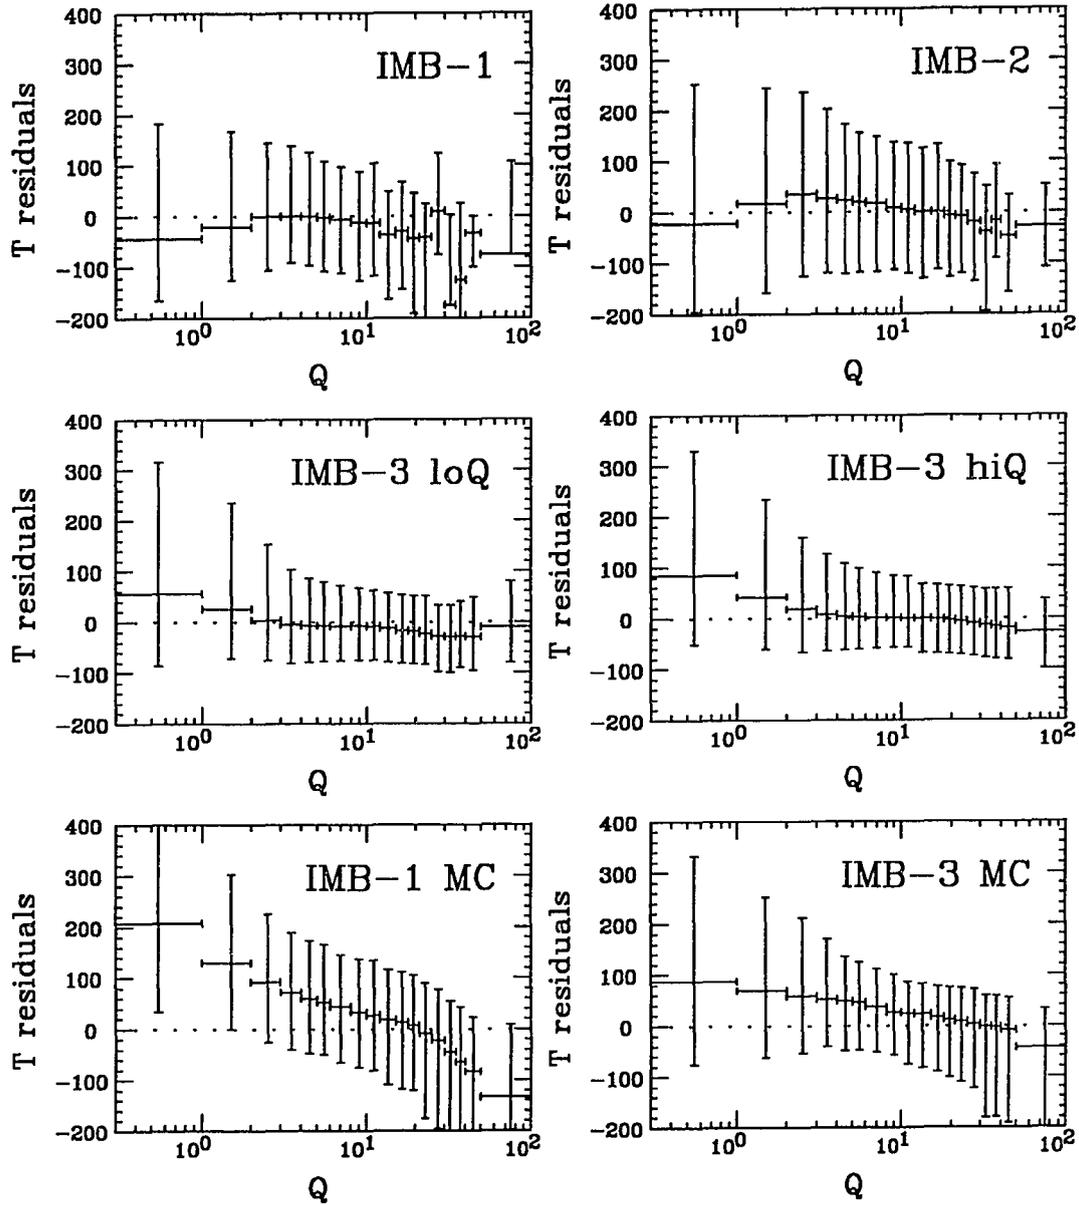


Figure 4.6: Timing residuals: mean value and early/late resolution, as a function of pulseheight, separately for the four phases of the detector and the monte-carlo simulations. The horizontal lines are the mean residuals, their width indicates the pulseheight bins; the vertical lines show the RMS resolution.

Several features are visible. The horizontal lines show the position of the mean timing residuals; their width shows the width of the bins in pulseheight. With the exception of the monte-carlo simulation, the mean values are fairly small, nearly all within 40 cm* which corresponds to less than 2 ns. This shows that the calibration procedure has already done the slewing correction very well. However, this is not the case for the monte-carlo simulated data; that doesn't cause a problem, since the procedure outlined above acts as a pulseheight-dependent slewing correction.

The two vertical lines (which look like error bars, but aren't) show the RMS resolution (the sigma) of the residuals. The positive resolution (late light) is larger than the negative one (early light), in particular at low pulseheights, exactly as predicted above. This may indicate that a fair fraction of the hits registered at low pulseheights are due to scattered light. Also note how the timing resolution worsened from IMB-1 to IMB-2, due to the addition of the wavelength-shifter plates. Clearly visible is how the resolution improves as the pulseheight increases; in particular, the positive (late light) resolution improves markedly, since at higher pulseheights the influence of both late scattered light and light delayed by the wavelength-shifter decreases. Prepulping at very high pulseheights is not visible; note that it cannot be found in the IMB-1 and IMB-2 samples, since the measured pulseheights are limited.

In summary, the slewing of the photomultiplier hits can be corrected using the data from this section (not that there is much to begin with), and the RMS timing resolution ranges from typically (+12 ns, -6 ns) at 1 PE to ± 3.5 ns at above 10 PE. Using these tabulated values, one can calculate an approximately normally distributed variable δ_i , which expresses the timing residuals for photomultiplier hit i in a statistically meaningful way:

$$\delta_i = \begin{cases} (t_i - t_i^e - m(q))/\sigma_+(q), & t_i - t_i^e > m(q) \\ (t_i - t_i^e - m(q))/\sigma_-(q), & t_i - t_i^e < m(q) \end{cases} \quad (4.1)$$

where t_i is the measured time of the hit, t_i^e the expected time, q the measured pulseheight, $m(q)$ the mean residuals and $\sigma_{\pm}(q)$ the two residual resolutions (for late and early light) as a function of pulseheight, as determined above.

*Remember, time is measured in units of light travel, see section 4.1.

4.3.2 Pulseheight Resolution

As was already pointed out in section 2.1, in principle one should be able to predict from first principles and the known properties of the photomultiplier the observed pulseheight for any situation; however, any attempt to do so has failed, since the properties of the photomultiplier and wavelength-shifter assembly are not well enough known. On the other hand, knowledge of what pulseheight to expect for a certain muon is necessary for the fitting procedure. For the expected timing and timing resolution shown above the situation is much easier; the expected time of a hit can be easily calculated, and is (as shown above) systematically very close to the measured time. For the pulseheight, I will in the following use a different approach, which involves just recording the expected distribution of pulseheights for each possible geometry, and in the fitting procedure simply compare the measured pulseheight to the distribution of expected values.

To begin with, the expected pulseheight depends mostly on three factors:

- The distance the Čerenkov photon has to travel from the point it was emitted by the muon to the photomultiplier assembly. I express this distance in the distance d of closest approach between muon and photomultiplier; the actual distance the photon has travelled is larger by a factor of $1/\sin \vartheta_c$). The reasons for the decrease of pulseheight as d increases were already outlined in section 2.1: the photons are attenuated by the water, with the attenuation coefficient depending on the wavelength, and they are diluted by the $1/r$ phase space increase.
- The angle ζ ; at which the photons reach the photomultiplier assembly; it is measured from the normal on the photomultiplier assembly, with 0° corresponding to a photon reaching the photomultiplier frontally, 90° to glancing the wavelength-shifter plate, and 180° to striking the photomultiplier housing from the back. The overall behaviour of this factor is easily explained. First of all, the wavelength-shifter plate has a projected area of $A|\cos \zeta|$, which clearly drops to zero at glancing incidence. The front half of the photomultiplier is close to a hemisphere and therefore fairly independent of angle; in the backwards direction the photocathode is mostly shadowed by the photomultiplier housing.

- The energy of the muon. The amount of light emitted by a muon depends logarithmically on its energy; the main reason is that at higher energies the muon is accompanied by a cloud of particles, some of which can radiate Čerenkov light. The dominant mechanisms for that increase are knock-on electrons, production of e^+e^- pairs by Weizsaecker-Williams photons accompanying the muon, and creation of photons by bremsstrahlung off a nucleus; in either case the electromagnetic cascade created from the additional particles can radiate Čerenkov light. The light output of a muon also fluctuates heavily, since some of these three processes are rare compared to the length of the detector, but they contribute heavily to the total light output. The amount of light due to these three processes has different energy dependence, and they also differ in how much they contribute to the fluctuations. This effect will be employed in section 4.5 to estimate the energy of a muon from its light output, but at this point, when trying to understand the behaviour of the “standard” muon it adds a complication.

The procedure used to remove the difference in total light output between muons is to normalize the measured pulseheights q_i to the ones expected for the “standard” muon, which is taken to be the average of the calibration muons. In more detail this is done in the following way: for each functioning photomultiplier in the Čerenkov light cone (whether it has been hit or not), separately sum up the observed pulseheight q_i , and the pulseheights expected for this geometry $\bar{q}(d_i, \zeta_i)$. Not included in these sums are photomultipliers which are not functional (these “dead” photomultipliers are marked as such in the calibration file), and photomultipliers which have registered hits before any light is expected to hit them (more accurately, photomultipliers for which the scaled timing residual $\delta_i < -3\sigma$), which are assumed to be due to random firing before the light from the muon reached them; since the electronics can only register one hit per photomultiplier they are not available to register real hits. Similarly, hits with $\delta_i > 3\sigma$ are assumed to be due to late random or scattered light; in this case the expected pulseheight is summed but not the observed pulseheight, since the photomultiplier should have seen that amount of light, but didn’t. Note that for photomultipliers which were not hit the expected amount of light is still summed up. The ratio of total observed pulseheight $\sum_i q_i$ to total expected pulseheight $\sum_i \bar{q}(d_i, \zeta_i)$ is the relative light output ε ; its use will

be discussed in greater detail in section 4.5 below. The term *relative* indicates that this quantity is measured relative to the average light output of the calibration muons (which were used to measure the expected pulseheight).

This relative light output ε is then used to correct the differences in total light output between events by simply scaling all observed pulseheights for each event by its inverse: $q'_i = q_i/\varepsilon$. Using all the calibration events, the distribution of scaled pulseheight can now be determined for each geometry, that is as a function of d and ζ ; this is done in 36 bins in impact angle ζ from 0° through 180° and in 42 steps in impact distance d between photomultiplier and muon, from 0 m through 21 m. From the calibration events three quantities are determined: the mean scaled pulseheight expected for this geometry $\bar{q}'(d, \zeta)$, the probability of observing a hit in this geometry $p_h(d, \zeta)$, and the distribution of observed scaled pulseheights $D(q', d, \zeta)$, normalized such that $\sum_{q'=1}^{\infty} D(q', d, \zeta) = 1$.

At this point there is a little problem: to determine these means and distributions of *scaled* pulseheights one has to first determine the scaling factor ε ; but to determine that knowledge of the mean pulseheights $\bar{q}'(d, \zeta)$ is already required. Breaking out of this circle is accomplished by simply assuming an initial value of $\bar{q}'(d, \zeta)$, and iterating the procedure which determines them. It is found that the values become stable already after two iterations, and don't change significantly with more iterations. The mean value of ε for the calibration muons should come out to be one after the iterations, since ε is defined relative to those events, and indeed after two iterations it is unity within the statistical errors.

Given these distributions and probabilities (there are $36 \times 42 = 1512$ of them) the probability of observing any pulseheight or the absence thereof can be determined: if the photomultiplier was not hit, the probability of observing this non-hit is $p_i = 1 - p_h(d, \zeta)$; otherwise the probability of observing such a large pulseheight is $p_i = \sum_{q'=1}^{q_i} D(q', d_i, \zeta_i)$. These probabilities can then be turned into the equivalent number of standard deviations κ_i , using the inverse integrated Gaussian distribution (which is tabulated or available as a computer subroutine). Since κ_i is generated by the inclusive inverse Gaussian distribution of a properly normalized probability density function, κ_i is automatically

normally distributed, and can be used as the “standard deviations” of observed scaled pulseheight.

Unfortunately, these distributions together contain more information than can be conveniently displayed in a few graphs; therefore I am restricted to displaying the mean expected scaled pulseheight in figure 4.7 in units of photoelectrons, the probability of observing a photomultiplier hit in figure 4.9, and the RMS pulseheight resolution in figure 4.8. The left scale, which goes from 0 m to 21 m, is the impact distance d , the right scale is the impact angle ζ from 0° through 180° . The resolution in the third set of figures is determined from the scaled pulseheight distribution as half the distance between the $\pm 1\sigma$ points. The six figures again correspond to the six separately calibrated periods: IMB-1, IMB-2, IMB-3 before the extended Q-scale (labelled IMB-3 loQ), IMB-3 after the extended Q-scale (labelled IMB-3 hiQ), and the monte-carlo simulation of IMB-1 and IMB-2 (labelled IMB-1 MC) and of IMB-3 (labelled IMB-3 MC). Clearly visible is the increase in pulseheight from IMB-1 to IMB-2 and even more the two parts of IMB-3. Noteworthy is also the large increase in hit probability for IMB-3; in IMB-1 and IMB-2 the probability of a photomultiplier being hit was significantly less than one already a few meters away from the muon track, whereas for IMB-3 it stays close to one out to the largest distances and angles possible. With the high statistics available for the IMB-3 data, the drop in sensitivity at an impact angle of 90° is clearly visible; this is the angle at which the wavelength-shifter becomes useless. Again, the two monte-carlo simulations don't match the data too well. Fortunately the slight difference in behaviour between the two does not cause problems, since in the fitting procedure each event (whether monte-carlo or data) is reconstructed by comparing the hits to the correct resolutions and distributions.

The admittedly complex process described here has yielded a final quantity κ_i which is a normally distributed indicator of the relationship between observed pulseheight to the pulseheights spectrum expected from the calibration events for this particular geometry. In the previous section we had already obtained δ_i , which is a similar indicator for the relationship of observed to expected timing. Note that a photomultiplier which has not been hit has no valid δ_i since there is no timing information, but it does have a κ_i indicative of its observed pulseheight being zero. These two innocent-looking greek

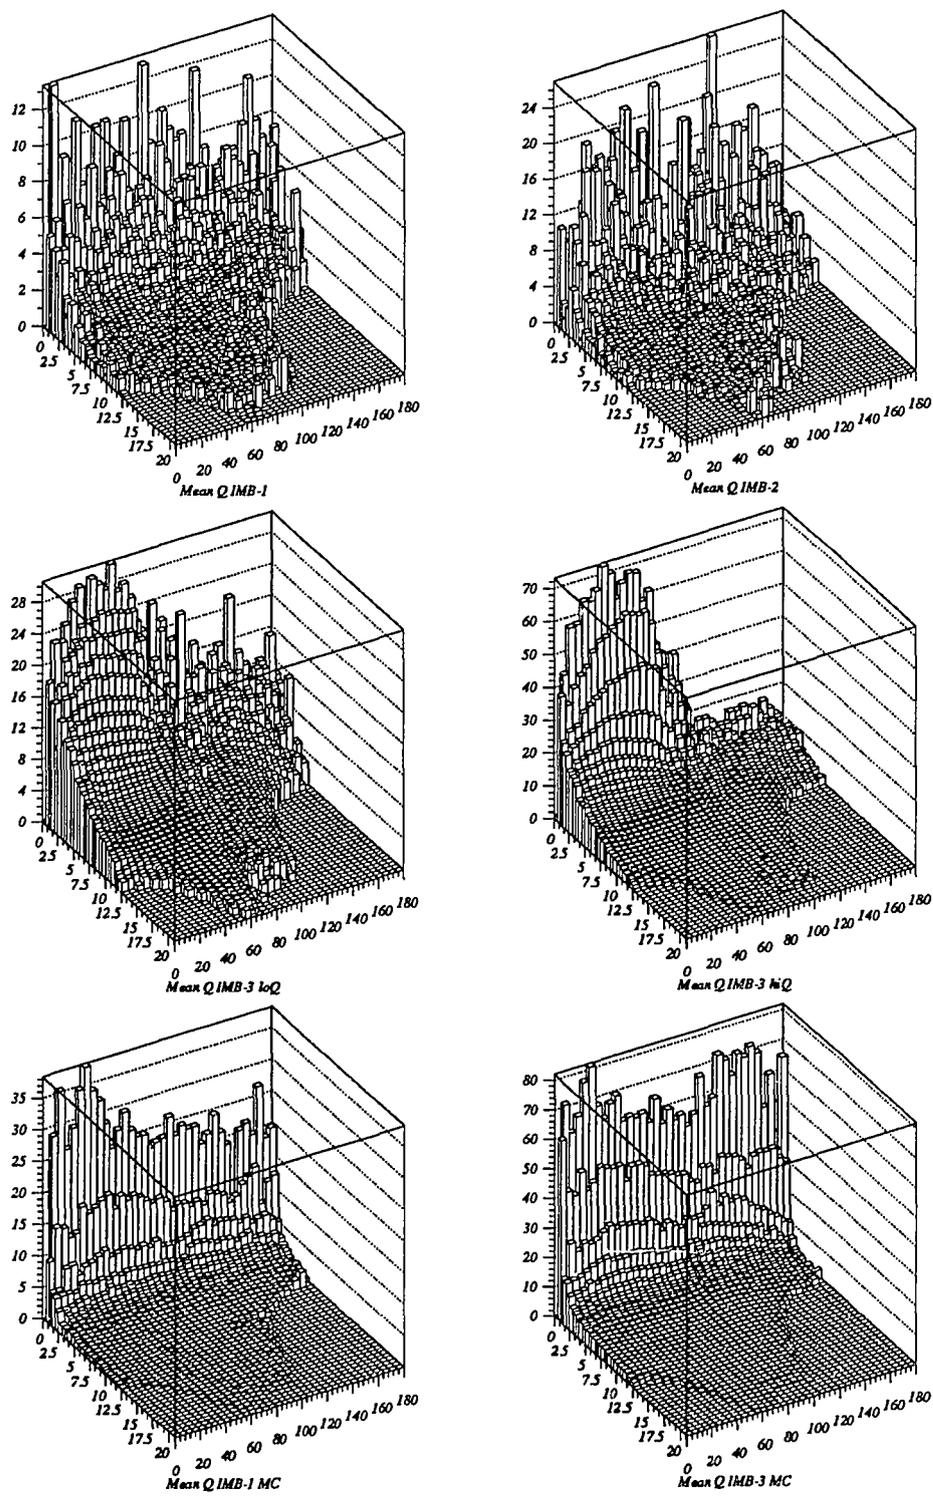


Figure 4.7: Mean expected pulseheight as a function of impact angle and distance

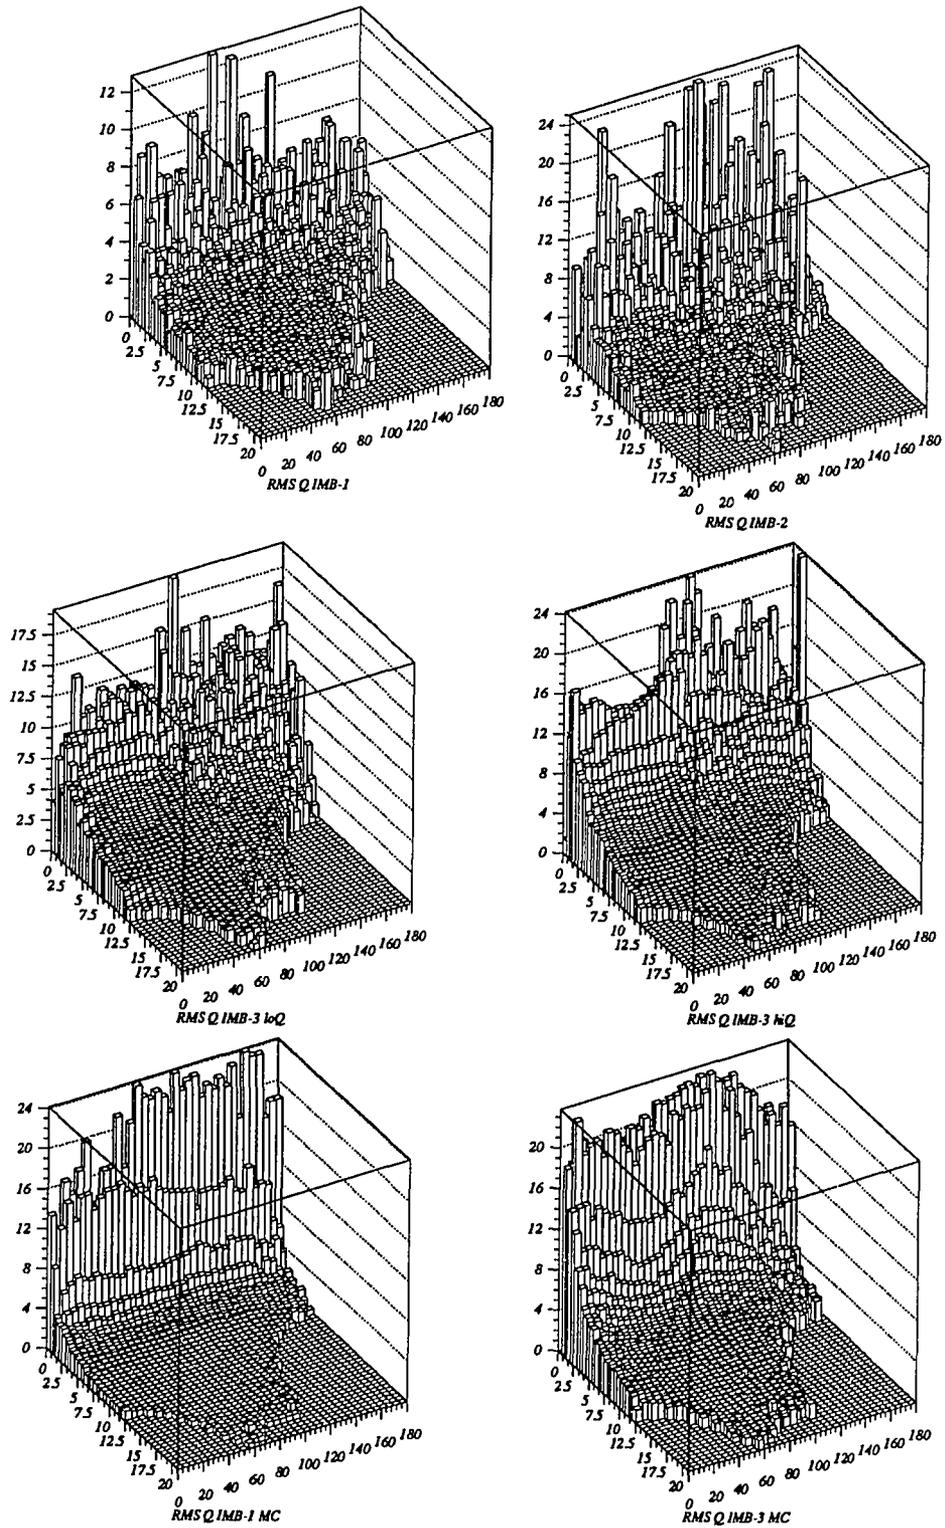


Figure 4.8: Pulseheight resolution as a function of impact angle and distance

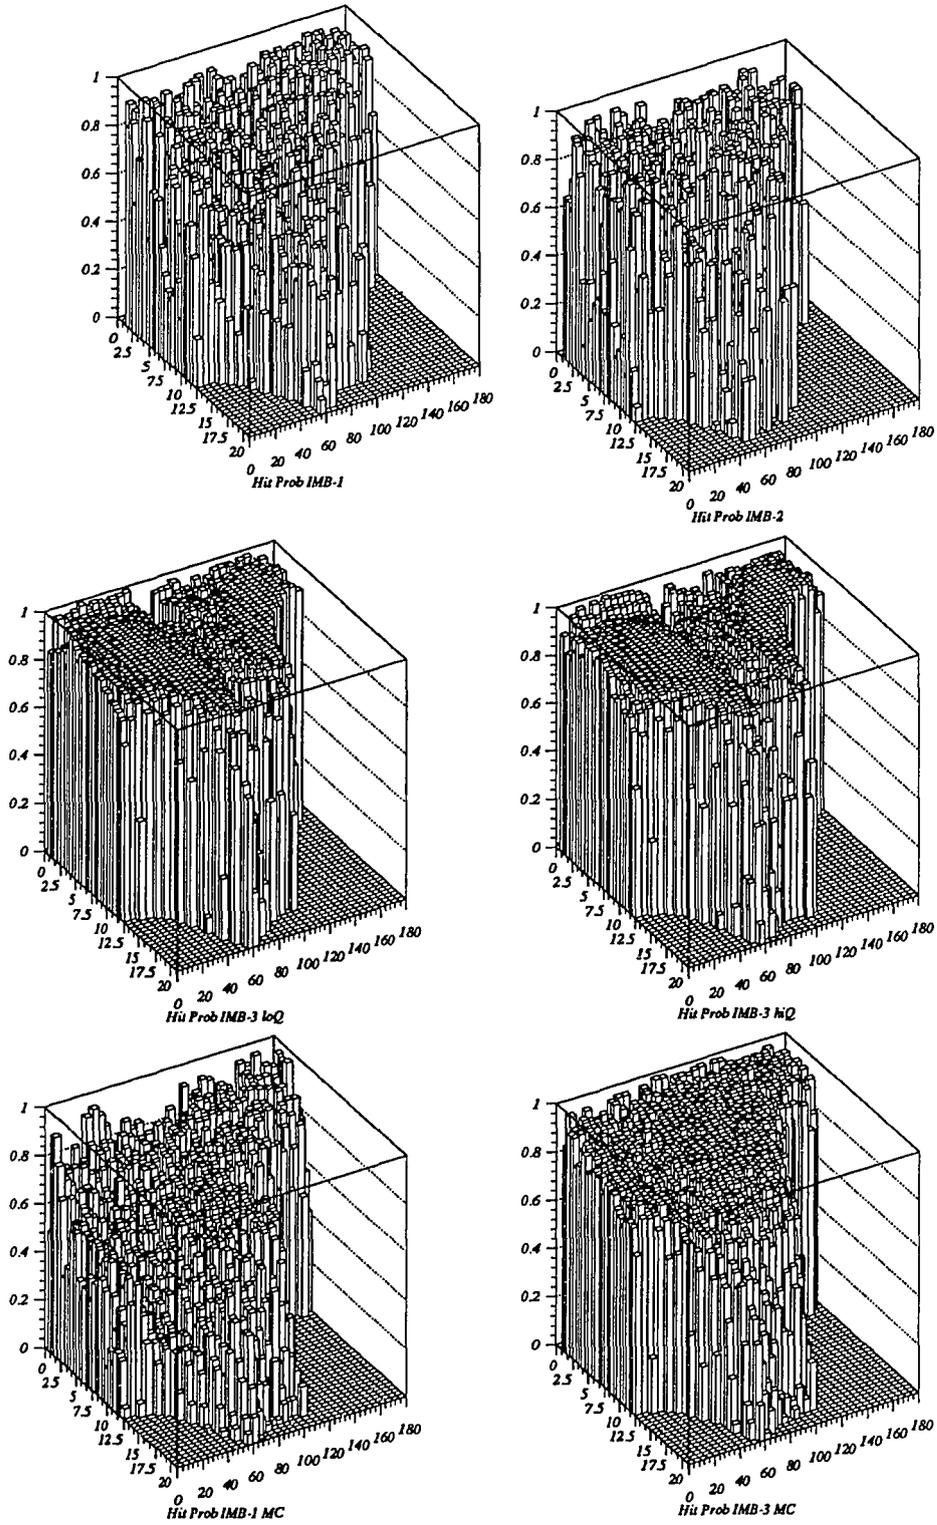


Figure 4.9: Hit probability as a function of impact angle and distance

letters are the thumb screws we will employ in the following section, where muons will be tortured until they reveal the direction they are going in.

4.4 Muon Direction Fitting

The preceding sections of this chapter have described how the data has been gathered and filtered, and what resolution of the detector is to be expected. Now a method to determine the direction the muons travel in is required. Not only are nearly all the results of this work based on this capability, it is also required to be able to select *upward-going* muons from the background.

4.4.1 Initial Estimate

An initial estimate of muon position and direction is required for several reasons. The online software will save events as upward-going muon candidates based on a quick estimate of muon direction. It also simplifies the work of the human scanner to have an initial guess of the fit. Last but not least, the high-precision fitter described later requires a fairly good initial estimate of the muon position and direction.

Two of the methods used to obtain that initial estimate will be described here: FastMu, the muon-fitting algorithm which is performed by the online computer, and the typical method a human scanner uses. Both methods rely heavily on the *entry and exit points* of the muon: one just looks for the earliest light as an estimate of the entry point, and the largest pulseheight hits as an estimate of the exit point.

In principle the very first Čerenkov light seen in the detector should be close to where the muon enters the water from the surrounding rock or from the air above. In practice there may be earlier hits, due to random firing of the photomultipliers. Two factors help to discriminate this random early light from Čerenkov light caused by the muon. Random hits typically have low pulseheights (equivalent to one or two PE), hits caused by a close-by muon typically have large pulseheight; see for example in figure 4.7: at a distance of about 1 m the expected pulseheights in IMB-3 are always above 10 PE. If the

muon traverses the wavelength shifting plate a lot of fluorescence light is generated in it; the geometrical probability of passing through a wavelength shifting plate is about 36%. Furthermore, the muon has to pass within about 50 cm of at least one photomultiplier, and even in the backwards direction (where the photomultiplier is mostly shadowed by its housing) the amount of light collected by a photomultiplier/wavelengthshifter assembly at such close a distance is fairly large, see figure 4.7 at angles above 90° . Also there are at least 50 cm of water outside the photomultipliers on all walls (except the top face of the detector, but in this work I am not concerned with muons entering from the top); therefore the muon will in most cases illuminate more than just one photomultiplier close to the entry point. A little complication is caused by malfunctioning photomultipliers; in principle one could loose the entry point of a muon if it illuminates only one photomultiplier, and that one happens to be dead. In practice this does not cause a lot of problems, since the entry point is usually larger than just one photomultiplier hit. In summary, the entry point of a muon into the detector will have some of the earliest light, with high pulseheights, and often in a tight cluster of photomultipliers.

Finding the exit point is easier, but more ambiguous. At the exit point the muon must again come to within 1 m of a few photomultipliers, and with the Čerenkov light cone facing the sensitive surface of the photomultipliers at such a short distance, they will register very high pulseheight, in fact usually the highest of the whole event. See figure 4.7, where the highest pulseheights are to be found at small distances and angles between 0° and 90° . So to find the exit point one just looks for the group of neighboring photomultipliers with the highest pulseheights.

Ambiguity on both ends of the fit can be caused by muons which stay very close to a plane of photomultipliers for a large distance (they are known as “wall glancers”), leaving a long trail of hits with large pulseheight. In this case the entry and exit point may be mis-identified, but are usually close since they are always within the trail of high-pulseheight hits. A larger problem is caused by events which cut across the edge of the detector with entry and exit point within a few meters of each other (so called “corner clippers”). In this case most of the light visible in the detector is generated outside the fiducial volume of the detector or is scattered light; the only useful part of the Čerenkov light cone is a little area at an edge or corner of the detector. In this case the entry

and exit spot are often inseparable, and the algorithms used to determine the first estimate of a fit may fail. This is not unexpected, both the efficiency and the accuracy of the fitting process fall off rapidly as the pathlength of the muon through the detector reaches a few meters.

FastMu is the most accurate and reliable of the muon fitting algorithms which is executed on the online computer. It is described in detail in reference [16]. It first looks for clusters of neighboring early and of neighboring large pulseheight hits. It selects the largest early cluster as the entry point. To find the exit point it tries all clusters of large hits and selects the one which agrees best with the assumption of the muon travelling at the speed of light. As a last stage it does a variational fit of the muon position and direction.

Obviously this algorithm employs several assumptions, which can be summed up by treating the event as being a single non-showering particle entering and penetrating the detector in two distinct points. There are cases where **FastMu** breaks down: contained events (it does not find the right entry point, there is none), stopping muons (there is no exit point) and multiple muons (there are too many choices for early and hot clusters). None of these cases are of interest here. For stopping muons, **FastMu** usually picks a point on the inner edge of the Čerenkov light cone as the exit point, from which it improves the fit in the last (variational) stage of the fit, often resulting in reasonably good fits as long as the pulseheights on the inner edge of the Čerenkov light cone are high enough. Since stopping muons are nearly all below 900 photomultiplier hits they will be saved in the data independent of the fit result, see section 2.2, so the fit quality of **FastMu** is not of great importance for them.

For throughgoing muons, **FastMu** is typically accurate to 5° (estimated both by monte-carlo events, and by comparing the best manual fits to **FastMu**), and even for “corner clippers” it performs well. In summary, **FastMu** is a very good general purpose fitter for muons, and is typically used as a starting point for the hand-fit by a human scanner.

The person scanning the event usually uses a method similar to computerized muon fitting algorithms. Typically one first looks for the earliest hits of reasonable pulseheight, ignoring single hits of small pulseheight occurring at random positions in the detector. Then one looks for the highest pulseheight hits, which should form a pattern with the

highest pulseheights in the center at the exit point (a “bull’s eye”). Connecting this entry point guess with the exit point guess gives a reasonable first estimate. To proceed efficiently at this stage a good initial guess of the fit is very helpful.

Now comes the crucial step: for a given muon trajectory one can determine which photomultipliers should have been hit; the event display programs typically either draw the outline of the Čerenkov light cone, or highlight photomultipliers which should have been hit. The important part is now that the contrast between the high light level in the Čerenkov light cone and the low light level outside is very good. In IMB-1 there is very little light scattering in the detector, so most of the hits in an event are in the Čerenkov light cone. The probability of a photomultiplier in the Čerenkov light cone being hit is significantly non-zero, even at the edge of the largest Čerenkov light cones it is still on the order of 25%, see figure 4.9; so by just drawing the outline of the area in which most of the photomultipliers are hit gives a good estimate of the true outline of the Čerenkov light cone.

In the IMB-3 detector the probability of a photomultiplier in the Čerenkov light cone being hit is very close to one, see the middle row of figure 4.9, so the cone appears as an area where all photomultipliers are hit. Unfortunately the amount of light scattered is significant, see figure 2.1, so even outside the Čerenkov light cone there are many photomultiplier hits. This scattered light can be effectively removed by using pulseheight, timing and clustering. In IMB-3 the expected pulseheight for hits inside the Čerenkov light cone is 2 PE or larger, whereas the typical pulseheight of scattered light hits is 1 or 2PE. Scattered light is typically later than the direct Čerenkov light, since it had to travel an extra distance to leave the Čerenkov light cone. The most important point is that the light generated by the muon is very coherent, the timing difference between adjacent hits is small, typically on the order of the spatial distance between those photomultipliers. Scattered light is less coherent, so one can remove it by removing hits which don’t have neighboring hits in close timing coincidence from the event. With a combination of these three methods (cuts to remove late hits and hits with low pulseheight, plus removing hits without time-coincident neighbors) one can usually very well remove the scattered light from an event. After that, the contrast between the high light level in the Čerenkov light cone and the lack of light outside is very good.

The crucial step in manually fitting an event is now to match the outline of the expected Čerenkov light cone with the observed outline of the area where photomultipliers have been hit, achieved in practice by moving the assumed muon track around until most of the hit photomultipliers are inside the Čerenkov light cone this muon would have generated, and as few as possible are outside it. This usually determines the muon track already fairly well, to an accuracy which is determined by the contrast at the edge of the Čerenkov light cone. Since this contrast is fairly high if one employs the pulse-height information, the resolution is typically one photomultiplier spacing, which leads to an accuracy of about $\pm 3^\circ$ for large events.

4.4.2 Precision Fitting

To further improve the fit one now starts using more and more information about the photomultiplier hits; we have available the times and pulseheights of the hits. In order to be able to use powerful statistical tests, such as the enhanced Rayleigh test described in section 3.2 these quantities have already been converted into indicators δ_i and κ_i , which are approximately normally distributed with mean zero and variance one *if and only if the muon is correctly reconstructed*.

First some notation. As outlined above, we characterize the path of the muon by \vec{v} , \hat{d} and t_0 . The photomultiplier hits are described by the position of the hit photomultiplier \vec{x}_i , the time t_i and the pulseheight q_i of the hit. Using this notation, the expected time of a hit is

$$t_i^e = t_0 + \frac{\vec{x}_i - \vec{v}}{n} \hat{d} + \frac{\sqrt{n^2 - 1}}{n} |(\vec{x}_i - \vec{v}) + ((\vec{x}_i - \vec{v}) \cdot \hat{d})\hat{d}| \quad (4.2)$$

This formula is derived in appendix B. The timing residuals are the difference between measured and expected time $t_i - t_i^e$, and are scaled to produce approximately normally distributed quantities δ_i according to the procedure of equation (4.1). Similarly, as described in section 4.3 the pulseheights q_i are transformed into normally distributed quantities κ_i by comparing them to the distribution of expected pulseheights, with the relative light output ε as a byproduct.

The general fit procedure has already been outlined at the beginning of section 4.3: assume a muon geometry (\vec{v}, \hat{d}) , calculate δ_i and κ_i , compound all of them into a single estimator of fit goodness, which is maximized by varying the geometry (\vec{v}, \hat{d}) . A little complication arises here: to completely describe the muon, knowledge of t_0 is required; but its value is not fixed and fluctuates from event to event, depending on when exactly the trigger conditions were fulfilled. Therefore t_0 is determined by an iterative process such that the mean timing residuals are zero. To prevent extraneous hits (early random light or late scattered light) from influencing this mean heavily, the determination of t_0 is repeated thrice, each time applying harder and harder cuts on the timing residuals to be included in the mean. This robust process estimates t_0 such that the mean timing residuals (and therefore the mean of the δ_i) is made zero, without regard to the tails of the timing distribution. The determination of t_0 has to be repeated any time the muon geometry (\vec{v}, \hat{d}) changes.

What is required now is a procedure to turn the δ_i and κ_i into a single estimator of fit quality. In the language of statistics, we need a test to determine whether they are normally distributed with zero mean and unit variance; if they are we have reached a “perfect” fit; if they are not, we can improve the fit. The null hypothesis in this test is *“the δ_i and κ_i are distributed normally”*, and the fit quality is the complement of the confidence level at which we reject the null hypothesis (a high confidence level of rejection means the null hypothesis is probably false, which in this case means the fit is probably not correct, so the fit quality is low). Note that in the case of the fitting procedure we are interested in tuning the parameters (\vec{v}, \hat{d}) until the null hypothesis is accepted (that is we have a perfect fit), which is the opposite of the usual situation, where one is trying to refute the null hypothesis.

The usual statistical test to use in such a situation is Parson’s χ^2 test; in fact most event fitting procedures use it exclusively. It however ignores valuable information in this case: the fact that the quantities δ_i and κ_i have been measured at distinct photomultiplier positions \vec{x}_i . As an example, assume that the muon fit is nearly correct, in the sense that the exit point of the muon is correctly reconstructed, but the reconstructed direction of the muon is wrong by a few degrees. Then the timing residuals on one side of the Čerenkov light cone will all be slightly biased towards positive, on the opposite side

slightly biased towards negative; similar the κ_i are offset. The χ^2 test treats all the δ_i the same, independent of the position they were measured at, and only uses the square of the indicators δ_i and κ_i , thereby ignoring their different sign on either side. If the bias is slight (as is the case if the reconstruction is close to correct) the χ^2 test is very insensitive to the slight increase in the mean of the absolute values of δ_i and κ_i . The enhanced Rayleigh test outlined in section 3.2 is sensitive to both the spatial distribution and to the sign of δ_i and κ_i . I will therefore make extensive use of it. In essence, this test measures the *dipole moment* of the distribution of δ_i and κ_i , and checks whether the magnitude of the dipole moment is statistically significant. The actual procedure is to choose an upward reference direction perpendicular on the muon direction (for convenience I choose the component of the $+z$ -axis perpendicular to the muon direction), project all the photomultiplier hit positions \vec{x}_i into a plane perpendicular to the muon direction, and measure the angle ϕ_i of each hit relative to the upward reference direction.

Another complication arises here. Just using the dipole moment of the spatial distribution of δ_i is not sufficient when fitting using the timing information contained in the δ_i . As an example, let's again assume we have fit a muon, the reconstructed exit point is correct but the direction is slightly wrong. The spatial distribution of δ_i will show a dipole moment towards one side. Similarly, if the direction is correct but the exit point is slightly wrong there will be a dipole moment in the spatial distribution of δ_i . Now the following can happen: we start out with a fit with the correct exit point but a slightly wrong direction. The fitting program will modify the parameters (\vec{v}, \hat{d}) , trying to maximize the fit quality, which in this case means minimizing the dipole moment. It can do so by keeping the wrong direction and moving the exit point; in essence one can compensate the dipole moment caused by a wrong direction using the dipole moment of a wrong exit point. From looking at the dipole moment it seems that the fit is indeterminate, that any error in the direction can be compensated by a corresponding error in the event position. To break this indeterminacy two remedies are possible. One is to use the dipole moment of the κ_i , thereby employing the pulseheight information. However, the κ_i are not very sensitive to small changes in the fit position, since the pulseheight resolution contributes much less to the fit resolution than the timing resolution. The better approach based on the following observation: when compensating

the dipole moment caused by a misfit direction with the dipole moment from the event position, a quadrupole moment is introduced into the spatial distribution of the δ_i . To employ the enhanced Rayleigh test for the quadrupole moment one just uses it with all the angles ϕ_i doubled, since it is not sensitive to the distribution of ϕ_i anyhow.

Summarizing, the procedure used to construct a single indicator of *fit quality* from the δ_i , κ_i and ϕ_i is to use the enhanced Rayleigh test, on the dipole moment of the δ_i , on the quadrupole moment of the δ_i (with the ϕ_i doubled), and on the dipole moment of the κ_i . The three results are then merged into one number using equation (3.2). In principle one could also use the quadrupole moment of the κ_i , but since they are not very sensitive to small changes in the fit that doesn't improve the fitting procedure very much. On an event-by-event basis, each of the three tests can be disabled; this capability is used mostly for debugging. One more problem remains: the fitting procedure can rotate the Čerenkov light cone in such a way that there are hardly any photomultiplier hits in it (for example by turning the fit direction by 90°). In this case the moments of the few photomultiplier hits remaining in the Čerenkov light cone are very small, since there are few hits, so this looks like an excellent fit. To prevent that, an additional penalty is included in the fit quality which measures the fraction of all photomultiplier hits in this event which are in the current Čerenkov light cone; this penalty is made very small as long as that fraction is reasonably close to 1, and only increases markedly at very small fractions. This prevents the fitting procedure from wandering off into unphysical regions.

The last step required to complete fitting an event is the algorithm to vary the parameters of the fit. In this case the parameters are the position and direction of the muon \vec{v}, \hat{d} . As was explained above, the two together contain four independent parameters; the fifth one, t_0 is recalculated for each change in \vec{v}, \hat{d} . It is however quite difficult to express \vec{v}, \hat{d} as four numbers; any attempt to parametrize them in terms of spherical coordinates suffers from the problem of indeterminacy at the poles of the coordinate system, which leads to a divergence in the fitting procedure. I circumvent this problem by expressing muon position and direction each in two angles *in spherical coordinates relative to the initial fit*, and re-expressing those angle-pairs using cartesian coordinates. This has several advantages: all the four parameters used in the fit have the same units of angle; they are

scale-free (that means their typical size doesn't depend on the muon geometry), there is no indeterminacy in the coordinate system (except at the pole opposite to the current direction, but a fit should never move 180° from the current fit).

The problem has now been reduced to varying four angles (which completely describe the geometry of a through-going muon) in such a fashion that the fit quality described above is maximized. For reasons of numerical stability I don't maximize the fit quality, but instead minimize its negative logarithm. There are many standard solutions to this *minimization problem*; I use two. The first one is a very simple parabolic interpolation algorithm, which successively minimizes the goal function by parabolic interpolations in each of the four angles; the second one is the Amoeba simplex-algorithms from reference [72]. Both methods usually yield the same result, but the second one is far superior, as it is much faster, terminates reliably when further improvement is impossible, and does not get fooled easily by multiple minima of the goal function.

This method of fitting an event is involved, and therefore also time-consuming. To calculate the set of δ_i and κ_i for a certain geometry requires several passes through the photomultiplier hit data, determining t_o and ε as by-products, and typically uses between 1 and 3 seconds of CPU time on an 11/780-class computer, depending on the number of photomultiplier hits. To completely fit a muon the geometry has to be varied many times until the highest possible fit quality is achieved; that typically consumes 10 to 30 minutes of CPU time, again on an 11/780-class machine. Furthermore, the success rate of the fitter is not 100%. Sometimes it finds a solution with a very high fit quality, but that solution is unphysical. Therefore all the fits have to be checked manually; for about 20% to 30% of all events I rejected the automatic fit, either because the fit quality finally achieved is too low, or because "it looks wrong". That success rate of 70% to 80% applies both to real data and to monte-carlo simulation.

4.4.3 Direction Resolution

As outlined above, this method of determining the direction of a muon should in principle yield better results than methods which are based just on the χ^2 test. It employs much more information, the spatial pattern of the timing and pulseheight residuals, and is

thereby sensitive to small changes in the indicators δ_i and κ_i , if these changes are spatially coherent, as they indeed should be for a mis-fit muon. Now a determination of the actual resolution of this fitting procedure is required. We are only interested in the resolution of the muon direction, since the position of the muon is of no consequence for the following sections dealing with neutrino astronomy.

This however is a large problem. This fitting procedure uses most of the information available from this detector for an event. In the best of all worlds we would have some mechanism independent of the detector which selects muons of a known position and has an intrinsic accuracy much better than what we want to achieve here, which is on the order of degrees. Such a device could for example be a set of well-calibrated and surveyed drift chambers covering at least two sides of the detector; but it doesn't exist, and would cost millions to install. There is an approximation: some data early in IMB-1 was taken with a 0.35 m^2 scintillator (a 66 cm diameter circle) installed in the bottom of the pool. This is unfortunately not very helpful; first of all it only fixes the exit point of downgoing muons, and doesn't help at all in determining their direction. Furthermore, even if another point on the muon track at the top of the detector were known with infinite precision, the intrinsic resolution of this device would be only about $66 \text{ cm}/17.5 \text{ m} = 2.2^\circ$; so there is no point trying to study angular resolution in the degree-range with this device.

It seems that the direction resolution cannot be determined from other parts of the detector, nor from devices installed externally to it. There is only one class of "events" left for which the true path of the muon is known: monte-carlo simulated events. Naturally it is dangerous to determine the resolution from simulated data; as was already seen in section 4.3, the exact behaviour of timing and pulseheight resolution for it is different from the real data. Although these differences should be accounted for in the correct definition of δ_i and κ_i ; it is still not certain that the properties of the fitting procedure will be the same for data and monte-carlo. However, the use of the simulated data is the only recourse left; only here can I actually calculate the difference in angle between the known path of the muon and the fit obtained by the method described above.

Initially, there was hope to determine the fit resolution on an event-by-event basis. The idea is the following: assume an event has been fit carefully by the above procedure, and the fit quality at the best fit is about unity (that means the fitter has achieved a near-perfect fit). If one moves the fit direction away from the best fit (that is if one tries a slightly “wrong” fit”), the fit quality will decrease. Then one can draw a two-dimensional map centered around the best fit, showing how the fit quality depends on the deviation from the best fit. The angular resolution could then be found as the area in which the fit quality is not significantly different from the best fit. There are several problems with this approach. It is very difficult to define where the fit quality is *significantly* different from the best fit; there is no reasonable concept of statistical error on the fit quality. The area of “acceptable fit quality” can be shaped very oddly; sometimes it is a nice ellipse, other times it has long pseudo-pods extending to parasitic unphysical maxima of fit quality. Lastly but most importantly, this method doesn’t work: using monte-carlo events I found that the estimate of fit resolution obtained from these two-dimensional maps of fit quality are not correlated at all with the true deviation between the “real” monte-carlo muon and the reconstructed fit. Another smart idea killed by an ugly fact.

Then many other possible indicators of fit resolution were tested after this failure; the intention being to find some quantity which is strongly correlated with the true deviation between the simulated and the reconstructed muon. Among the ones considered are total area of the Čerenkov light cone, number of photomultiplier hits in the Čerenkov light cone, total amount of light observed, pathlength of the muon, distance of either entry or exit point from the nearest edge of the detector, muon angle with respect to the wall through which it enters or exist, final fit quality at the best fit, difference between fit results using only timing and using only pulseheight. None of these showed any substantial correlation with the angular deviation between the generated muon track and the fit.

At this point, the only option left is to not try to estimate the angular resolution on an event-by-event basis, instead use one global resolution for all events. Figure 4.10 shows the distribution of angle between the “real” monte-carlo muon and the direction reconstructed by the fit (the automatic fit described above, plus manual corrections as described above for cases where the automatic fit fails to find a good enough fit quality),

for 251 IMB-1 and IMB-3 monte-carlo events. The RMS deviation is 2.55° , and is shown as a dotted line. Note that the RMS angle is larger than either the most probable angle or the mean deviation; although it makes the result look worse, it is the proper thing to use, since in the following chapters the events will be smeared out using a Gaussian distribution with this angle as variance.

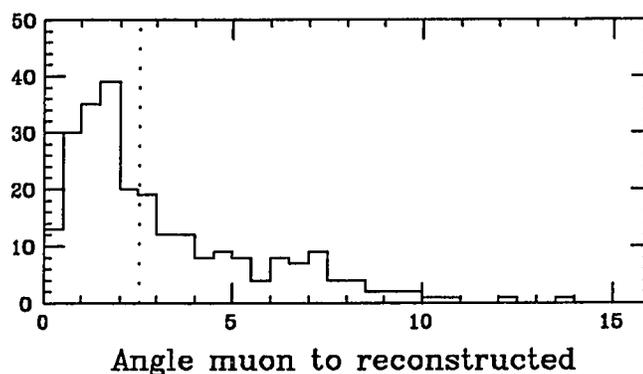


Figure 4.10: Distribution of angle deviation between generated monte-carlo muon and reconstruction, for 251 simulated events. The dotted line shows the RMS of 2.55°

There are other small sources of uncertainty in the determination of angle. The detector is surveyed to better than 1° . The accurate time for most IMB-3 data is read out from the radio- or satellite-clock; for some events where that clock failed, the time-of-day clock of the online computer is used. We can safely assume that the event time is accurate to 5 minutes for all events (in fact, for most events the accuracy is more on the order of milliseconds), which adds at most another degree of uncertainty to the direction (due to the earth's rotation). The position of the photomultipliers is slightly uncertain; even if we assume that worst case error of uniformly 20 cm, that only adds another degree of uncertainty across the detector. To account for all these possible small sources of error, I choose to quote the event reconstruction angular resolution as 3° .

In summary, a method of determining the direction of throughgoing muons has been developed which is sensitive to the pattern of timing and pulseheight, and yields a statistical estimate of fit quality. After spending months to develop and implement this method, hundreds of hours of CPU-time to fit events under study here plus a similar

number for monte-carlo simulated events, and weeks of event scanning to manually check and sometimes correct all the fits, an angular resolution for the fit of 2.55° is achieved. However, all attempts to determine that resolution on an event-by-event basis have failed.

4.5 Muon Energy Estimation

It was already pointed out in section 4.3 that the mean Čerenkov light output of a muon increases with energy. One can therefore use the measured light output to estimate the energy of a muon. Unfortunately, the Čerenkov light output also fluctuates very much. Any estimate like this will not be an accurate measurement of muon energy, but rather a *trend indicator*. Obviously, the total amount of light observed in the detector cannot be used directly to estimate the amount of light emitted by the muon *per unit tracklength*; even scaling the observed number of PEs to the total pathlength will not do, since the detector is highly non-linear in measuring pulseheight. Fortunately, the calculation of the pulseheight residuals κ_i yields the quantity ε , which is the ratio of observed to expected pulseheight, averaged over all photomultiplier hits. This ε will be used here to estimate the muon energy.

An estimation of the muon energy is interesting for two reasons. In itself, it might allow a measurement of the energy spectrum of upcoming muons and therefore of atmospheric neutrinos; there has never been an accurate measurement of their energy spectrum. However, as we will see below, such a measurement is not possible in this case. Less fundamental but technically more useful, the estimate of muon energy immediately yields an estimate of the energy of the parent neutrino; and the neutrino energy is strongly correlated with the angle between the neutrino and the muon. Also, the multiple scattering of the muon depends on its energy. At the energies of interest here, this angle is usually larger or of the same order of magnitude than the event reconstruction uncertainty of about 3° derived above. At the energy spectrum typical of atmospheric neutrinos the angle deviation between muon and neutrino is the dominating component of the uncertainty in measuring neutrino directions using this detector.

The only way to estimate the muon-to-neutrino angle from the observed scaled light output ϵ is to completely rely on the monte-carlo simulation. The kinematics of the $\nu_\mu + \mathcal{N} \rightarrow \mu^\pm + X$ process and of the angle between ν_μ and μ^\pm is well understood, and one can assume these processes are correctly implemented in the monte-carlo. The mechanisms by which the light output increases is less well documented; since there is no artificial source of muons of these energies this cannot be observed in the laboratory under controlled circumstances. In fact, it is known that the monte-carlo program does not correctly simulate bremsstrahlung photons above a muon energy of 1 TeV; however, there are very few events above that energy, so that problem should be only a very small effect.

As was already pointed out in section 4.3, the quantity ϵ is the light output of an event *relative to the average of the calibration muons used to create the pulseheight distribution for a particular part of the data*. Therefore, any estimate of muon energy based on ϵ only yields information relative to these calibration muons. Unfortunately, different sets of events with different energies had to be used for the different sets of data. For IMB-3 the data was calibrated using downgoing muons only, for IMB-1 and IMB-2 a mix of downgoing and upward-going muons was used, and for the two monte-carlo simulations only upward-going muons were used. The mean energy of downgoing and upward-going muons are quite different, although their spectral indices are similar; and therefore the values of ϵ obtained from the monte-carlo cannot be used directly as a calibration for the energy estimation of the data. To correct for this difference, I introduce a scale factor in the ϵ measurement for the monte-carlo: since the energies of the monte-carlo events are lower, their light output is lower; therefore when comparing them to the data, one would overestimate the energy for data events, therefore the light output measurement of the data has to be artificially increased too. This correction factor is determined by matching the ϵ distribution for data and monte-carlo, using the Kolmogoroff-Smirnov probability for the two distributions to be identical. The resulting distributions are shown in figure 4.11; IMB-1 and IMB-2 are the top graph, IMB-3 is the bottom graph. The resulting scale factors and Kolmogoroff-Smirnov (K-S) probabilities are listed in table 4.3. As one can see, and as the K-S probability proves, the match between the distribution of ϵ for data and monte-carlo after applying this correction is very good.

This is the best indication that the handling of energy loss mechanisms in the monte-carlo program is adequate for the purposes of this study.

An important point is that at this point, the chance of measuring the energy spectrum of the muon events has been lost; in fact, the assumed spectrum has been used to calibrate this method of estimated energy. All we are left with is means to estimate the energy of individual muons.

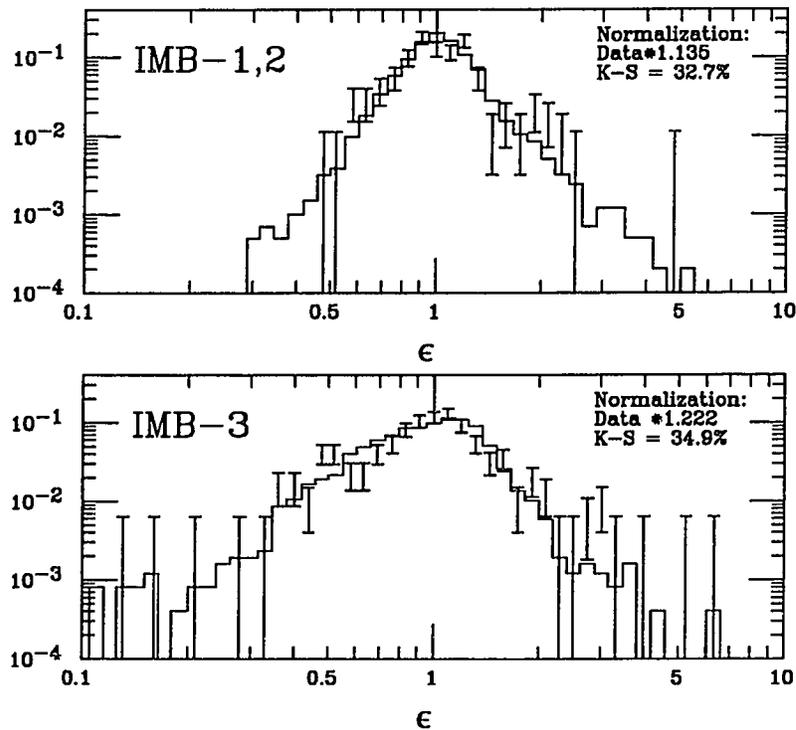


Figure 4.11: Distribution of relative light output ϵ . Top figure is for IMB-1 and IMB-2, bottom is for IMB-3. Solid line is monte-carlo, error bars are data. The data has been scaled by the factor indicated in the top right corner; also indicated is the Kolmogoroff-Smirnov probability.

Period	Correction factor	Resulting K-S Probability
IMB-1 and IMB-2	1.135	32.7%
IMB-3	1.222	34.9%

Table 4.3: Correction factors used for the estimation of muon energy from observed light output, and the corresponding Kolmogoroff-Smirnov probabilities

The next step is to obtain the relation between the relative light output ϵ , and the actual muon energy and the angle between neutrino and observable muon. For that purpose, the monte-carlo events are split into 11 bins in ϵ , with the lower limits of the bins being 0, 0.615, 0.735, 0.81, 0.87, 0.93, 0.975, 1.035, 1.125, 1.26 and 1.8. These bins were chosen such that the number of events in each bin is about the same, except the last one, which is designed for events with extremely high light output and presumably extremely high energy. Since the light output of a muon fluctuates so much, one expects that the energy estimation will work better for longer pathlengths of the muon in the detector, where the fluctuations are better averaged out. Therefore it is advantageous to keep events with different pulseheight separate, to prevent the ill-measured ones at low pathlength from introducing too much scatter into the better measurements for longer paths. The monte-carlo events are further split into five bins in pathlength, with lower limits of 0 cm, 930 cm, 1350 cm, 1740 cm and 2010 cm; again the limits were selected such that the population per bin is about the same. For each of the 11×5 bins in ϵ and pathlength a distribution of muon energy and one of angle between observable muon and neutrino was produced. From each of the distributions, the mean muon energy and the RMS angle is then calculated. Figure 4.12 shows one fifth of these distributions, for pathlength between 1740 cm and 2010 cm. The top graph shows the distributions of muon energy in GeV; there is one distribution for each bin in ϵ . The mean muon energy is marked by a star. Similarly, the bottom graph contains the distributions of angle between muon and neutrino in degrees, with the RMS angle marked. A trend is visible: The mean muon energy increases with increasing ϵ , and the RMS angle decreases. The same trend is shown more clearly in figure 4.13 for the muon energy as a function of the bin in ϵ , and similarly in figure 4.14 for the RMS muon-to-neutrino angle. The five lines in each graph correspond to the 5 bins in pathlength, with "A" being pathlength below 930 cm through "E" for pathlength above 2010 cm.

These RMS muon-to-neutrino scattering angles are tabulated in table 4.4. They will be used together with the reconstruction angular resolution of 3° in estimating the angular spread of each event in sections 6.3 through 6.5; the final angular resolution σ ; for each event is determined by adding the energy- and pathlength-dependent RMS resolution and the 3° reconstruction resolution in quadrature.

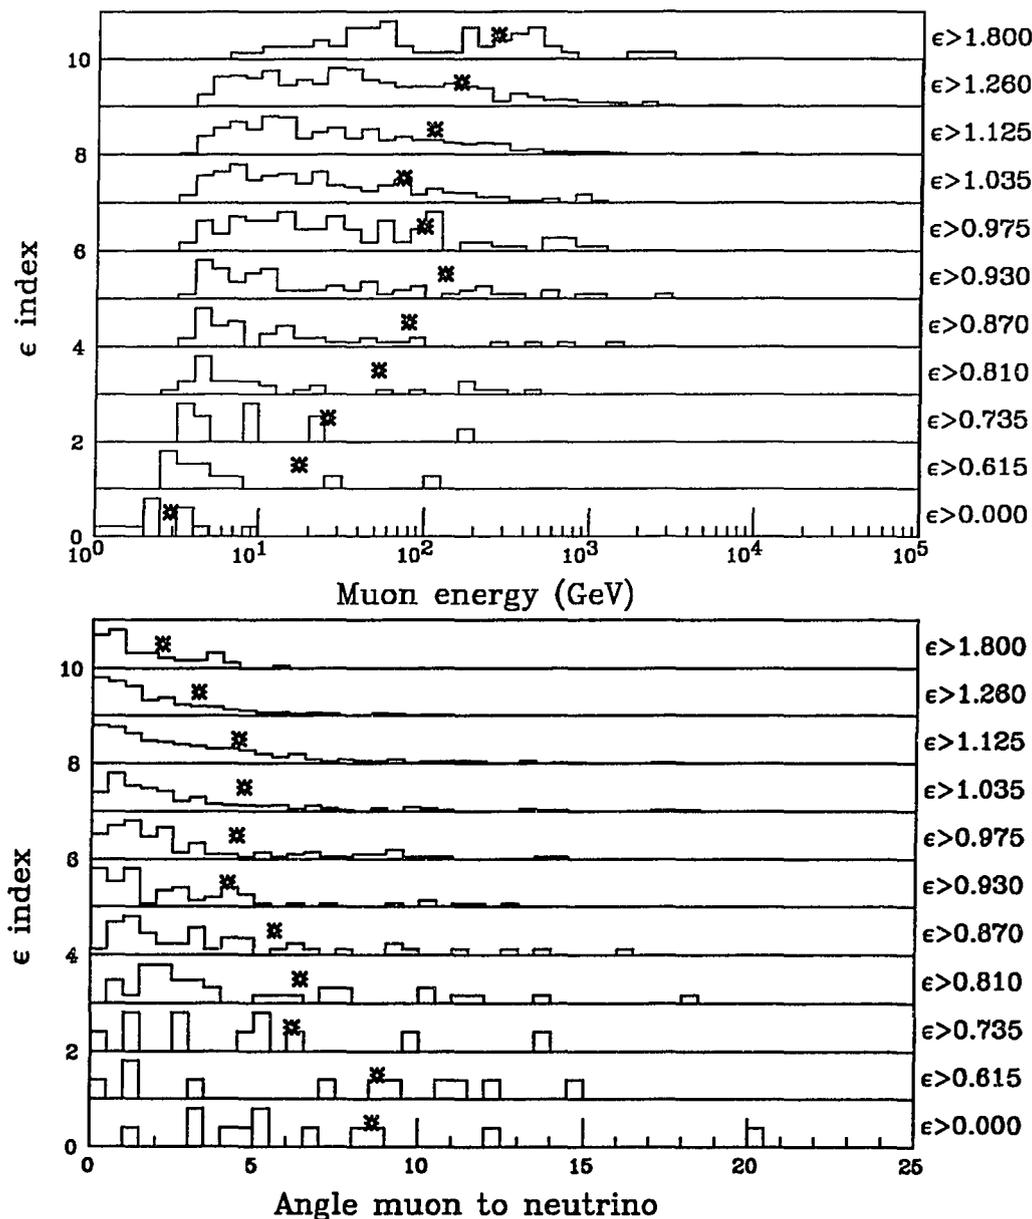


Figure 4.12: Change in distribution of muon energy and muon-to-neutrino angle with ϵ . The two graphs are for muon pathlengths > 1740 cm and < 2010 cm. The top figure shows the evolution of muon energy; the mean muon energy for each bin in ϵ is marked by a star. The bottom figure shows the evolution of the angle between muon and neutrino in degrees, the RMS angle is marked by a star.

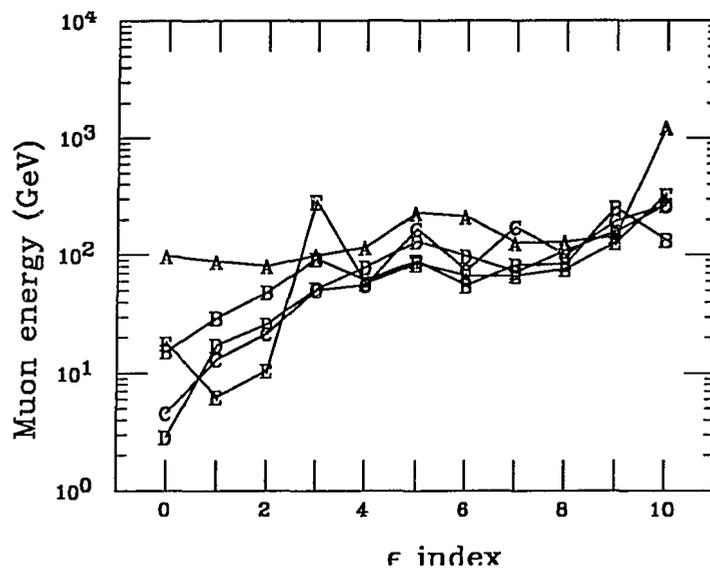


Figure 4.13: Mean muon energy as a function of ϵ and pathlength. The horizontal axis is the index for the bin in ϵ . The five lines marked "A" through "E" correspond to the five bins in pathlength.

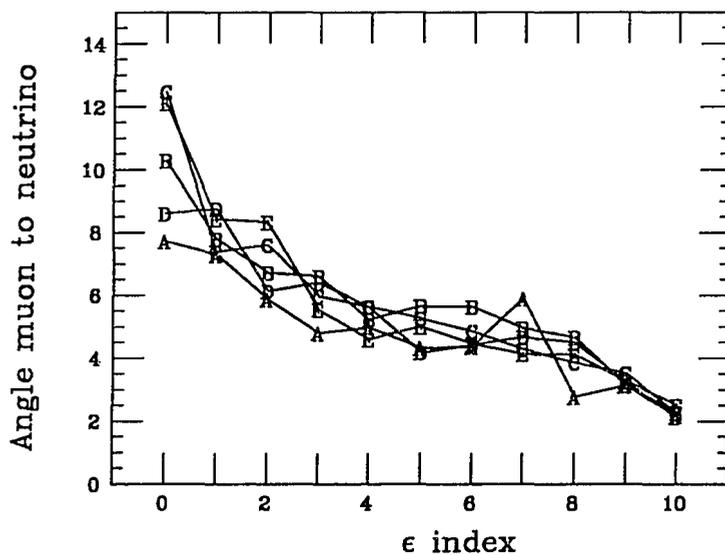


Figure 4.14: RMS muon-to-neutrino angle in degrees as a function of ϵ and pathlength. The horizontal axis is the index for the bin in ϵ . The five lines marked "A" through "E" correspond to the five bins in pathlength.

Pathlength:	< 930 cm	<1350 cm	<1740 cm	<2010 cm	>2010 cm
$\epsilon > 0.000$	7.73	10.29	12.47	8.61	12.12
$\epsilon > 0.615$	7.30	7.80	7.40	8.77	8.44
$\epsilon > 0.735$	5.89	6.72	7.61	6.14	8.34
$\epsilon > 0.810$	4.81	6.61	5.99	6.41	5.54
$\epsilon > 0.870$	4.99	5.24	5.65	5.62	4.60
$\epsilon > 0.930$	4.33	5.67	5.27	4.19	5.01
$\epsilon > 0.975$	4.37	5.65	4.87	4.44	4.47
$\epsilon > 1.035$	5.91	4.97	4.29	4.67	4.13
$\epsilon > 1.125$	2.78	4.66	3.86	4.50	4.12
$\epsilon > 1.260$	3.14	3.15	3.53	3.26	3.30
$\epsilon > 1.800$	2.18	2.29	2.25	2.13	2.51

Table 4.4: RMS muon to neutrino scattering angles, as a function of pathlength and relative light output ϵ

4.6 Detector Effective Area

In section 6.1 the total flux of atmospheric neutrinos will be calculated. That requires knowledge of the detector live time, which was already covered in section 4.2.1, and knowledge of the effective area of the detector, which this section will provide.

Again, the help of the monte-carlo simulation is required. Muon events are generated in a large volume around the detector (a sphere of 2 km radius, cut off at the surface of the earth), with the correct angular and energy distribution for atmospheric neutrinos. Events are generated down to very low energies, and also including downward going events. Most of these events do not even get close to the detector, so the first cut applied to them is that the muon has to intersect a box as big as the nominal water volume of the detector, whose half-sizes are $1204 \times 893 \times 926 \text{ cm}^3$. The muons which intersect this box in any fashion are saved as the raw muons. These muons are then processed through the detector simulation, and the same cuts are applied to them as for real data events. The ones which survive are then scanned by humans (just like real events), and the human scanner saves the events he considers to be upward-going muons. The only concession made is that the human scanner is not required to fit an event (that would be extremely time-consuming) as long as he can decide without fitting whether the event is an upward-going muon. The resulting energy distribution of scanned monte-carlo events is shown in figure 4.15; the dashed line is for IMB-1 and IMB-2, the solid one for IMB-3.

Using the scanned monte-carlo events, one can now calculate the efficiency of the detector by comparing them to the raw events as generated. Clearly that efficiency will be fairly small; many of the raw muons which just intersect the detector do so at very low energy, not enough to generate the amount of light required by the data selection, or their pathlength in the detector is very small. However, the absolute size of this efficiency is of no particular interest, since it is measured relative to an arbitrary sample of events.

The problem in calculating the efficiency is that it consists of two components which are interrelated: the probability of an event being saved depends on its energy and on the geometry of how it reaches the detector. Both factors are expressed as the effective

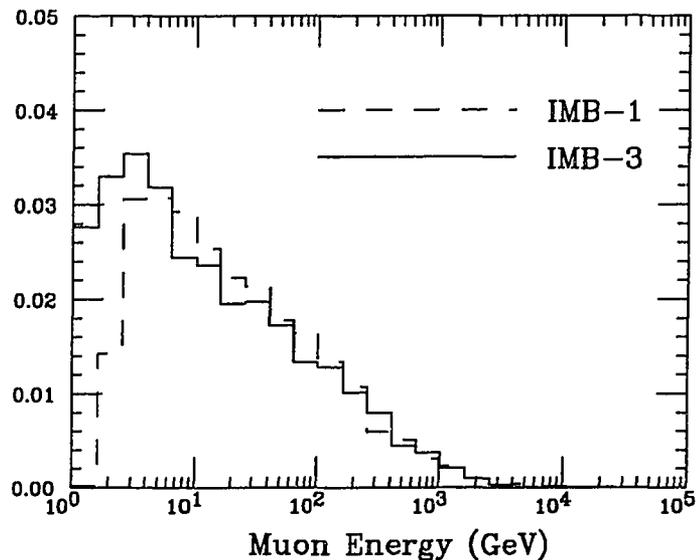


Figure 4.15: Energy distribution of simulated atmospheric neutrino events, after scanning.

area of the detector. It is a strong function of energy at low energies, since muons have to travel a minimum pathlength in the detector to pass the reconstruction cuts. The effective area also clearly depends on the zenith angle at which the muon reaches the detector; for example, it is identically zero for downgoing muons (since those have been rejected). I will express the effective area as a function of zenith angle only. That involves averaging it over azimuth angle; it does clearly depend on azimuth angle (for example, the east and west walls of the detector are considerably smaller than the north and south walls).

The first step is to measure the energy-dependent recovery efficiency, by just determining the fraction of raw events which survive selection and scanning. This is done in 50 bins in $\log E_\mu$ from 1 GeV to 100 TeV. The result is expressed as an effective area in figure 4.16, with the top graph corresponding to IMB-1 and IMB-2, and the bottom graph to IMB-3. In both cases the effective area increases slowly, reaching a plateau at about 60 GeV; since at high energies statistics becomes very poor, all the data above 1 TeV has been lumped together in the last data point. The dotted line in the two graphs is the plateau value, calculated as being the average effective area above 60 GeV; the values correspond to an

efficiency of 50.1% for IMB1 and IMB-2, and 48.9% for IMB-3. That plateau is not at an efficiency of one, since the comparison is of accepted events to muons which intersect an unrealistically large volume. This method guarantees that no events are lost in the monte-carlo simulation; and after all, the absolute value of the efficiency is of no interest, as long as it is used consistently when comparing data and monte-carlo.

The lower end of the energy spectrum of observed events is set by the requirement of at least 400 photomultiplier hits in the detector (250 for IMB-1 and IMB-2). That requires the pathlength of the muon in the detector to be at least about 2 m and therefore imposes a lower energy cutoff of about 1.5 to 2 GeV.

The second step is the calculation of effective area as a function of zenith angle. This should be done using the energy-dependent recovery efficiency calculated above; to implement that, every event is weighted by the efficiency appropriate for its energy, divided by the plateau efficiency. The reason to remove the plateau efficiency from the calculation of effective area is to not have to multiply the efficiencies later on (the effective area should describe the detector efficiency averaged over the correct muon energy spectrum, without having to explicitly include the energy-dependent recovery efficiency).

The raw events were selected by the criterion that they have to intersect the box described above, therefore the area over which we accepted raw events is the “shadow” of that box, projected in the direction the muon travels. If $l_i = (1204, 893, 926)$ cm are the half-sizes of the detector walls and \hat{d} is the unit vector describing the direction of the muon, the total projected area of the detector box is $A_s(\hat{d}) = 4 \cdot \sum_{\{ijk\}} |d_i| l_j l_k$, where $\{ijk\}$ is a permutation of 123. This shadow area is then averaged over the raw events, using the proper energy-dependent weight. From this mean shadow area the effective area is then calculated as the fraction of events (again properly weighted) which are saved, as a function of zenith angle. In summary, if \hat{d}^i , ϑ^i , and E^i are the parameters of the monte-carlo events, Δ^i is one if the event was saved in the scanning, zero otherwise, $w(E)$ the recovery efficiency described above, w_0 the plateau efficiency, then the shadow area A_s , the mean shadow area $\langle A_s \rangle$ and the effective area A_e can be expressed as

$$A_s(\hat{d}) = 4 \cdot \sum_{\{ijk\}} |d_i| l_j l_k$$

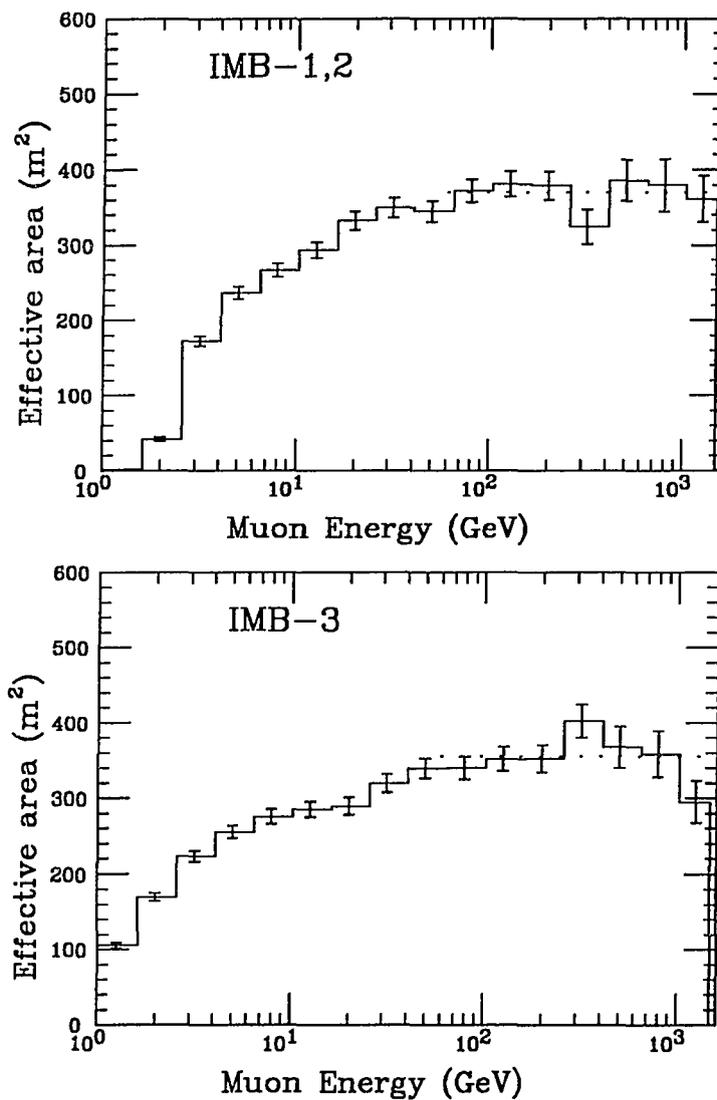


Figure 4.16: Energy-dependent recovery efficiency expressed as effective area, calculated by comparing raw monte-carlo muons to scanned events; top figure for IMB-1 and IMB-2, bottom for IMB-3. The dotted line is the average effective area above 60 GeV.

$$\langle A_s \rangle(\vartheta) = \frac{\sum_{\vartheta} A_s(\hat{d}^i) \cdot w(E^i)/w_0}{\sum_{\vartheta} w(E^i)/w_0}$$

$$A_e(\vartheta) = \frac{\sum_{\vartheta} \langle A_s \rangle(\vartheta) \cdot \Delta^i}{\sum_{\vartheta} w(E^i)/w_0}$$

where the sums are functions of ϑ , that means events are included in the sum only if they fall into the correct bin in ϑ .

The resulting effective area is shown in figure 4.17, again the top figure for IMB-1 and IMB-2, the bottom figure for IMB-3. The dotted line is the effective area averaged over the correct zenith angle distribution of atmospheric neutrino events, for events above an energy of 60 GeV. This mean effective area comes out to be 370.1 m² for IMB-1 and IMB-2, and 355.7 m² for IMB-3.

A note of caution is required here. The effective area can be quoted in many different ways, depending to what one compares it. Due to the way the events have been weighted, the effective area used here is normalized to *events above 60 GeV with the angular and energy distribution of atmospheric neutrinos*, and it rolls off towards a minimum energy of about 2 GeV.

For the purpose of neutrino astronomy we will need the effective area of the detector not as a function of event zenith angle ϑ , but instead as a function of source declination δ . This is calculated by folding the effective area calculated above over the path in ϑ a source at a constant δ traces in the sky (or on our case, in the earth). The result is shown in figure 4.18.

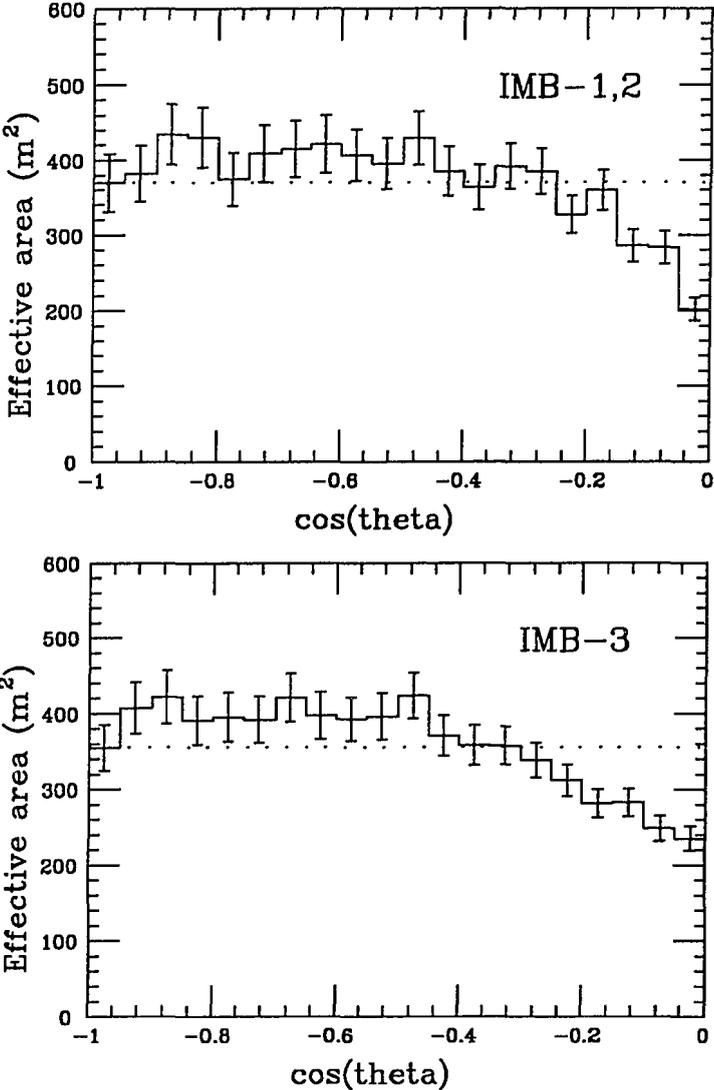


Figure 4.17: Effective area in m² as a function of zenith angle; top figure for IMB-1 and IMB-2, bottom figure for IMB-3. The dotted line is the average over the angular distribution of atmospheric neutrino events.

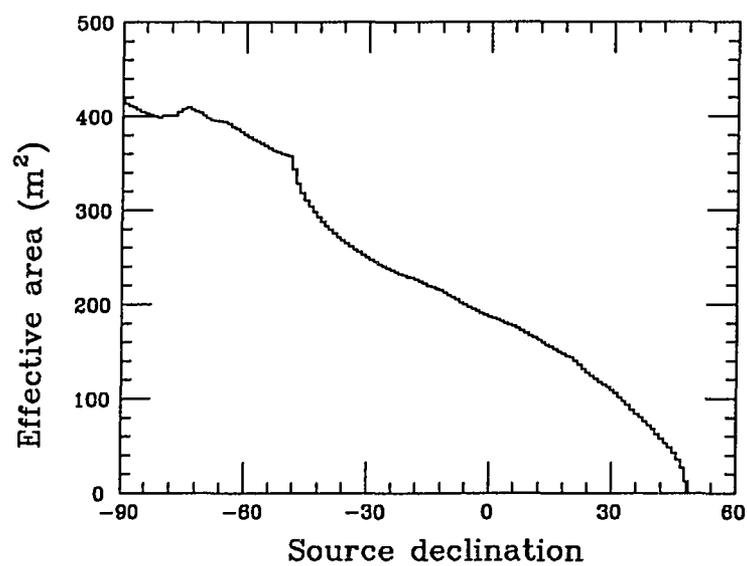


Figure 4.18: Effective area in m^2 as a function of declination in degrees, for a source at constant declination.

Chapter 5

Statistical Interlude II

In this work I ask three questions about the distribution of neutrino arrival directions:

1. Are the neutrino arrival directions statistically significantly distributed non-randomly, in the sense of being spatially correlated with each other ?
2. How many points in the sky show a statistically significant excess of neutrinos, and if there are any, where are they ?
3. Of the suspected sources of neutrinos listed in section 1.4, which ones if any show a statistically significant excess of neutrinos ?

The null hypothesis in all cases is the answer "no" or "none".

The first question is *global*: it is only concerned with the ensemble of all neutrinos. If it finds a significant correlation (for example by finding tight clusters of events), it will not yield an answer on how many clusters there are, nor where they are. The twopoint correlation technique described in section 5.1 will deal with this question.

The second and third question are directly connected and are *local* in that they look for an enhancement of neutrino flux from some direction. The second question asks whether such directions of enhancement exist, and if so where they are, while the third question asks whether there are any such excesses at points which are *a priori* suspected to show an enhancement in neutrino flux. The method used for these questions are skymap ideograms, described below in section 5.2.

Common to both methods is that they require an estimate of the expected background of muons under the assumption of no signal. The technique used to estimate that background is common to the two, so let me describe it here.

The observed data consists mostly of atmospheric neutrino events, and in section 6.1 we will see that the overall flux of muons can be well explained by just atmospheric

neutrinos. We are trying to determine whether the data contains additional neutrinos from astronomical sources, and we need another set of simulated data without such an admixture to compare the measured data to. The observed events have many observational biases, grouped into two categories:

- *Bias in local angle:* The most obvious one is the restriction that only upward-going events have been kept through the analysis. The effective area of the detector in this analysis varies greatly with the event altitude, as figure 4.17 shows. It is also obvious that the effective area is not constant over azimuth either, the south and north sides of the detector are nearly one third larger than the east and west sides. More complicated variations of the efficiency as a function of both altitude and azimuth may also exist; for example at certain angles of the muon relative to the detector walls the Čerenkov light cone strikes the photomultiplier assemblies edge-on, depositing less light, which might lead to a slightly smaller reconstruction efficiency around the diagonals of the detector.
- *Bias in local time:* Again large biases are obvious; the detector has been turned off for long periods, sometimes close to one year; and the detector has undergone several large changes, which modify all efficiencies. Even during normal operation the efficiency (or live fraction) of the detector varies with time. There are several small effects: during the summer months (when many of the collaborating physicists don't have to teach) repair and improvement projects are more frequent, with a corresponding downtime of the detector. Furthermore there is a daily variation of detector performance. If something goes seriously wrong during the night the detector just stops gathering data, and is usually restarted in the morning (around 7:30AM) when the technicians and physicists arrive. This gain in efficiency is compensated by the tendency to do repairs or improve the detector (or just mess around) during the day, leading to a loss of live time.

One could simply say that the two rotations of the earth (the daily and yearly cycle) will effectively smooth out all these biases. However, this approximation is too crude for the small effects being studied here. For example, the combination of the daily bias (the morning fixes all problems effect), the yearly bias (the "summer is time to play" effect) with the spatial nonhomogeneity of the detector will create a non-uniform distribution

of events in astronomical coordinates. Fortunately there is no indication of effects which correlate a spatial with a temporal bias. Such a bias would for example be caused by a part of the detector being switched off for part of the day. The spatial biases are not time-dependent, nor do the temporal biases vary with the event direction. This statement is not strictly true; the spatial biases of the several modified phases of the detector are not necessarily the same. But the differences between them are small, and each modified version of the detector has operated for a long time, so the effect of these three major changes are restricted to a fraction of a day, which is insignificant compared to the total data of over 1200 days.

In summary, the data shows strong efficiency biases as a function of local coordinates and of local time, but these biases should not be correlated. The question is now how to best model these biases. In principle the detector simulation monte-carlo program (described in section 4.6) should correctly simulate the spatial biases. The temporal biases could be measured by studying the distribution of detector live time fraction over real time for all the data analyzed here. However, I will not follow this approach here. First of all, the accumulation of live time would be very complex and error-prone: for IMB-1 there is no reliable real-time clock, only a time scaler relative to the beginning of the run, so the times would have to be reconstructed from the detector logbooks; even for IMB-3 the clock systems sometimes fail. Also, this procedure would rely completely on the monte-carlo simulation of the detector being correct, which does not seem a prudent choice.

However, there is a more natural way of solving this problem. Why not just use the measured data to estimate the bias? For example, if the monte-carlo detector simulation were to be used, one would fine-tune it until it agrees with the distributions of the measured data. So one might as well just use the distributions of the measured data.

The standard technique is the following. Assume that the signal we are looking for is small, so the data set is dominated by the background of atmospheric neutrino events. As explained above, one also assumes that the biases of the data are in local angles and in local time, but not correlated between the two. Then just use the observed distribution of measured angles and measured local times to generate simulated background events.

In principle one could generate those distributions (for example a histogram of local times, then fit a smooth function to that histogram). However, the best estimator of the distribution is the actual input data. So the procedure is: pick the local angles from event i and the local time from event j , and use those to calculate the astronomical coordinates. In practice one generates the second index j as a random permutation of the numbers $1 \dots N_e$.

This technique generates new data sets, in which the biases inherent in the distribution of local angles and local time of the events have been preserved, but any bias from a correlation between the two are removed. If there were a neutrino point source signal in the data one could think of it as a correlation between spatial and temporal bias: there is this hot spot in the detector, which rotates around the earth axis once per sidereal day. Such a correlation would be removed by this method; therefore such an effect could be found by comparing the real data with many sets of such simulated “scrambled” data. This technique is conservative in that it may remove too many correlations. For example, assume that there were a source of neutrinos which is a uniform ring at a constant declination (along the equator in equatorial coordinates). In detector coordinates it would be mapped into a ring (diagonal in local coordinates). However, using this method of estimating the background, the excess along this ring would persist into the simulated data, so by comparing data and simulation such a source could not be found. Fortunately, there is no reason to expect such a source of neutrinos. In general, if there is a point source of neutrinos, it will increase the flux at that particular declination, and thereby the background estimate for that declination will increase slightly. This results in the flux excess from the point source being *underestimated*; but this is a small effect, and an error in the conservative direction, so not a cause for great concern. However, there is one case where the underestimation is complete: the celestial poles; this background simulation technique will always make data and background match at the poles. Therefore no point sources of neutrinos can be found at each pole. The simplicity and safety of this method is worth this loss.

This statistical technique, where the expected distribution of the data is not inferred from first principles but instead the observed data is used to generate its own distribution is

known as bootstrap statistics, since it allows bootstrapping the required distributions without exact *a priori* knowledge.

5.1 Two Point Correlation

For studying the first question and search for global correlations I use the twopoint correlation technique. Its discussion will begin with a detailed description of the algorithms involved, followed by an interpretation.

Given is a set of neutrino events i , $i = 1, N_e$, characterized by their local angles (the result of the fit to the muon track) (ϑ_i, φ_i) , the local sidereal time τ_i , and the angular resolution estimate σ_i , determined as described in sections 4.4.3 and 4.5. Note that the muon fit direction is the direction in which the muon travelled, whereas for astronomical purposes one is interested in the direction from which the neutrino came, which is obviously the inverse. Also note that the sidereal time can be trivially determined from the local time and date of the event. From these quantities one calculates the direction from which the neutrino arrived, in astronomical equatorial coordinates, expressed in right ascension α_i and declination δ_i . Right ascension (often abbreviated RA) is expressed in hours and minutes; in the following I will assume that the units are converted to degrees or radians when appropriate. Declination is measured in degrees, 0° at the equator, positive values on the northern hemisphere.

The core of the twopoint correlation technique is to study the distribution of relative distances between all possible event pairs

$$\cos \epsilon_{ij} = \sin \delta_i \sin \delta_j + \cos \delta_i \cos \delta_j (\cos \alpha_i \cos \alpha_j + \sin \alpha_i \sin \alpha_j).$$

There are $N_e(N_e - 1)/2$ such distances, a reasonably large number. Let $D_d(\epsilon)$ be the distribution of pair distances, where the subscript indicates that this is distribution for the observed data, normalized such that $\int_0^{180^\circ} D_d d\epsilon = 1$ (in reality the distribution is approximated by a histogram, and the integral becomes a summation). Equivalently, one can express the same data as a distribution of $\cos \epsilon$: $D_d(\cos \epsilon)$, where the equivalent normalization condition is $\int_0^1 D_d d \cos \epsilon = 1$ (using the same symbol for the distribution

of ϵ and of $\cos \epsilon$, since there is no danger of confusing them). In the following I will express everything in terms of ϵ and not of $\cos \epsilon$.

It seems obvious that this distribution of pair distances contains some information about correlation between the arrival directions of neutrinos. For example, in the extreme case where many events come from just one place in the sky (so their directions are very correlated), all pair distances will be very small, so $D_d(\epsilon)$ will show a large spike at very small distances. But to quantitatively study $D_d(\epsilon)$ one has to determine what to expect for it in the absence of any extraterrestrial sources of neutrinos. The method used here is to compare the distribution with one obtained from a data sample which has the effects of any astronomical sources of neutrinos removed.

To do this, I use the background generation technique described above; it generates new data sets (α'_i, δ'_i) which preserve the biases inherent in the distribution of (ϑ_i, φ_i) and τ_j , but removes any bias from a correlation between the two. The distribution of pair distances expected from background $D_b(\epsilon)$ is then calculated by generating many such “scrambled” sets of simulated data. Note that $D_b(\epsilon)$ has much better statistics than $D_d(\epsilon)$. Again, the normalization is $\int_0^{180^\circ} D_b d\epsilon = 1$.

The twopoint correlation function $C(\epsilon)$ and the cumulative twopoint correlation function $I(\epsilon)$ are then defined as

$$C(\epsilon) = \frac{D_d(\epsilon)}{D_b(\epsilon)} - 1$$

$$I(\epsilon) = \int_{\epsilon'=0}^{\epsilon} C(\epsilon') d\epsilon'$$

The normalization of the correlation function follows from those of D_d and D_b : $\int_0^{180^\circ} C d\epsilon = 0$, which is equivalent to $I(180^\circ) = 0$. Furthermore, the statistical error of $I(\epsilon)$ can be determined easily; the statistical errors of D_d and D_b are just poissonian (for each bin in angle, one counts event pairs where ϵ is in that angle bin), and from that the statistical errors $\Delta^C(\epsilon)$ and $\Delta^I(\epsilon)$ of $C(\epsilon)$ and $I(\epsilon)$ can be determined by simple error propagation. Using these statistical error, either of the two can be expressed as a (positive or negative) number of standard deviations, for example

$$\sigma^I(\epsilon) = I(\epsilon)/\Delta^I(\epsilon)$$

A peak in $C(\epsilon)$ with a large positive value indicates that this distance between events is preferred in the data compared to the background; that means, the event directions are correlated, with a correlation length being the peak location. Similarly, a positive value of $I(\epsilon)$ indicates that events are correlated at angles up to the peak angle. It is often helpful to not use $I(\epsilon)$ itself in this, but $\Delta^I(\epsilon)$.

The problem with this twopoint-correlation technique is that it may indicate a correlation, but it will not indicate where in the sky it occurs, which events are correlated, or what the spatial shape of the correlation is. The only way to intuitively understand the results of the twopoint-correlation test is to generate “fake” data sets which may have signals injected into them.

As an example, a fake data set was generated. The real events were scrambled as described above; additionally, they were smeared by their appropriate resolutions. Then four “point sources” of neutrinos were added, each with 5 properly smeared out events, and placed at positions (2:39, +33°), (14:14, -9°), (0:08, -22°) and (12:46, -52°). These four positions were chosen randomly. For each event at those point source one event was removed from the data, to keep the total number constant; special care was used to make sure that this replacement does not significantly alter the sidereal time or zenith angle distribution of the data. The background simulation was carried out independently from the one for the real data. The results are shown in the following figures. The first one, 5.2 shows the distribution of event pair separation angles. The data points are the “fake real” data, the histogram is the background expectation. In this form, no effect is immediately obvious. Figure 5.2 shows the twopoint correlation function as a function of angle, that is the ratio of measured to expected pair separation angle distribution minus one. It is obvious that the first few bins are significantly high compared to the expectation of zero. In the next figure, 5.3, the cumulative correlation is shown. The significance of the first few bins is integrated, and there is quite a bump between angles of about 10° to 20' *deg*. The last figure, 5.4 shows the same cumulative twopoint correlation function, now expressed in standard deviations from the expected value of zero. The bump between 10° and 20° is quite significant; the highest significance is 3.4σ at 19°. The interesting feature is the location of that highest correlation (that is the correlation length): the correlation length should be about the typical event resolution, which is on

the average between 5° and 7° . The highest correlation is indeed found at 7° , but due to the integration and the increasing statistics, the highest *significance of the cumulative correlation* occurs at over twice that angle.

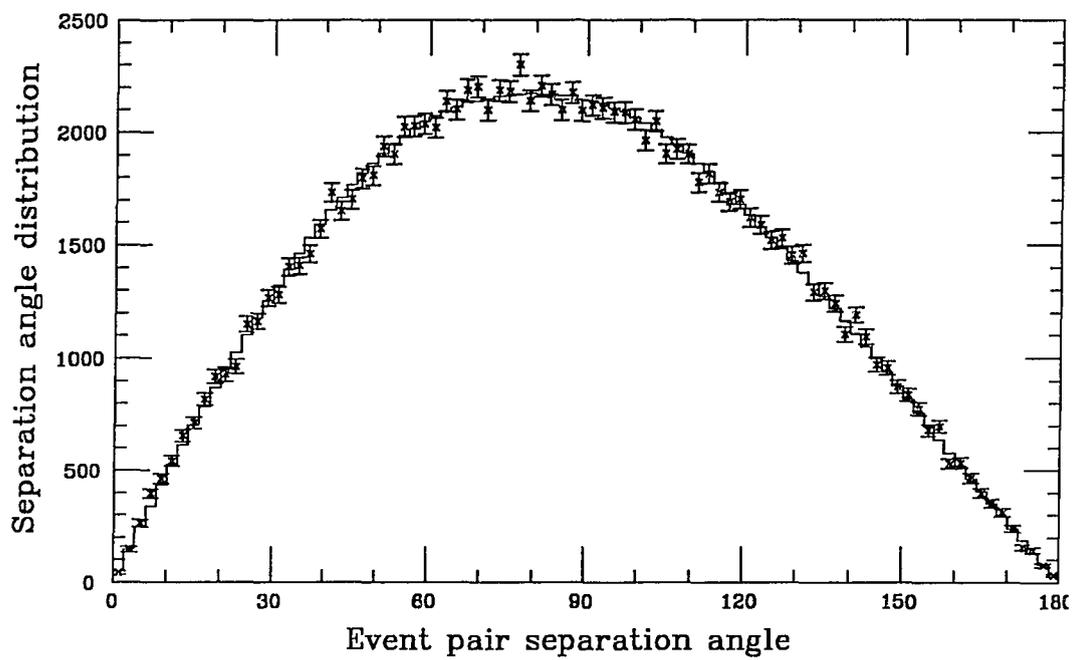


Figure 5.1: Distribution of pair separation angle

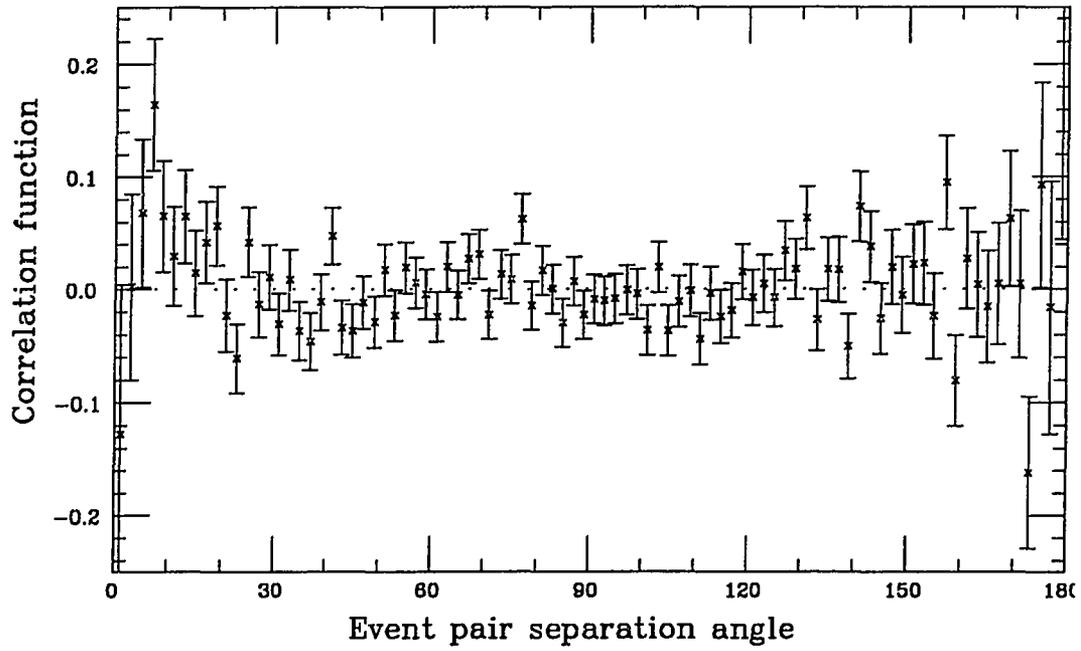


Figure 5.2: Twopoint correlation as a function of pair separation angle

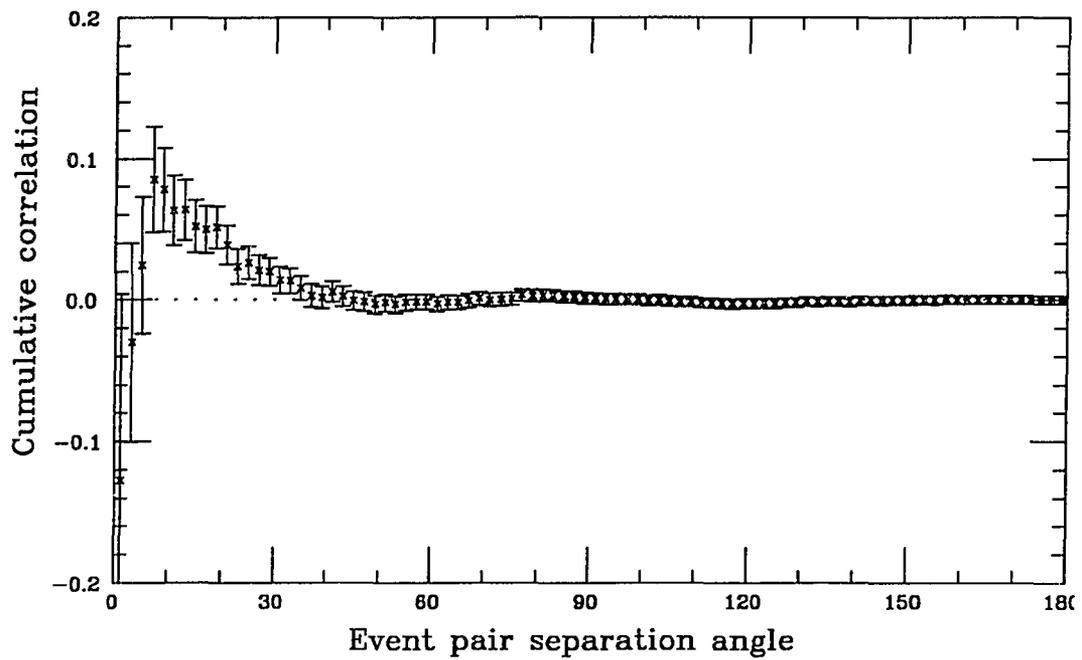


Figure 5.3: Cumulative twopoint correlation as a function of pair separation angle

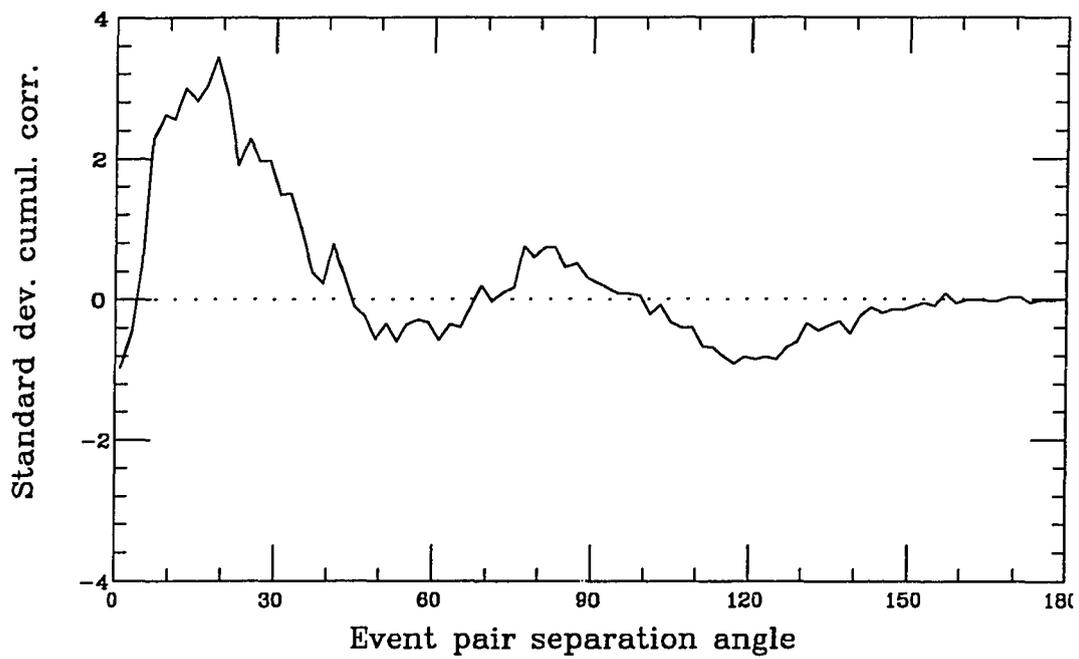


Figure 5.4: Standard deviation of the cumulative twopoint correlation

5.2 Skymap Ideograms

The technique explained in this section estimates whether there is an excess flux of events for each point in the sky. In essence, it performs a hypothesis test for each point in the sky, with the null hypothesis \mathcal{H}_0 being “the flux at this point is not larger than the one expected from background simulation”, and the signal hypothesis is “the flux at this point is higher with the one expected from background simulation”. An appropriate definition of *flux* will be provided later on. Note that the test is performed many times, once for each point in the sky. It yields a confidence level for refuting \mathcal{H}_0 at each point, which is the complement of the chance probability of \mathcal{H}_0 being correct. Since \mathcal{H}_0 and \mathcal{H}_1 are exclusive and their union is complete (that means if one is true the other one must be false), in this rare case we can actually employ equation (3.1) to calculate the probability of \mathcal{H}_1 , which is the complement of the chance probability of \mathcal{H}_0 , and is therefore directly the confidence level for refuting \mathcal{H}_0 . All these convolutions boil down to the fact that in this fashion we have obtained a map of the probability of having observed a flux excess from a particular point in the sky.

What can such a map be used for? The first use is to look for places where the probability of observing a flux excess is particularly low (that means, the observed flux excess is very unlikely). These locations may be point sources of neutrinos. However there is a complication in determining the statistical significance of such a search, which will be described below. The second use is to check an expected source of neutrinos whether it shows a flux excess. This case is fairly simple; the probability of having observed a flux excess at this point is directly given by the confidence level at which \mathcal{H}_0 is refuted at this point.

The first case is more complicated, since in looking for a particular place in the sky where the flux is specially unlikely, we have made a choice with knowledge of the data. The chance probability of observing such a high flux excess *anywhere in the sky* is far larger than the chance probability of observing it in a particular place. The issue here is the difference between knowing where to expect a flux excess, that is doing a statistical test with *a priori* knowledge of its parameters, and choosing a place where one looks for a

point source, that is doing a statistical test with *a posteriori* selection. In this case, the chance probability has to be determined in a separate process.

The concept of the method used to obtain the improbability map is the following. First take the observed event directions and turn them into a flux map, using as much information about each neutrino-induced event as possible; in particular making use of the reconstruction resolution and energy-dependent angular resolution estimate obtained in section 4.5. This flux map shows the muon flux (events per unit area per unit time) coming from each particular direction. Clearly it is a 2-dimensional map, most conveniently projected onto the surface of a unit sphere. Second, estimate the expected background to that flux map, calculated under the assumption that there is no signal. Third, compare the measured flux to the background expectation at each point, which yields a flux excess (observed flux minus mean expected flux) at that point, along with an estimate of the statistical significance of that excess: the chance probability of the the background fluctuation up to the value of the measured flux at that point.

To search for *a priori* unknown places with a flux excess two more steps are required. Fourth, look for peaks in the map of flux excess statistical significance, using some suitable definition of peak. Lastly, compare the peaks found in the data with the background expectation for peaks.

Before discussing the intricacies of these analysis steps, one preliminary problem has to be cleared up. All the maps mentioned above are functions of two parameters on the unit sphere, the declination δ and the right ascension α . In this chapter I will express everything in those parameters; one can easily use the same method in local coordinates (ϑ, φ) , or galactic coordinates. Clearly we will not analytically calculate all the maps as continuous functions of two parameters, but instead sample them on a finite set of points. In principle it would be advisable to do this *bin-free*: one should not impose an *a priori* selection of points on which to calculate these quantities, instead let the measured data determine at which positions to calculate the quantities involved. Unfortunately, I am not aware of any bin-free method of doing sky maps which is similar in power to the method used here, and not prohibitively complicated or computer-time intensive.

Therefore these quantities have to be calculated on some reasonably selected set of points, which we shall call a grid. This grid shall fulfill the following requirements:

- The position of each grid point should be easy to calculate.
- The number of grid points should not be unnecessary large, since the amount of computer time spent on these calculations is about proportional to the number of grid points.
- It should be fine compared to the smallest scale of expected or observable phenomena. The smallest length scale in the problem measurable in this work is given by the events with the smallest angular resolution, which is about 3° , see section 4.4.3. I will therefore impose that the largest distance between any neighboring grid points is one degree or smaller. Not only is that fine enough that there is no danger of large variations in measured quantities over the distance between grid points, it is also a round number and easy to remember.
- The area of each cell (the little polygon or “square” around each grid point formed by drawing lines between adjacent points on the grid) should be fairly constant. Otherwise anytime an algorithm averages or compares quantities calculated at different grid points they would have to be weighted in some fashion, which makes the process time-consuming and error prone. Since no grid distance shall be larger than 1° and the number of cells should be as small as possible, this requirement means that the cell area should be about $(1^\circ)^2$, and ideally there should be 41253 cells.

These requirements already rule out most simple gridding algorithms. For example, a rectangular grid of 360 equidistant points in α and 180 equidistant points in δ does have no distance larger than one degree, but the area of cells varies by more than a factor of 100. A grid of 360 equidistant points in α and 114 equidistant points in $\sin \delta$ has cells which are all $(1^\circ)^2$ in area, but the largest distance between grid points is at the poles exceeds 10° . A similar grid but with 13132 points in $\sin \delta$ fulfills most of the requirements, but would be extremely too fine grained, having nearly 5 million grid points. The optimal solution would be a regular pythagorean body (tetrahedron, cube, dodecahedron, icosahedron), but no regular body with more than 20 faces exists.

The solution I adopt here is a doubly linear grid in α and δ , but with a varying number of grid points in α . The grid points are on 179 rings of equal declination at $\delta = -89.5^\circ, -88.5^\circ \dots -0.5^\circ, 0.5^\circ \dots 89.5^\circ$; therefore the north and south edges of the cells around each grid point lie at integer values of δ . The number of grid points in the α direction varies from ring to ring, and is chosen to be the smallest integer such that no edge of a grid cell is longer than 1° . That results in number of grid points on each ring varying from 7 at the poles to 360 at the equator. The grid points are then spread equidistant in α between 0° and 360° . This scheme yields a grid where each edge of grid cell is parallel to lines of constant declination or ascension, but the east and west edges don't line up vertically. The total number of grid points and therefore cells is 41702; since the study of upward-going muons in this work is limited to declinations below $+60^\circ$, only 38832 of these grid points are used. The grid cell area is fairly uniform, as one can see from the ratio of 41253 expected cells to 41702 cells in this scheme. The area does not deviate by more than 10% from the ideal value, except for the 5° closest to the poles. It is an unfortunate disadvantage of the celestial sphere that it has two poles, at which spherical geometry becomes indeterminate.* Since the number of grid points in α can be precalculated for each of the 179 possible values of δ , and since α can be found by linear interpolation, this method does not use too much computer resources, and it has the advantage of having a nearly optimally low number of grid points.

One unfortunate disadvantage of this scheme is that the concept of "nearest neighbor" of a grid point becomes more complicated. If the grid were strictly rectangular, each grid point would have 8 neighbors. In the scheme used here only the two neighbors to the left and right can be defined trivially. Therefore a different definition is adopted: a nearest neighbor is any grid point closer than 1.82° , except within 4° of the poles, where that limit has to be adjusted down to 1.8° , 1.65° and 1.2° . These limits are hand-crafted to make sure that no grid point has less than 8 neighbors, and have to be corrected at the poles to prevent any grid point from having more than 10 neighbors. That definition leads to each grid point having 8.11 nearest neighbors on the average.

*Astronomers should have long flattened the cosmos to fit on a square sheet of paper more conveniently.

Figure 5.5 shows a little part of the grid around the south pole. The left figure is a rectangular projection (in declination and right ascension, expressed in degrees), the right figure a circular projection around the south pole. One bin which happens to have 10 nearest neighbors has been highlighted.

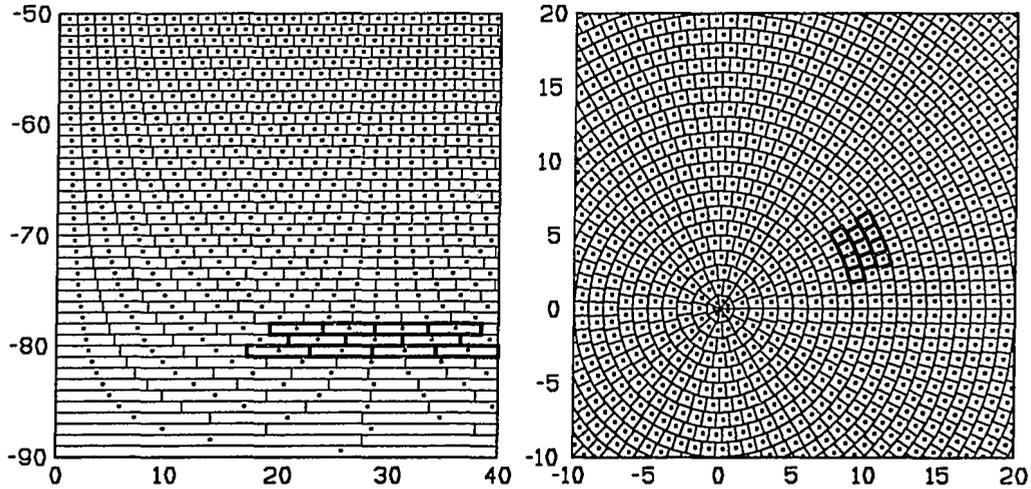


Figure 5.5: Sample of the grid used to calculate ideograms, showing a small area around the south pole. The left figure is a rectangular projection, with declination along the vertical axis and right ascension along the horizontal axis, both in degrees. The right figure is a polar projection. A sample bin with 10 nearest neighbors is highlighted.

With that grid in hand, the muon flux at each point is now calculated by smearing each event over a two-dimensional Gaussian for the appropriate resolution. If $\epsilon_i(\delta, \alpha)$ is the angular distance between a grid point and event i , the smeared muon flux F is defined as

$$F(\delta, \alpha) = \sum_i \frac{5}{\sigma_i} e^{-\epsilon_i(\delta, \alpha)^2 / 2\sigma_i^2}, \quad (5.1)$$

where σ_i is the angular resolution of event i , determined as described in section 4.5. The scale factor of 5 was chosen arbitrarily to make the flux similar in magnitude to the observed number of events.

The next step is the generation of background flux. The scrambling technique for generating new samples of events which keep the biases in local angle and local time has

been described above. When applying this technique here, the question arises whether the event resolution σ_i should be kept together with the event local angles (ϑ_i, φ_i) or whether it should be permuted together with the sidereal times τ_i of the events. I selected the first option; the justification being that the angular resolution of events should not be time-dependent, but since it depends on the energy measurement (which is coupled to the pathlength of the muon) it may be correlated with the local angles of the events. Using the scrambled events, many background flux maps can be calculated using equation (5.1), let me call them $F_j^b(\delta, \alpha)$, where j is the index of permuted background maps.

The next step is to combine all these background flux maps into the distribution of flux expected from background $D(F, \delta, \alpha)$, which is normalized to unity, and from it the mean expected flux $\tilde{F}(\delta, \alpha)$:

$$\int_0^\infty D(f, \delta, \alpha) df = 1$$

$$\tilde{F}(\delta, \alpha) = \int_0^\infty f \cdot D(f, \delta, \alpha) df.$$

Each of the distributions in $D(f, \delta, \alpha)$ is known as an ideogram. Obviously $D(f, \delta, \alpha)$ is not a continuous function of F either; I use a histogram with 50 bins for it. The limits of the histogram were set such that the largest flux observed in the data (which conveniently happens to be a little less than 10) can be accommodated. In the background estimation, flux maps cases of $F > 10$ do occur; such overflows in the flux are simply lumped together in the last bin. As will be shown below, this clipping does not do any harm. Note that $D(f, \delta, \alpha)$ contains 38832 histograms, each with 50 bins, so it is quite a large set of data and cannot be conveniently displayed graphically. To prevent roundoff-errors, the mean flux is not actually calculated by integrating $D(f, \delta, \alpha)$ but directly from the $F_j^b(\delta, \alpha)$.

At this point, we are nearly done. The flux excess $E(\delta, \alpha)$ is just the difference $F(\delta, \alpha) - \tilde{F}(\delta, \alpha)$; note that it may be negative (corresponding to a deficit). The chance probability $C(\delta, \alpha)$ of the measured flux being an excess can be determined by comparing the measured flux to the distribution of background flux:

$$C(\delta, \alpha) = \int_{f=F(\delta, \alpha)}^\infty D(f, \delta, \alpha) df$$

Since $D(f, \delta, \alpha)$ is actually a histogram the chance probability is calculated by summation. The bin in which $F(\delta, \alpha)$ lies is *included* in the sum; that tends to *overestimate* the chance probability, thereby *underestimating* the significance of a possible excess; this is the conservative approach. Note that in the same fashion for a flux deficit could be calculated (by integrating from 0 to $F(\delta, \alpha)$); due to the inclusion of the *whole* bin in which $F(\delta, \alpha)$ lies, it would be a little larger than $F(\delta, \alpha)$. But we are not interested in flux deficits.

To reliably estimate $C(\delta, \alpha)$ one has to generate enough iterations of the background to make the above integral non-zero; that means that to achieve chance probability levels of 10^{-4} at least 10000 iterations are required.

Since the excess improbabilities $C(\delta, \alpha)$ can be fairly small numbers, it is convenient to express the logarithmically, in a quantity called significance $S(\delta, \alpha)$:

$$S(\delta, \alpha) = -\log_{10} C(\delta, \alpha)$$

This process of calculating a smeared event flux map, comparing it to a distribution of background maps, and obtaining the excess significance of each point, will be demonstrated in a simpler one-dimensional fashion in section 6.3. This *ideogram* technique, named after the distributions contained in $D(f, \delta, \alpha)$, has been independently developed by the Fly's eye group.¹⁷

Chapter 6

Results

6.1 Atmospheric Neutrinos

The first and easiest result is the calculation of the flux of muons generated by atmospheric neutrinos as described in section 1.2. In this section I assume that all observed events are due to atmospheric neutrinos, and none (or only a small fraction) are created from astronomical point sources of neutrinos.

All the information required for this simple calculation has already been presented: the live time T in section 4.2.1, the efficiency and effective area $A_e(\vartheta)$ in section 4.6. As mentioned in section 4.2.1 the data from the 8mm tapes is excluded from this calculation. With these results, the flux of muons is

$$\mathcal{F}_\mu(\vartheta) = \frac{N_{\text{evt}}(\vartheta)}{2\pi \text{ sr} \cdot A_e(\vartheta) \cdot T}$$

where the $2\pi \text{ sr}$ is the solid angle observed by the detector, and the units have to be converted to $\text{cm}^{-2}\text{s}^{-1}\text{sr}^{-1}$. The resulting differential flux as a function of ϑ is shown in figure 6.1; the top graph is for IMB-1 and IMB-2, the bottom one for IMB-3. The data are the points with error bars, the monte-carlo prediction is the histogram. The total flux (integrated over ϑ) is given in table 6.1.

Two different theoretical calculations of the expected flux have been used, the one by Volkova⁹⁵ and the one by Lee and Koh⁵⁸; the histogram in figure 6.1 uses the Volkova fluxes. Good agreement between data and the monte-carlo using the Volkova flux is found, whereas the Lee flux expectation is too low. The errors given are only the statistical errors; they include the statistical error on the determination of the effective area. The absolute normalization of the monte-carlo prediction is uncertain by at least $\pm 10\%$, probably by $\pm 20\%$, as the comparison of the calculations of Volkova and Lee

shows. The systematic error of the data is small compared to that uncertainty; the largest factor is the determination of the effective area and its dependence on the energy efficiency, which is estimated to be uncertain by $\pm 5\%$. Since that uncertainty of the effective area affects the measured and the expected flux in the same way it causes no problem when comparing the two. The largest relative uncertainty between the two stems from the live time determination, which is estimated to be better than $\pm 2\%$.

Period	Events	Livetime hours	Measured flux in $10^{-13} \text{ cm}^{-2} \text{ s}^{-1} \text{ sr}^{-1}$	Expected	
				Volkova	Lee and Koh
IMB-1 and IMB-2	200	10920	2.19 ± 0.16	2.14 ± 0.06	2.04 ± 0.06
IMB-3	268	11004	3.02 ± 0.19	2.80 ± 0.07	2.53 ± 0.06

Table 6.1: Total flux of muons from atmospheric neutrinos

By comparing the expected to the measured flux, taking into account the 20% systematic uncertainty of the expected flux, one finds that at 90% confidence level, a flux of less than $0.48 \cdot 10^{-13} \text{ cm}^{-2} \text{ s}^{-1} \text{ sr}^{-1}$ for IMB-1 and IMB-2, and less than $0.79 \cdot 10^{-13} \text{ cm}^{-2} \text{ s}^{-1} \text{ sr}^{-1}$ for IMB-3 can be caused by neutrinos other than atmospheric neutrinos. These limits are fairly bad, due to the large systematic uncertainty in the expected flux.

6.2 Neutrino Direction Correlation

The twopoint correlation technique has been introduced in section 5.1, and was tested on a “fake” data set, which had four artificial point sources injected. In this section, I will present the results of using it on the observed data. The algorithm described in section 5.1 was used on the 496 throughgoing events, and 1000 iterations were done to generate the background. Note that there are 122760 event pair distances between 496 events, so the statistics in the background are quite good.

The results are shown in the four figures below. Figure 6.2 is the comparison of the distribution of event pair distances between data and background estimation. In this form, no gross difference is visible between the two, although the first few bins are systematically high. Figure 6.3 shows the twopoint correlation function, that is the

ratio of measured to expected event pair separation angle, minus one. Except for the second bin (where statistics is low) the first ten bins are all above the expected value of zero, most of them by about one standard deviation. Among them, there is no clear candidate for a largest correlation function, the maximum may occur at 7° , 9° or 17° . The cumulative correlation function in figure 6.4 doesn't clarify the situation either; the cumulative correlation rises to its maximum value of about 4% at 9° and stays there up to about 20° . Lastly, figure 6.5 shows the cumulative twopoint correlation expressed in standard deviations; the peak significance of 2.6σ is found at a correlation length of 19° .

These four figures can be compared to figures 5.1 through 5.4, the equivalent ones for data which had four point sources of five events each injected artificially. A more detailed comparison of the shape and height of the two cumulative correlation functions however shows some differences. In the fake data, the cumulative correlation rises quickly, reaches its maximum at 7° and then drops off slowly; in the real data it doesn't rise that fast, and reaches a plateau between 9° and 20° . This might suggest that the correlation length in the data may be larger than the one used for creating the "fake" data set, although that conclusion is by no means significant; and as we will see below, the two functions statistically agree very well. The events injected at the four point sources in the fake data set were all smeared by the resolution appropriate for them, as determined in section 4.4.3 and section 4.5.

The magnitude of the cumulative correlation of the data, when compared to the fake data, suggests that the number of events which are correlated is smaller than the 4×5 events used in the fake data. To test that, the cumulative correlation function for the fake data was scaled down by a multiplicative factor, and that scale factor was determined by a χ^2 fit between the two cumulative correlation functions. The fit results in a scale factor of 0.843 for the fake data, indicating that the equivalent number of correlated events in the real data is about 17. The statistical agreement between the two cumulative correlation functions using that scale factor is excellent, yielding a χ^2 probability of 99.6%. This near-perfect match between the observed data and the fake data set when properly scaled down may indicate that the correlation in the data is indeed due to some point sources with an angular distribution of events similar to the one used for the fake data.

In summary, the data shows a correlation between event locations, statistically significant at the 2.5σ level. The shape of the correlation function agrees fairly well with what is expected from a small set of point sources with a few events each. The correlation length of the data seems to be a little too high, but that difference to the correlation function expected from the fake data set is not statistically significant.

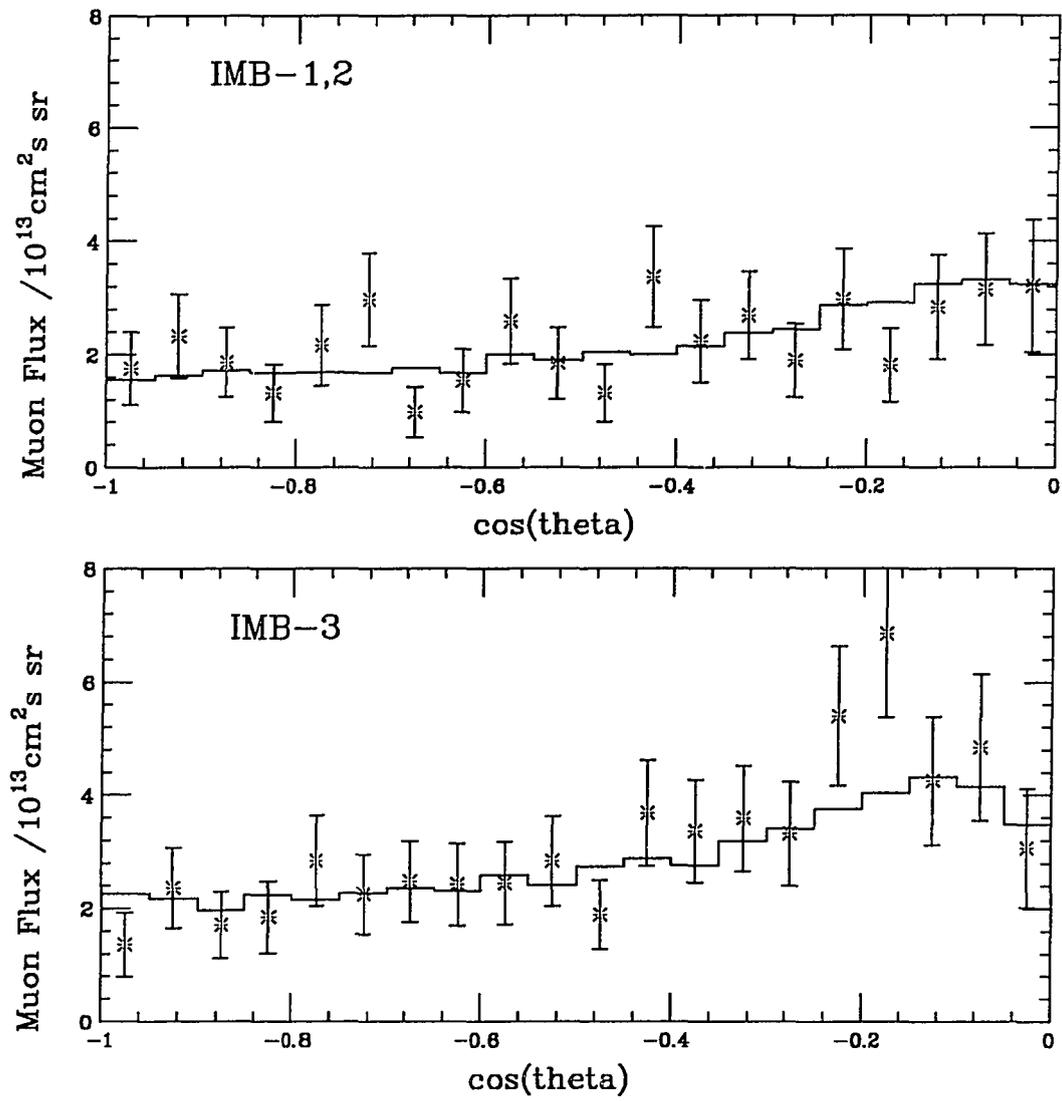


Figure 6.1: Flux of muons from atmospheric neutrinos. Top figure for IMB-1 and IMB-2, bottom for IMB-3. The histogram is the monte-carlo prediction.

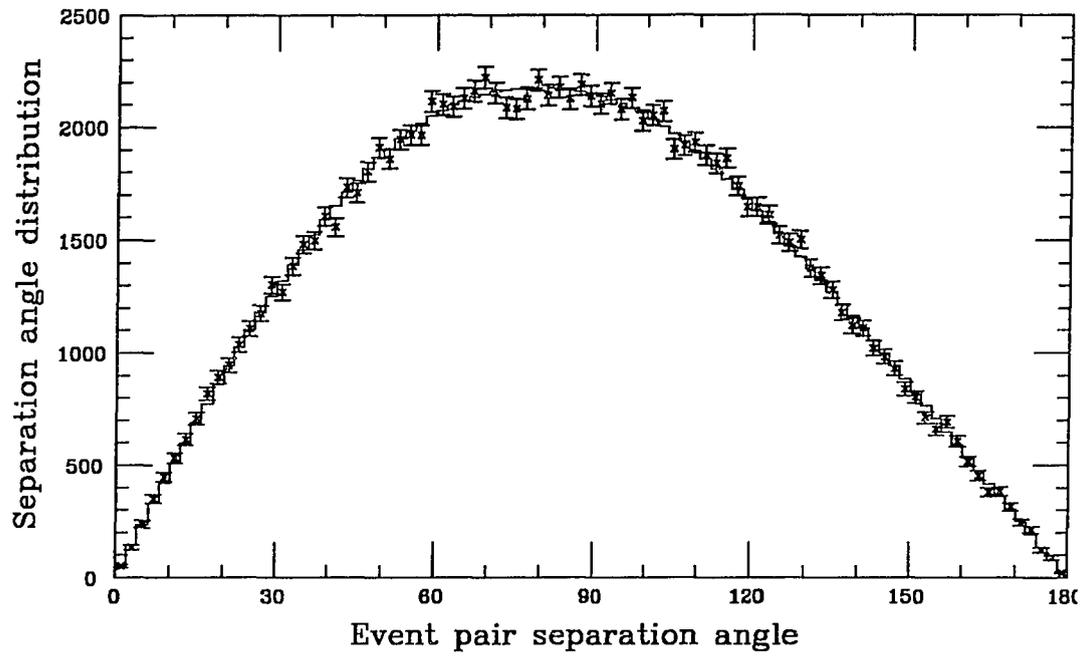


Figure 6.2: Distribution of pair separation angle

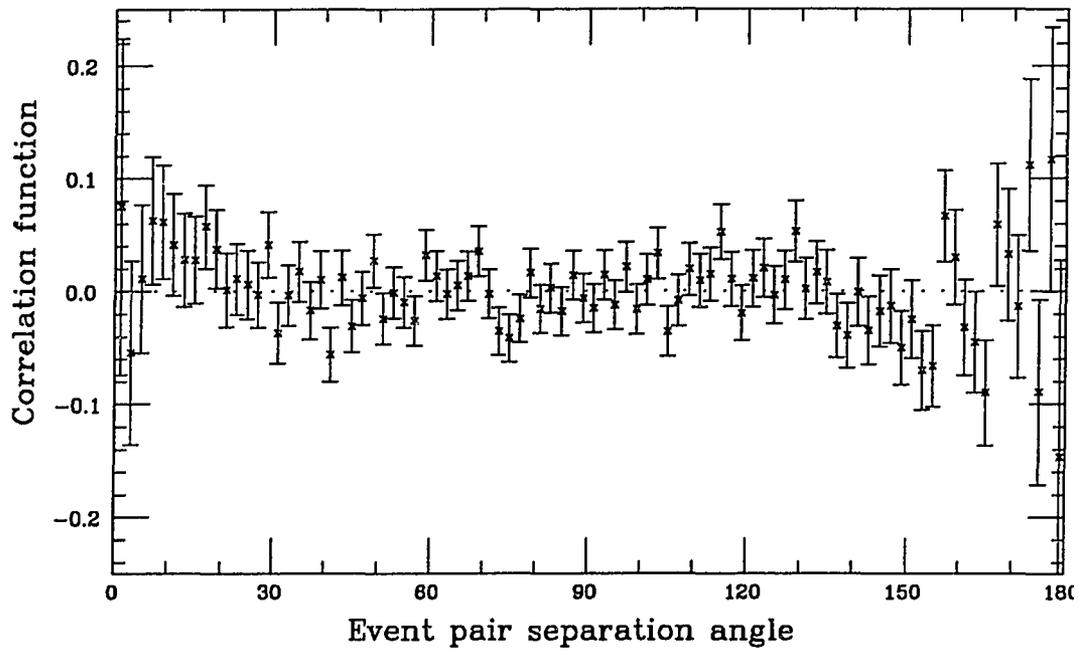


Figure 6.3: Twopoint correlation as a function of pair separation angle

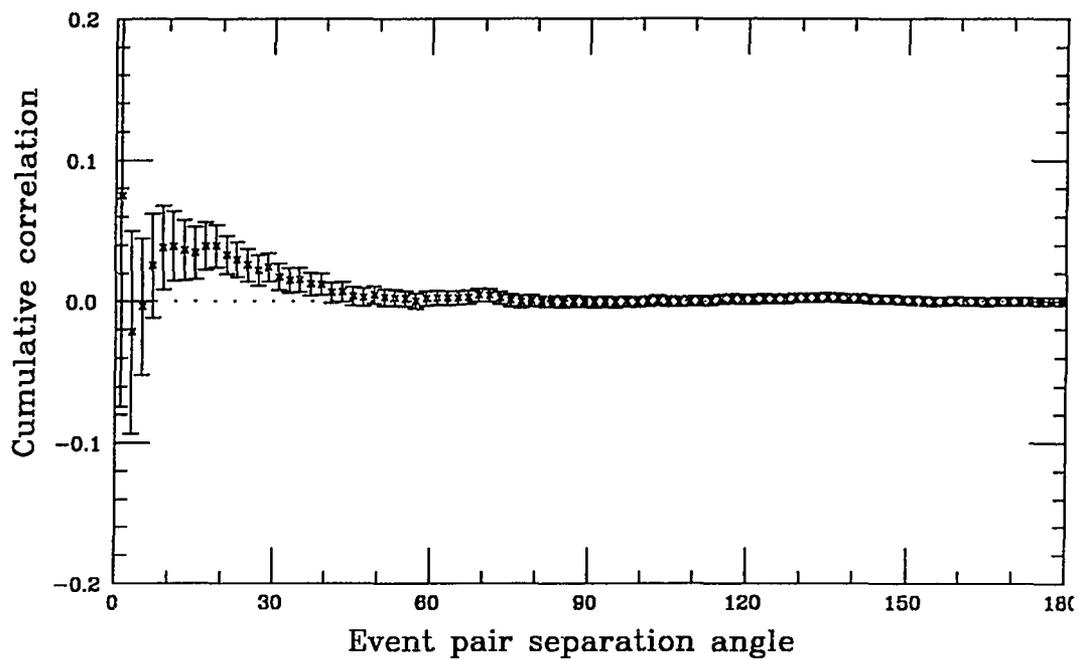


Figure 6.4: Cumulative twopoint correlation as a function of pair separation angle

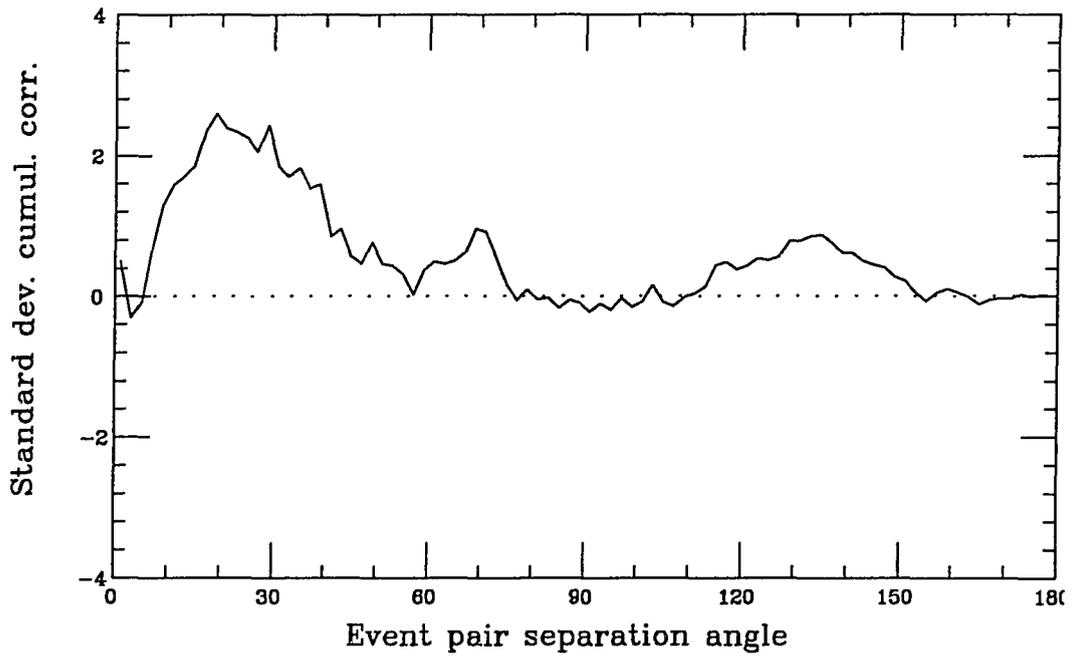


Figure 6.5: Standard deviation of the cumulative twopoint correlation

6.3 Flux Excess from the Galactic Plane

The first application of the ideogram technique to search for a *localized* excess of events from a certain direction is a simplified version. The aim here is to study whether there is an excess of events from the galactic plane. Most observable sources of neutrinos are expected to be in our galaxy, and indeed most sources of high-energy photons observed so far are in our galaxy. The reason is essentially the same as the reason for Olbers' paradox to be wrong: the galaxy is an upward fluctuation in matter density. There is a lot more matter (some forms of which are expected to emit neutrinos) in our neighborhood up to distances of tens of kiloparsecs than one would expect from the mean density of the universe. Furthermore, in Olbers' paradox the uniform radiation background is of interest, but in this work a uniform background of neutrinos created by many weak sources at large distances would not be detectable, due to the self-normalizing search technique used here.

The feature which makes a search for an excess from the galactic plane so simple is that it is a *one-dimensional* search, using the technique described in section 5.2: the events are binned in galactic latitude l . The other simplification is that the binning is made much coarser; we are interested in integrating all of the galactic plane in one bin, so I choose to use 15 bins in l , each 12° wide. The number of bins should be odd, so there is one bin centered around $l = 0^\circ$, and the bin-width should be about twice the typical angular resolution of events; therefore this choice is natural. Let me repeat the formulas from section 5.2, rewritten for this one-dimensional case. The flux of muons is $F(l)$; since the binning in l is very coarse, I don't calculate it by sampling each event at the bin center, but instead properly smearing each event in galactic latitude:

$$F(l) = \sum_i \int_{l-6^\circ}^{l+6^\circ} \frac{5}{\sigma_i} e^{-(l_i-l')^2/2\sigma_i^2} dl'$$

where l_i are the galactic longitudes of the events, and this formula is to be evaluated at the 15 bin centers, located at $-84^\circ, -72^\circ \dots 0^\circ \dots 84^\circ$. The next step is generating 30,000 background sky maps, and from them assembling the 15 distributions of expected flux $D(F, l)$ and the mean expected flux for each bin $\bar{F}(l)$. The results are shown in figure 6.6. The horizontal axis is the measured flux, which roughly corresponds to the

smear number of events in each bin. The vertical axis shows the 15 bins in galactic latitude. The bold black line in each bin shows the value of the measured flux, and the 15 shaded histograms are the distributions of expected flux for each bin. To guide the eye, three white lines are drawn across each histogram; the center one is the mean expected flux, the other left and right of it are the fluxes at which the distribution has dropped to correspond to $\pm 1\sigma$. The technique of comparing the measured flux to the expectation histogram can be demonstrated very well here. For example, in the center bin (latitude 0°) the measured flux is a little higher than the -1σ mark; therefore the flux observed here is a deficit of about 1σ . The chance probability of observing such a large flux “excess” is fairly high, about 75% of the background flux distribution is above the observed flux; therefore the flux excess observed in this bin is not significant at all.

The result of this section is that there is no observed flux excess from the galactic plane; nor is there any statistically significant excess from any of the other 14 bands in galactic latitude.

6.4 Point Source Search

Unfortunately, a full two-dimensional search for point sources of neutrinos cannot be displayed as easily as the one-dimensional search of the previous chapter. Again, I will be using the method outlined in section 5.2, this time on the complete two-dimensional sky, using equatorial coordinates. Again, this search uses the 496 throughgoing events; the positions of the events are shown in figure 6.7. The horizontal axis is right ascension (from 0 hours through 24 hours), the vertical axis is declination (from -90° through 60°). Since the flux of muons drops to zero at a declination of about 55° , the northern part of the sphere is omitted. The solid and dotted white grid lines mark every hour in right ascension and every 15° in declination.

The events are smeared using the appropriate resolutions, which results in the flux map shown in figure 6.8. In this and the following maps, flux is shown as darkness, with darker areas corresponding to higher flux. Areas of different “coloration” are separated

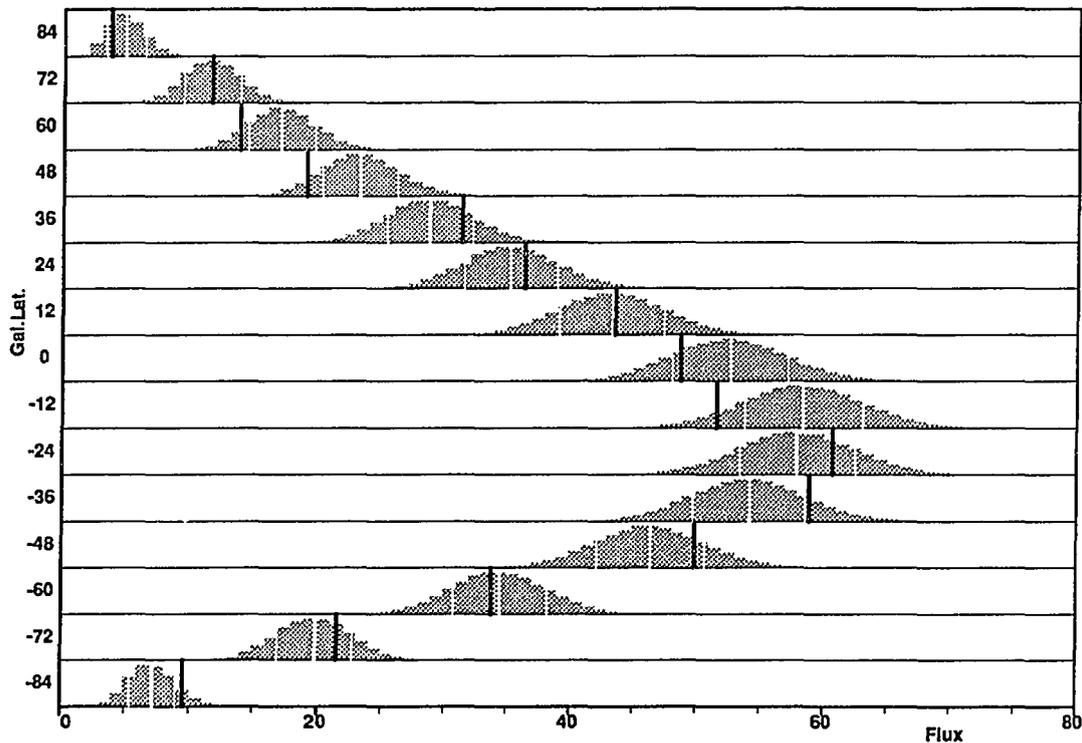


Figure 6.6: One-dimensional muon flux excess search in galactic coordinates. The bins in the vertical direction are the bins in galactic latitude, the horizontal axis is the muon flux in each bin. The black line shows the observed flux in each bin. The histogram is the distribution of expected flux in each bin. The three white lines across the histogram are the mean expected flux, and the $\pm 1\sigma$ points in the histogram.

by black contour lines. The scale at the right shows the conversion from shading to the flux value; the numbers in each box indicate the flux *at the center of the bin in shading*.

Then 32,500 iterations of the scrambled background are generated, using the procedure described in sections 5 and 5.2. The mean flux expected from the background is shown in figure 6.9. Note that the mean background is nearly but not completely uniform over right ascension; as explained above, the small non-uniformity stems from the slight correlation between the different temporal biases of the data. By comparing the measured flux to the ideogram for each bin, the chance probability of seeing such an excess can be determined. Figure 6.10 shows the result, with the chance probability displayed on

a logarithmic scale as the significance $S = -\log_{10} P$. The same excess significance data is displayed again in figure 6.11 in galactic coordinates, with the horizontal axis being galactic longitude and the vertical axis galactic latitude. In both figures, the positions of the candidate sources are shown; the exact position of each is at the center of the identifying text for each. The sine-wave in figure 6.10 is the galactic plane. Note that the information contained in both figures is identical, just displayed in a different coordinate system. The region of the sky which cannot be observed (above a declination of $\delta = 60^\circ$) forms a round region centered around galactic longitude 123° and latitude 23° .

This excess significance map will be used in the next section to determine the chance probability of flux excesses *from sources known a priori*; but in this section we are interested in finding new sources. The map is not perfectly smooth and uniform, but does show structure; the question is how to quantify that structure. Clearly, one cannot just pick out the place in the skymap which shows the largest excess significance (that is the smallest chance probability of the flux excess being due to background); there is a large statistical penalty involved in the process of choosing a particular location which one wants to call a point source. Furthermore, the number of possible choices of point source positions on the map is not easy to determine; it is clearly much smaller than the number of grid points, since the grid was purposely crafted to be much finer than any feature that can be observed in the data.

So how does one go about defining a point source candidate in a quantitative way? It should form a peak in the excess significance map; so the algorithm is to define the term "peak" in a sensible fashion and determine the properties of the peaks in the data; repeat this for the 32,500 iterations of the background, and compare. The definition of peak I use here is any grid point in the sky which has a higher excess significance than all of its nearest neighbors, and which has an excess chance probability of 3% or less. Why three percent? Because it yields a reasonable number of peaks to work with; if I had made that cut either 50% or 0.01%, the number of peaks would have been overwhelming or near zero, and in either case statistical comparisons become inefficient. There are order of 1000 statistically independent points in the observable portion of the sky (using an angular resolution of about 5°), so a threshold level of 3% should yield about 30 point source candidates.

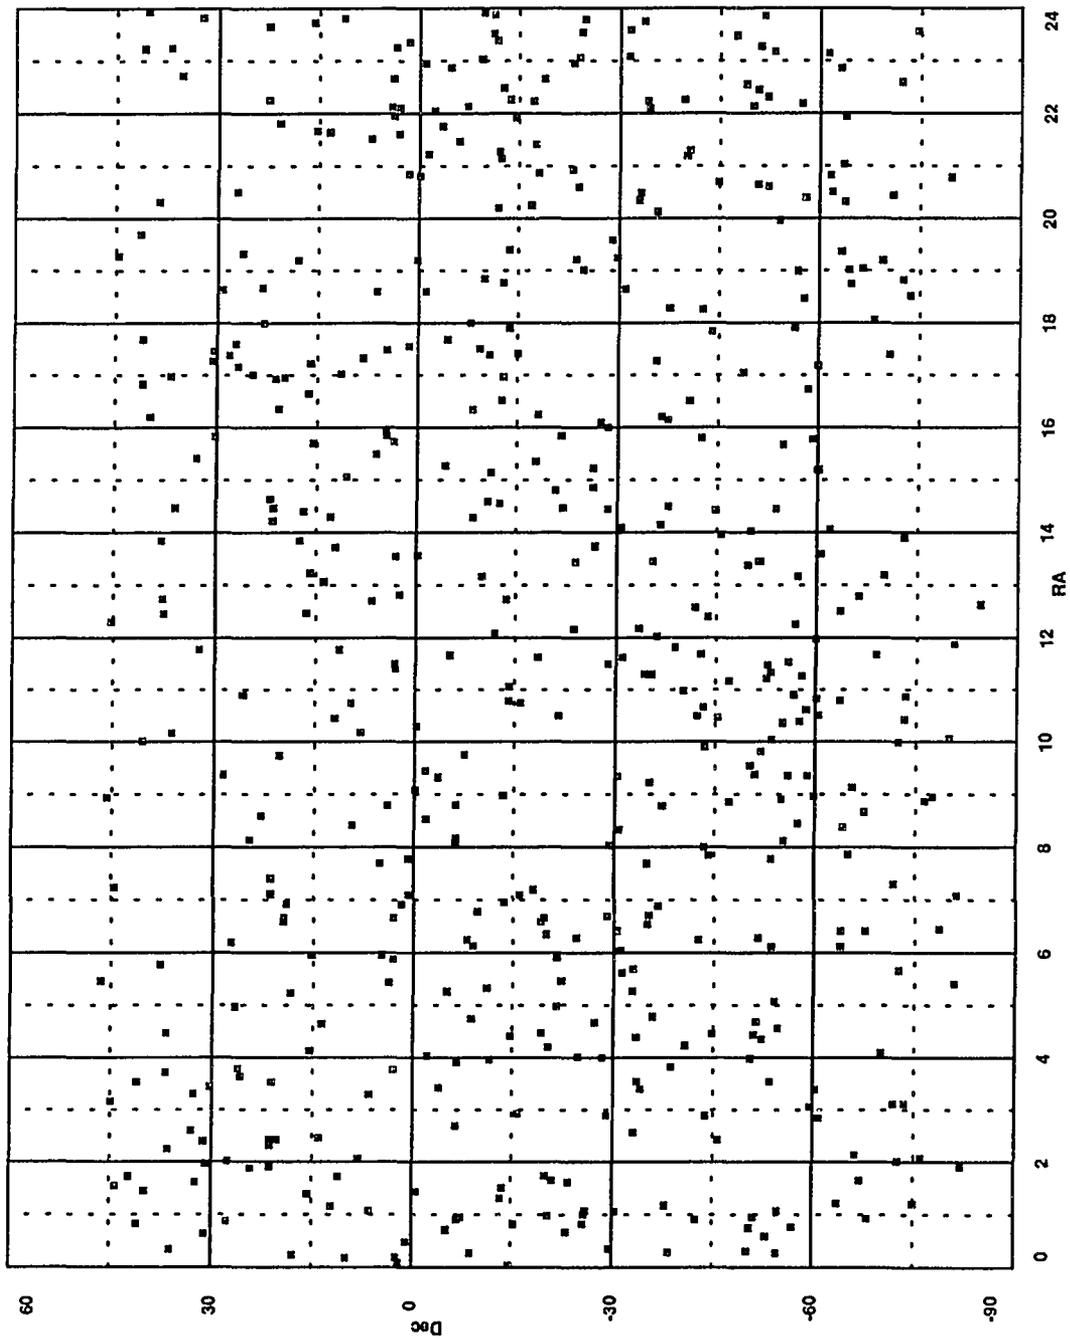


Figure 6.7: Event positions in equatorial coordinates

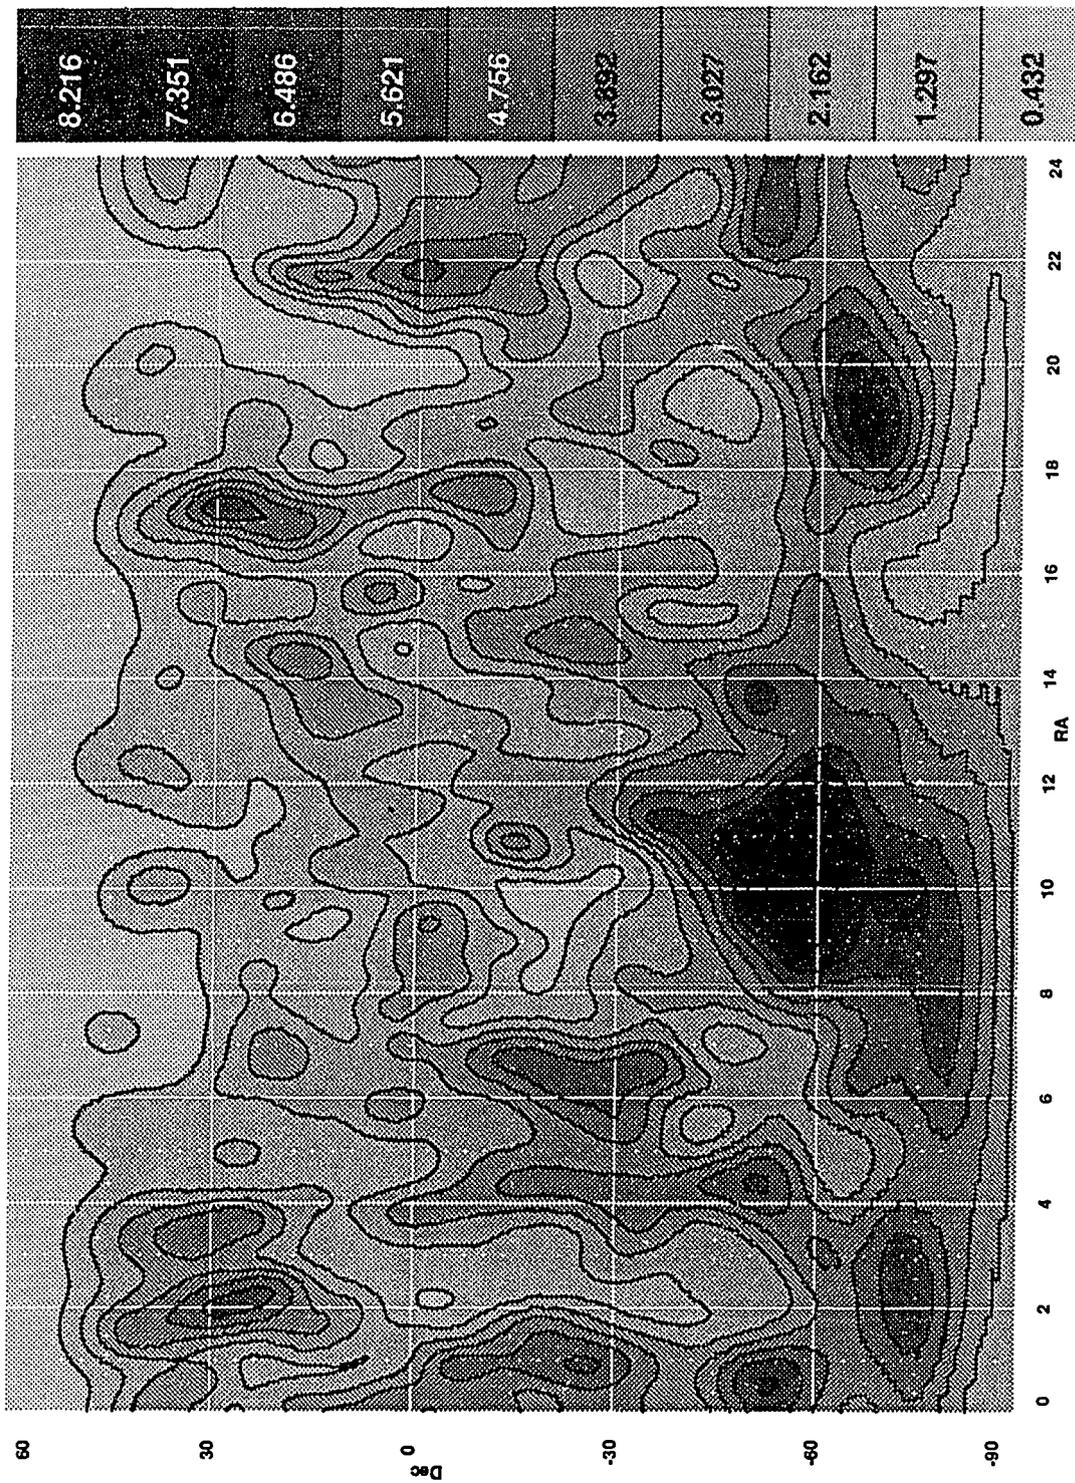


Figure 6.8: Smeared event flux

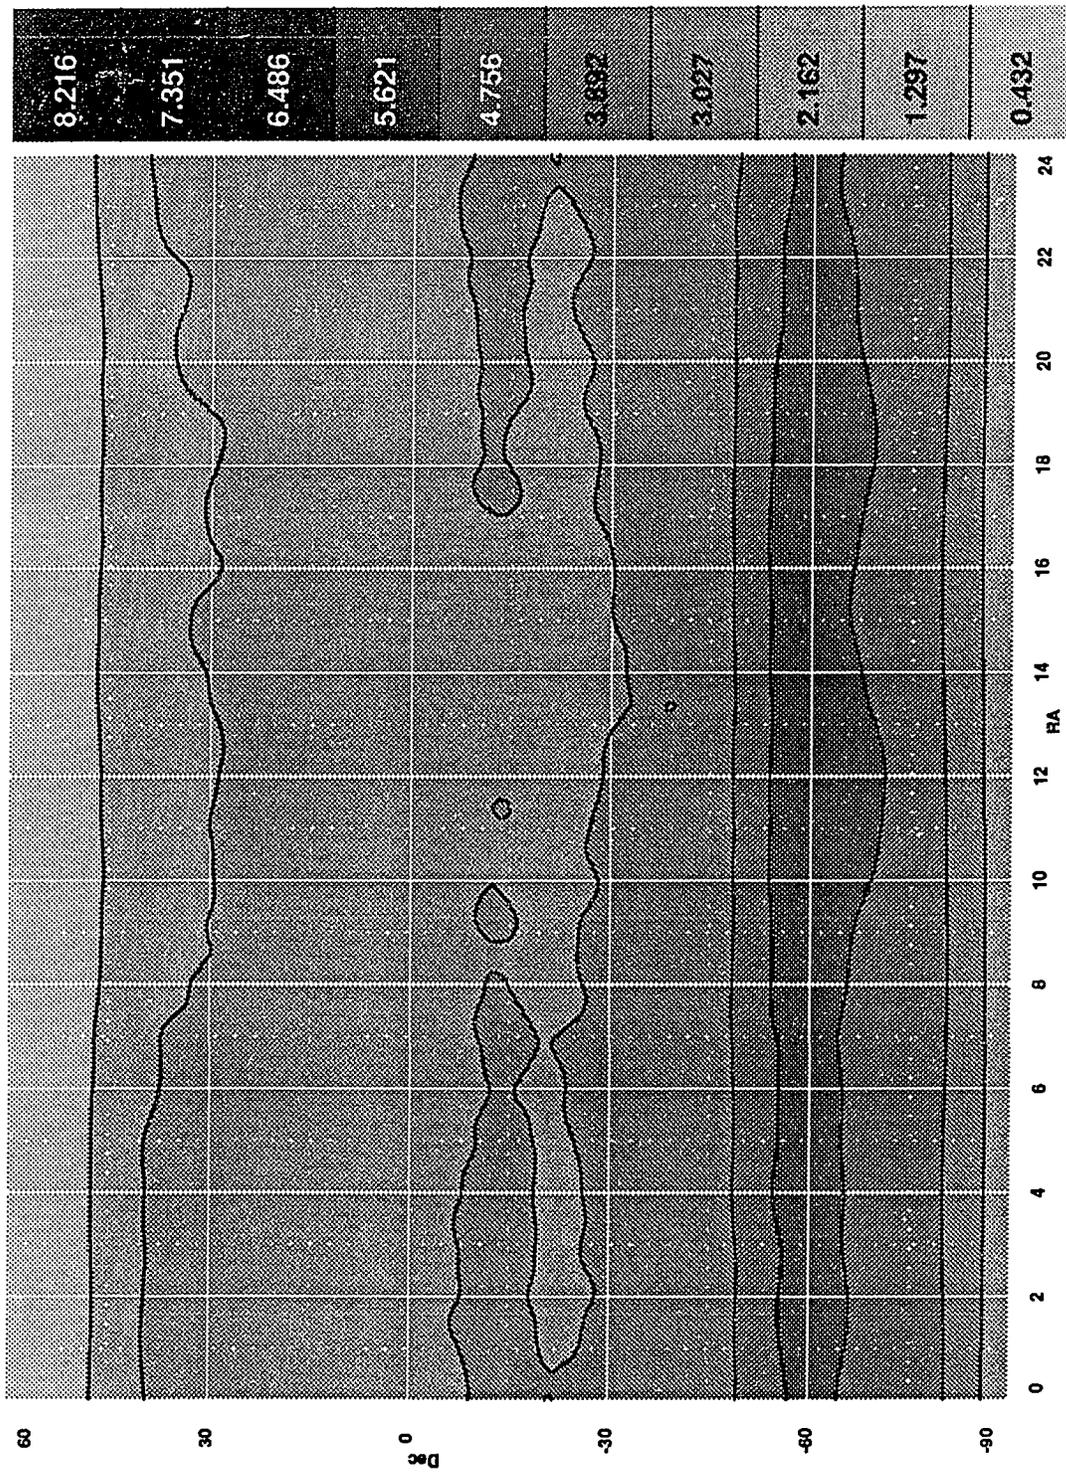


Figure 6.9: Mean expected event flux

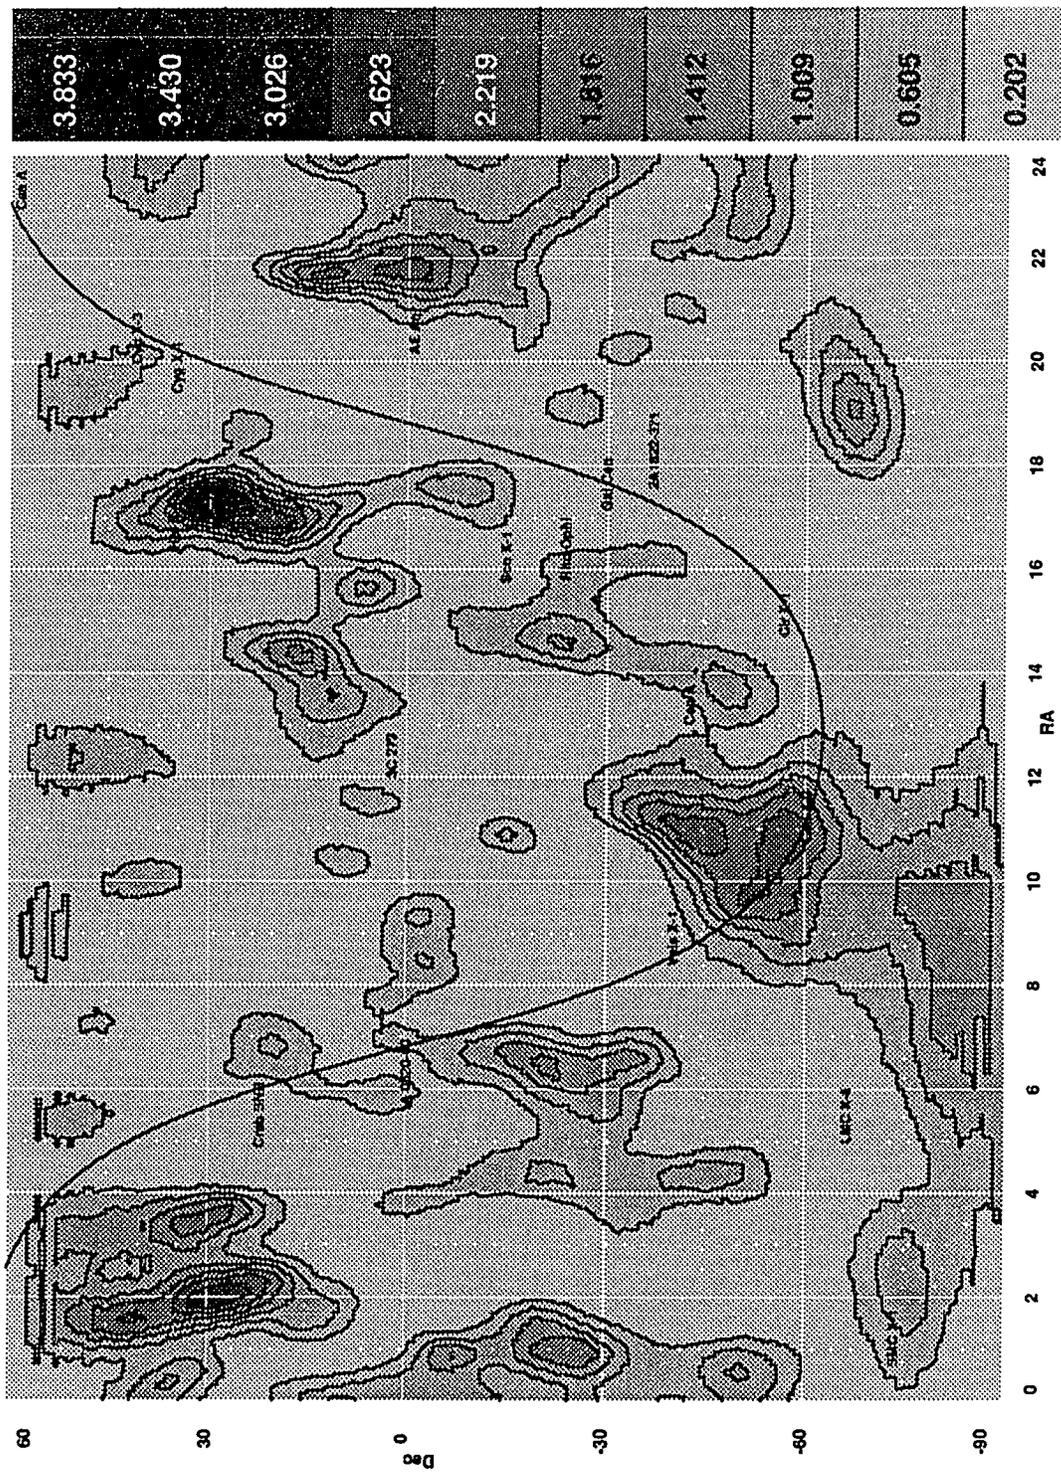


Figure 6.10: Significance of flux excess

There is one more technical issue which has to be discussed. The excess probability or significance has been obtained by comparing a measured value to a *histogram*. Due to the binning inherent in that comparison, the measured excess chance probabilities show a very small spatial ripple from bin to bin, which is too fine to be seen on the sky map or to influence the results; however it can create spurious peaks. To remove any such effect from the data, the excess significance map is smoothed slightly, using a Gaussian smearing function with about 0.5° resolution. Note that the smearing resolution is actually smaller than the grid spacing, so this smearing step will not lead to an appreciable loss of resolution.

Using this definition, 21 peaks are found in the data. Their positions are shown in table 6.2. They are ranked by excess significance, that is by the negative logarithms of the chance probability, with the most unlikely peak first. The first column of the table lists the rank, the second the position in equatorial coordinates. The third column gives the significance and the fourth one the flux observed at that position. Note that the chance probability of the flux excess at the first peak being due to a random fluctuation of the background *at this point in the sky* is $10^{-4.035} = 9.2 \cdot 10^{-5}$.

To be able to assess the statistical significance of these peaks, they have to be compared to the peaks expected from the random background. Therefore, the same process is repeated for the simulated background data, resulting in 32,500 such sets of peaks. The first check is whether the total number of peaks observed in the data is significantly different from the expectation. Figure 6.12 shows the distribution of the number of peaks observed in the background runs; the histogram is the distribution itself (use the left scale for it), the smooth curve the chance probability in percent of observing *that number or more* peaks (use the right scale for the smooth curve). The dotted line indicates the number of peaks observed in the data; the chance probability of the observed 21 (or more) peaks being due to background is 31.865%, therefore compatible with the expectation from background. There is no evidence for an excessively large number of peaks in the data. Such an excess could be caused by “lumpiness” in the data, for example if the 496 events studied here were grouped into 100 very tightly spaced clusters of about 4 events each. Such a clustering would also have caused a strong correlation at small angles to be observed in the twopoint correlation study above.

Rank	RA	Dec	Signif.	Flux	C.P.
1	17:15	+29.5	4.0350	5.8234	12.043%
2	2:00	+28.5	2.7638	5.4532	26.492%
3	1:57	+30.5	2.6669	5.2347	29.557%
4	10:50	-45.5	2.4150	6.2803	32.778%
5	10:58	-42.5	2.3753	5.8214	34.222%
6	10:54	-57.5	2.3216	8.6338	33.357%
7	10:43	-55.5	2.2791	8.4542	31.766%
8	21:43	+12.5	2.2109	4.4724	30.397%
9	21:46	+2.5	2.1656	5.2058	28.369%
10	1:32	+41.5	2.1484	4.0615	26.135%
11	21:42	-1.5	2.1337	5.2985	23.671%
12	9:47	-52.5	2.1228	7.8206	21.077%
13	21:46	+0.5	2.0822	5.4052	18.892%
14	1:35	+39.5	2.0650	4.2241	16.865%
15	0:59	-24.5	1.9856	5.2778	15.218%
16	3:22	+33.5	1.7725	4.0901	14.775%
17	14:22	+15.5	1.7369	4.0704	14.258%
18	18:58	-67.5	1.6650	7.2474	13.495%
19	6:21	-22.5	1.6237	5.0019	12.892%
20	8:31	-84.5	1.5709	4.2595	12.031%
21	0:50	-6.5	1.5472	4.8044	11.228%

Table 6.2: Peaks found in the flux excess significance skymap. The first column is the rank of the peak; the second and third its position in equatorial coordinates. The fourth and fifth column are the flux excess significance (negative logarithm of the chance probability) and the flux at the peak. The last column is the chance probability of aggregate significance for this peak being greater than or equal to the background.

Checking the total number of peaks is only the initial step in the comparison between the observed data and the background. More interesting is a comparison of peak “height”: are the peaks in the data “higher” than those in the background? The “height” should be somehow related to the excess significance: are the peaks in the data less probable than the ones in the background? Comparing the event flux would not do, since the absolute flux depends drastically on the declination, but comparing excess significance (or equivalently excess chance probability) is fair, since the significance is a good measure for how “interesting” an excess is. Using this test, the chance probability of the highest peak in the data being of a significance at or above 4.0350 is 12%; this is obtained

from comparing the least likely peak in the data to the least likely peaks in each of the 32,500 background maps.

Performing this test for the other peaks is not this easy. It is not fair to compare the second highest peak in the data to all the second highest peaks in the background. The reason is in the philosophy of hypothesis testing, that is checking whether the observed data is compatible with being caused by background. Assume that the highest peak in the data were incredibly unlikely; we would then call that peak a *point source detection*, and remove it from further consideration. Now the second highest peak should be tested *against the highest peak in the background*, to check whether it is caused by background. Since it is not possible to individually single out some of the peaks as being point source detections and then just move the others up in rank, I use a different approach here: to compare the aggregate chance probability of the first n peaks with the same aggregate for all the background maps, where n is the peak rank we are interested in. The formula (3.3) for compound hypotheses tests is used to calculate the compound chance probabilities; but since the only use of this compound chance probability is comparing it to other compound chance probabilities for the same number of tests, one can as well just multiply the probabilities. The result is shown in the last column of table 6.2, and in figure 6.13 showing the chance probability as a function of rank. The chance probabilities for the first few peaks (excluding the first one) are not significant on their own. Note that the chance probability of the last peaks decreases steadily, reaching 11% for the last one. That indicates that although the total number of peaks in the data is not significantly higher than what is expected from the background, the aggregate chance probability of the peaks in the data is lower than expected from the background. That could for example be caused by a slight correlation at small angles between events at a few locations; since only a few locations are involved, it does not contribute heavily to the total number of peaks, but it makes a few of the peaks less likely than expected, which shows up in the aggregate chance probability. Note that this model is compatible with the observation of a correlation in section 6.2 above.

Note that these 21 peaks are not at all distributed uniformly over the sky; peaks 2 and 3, peaks 4 and 5, peaks 6 and 7, peaks 9, 11 and 13, and peaks 10 and 14 all form groups which are within 5° of each other. Such a behaviour is not unexpected, since the slopes

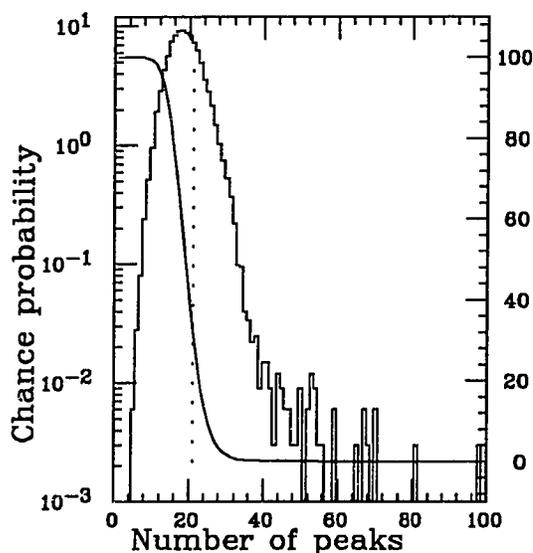


Figure 6.12: Distribution of peak number in the background, shown in the histogram using the left scale. Curve is the chance probability in percent of exceeding a given number of peaks, use right scale. The dotted line indicates the observed number of peaks.

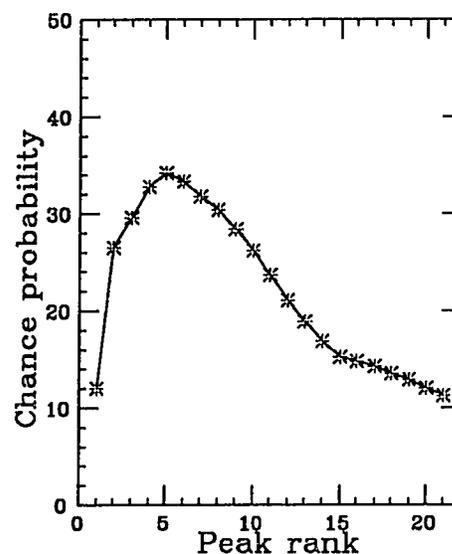


Figure 6.13: Chance probability of aggregate peak flux excess significance as a function of rank

on the side of the excess significance shape can have little ripples, which may be counted as peaks on their own right. Naturally, these multiple peaks might also be caused by tightly spaced point sources, but in this case they could not be reliably resolved by data with our angular resolution anyhow. In these cases, I will only use the highest peak of each group; table 6.3 is a revised version of table 6.2 where these secondary peaks have been removed.

The example of the first peak points out the difference between searching for an *a priori* known source of neutrinos, and the attempt of finding one and selecting its location *a posteriori*. The chance probability of the excess observed at the location of that peak is $10^{-4.035} = 9.2 \cdot 10^{-5}$, whereas the chance probability of seeing such a large peak in excess *anywhere in the map* is 12.043%. The ratio of the two probabilities is 1305; this could be

Rank	RA	Dec	Signif.	Flux	C.P.
1	17:15	+29.5	4.0350	5.8234	12.043%
2	2:00	+28.5	2.7638	5.4532	26.492%
4	10:50	-45.5	2.4150	6.2803	32.778%
6	10:54	-57.5	2.3216	8.6338	33.357%
8	21:43	+12.5	2.2109	4.4724	30.397%
9	21:46	+2.5	2.1656	5.2058	28.369%
10	1:32	+41.5	2.1484	4.0615	26.135%
12	9:47	-52.5	2.1228	7.8206	21.077%
15	0:59	-24.5	1.9856	5.2778	15.218%
16	3:22	+33.5	1.7725	4.0901	14.775%
17	14:22	+15.5	1.7369	4.0704	14.258%
18	18:58	-67.5	1.6650	7.2474	13.495%
19	6:21	-22.5	1.6237	5.0019	12.892%
20	8:31	-84.5	1.5709	4.2595	12.031%
21	0:50	-6.5	1.5472	4.8044	11.228%

Table 6.3: Peaks found in the flux excess significance skymap, excluding secondary peaks. The format is identical to table 6.2.

taken to be the statistical penalty one incurs for searching the whole sky. In a very crude approximation, that statistical penalty should be the number of statistically independent points one can search in the part of the sphere observable to this study (from the south pole to a declination of $+60^\circ$). So the size of the search circle which would be equivalent to this study can be determined from this statistical penalty, it comes out to be a circle (or cone) of radius 3° , which is reasonable compared to the angular resolution which was used in this study.

6.5 Flux from Point Source Candidates

This section is devoted to the question whether any of the astronomical objects discussed in section 1.4 is a source of neutrinos visible above the background of atmospheric neutrinos. As a reminder, in that section I had made up two lists of candidate sources. The objects on the first list can be assumed to emit neutrinos, albeit maybe not at a level detectable in this experiment. They are the ones which will be studied seriously here. The second list is a sample of other sources which may emit neutrinos; they

will just be listed for comparison in this section. There are 9 sources in the first list: 4U0115, Centaurus X-3, Circinus X-1, Crab pulsar, Cygnus X-3, Hercules X-1, LMC X-4, Scorpio X-1 and Vela X-1. Of those, 4U0115 cannot be observed in this work, it is too far north to ever drop below the horizon; it will be omitted from the tables and figures of this section. Two other sources, Cygnus X-3 and Hercules X-1 are very far north, and the efficiency of observing them is fairly small.

As seen in the previous section, there is no convincing evidence for new point sources of neutrinos. Nevertheless, I will add the three most significant peaks from table 6.3 to the list of sources, naming them “IMB ν -1” through “IMB ν -3”; not because they are claimed as point sources of neutrinos (which would be quite ridiculous), but only to conveniently list the muon and neutrino fluxes for them, which will be calculated in this section.

All the three categories of sources are listed in table 6.4. The first two columns of that table are the name of the source and its position in equatorial coordinates. The third column is the effective area of the detector for a source at this position, obtained in section 4.6.

The chance probability of the observed flux at a certain position in the sky being in excess of the background expectation has already been discussed in the previous sections; in a similar fashion, the deficit chance probability can be estimated. The fourth column of table 6.4 lists the significance $S = -\log_{10} P$ of the chance probabilities for flux excess and deficit. None of the sources shows a significant deficit of events. Two of the sources in the first category show a significant flux excess: Centaurus X-3 at a chance probability of 1.48%, and Hercules X-1 at a chance probability of 1.84%. The chance probability of observing these two significances in eight trials (4U0115 doesn't qualify as a trial) is hard to determine. One possibility is to assume a statistical penalty of a factor of 8 for the first source, and a factor of 7 for the second one (there is one less to choose from); that leads to a chance probability after trials of 11.8% for Centaurus X-3 and 12.9% for Hercules X-1. Another possibility is to lump the two together using equation (3.2) and apply a correction for 56 trials (8×7), which leads to a chance probability after trials of 11.0%. From those probabilities, all in the neighborhood of 12%, follows that

the observation of a flux excess from two of these eight sources is marginally statistically significant. In the following, I will assume that Centaurus X-3 and Hercules X-1 have been “detected”. I will calculate muon and neutrino flux *limits* for all sources in the three categories; for the two sources which show a significant flux excess and for the three peaks found in the point source search, I will also estimate the muon and neutrino flux “observed” from these sources.

The calculation of flux limits and expected fluxes is unfortunately complicated by the fact that there is no direct measurement of the *integer* number of events contributing to the flux at a certain point. If the number of events could be determined by simple counting, the flux limit could be established using the statistics of a poisson distribution. Since the event “flux” is not just the observed number of events but a weighted and smeared quantity that approach is not feasible.

Instead I once again compare the measured data with data which would be expected if any of these sources emit neutrinos. The process is the following: the observed events are scrambled, and the event flux at the position of the source under consideration *assuming no neutrino events from it* is recorded. This has already been done 32,500 times, when generating the background for the point source search. Therefore, we already have in hand the distribution of event flux expected for the case of no signal from the source. What has to be added now is the distribution of event flux in the case that there a given source emits neutrinos. A new scrambled map is generated, but this time n_i randomly selected events are shifted so they point at the source we are interested in; n_i is the number of *injected* events. When moving the events, their angular resolution is taken into account properly, so they don’t always point straight at the source.

Several new distributions of event flux are made, this time as a function of the number of injected events n_i . In figure 6.14 this is shown for the case of Hercules X-1; the horizontal axis is the measured event flux, the vertical axis is the number of events injected. For each number of injected events, there is a distribution of event flux. The dashed line shows the observed event flux on this source, and the diagonal solid line joins the mean values of each distribution. The bottommost distribution (for no injected events) is the

background distribution calculated for the point source search above; it has much better statistics, since injecting these events is a time-consuming process.

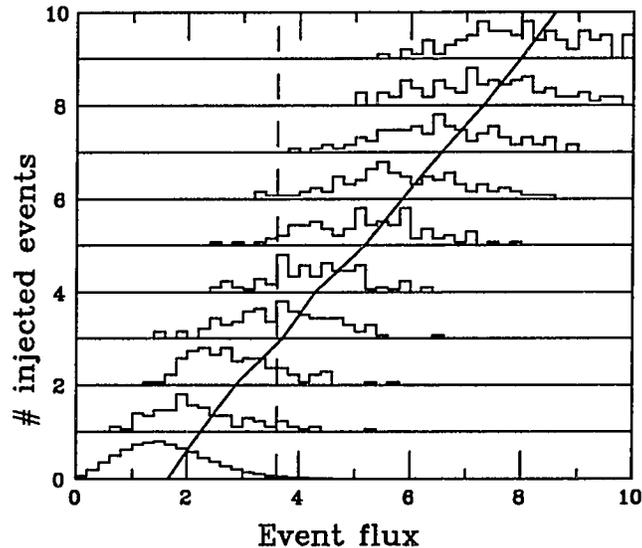


Figure 6.14: Event flux distributions for background data with injected signal events, shown for the location of Hercules X-1. The horizontal axis is the observed event flux, the vertical axis the number of signal events from the source which were injected. The dashed line is the observed flux from that source; the straight diagonal line joins the means of all the distributions.

The method used to determine the event flux limit (and the expected flux) is now to compare these distributions with the observed flux. For a certain number of on-source muons N , the real number of observed events (which I am trying to simulate with the n_i injected events) will have a Poisson distribution:

$$P_N(n_i) = \frac{e^{-N} N^{n_i}}{n_i!}$$

To incorporate this Poisson distribution into the process, the distribution of event flux for N muons from the source is made up from a weighted sum of the distributions as a function of n_i , each weighted by the poisson probability $P_N(n_i)$. Note that N does not have to be an integer. The resulting weighted distributions are shown the left half of figure 6.15, again for the case of Hercules X-1; again the horizontal axis is the event flux, the vertical axis is the number of muons N . The distributions are shown only for

integer number of muons, but could as easily be calculated for a fractional number of muons.

The 90% confidence level upper limit on the number of events can now in principle be read off that graph, by moving up in the number of muons, until 90% of the event flux distribution is above the observed flux. Note that the number of muons injected does not have to be integer. In the same manner the expected flux can be determined for those sources where there is some indication for the flux being larger than just background: adjust the real number of events to the point where half the distribution of event flux is below, half above the observed flux. This process is illustrated in the right side of figure 6.15, which shows the the fraction of the event flux distribution which is above the observed flux (on the horizontal axis) as a function of the real number of muons (on the vertical axis).

One more complication arises in cases where the observed event flux is lower than what is expected from just background alone. In these cases, the number of muons is less than what is expected, which is unphysical. Clearly, we cannot start injecting negative muons into the data to account for that. In these cases using the above procedure would yield too good a limit; the conservative way of estimating a flux limit (which may be too high, but that is on the safe side) is to replace the observed flux with the mean flux expected from background. In essence this means that if the observed flux is lower than the background expectation, this is not to be taken as a flux *deficit*, but instead is discounted as a statistical fluctuation, and the expected flux is used.

One note of caution is required here. When injecting the extra events into the scrambled background, their position was smeared *using the resolution determined in section 4.5, which applies to atmospheric neutrinos*. The problem here is that a point source of neutrinos is expected to have a much harder neutrino spectrum (much lower spectral index) than atmospheric neutrinos; at these typically higher energies, the scattering angle between neutrino and muon is smaller. In principle the events should be injected with an angular resolution more typical for point source neutrino induced events, which would be lower. However, we have only observed atmospheric neutrinos so far, and we know that most of the data is well described by them. Furthermore the exact value of the

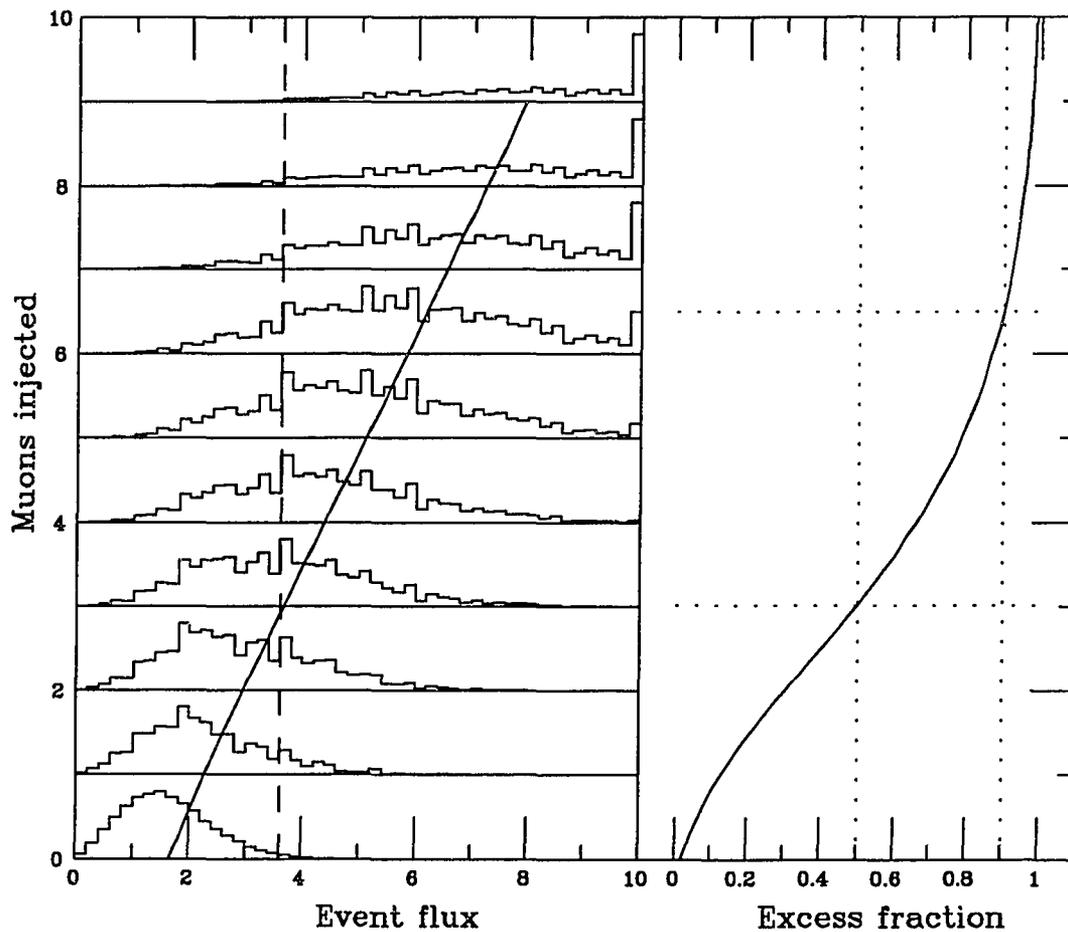


Figure 6.15: Event flux distributions for background data with injected Poisson-distributed muons, shown for the location of Hercules X-1. The vertical axis is the expected number of muons from the source which were injected; note that it does not have to be integer. The horizontal axis of the left graph is the observed event flux. The dashed line is the observed flux from that source; the straight diagonal line joins the means of all the distributions. The right graph shows the fraction of the distribution which is below the observed flux.

neutrino spectral index is not known. Last but not least, the choice of angular resolution is *conservatively pessimistic*, that means it will yield flux limits which are higher, but safe and assumption-free.

Calculating the muon flux or the muon flux limits from the number of muons or muon number limit obtained above is quite simple: the flux is just the number of muons, divided by the effective area of the detector for the given declination of the source and the live time, with the units converted to $\text{cm}^{-2}\text{s}^{-1}$. The 90% confidence level limits on number of muons and on muon flux are shown in the fifth column of table 6.4; all the muon fluxes are in units of $10^{-15} \text{cm}^{-2}\text{s}^{-1}$. For those sources where a flux estimate is meaningful, it is given in the following column of that table.

Calculating the neutrino flux or flux limit from the muon flux is more involved. First the cross-section of the neutrino interaction and the range of the muon have to be taken into account; both are strong functions of energy. For example, reference [40] calculates the probability for a neutrino of a given energy to create an observable muon as a function of neutrino energy, see for example figure 1 in that reference. Since the detection efficiency for a muon depends on its energy, the relation between muon and neutrino energy has to be folded into the calculation too; this requires knowledge of the y -distribution of neutrino interactions, which can for example be found in reference [84], or in a simple parametrized form in reference [91]. The detection efficiency for muons as a function of energy has already been explained in section 4.5, and is shown in figure 4.16. Since all the factors in this calculation are strong functions of energy, the result will depend strongly on the assumed neutrino spectral index. Lastly, the upper cutoff energy of the neutrino source is required; I am assuming 10^7GeV , but the result doesn't depend strongly on that assumption, since the flux at these extreme energies is very low. Figure 6.16 shows the resulting conversion factor from muon flux to neutrino flux as a function of spectral index, for spectral indices in the range from 2.0 to 3.0.

From the neutrino flux or flux limit the neutrino luminosity of an object can be calculated using the known distance to it, with the units converted to erg/s . I assume a neutrino spectral index of 2.1 for all of the sources. The results on neutrino fluxes and luminosities for the sources which yield an observable flux estimate are listed in table 6.5. For all

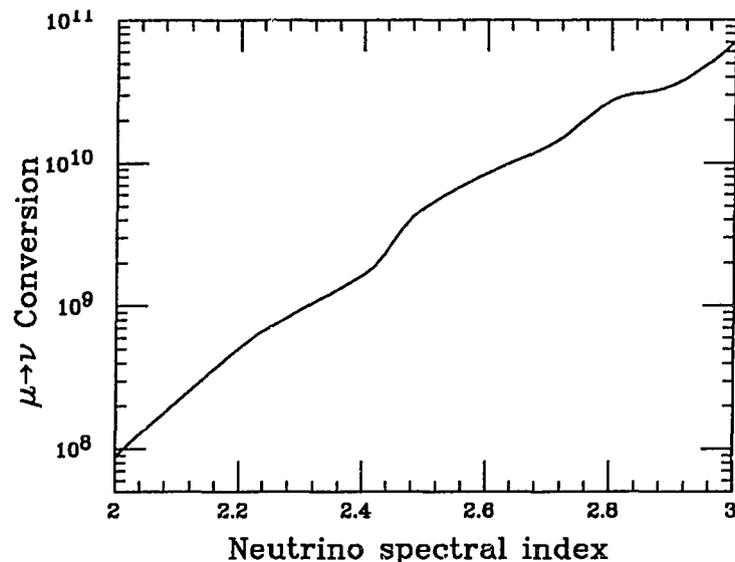


Figure 6.16: Conversion from μ to ν flux, as a function of neutrino spectral index.

sources neutrino flux and luminosity *limits* are given in table 6.6. Also listed are previous experiments which have set competitive limits on the neutrino flux of astronomical objects: IMB-1,⁹³ Kamiokande⁶⁷ and Fréjus.⁷ Note that reference [67] and [7] use a neutrino spectral index of 2.1, just as I do, reference [93] uses 2.3 (but also quotes figures for 3.0). Also note that all results except this work use the angular resolution expected for the hard spectrum from neutrino point sources; therefore they are often background-free (no events within the search cone around the source), whereas I use the angular resolution for the observed events, which is assumption free (and pessimistic).

As was pointed out in section 1.3, in most models of the production of neutrinos in point sources there is a connection between the production between the neutrino flux and the gamma-ray flux: the neutrino flux is typically larger than the gamma-ray flux, by a factor known as λ , which is expected to be between about 1 and 1000, with typical values of 10 to 30 for binary pulsars. Figures 6.17 and 6.18 show a comparison of the neutrino flux limit (for the case of the Crab pulsar in figure 6.17) and the observed neutrino flux (for the case of Hercules-X1 in figure 6.18) to typical gamma-ray observations of those objects. The gamma-ray data are not a complete survey of all observations but a representative sample; they are taken from reference [51]. All fluxes shown are integral

fluxes above the detection threshold, and are plotted at the detection threshold (which nearly coincides with the median energy of the observed events for the gamma-ray data). As can be seen, the gamma-ray data can be well described by a power-law spectrum with a differential spectral index of 2.1 for the Crab pulsar and 2 for Hercules X-1. The low-energy measurements below 100 MeV are observations of X-rays and low-energy γ -rays from balloon or space detectors; the observations around 100 GeV are typically ground-based air Čerenkov telescopes, and the observations around 1 TeV are air shower experiments. Unfortunately there are no gamma-ray observations in the energy typical of this neutrino experiment, since the flux at these energies¹ is too low for airborne detectors, but the energy is too low to operate ground-based experiments. Since the gamma-ray data is well described by a power-law spectrum one can probably interpolate the expected flux at these energies safely.

This yields a neutrino flux over gamma-ray flux enhancement (the λ factor) of less than 18 for the Crab pulsar (for which no neutrino flux is observed, so this is an upper limit), and about 45 for Hercules X-1. Those two figures are well in the range of expected values for λ . One can conclude that the observed neutrino flux from Hercules X-1 (if indeed it is a neutrino flux, and not just a statistical fluctuation) is about compatible with the flux expected by interpolation of gamma-ray observations of the same source.

In conclusion, two out of the eight sources which are seriously studied show a flux excess, each with a significance of above 98%. After taking the trials factor into account, the chance probability of observing a flux excess at that significance for two out of eight sources is about 12%. These two sources are Centaurus X-3 and Hercules X-1. For all these sources and for the three most unlikely point source candidates observed in the previous section muon and neutrino fluxes and neutrino luminosities are calculated; for all sources flux limits are given.

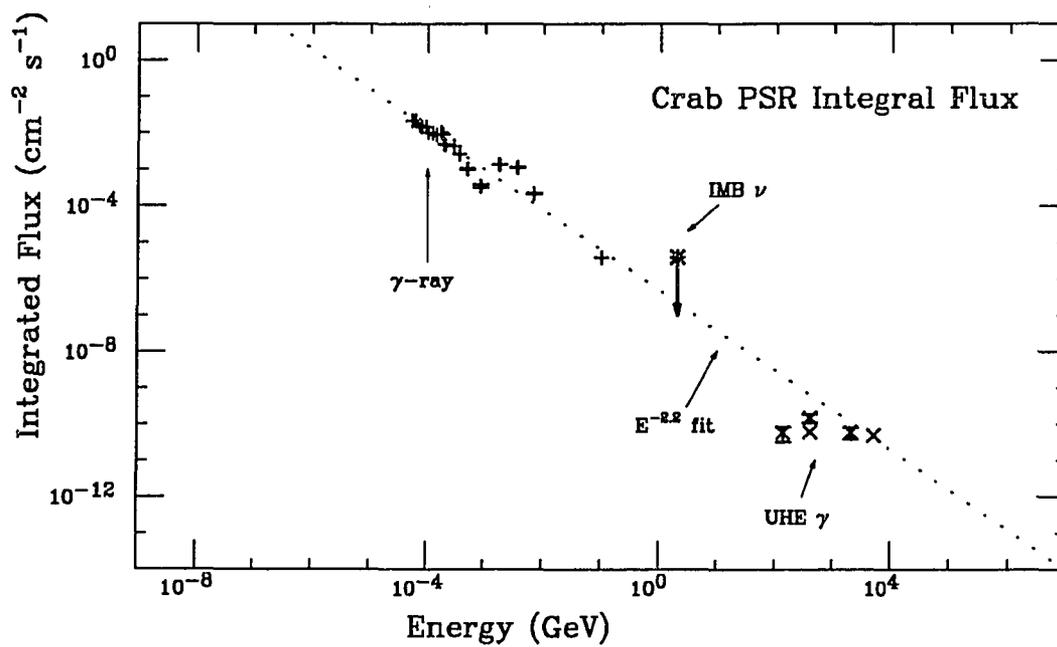


Figure 6.17: Comparison of neutrino flux from the Crab pulsar to gamma-ray measurements. All fluxes are integral fluxes above the detection threshold energy, and are all plotted at the threshold energy. The neutrino flux is an upper limit.

Source Name	Position RA	Dec.	A_{eff} m ²	Significance Exc.	Def.	Upper limit Evt.	Flux	Observed Evt.	Flux	Comment
Category 1: Sources which should emit neutrinos										
Cen X-3	11:19	-60.3	380.2	1.830	0.005	8.15	19.63	4.15	9.99	
Cir X-1	15:16	-56.9	373.3	0.135	0.495	2.26	5.54			
Crab PSR	5:32	+22.0	136.2	0.050	0.773	2.74	18.42			
Cyg X-3	20:32	+41.0	62.4	0.263	0.254	2.31	33.89			
Her X-1	16:57	+35.3	88.5	1.736	0.005	6.50	67.24	3.02	31.24	
LMC X-4	5:32	-66.4	394.7	0.056	0.789	3.50	8.12			
Sco X-1	16:19	-15.0	219.9	0.305	0.236	3.44	14.32			
Vela X-1	9:02	-40.5	283.2	0.234	0.306	2.33	7.53			
Category 2: Sources which may emit neutrinos										
1A0620	6:22	-0.3	187.2	0.315	0.225	3.54	17.31			
2A1822	18:22	-37.1	271.3	0.324	0.225	3.48	11.74			
3C273	12:28	+2.1	184.1	0.159	0.411	3.02	15.02			
AE Aqr	20:38	-1.0	189.0	0.208	0.333	3.66	17.73			
Cen A	13:25	-42.8	297.9	0.304	0.236	2.60	7.99			
Cyg X-1	19:58	35.0	88.5	0.211	0.317	2.49	25.76			
Gal. Cen.	17:42	-30.0	249.9	0.007	1.477	2.90	10.62			
LMC X-3	5:38	-64.1	392.3	0.096	0.600	2.73	6.37			
Rho Ophi	16:27	-24.0	235.9	0.324	0.218	2.66	10.32			
SN1987a	5:35	-69.0	397.9	0.112	0.547	2.94	12.65			
SMC X-1	1:15	-73.7	407.2	0.715	0.074	6.58	14.79			
Category 3: Peaks found in the point source search										
IMB ν -1	17:15	+29.5	109.2	4.035	0.001			5.96	49.97	Peak C.P. = 12.0%
IMB ν -2	2:00	+28.5	112.1	2.716	0.001			4.98	40.67	Peak C.P. = 26.5%
IMB ν -3	10:52	-45.5	318.4	2.415	0.001			4.28	12.31	Peak C.P. = 32.8%

Table 6.4: Significance and muon flux of source candidates. All muon fluxes are in units of $10^{-15} \text{ cm}^{-2} \text{ s}^{-1}$.

Source	Distance in kpc	\mathcal{F}_ν limit $10^{-6}\text{cm}^{-2}\text{s}^{-1}$	Luminosity limits in erg/s			
			This work	Kamiokande	IMB-1	Fréjus
Cen X-3	5-10	4.24	$2.82 \cdot 10^{38}$ – $-1.13 \cdot 10^{39}$			
Cir X-1	?	1.20				
Crab PSR	2	3.98	$4.24 \cdot 10^{37}$	$1.3 \cdot 10^{38}$	$1.0 \cdot 10^{39}$	$3.4 \cdot 10^{38}$
Cyg X-3	11	7.32	$2.36 \cdot 10^{39}$	$6.3 \cdot 10^{39}$	$8.5 \cdot 10^{40}$	$1.7 \cdot 10^{40}$
Her X-1	5	14.51	$9.67 \cdot 10^{38}$	$1.6 \cdot 10^{39}$	$4.9 \cdot 10^{39}$	$1.9 \cdot 10^{39}$
LMC X-4	44	1.75	$9.05 \cdot 10^{39}$	$3.6 \cdot 10^{40}$	$2.5 \cdot 10^{41}$	
Sco X-1	0.3-1	3.09	$7.42 \cdot 10^{35}$ – $-8.24 \cdot 10^{36}$			
Vela X-1	1.4	1.63	$8.49 \cdot 10^{36}$	$4.0 \cdot 10^{37}$	$1.9 \cdot 10^{38}$	
1A0620	?	3.74				
2A1822	?	2.53				
3C273	$(6-9) \cdot 10^5$	3.24	$3.11 \cdot 10^{48}$ – $-7.00 \cdot 10^{48}$		$8.8 \cdot 10^{49}$	
AE Aqr	0.053	3.83	$2.87 \cdot 10^{34}$			
Cen A	4400	1.72	$8.90 \cdot 10^{43}$		$2.5 \cdot 10^{45}$	
Cyg X-1	?	5.56				
Gal. Cen.	10	2.29	$6.11 \cdot 10^{38}$	$3.1 \cdot 10^{39}$	$1.1 \cdot 10^{40}$	
LMC X-3	55	1.37	$1.11 \cdot 10^{40}$			
Rho Ophi	0.16	2.23	$1.52 \cdot 10^{35}$			
SN1987a	50	2.73	$1.82 \cdot 10^{40}$	$8.2 \cdot 10^{40}$		$2.8 \cdot 10^{43}$
SMC X-1	65	3.19	$3.60 \cdot 10^{40}$			
IMB ν -1	?	13.28				
IMB ν -2	?	11.93				
IMB ν -3	?	3.322				

Table 6.6: Neutrino flux and luminosity limits from point sources. The results for Kamiokande are from reference [67], the IMB-1 results from reference [93], and the Fréjus results from reference [7]

Source	Distance in kpc	ν Flux $10^{-6}\text{cm}^{-2}\text{s}^{-1}$	ν Luminosity in erg/s
Cen X-3	5-10	2.16	$1.44 \cdot 10^{38}$ – $5.75 \cdot 10^{38}$
Her X-1	5	6.74	$4.49 \cdot 10^{38}$
IMB ν -1	?	10.78	
IMB ν -2	?	8.78	
IMB ν -3	?	2.66	

Table 6.5: Neutrino flux for sources with an observable flux estimate

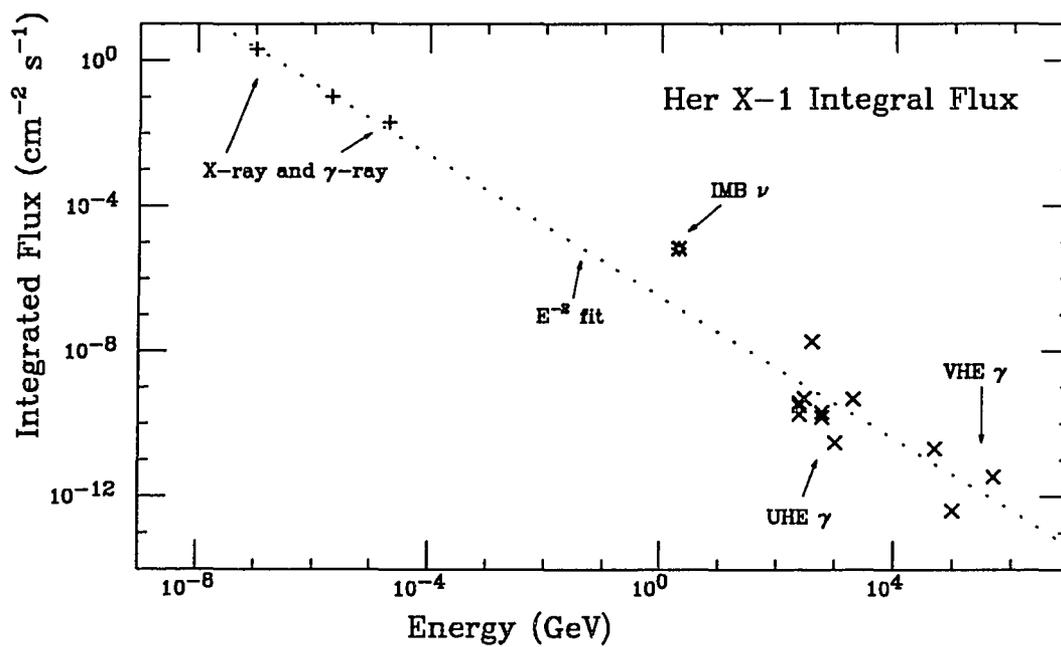


Figure 6.18: Comparison of neutrino flux from Hercules X-1 to gamma-ray measurements. All fluxes are integral fluxes above the detection threshold energy, and are all plotted at the threshold energy.

Chapter 7

Conclusions

The upward-going muon data from a 3.5 live year exposure of the 8000 ton IMB detector has been analyzed. 640 events are found above an energy threshold of about 2 GeV, 496 of them throughgoing, 144 of them stop in the detector. These events are assumed to be the product of neutrino interactions in the rock below the detector, and are used to study the properties of those neutrinos. The effective area of the detector is determined to be about 360 m^2 averaged over the angular distribution of atmospheric neutrinos.

The event directions have been determined using a very sensitive fitting technique, which minimizes the multipole moments of the pattern of timing and pulseheight residuals. The angular resolution of that fitting technique is determined from monte-carlo simulation to be 2.55° ; adding other direction uncertainties, the event reconstruction resolution is estimated to be 3° . Unfortunately, the angular resolution of the fitting procedure cannot be determined on an event-by-event basis.

The observed light deposition of the muon is used for a crude measurement of muon energy, and from that the expected angular deviation between the parent neutrino and the muon is determined. Quite a strong statistical correlation between the observed light output and the muon energy is found in monte-carlo simulation, even for a detector as small as IMB. This measurement does unfortunately not yield an absolute energy calibration, but has to be normalized to the expected energy distribution of atmospheric neutrinos.

The flux of atmospheric neutrinos and its angular distribution is studied, under the assumption that nearly all observed muon events are due to atmospheric neutrino parentage. The agreement between the observed total flux and its angular distribution, and the expectation for atmospheric neutrinos is very good. The number of observed muons,

their energy and angular distribution is compatible with most of them being due to atmospheric neutrinos.

To test whether a fraction of the observed muons may be due to point sources of neutrinos, a twopoint-correlation study is carried out. Correlations are found at small angles, significant at the 2.5σ level. The shape and size of the correlation function is very similar to the one expected for several point sources of neutrinos with a handful of events each. The observed correlation is not compatible with the direction of the neutrinos being completely random with a significance of 2.5σ , and is compatible with a small number of events (about 15 to 20) coming from a few point sources. However, the twopoint-correlation technique cannot prove that the observed correlations are indeed due to such point sources, nor can it determine the location of these point sources.

A skymap of flux excess chance probability is generated, using an ideogram method, with the expected background generated from the observed data using bootstrap statistics. This flux excess map is searched for peaks, as would be expected from point sources of neutrinos. The deviation in number of peaks and in the flux excess at the peaks from the background expectation is not statistically significant, with the chance probability of the highest peak being 12%. No convincing evidence for new point sources of neutrinos is found. Note that a search for point sources of neutrinos without *a priori* knowledge of the positions of the source incurs a large statistical penalty of about a factor of 1000, so the lack of significant results from a dataset this small is no surprise.

A list of sources expected to emit neutrinos at a level observable in this detector is made. Only sources which have been repeatably observed in high-energy gamma rays are considered candidates; any source which hadronically generates gamma-rays is also expected to generate neutrinos, at a flux comparable to or higher than the photon flux. Of the 8 sources, two are found to show a significant flux excess of below 2%, Hercules X-1 and Centaurus X-3; the chance probability of observing two out of eight sources at such a significance is about 12%. The muon flux from these two sources is $3.1 \cdot 10^{-14} \text{ cm}^{-2}\text{s}^{-1}$ and $1.0 \cdot 10^{-14} \text{ cm}^{-2}\text{s}^{-1}$. Assuming a neutrino spectral index of 2.1 and a detection threshold energy of 2 GeV, this translates into neutrino fluxes of $6.7 \cdot 10^{-6} \text{ cm}^{-2}\text{s}^{-1}$ and $2.2 \cdot 10^{-6} \text{ cm}^{-2}\text{s}^{-1}$, and neutrino luminosities of $4.5 \cdot 10^{38} \text{ erg/s}$ and $(1.4 - 5.8) \cdot 10^{38} \text{ erg/s}$

for Hercules X-1 and Centaurus X-3 respectively. This flux can be compared to observed gamma-ray fluxes, which are well described by a spectral index of 2.1 over several orders of magnitude in energy. The neutrino flux observed is high compared to the gamma-ray flux, but within what is reasonable using the expected conversion factor from gamma-ray to neutrino flux.

These results are not the proof that point sources of high-energy neutrinos exist; the observed events deviate from just being randomly generated atmospheric neutrinos, but only at statistical confidence levels which are between 2% and 10%. They are consistent with a small fraction of the observed events being generated by neutrinos from a few point sources; but consistency with a hypothesis and rejection of the null hypothesis is not proof. Further observations of neutrinos will have to be made to determine whether these results are the first hints of things to come, or whether we have observed statistical fluctuations.

The sensitivity of this search is determined by several factors: the effective area of the detector, the live time during which it is exposed to such neutrinos, the solid angle coverage of the detector, and in the case of the IMB detector (where this search is already pushing into the background of atmospheric neutrinos) the angular resolution with which muons can be reconstructed. The detector angular resolution is only of interest as long as it is comparable in magnitude to the muon-to-neutrino scattering angles at the energies expected for each detector threshold; in the case of the IMB detector, the angular resolution of the detector is not the limiting factor for neutrinos of relatively low energy any longer. At the high neutrino energies expected for point sources with hard spectra the detector resolution remains to be improved. Further observations will require much larger detectors with better angular resolution and larger solid angle coverage, but unfortunately such detectors also have higher energy thresholds. These large detectors should be complemented by detectors such as the IMB laboratory, which operate at significantly lower thresholds, and which should be upgraded to operate with significantly better event reconstruction resolution.

Appendix A

Flux Limits and Excesses

In section 6.5 the flux excesses from some sources and the flux limits from all sources were shown. This appendix contains the figures equivalent to figure 6.15 for all other sources. See that section for an explanation of these figures and the numeric results.

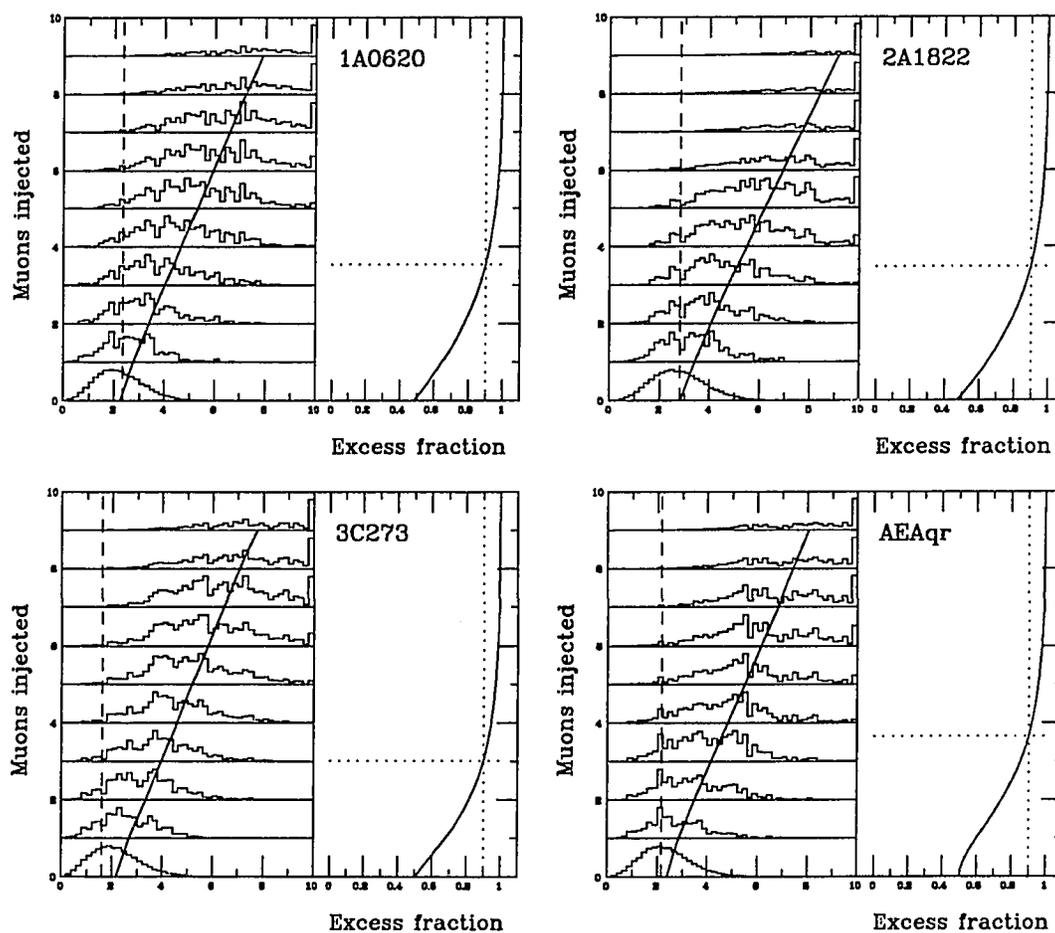


Figure A.1: Flux limits for sources 1A0620 through AE Aqr

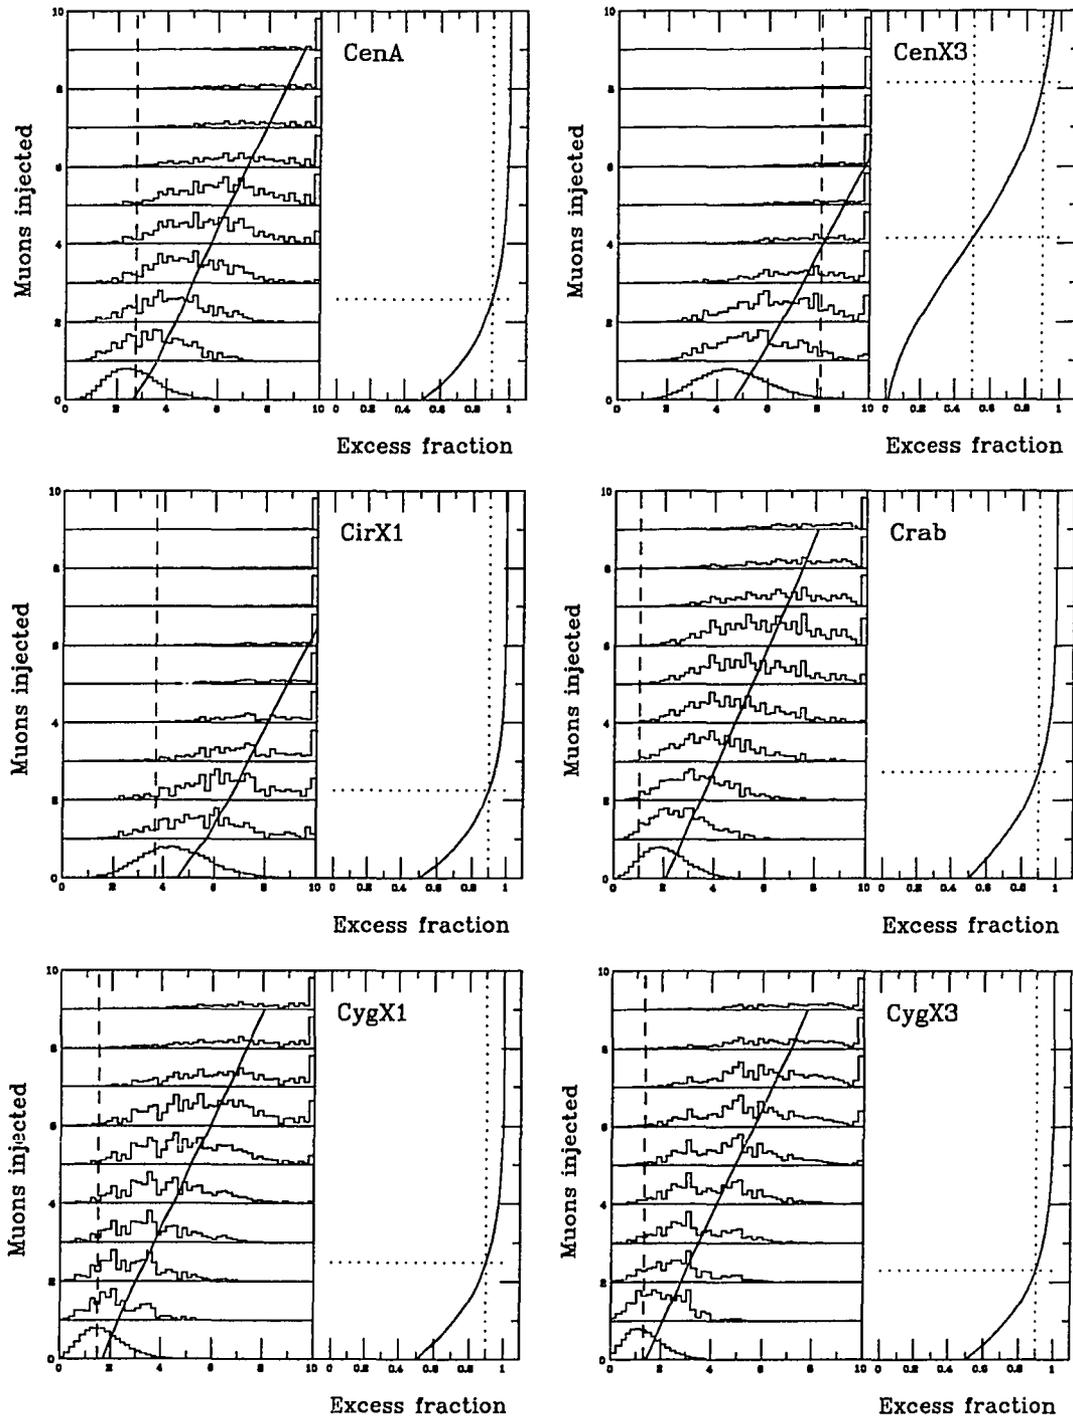


Figure A.2: Flux limits for sources Cen A through Cyg X-3

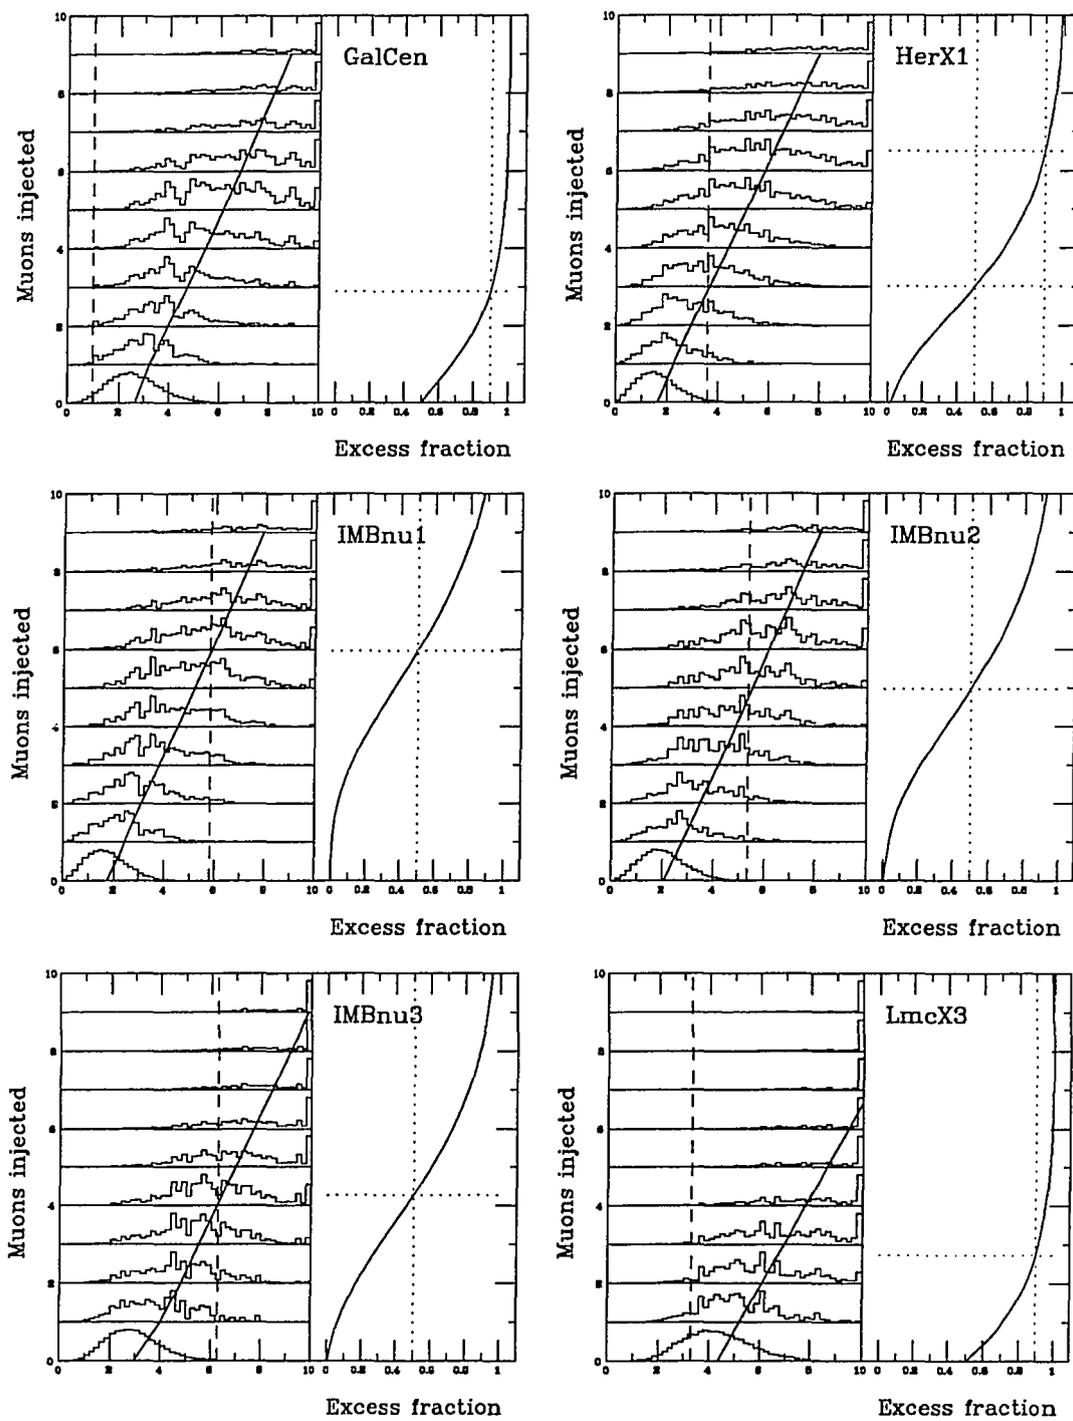


Figure A.3: Flux limits for sources Gal. Cen. through LMC X-3

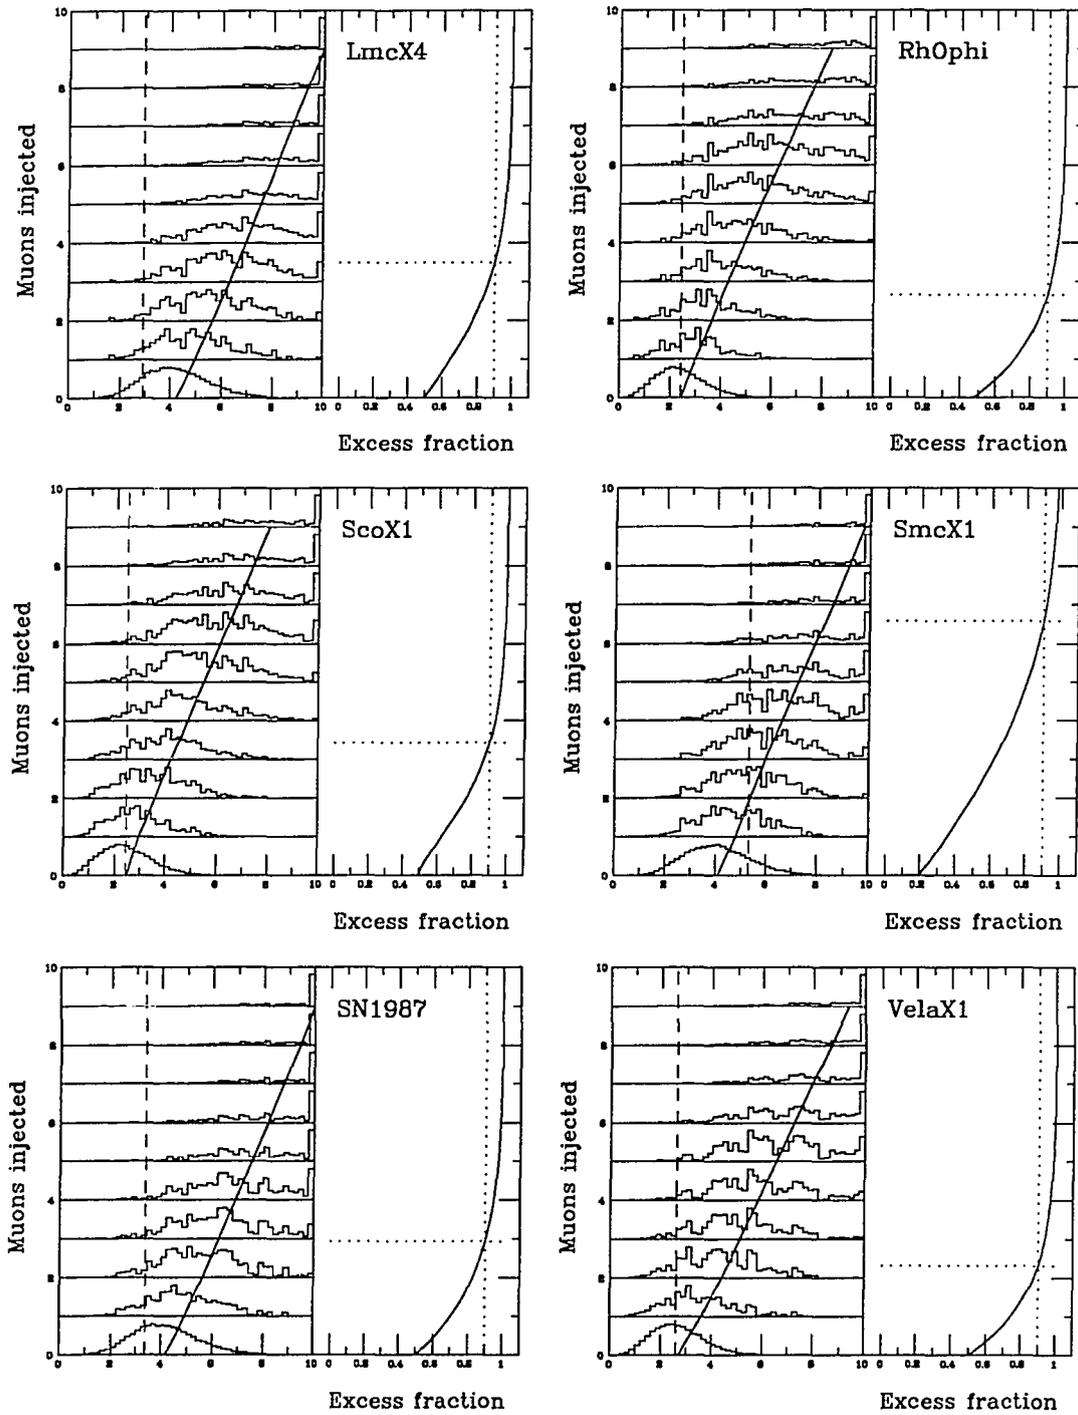


Figure A.4: Flux limits for sources LMC X-4 through Vela X-1

Appendix B

Geometry of Expected Timing

Here I will derive equation (4.2), which describes the expected time of a hit for a given geometry. To recapitulate the symbols: the muon passes the position \vec{v} at time t_0 travelling in direction \hat{d} . We observe a Čerenkov photon at a position \vec{x} .

The index of refraction of water is n , and the Čerenkov angle is ϑ_c , where $\cos \vartheta_c = 1/n$, $\sin \vartheta_c = \sqrt{n^2 - 1}/n$, $\tan \vartheta_c = \sqrt{n^2 - 1}$.

The point at which the photon was emitted is \vec{e} , and the distance the muon has travelled until there is $\lambda = (\vec{e} - \vec{v}) \cdot \hat{d}$. Note that it isn't just the magnitude of $\vec{e} - \vec{v}$, since λ can be negative. The point of closest approach between the muon and the photomultiplier is \vec{a} . Let μ be the distance the muon travelled from the emission point to the point of closest approach, γ the distance the photon travelled from being emitted to being observed, and i the impact distance between photomultiplier and muon at the point of closest approach. Figure B.1 shows the relation of all these vectors. Obviously $\mu = |\vec{a} - \vec{e}|$, $\gamma = |\vec{x} - \vec{e}|$ and $i = |\vec{x} - \vec{a}|$. Also, due to the geometry in the triangle $(\vec{e}, \vec{i}, \vec{x})$ the following two relations hold: $\gamma = i / \sin \vartheta_c$ and $\mu = i / \tan \vartheta_c$. Written explicitly the point of closest approach is $\vec{a} = \vec{v} + ((\vec{x} - \vec{v}) \cdot \hat{d})\hat{d}$.

The hit should have occurred at a time t^e , which is the time the muon is at $\vec{v} = t_0$, plus the time it takes the muon to get from \vec{v} to \vec{e} , plus the time it takes the photon to get from there to \vec{x} . Also remember that in our units a photon travels with a velocity of 1, and a muon with a velocity of n . Therefore:

$$\begin{aligned} t^e - t_0 &= \lambda/n + \gamma \\ &= \frac{\vec{e} - \vec{v}}{n} \cdot \hat{d} + \frac{n i}{\sqrt{n^2 - 1}} \\ &= \frac{\vec{a} - \vec{v}}{n} \cdot \hat{d} - \frac{\mu}{n} + \frac{n i}{\sqrt{n^2 - 1}} \end{aligned}$$

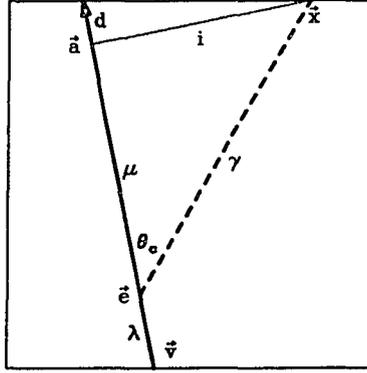


Figure B.1: Geometry of muon, photomultiplier and Čerenkov light

$$\begin{aligned}
 &= \frac{(\hat{d} \cdot (\vec{x} - \vec{v})) \cdot \hat{d}}{n} \hat{d} - \frac{i}{n \sqrt{n^2 - 1}} + \frac{ni}{\sqrt{n^2 - 1}} \\
 &= \frac{\vec{x} - \vec{v}}{n} \hat{d} + \frac{\sqrt{n^2 - 1}}{n} i \\
 &= \frac{\vec{x} - \vec{v}}{n} \hat{d} + \frac{\sqrt{n^2 - 1}}{n} |(\vec{x} - \vec{v}) + ((\vec{x} - \vec{v}) \cdot \hat{d}) \hat{d}|
 \end{aligned}$$

Note that this requires only a few simple operations to calculate; the only vector operations are one difference, one sum, two scalar products, two products of vectors with scalars, and one magnitude, which requires one square root. No other transcendental functions are used.

Appendix C

Event Format

The IMB Collaboration is quite extraordinary in that it probably has more different and incompatible ways of storing events than it has collaborators. It is a tradition that graduate students outline the current state in event storage formats in an appendix to their thesis.

C.1 Tape and Event Format

The fundamental structure of an event is

Header: General event-related information, for example the event identification, times, event-related CAMAC readout.

RN: User data which is added during analysis of this event, for example track reconstruction information.

MCINFO: For monte-carlo events a description of how this event was generated, in particular the physical particles involved.

Tube data: The tube hit information in a highly compressed format, known as packed tube data. There are two different methods of packing, one used in IMB-1 and IMB-2, and a new one used in IMB-3.

IA: User data which is added during analysis of this event, for example track reconstruction information.

IX: The decoded and calibrated tube data.

RN, MCINFO, IA and IX may be missing. RN is only used in the east-coast analysis, IA and IX only in the west-coast analysis. MCINFO is present only in monte-carlo events.

The header, RN, MCINFO and the tube data all reside in one long event storage buffer, known as InBuf, which is declared as an array of words.* IA and IX are stored in separate buffers. IA is an array of 350 longwords, which will be described in more detail later. IX is an array of 6×2048 longwords, also to be described later. MCINFO is an array of words of varying length, for its description please refer to reference [16].

The reason to have both the highly compressed tube data (in InBuf) and the decoded and calibrated tube data (in IX) is mostly historical: in the early days the analysis was

*The term *word* refers to 16-bit quantities, used usually as signed integer variables INTEGER*2. The term *longword* refers to 32-bit quantities, used usually as signed integer variables INTEGER*4 or floating-point variables REAL*4.

done on a CDC Cyber computer, which has 60-bit words. On this machine decoding the tube data was a very complex and time-consuming process; so it was considered more efficient to go through that process only once and store the decoded information with the event from than on. Unfortunately, this increases the size of each event by a factor of about five.

One more complication is the record format used on the computer. A "record" is a storage buffer of variable (but limited) length which contains one event, or one block of events. This usage of the term "record" pertains to a *logical* record. Data on tape may be stored on a labelled tape, which is structured into files, or on an unlabeled tape consisting just of the data blocks. In the latter case life is simple: one physical data block on tape corresponds directly to one logical record, and therefore also to one block containing one or several events.

On disk and on labelled tapes things are more complicated. Records can be stored in variable format, which means that each logical record is stored in one physical record of varying length. This is the simple case, and the one-to-one correspondence of event block, logical record and physical record is kept. In this format it is possible to determine the length of the next record to be read before actually reading it.

The second case are records stored in the *segmented* format: one logical record (which is what we consider one event or one block of event) is split into several physical records. This is done to circumvent some limitations of some operating systems, which might restrict a physical record length too much. Unfortunately, *segmented* records can be read directly by only by VAX FORTRAN programs, not by other computers, nor by other programming languages or system utilities. When reading *segmented* records the length of the next record to be read cannot not be determined in advance. To circumvent that the record length is sometimes stored in front of the actual record content; in that case it can be determined while reading. In the following we will use the notation LEN for this record length (LEN is a word). If this record length is not provided for a *segmented* record, it can still be read *if each record contains a unmistakable end-of-block marker* using the following time-consuming and error-prone method: first clear the data buffer used for the input record, and then just request to read the longest possible input record, ignoring errors caused by the input record being shorter than what was requested. Then search backwards from the end of the buffer for the end-of-block marker. The *segmented* record format is a holdover from times when computer power and software were very restrictive, but it seems to be firmly entrenched now.

Now to the possible event formats, listed in order of increasing complexity:

Mine format, also known as **online** format. This is how the data-taking computer assembles events. No data is actually written in this format, since on the raw tapes from the experiment the events are in blocked format (see below). Each event record consists of four integer word pointers, followed by the buffer (which

contains the header and the packed tube data), terminated by an end-of-event marker: `NxtEvt`, `RNaddr`, `MCaddr`, `TDaddr`, `InBuf(1...LEN)`, `EoE`.

Let me describe the meaning of the four pointers: `NxtEvt` points to the next event record in this block (see the description of the blocked format). It also indicates the length of this event record. `RNaddr` points to the beginning of the RN-array in the `InBuf` array and `MCaddr` points to the beginning of the MCINFO block in the `InBuf` array. Finally `TDaddr` points to the beginning of the packed tube data in the `InBuf` array. Data written at the experiment does not contain a RN array or an MCINFO block, so the last three pointers are all equal; but provision for these two additional features is already made in the event layout. The very last word in the event record is the end-of-event marker word `EoE`, decimal `-2` or hexadecimal `16FFFE`. Obviously the length of the record can be inferred from `NxtEvt`: the last word of the packed tube data is at word `NxtEvt-2` and the end-of-event marker is at word `NxtEvt-1`. This format is shown in some more detail in table C.1.

Blocked format is what the experiment's data tapes are written in. For efficiency reasons the data blocks written at the experiment are fairly long (on regular reel tape between 10240 and 16383 words, on 8mm tape up to 8192 bytes). To accomplish this several events are put into one data block. The first word in the data block is `NEvents`, the number of events in this data block. Then follow (`NEvents`) events, each in the format described above for the mine format. After the last event follows the end-of-block marker word, decimal `-28673` or hexadecimal `168FFF`. The `NxtEvt` pointer of each event points to the address of the next event or to the end-of-block marker. This blocked format is shown in table C.2.

This format is used only on data tapes from the experiment. One problem with this format is that you have to determine the length of the block while reading it. This is easily done on unlabeled tapes, for which this format was designed. The 8mm data tapes from the experiment use this blocked format but with *segmented* records, so the reading method outlined above has to be employed.

Special records are the exception to the blocked format. They are not actual events, but other information from the experiment. Their first word (which should have been `NEvents` if they were an event block) is a large number known as the special record type; that is how they can be recognized (for normal blocked events `NEvents` is always a small integer, at most about 20). Their length has to be determined while reading them, just as with the blocked event format, but they contain an end-of-block marker which helps in determining their length.

Raw format, also known as disk format, used in the west-coast analysis. This format is the event described above, but without the four pointers. Therefore the event cannot contain a RN array or an MCINFO block. The pointer to the beginning of the tube data is recorded in `InBuf(4)`. As the name says, it is used mostly for disk files, and contains no information except the raw event as it was read by the online computer. One *segmented* record per event. Format: `LEN`, `InBuf(1...LEN)`.

- Western MC** format is just like raw format, except that there is an MCINFO block between the header and the tube data. The format is again LEN, InBuf(1..LEN). Since there are no pointers, the position of the MCINFO block is determined by convention: it always begins at InBuf(301). Since the first longword of MCINFO is the number of bytes in MCINFO, you can determine its length after reading the event. To accommodate the MCINFO block the tube data is pushed further back in InBuf.
- Pass 0** format (this format is also used by the UpMu pass 0 and pass 1). Just like the raw format, except that after the event there are six additional longwords. The names and meanings for these six longwords vary, but they are always read into IA(201..206). Format: again one segmented record per event, composed of LEN, InBuf(1..LEN), IA(201..206). Note that LEN is not the length of the whole record, but only the length of InBuf.
- DST** format, also known as **pass 1** format. Careful, the UpMu pass 1 uses the pass 0 format described above. Contains all the information available in the west-coast analysis. One segmented record per event consisting of LEN, InBuf(1..LEN), IA(1..350), IX(1..6, 1..IA(15)). Note that the whole IA array is read, that IA(15) is the number of PMT hits in IX, and that the IX array is calibrated, so one accesses the tube hits in real format through the X array without having to unpack or calibrate the event. This format is fairly universal and easily interchanged among the west-coast analysis chain, but is very wasteful, with events typically being 30 KBytes long.
- SDST** format, which stands for shortened DST. A slight variation of the DST format: since IX is nothing but the calibrated and decoded PMT hits it is not actually required to read it; instead for this format one just reads LEN, InBuf(1..LEN), IA(1..350), and recreates the IX array from the packed tube data in InBuf. Note that (to my knowledge) no data is actually *written* or *stored* in this format; it is just a simplified (and faster) way of *reading* the DST format.
- Hawaii DST** format is the next logical step from the SDST format. The only difference to the is that now each event is one **variable** record, not a **segmented** record. That makes it possible to read these event easily using languages other than FORTRAN, and to use system utilities like SORT on this format. The format within the **variable** record is still the same as in the SDST format: LEN, InBuf(1..LEN), IA(1..350). Since no other collaborator is set up to process this format, I only use it internally for short-term storage.
- Disk** format, used in the east-coast analysis. Careful, this format is very different from the west-coast disk format. It is nearly the same as the mine format. Events are written in **segmented** records, one event per record. Instead of the NxtEvt pointer (which is not needed since by definition there is no next event) the length of the record is written, and by convention a zero word is added to the end of the record (after the end-of-event marker), making the format: LEN, RNaddr, MCaddr, TDaddr, InBuf(1..LEN-6), EoE, 0. This format can include a RN array (of any length, up to maximum of 256 words) and an MCINFO block. This

format is universal enough to be used for all stages of the east-coast analysis and monte-carlo work.

The different storage formats are summarized in table C.3.

Address in Buffer	Address in InBuf	Content
1		NxtEvt. Pointer to the beginning of the next event in Buffer, or to the EoE marker.
2		RNaddr. Pointer to the RN array in InBuf.
3		MCaddr. Pointer to the MCINFO block in InBuf.
4		TDaddr. Pointer to the tube data in InBuf.
5 ... RNaddr+4	1 ... RNaddr-1	Event header (sometimes incorrectly referred to as InBuf).
RNaddr+5 ... MCaddr+4	RNaddr ... MCaddr-1	Analysis information: the RN array.
RNaddr+5 ... MCaddr+4	RNaddr ... MCaddr-1	Monte-carlo information: the MCINFO block.
TDaddr+5 ... NxtEvt-2	TDaddr ... NxtEvt-1	Packed tube data.
NxtEvt-1	NxtEvt-6	EoE

Table C.1: Event format layout. EoE is the end-of-event marker $^{16}\text{FFFE}$. For the east-coast disk format replace NxtEvt with LEN-1, and add one word (content zero) at the end.

Word address	Content	
1	NEvents	
2	NxtEvt ₁	Event 1
3	RNaddr ₁	
4	MCaddr ₁	
5	TDaddr ₁	
6...NxtEvt ₁ -2	InBuf ₁	
NxtEvt ₁ -1	EoE	
NxtEvt ₁	NxtEvt ₂	Event 2
NxtEvt ₁ +1	RNaddr ₂	
NxtEvt ₁ +2	MCaddr ₂	
NxtEvt ₁ +3	TDaddr ₂	
NxtEvt ₁ +4...NxtEvt ₂ -2	InBuf ₂	
NxtEvt ₂ -1	EoE	
...		
NxtEvt _{n-1}	NxtEvt _n	Event n
NxtEvt _{n-1} +1	RNaddr _n	
NxtEvt _{n-1} +2	MCaddr _n	
NxtEvt _{n-1} +3	TDaddr _n	
NxtEvt _{n-1} +4...NxtEvt _n -2	InBuf _n	
NxtEvt _n -1	EoE	
NxtEvt _n	EoB	

Table C.2: Blocked event format. EoE is the end-of-event marker ¹⁶FFFE, EoB is the end-of-block marker ¹⁶8FFF.

Mine	NxtEvt, RNaddr, MCaddr, TDaddr, InBuf(1...LEN), EoE Never actually used directly.
Blocked	NEvents, Event ₁ ... Event _n , EoB In physical data blocks on reel tape In segmented records on 8mm tape.
Raw (west)	LEN, InBuf(1...LEN) segmented records.
MC (west)	LEN, InBuf(1...LEN) segmented records, MCINFO is at InBuf(301).
Pass 0	LEN, InBuf(1...LEN), IA(201...206) segmented records.
DST	LEN, InBuf(1...LEN), IA, IX(1...6, 1...IA(15)) segmented records.
SDST	LEN, InBuf(1...LEN), IA segmented records.
Hawaii DST	LEN, InBuf(1...LEN), IA variable records
Disk (east)	LEN, RNaddr, MCaddr, TDaddr, InBuf(1...LEN-6), EoE, 0 segmented records.

Table C.3: Summary of storage formats

C.2 Layout of the INBUF Data

InBuf contains a large number of items of general interest for an event, in particular the information which identifies the event. In the following listing, the first column identifies the word address in InBuf where an item begins, the second column defines the data type of the item, and the last one is a description. The following data types are used: I*1 is an unsigned integer byte, I*2 is a signed integer word, L*2 is a logical word (even values mean "false", odd values mean "true", typically 0 and -1 are used), I*4 is a signed integer longword, R*4 is a floating-point longword. Longer items are I*6 a 48-bit tripleword and I*8 a 64-bit quadword, neither of which can be represented by an atomic data type on most computers. $\vec{I*2}$ is a vector formed by three signed integer words, similarly $\vec{I*4}$ and $\vec{R*4}$ are vectors formed from I*4 and R*4 longwords. All distances and locations are measured in cm, all velocities are measured in cm/ns, scaled up by a factor of 100. Where an integer vector specifies a direction (the three elements of the vector are the three direction cosines) the integer value is scaled up by a factor of 1000.

The following list is hopefully accurate for the end of 1990; many items in it have changed (and presumably will always change) over time. A very good reference for the data in InBuf is reference. [16] for the current status, and reference. [93] for 1986; the ultimate reference is always the GETHDR routine of the online computer. Furthermore, a few items are interpreted differently for calibration runs. The entries marked with a (C) are valid only for calibration runs.

1	I*2	Data type: Normally 0, add 100 if RN is present, add 200 if MCINFO is present. Numbers > 4096 mark special records.
2	I*4	Event number on tape, also known as MBD number. This event number is assigned when an event is written to tape.
4	I*2	Location of the tube data in InBuf. If the TDaddr pointer is available use it instead of InBuf(4).
5	I*4	Trigger number in this run, also known as SPU number. This event number is assigned by the CAMAC system when the detector triggers, and may be different from the MBD number on a cut tape.
7	I*6	Total time since the beginning of this run, in units of 2^{-18} s.
10	I*6	Dead time since the beginning of this run, in units of 2^{-18} s.
13	I*4	Total number of NPATCH triggers in this run.

15	I*4	Accepted number of NPATCH triggers in this run. Counted only when the detector is live.
17	I*4	Total number of NTUBE triggers in this run.
19	I*4	Accepted number of NTUBE triggers in this run.
21	I*2	Buffer memory from which this event was read, or -1 if the read bypassed the buffer memory.
22	I*2	Hardware trigger pattern unit, one bit for each possible trigger.
23	I*2	Software pattern unit, one bit for each possible reason to save the event.
24	L*2	Did the CONE fitter run ?
25	I*2	Number of hits in the T1 timescale.
26	$\overrightarrow{I*2}$	Direction fitted by CONE.
29	I*2	Calibration laser filter setting (C).
30	I*2	Measured laser fire delay in ns (C).
31	I*2	Measured detector trigger delay in ns (C).
32	I*2	Measured trigger pulse delay in ns (C).
33	I*2	Photodiode Q charge (C).
34	L*2	True if the CAMAC system was read out for this event, false if only DREAM data is available.
35	I*2	Run number, also known as tape number.
36	32×I*1	Number of T1 hits for each patch.
52	32×I*2	Patch TDCs for each patch, in units of 1/2 ns.
84	32×I*2	Patch scalers, hits received in each patch in this run.
116	I*2	File number on the tape.
117	I*2	Digitization pulse fanout difference between the long and short cables, in units of 20ps. Often just known as the fanout TDC.
118	5× I*2	Scanning environmental ADC, among them water depth.
123	I*2	First reason to save event, or -1 if no cuts were applied.
124	$\overrightarrow{I*2}$	Center of mass of all hits.
127	I*2	Number of patches hit.
128	I*2	Muon velocity from the FASTMU fit.
129	I*2	Mean uncalibrated Q in T1.
130	I*2	Total calibrated Q in T1.
131	L*2	Did this event occur within 10ms of the previous one ?

132		Digitizer card occupancy flags, one bit for each of the 256 digitizer cards, set if any tube on this card was hit in the last 25 events.
148	I*2	Number of hits in the T2 timescale.
149	L*2	Did FASTMU run ?
150	I*2	Offset in ms to be added to the WWV clock time.
151	5× I*2	Scanning environmental ADCs, among them room temperature.
156	I*6	WWV clock in packed BCD format.
159	I*4	SPU number of last event which had CAMAC read out.
161	I*8	Time recorded by the online computer, in VAX format.
165	$\vec{I*2}$	Position fitted by FASTMU (not the vertex).
168	$\vec{I*2}$	Direction fitted by FASTMU.
171	I*2	Time since previous trigger in ms.
172	I*2	Block number on tape.
173	L*2	Did WAVER run ?
174	I*2	Muon velocity from the WAVER fit.
175	$\vec{I*2}$	Position fitted by WAVER (not the vertex).
178	$\vec{I*2}$	Direction fitted by WAVER.
181	I*2	Number of early clusters found by FASTMU.
182	I*2	Size of largest early cluster.
183	I*2	Size of second largest early cluster.
184	I*2	Number of hot clusters found by FASTMU.
185	I*2	Size of largest hot cluster.
186	I*2	Size of second largest hot cluster.
187	32×I*2	Sum of calibrated Q for each patch.
221	I*2	Number of events since DREAM last read.
222	I*2	Number of full DREAM buffers.
223	I*2	Next free DREAM buffer number.
224	I*4	DREAM counter: WWV offset, in ms.
226	I*4	Spare DREAM counter 1.
228	I*4	Spare DREAM counter 2.
230	I*4	DREAM counter: 2 ¹⁸ Hz clock.
232	I*4	4-bit pattern unit read from DREAM.
234	I*4	Difference in DREAM counter WWV since last trigger.
236	I*4	Difference in DREAM counter Spare 1 since last trigger.

238	I*4	Difference in DREAM counter Spare 2 since last trigger.
240	I*4	Difference in DREAM counter 2^{18}Hz since last trigger.

An unfortunate habit of the west-coast analysis is to overwrite items in the header during data analysis; this is caused by the fact that an event increases so much in size when going from raw format to DST format, so it is often more efficient to keep the event in raw format and store additional information in the header. The only additional information in the header of interest for this analysis are the pass 0 and pass 1 upward-going muon fits used in IMB-1 and IMB-2:

28	I*2	Pass 0 fit entry: tube number nearest to entry point.
29	$\vec{I*2}$	Pass 0 fit direction, scaled by 10000.
32	$\vec{I*2}$	Pass 1 fit entry position.
52	$\vec{I*2}$	Pass 1 fit direction, scaled by 10000.
55	I*2	Pass 1 fit quality estimate.

C.3 Layout of the IX, IA and RN Data

The IX and X arrays contain the decoded and eventually calibrated tube hit data. IX is defined as an array of 6×2048 integer longwords, and X as an array of 6×2048 floating-point longwords. They occupy the same storage space, so only one of them can be valid at a given time. When unpacking the tube data the IX array is filled, and when calibrating the calibrated data is moved in place into the X array.

After unpacking, $IX(1, i)$ contains the PMT number of hit number i , $IX(4, i)$ contains the uncalibrated time of the hit on the T1 scale in ns, $IX(6, i)$ contains the uncalibrated time of the hit on the T2 scale (in units of 15ns). If there is a T1 hit $IX(5, i)$ contains the uncalibrated Q associated with the T1 hit, otherwise it contains the uncalibrated Q associated with the T2 hit. The number of valid hits which can be accessed in IX (that is the range of the index i) is found in IA(15).

After calibrating, $X(1, i)$ through $X(3, i)$ contain the geometrical position of the hit PMT, $X(4, i)$ the calibrated time of the on the T1 scale (now measured in cm, the conversion factor is the speed of light in water, $c_W = 22.556\text{cm/ns}$). $X(6, i)$ contains the calibrated time of the hit on the T2 scale in ns. $X(5, i)$ contains the calibrated Q of the his in units of photo-electrons; if there is a T1 hit is corresponds to that hit, otherwise to the T2 hit. The range of the i index is a little more complicated: $i = 1$ through $i = IA(11)$ are the T1 hits which calibrated correctly, $i = IA(11) + 1$ through $i = IA(13)$ are the T2 hits which calibrated correctly, only for PMTs which did not have a T2 hit, and finally $i = IA(13) + 1$ through $i = IA(15)$ are the hits which could not be calibrated. Any of the three ranges may be empty.

The IA array contains additional reconstruction information from the west-coast analysis. It is defined as an array of 350 integer longwords, and some entries are overlaid to floating point longwords. I will only list entries used in this analysis:

11	I*4	Number of good calibrated T1 hits.
13	I*4	Number of good calibrated T2 without T1 hits.
15	I*4	Total number of hits in the IX or X arrays.
90	R*4	Year of the event, in universal time.
91	R*4	Month.
92	R*4	Day.
93	R*4	Hour.
94	R*4	Minute.
95	R*4	Second.
96	R*4	Source of the time information: 1=Detector logbook, 2=Computer system clock, 3=WWV clock.
97	R*4	Local sidereal time.
98	R*4	Right ascension of this event in equatorial coordinates.
99	R*4	Declination of this event in equatorial coordinates.
100	R*4	Event type flag: 0=through-going muon, 1=stopping muon.
101	R*4	Path length of this muon in the detector, in cm.
102	R*4	Sum of Q in the Čerenkov cone, in photo-electrons.
103	R*4	Anisotropy.
104	R*4	Julian day of this event.
201	$\overrightarrow{I*4}$	Upward-going muon pass 0 and pass 1 fit direction.
204	I*4	Upward-going muon pass 0 and pass 1 fit entry, specified as the tube number nearest to the entry point.
205	R*4	Upward-going muon pass 0 and pass 1 fit quality estimate.
300	I*4	Flag bits for each of the following fits: 1=red, 2=blue, 4=purple, 8=green. For upward-going muons the red fit shall be valid, the other three will be ignored.
305	$\overrightarrow{I*4}$	Position of the red fit.
309	$\overrightarrow{I*4}$	Direction of the red fit.
312	$\overrightarrow{I*4}$	Position of the blue fit.
316	$\overrightarrow{I*4}$	Direction of the blue fit.

319	$\overrightarrow{I*4}$	Position of the purple fit.
323	$\overrightarrow{I*4}$	Direction of the purple fit.
326	$\overrightarrow{I*4}$	Position of the green fit.
330	$\overrightarrow{I*4}$	Direction of the green fit.
350	$I*4$	Calibration tape used to calibrate this event.

Finally, the RN array is the east coast equivalent of the west-coast IA array: it contains additional reconstruction information. It consists of two overlaying arrays of 256 integer words and 128 floating-point longwords. The addressing convention used here is different from the one used in the event header and in IA: $IRN(i)$ refers to the i -th integer word of the storage area, and $RN(i)$ to the i -th real longword of the same storage space. Therefore $RN(i)$ occupied the same space as $IRN(2 \cdot i - 1)$ and $IRN(2 \cdot i)$. The only items of interest for this analysis are $RN(101 \dots 103)$, the muon vertex, and $RN(104 \dots 106)$, the muon direction.

Appendix D

Pauli's 1930 Letter

Open letter to the group of radioactives at their "regional chapter meeting" in Tübingen:

Physics Institute
of the Federal Institute of Technology
Zürich

Zürich, December 4, 1930
Gloriastrasse

Dear Radioactive Ladies and Gentlemen,

As the carrier of these lines, whom to listen to graciously I ask you, will explain in more details, in view of the "wrong" statistic of the N and ${}^6\text{Li}$ nuclei, and also in view of the continuous beta-spectrum, I have come to a desperate way out to rescue the "change law" of statistics and the energy law. Which is: the possibility that there exist new particles, which I shall name neutrons, which have spin $1/2$ and follow the exclusion principle, and are different from light quanta by not propagating at the speed of light. The mass of neutrons should be of the same order of magnitude as the electron mass, and in any case not larger than 0.01 proton mass. — The continuous beta-spectrum would then be explainable by assuming that in each beta decay an additional neutron is emitted with the electron, in such a fashion that the sum of energies of neutron and electron is constant.

Now the next question is which forces act on the neutrons. For reasons of wave mechanics (the carrier of these lines knows more) the most likely model seems to be one in which the neutron at rest is a magnetic dipole of a certain moment μ . The experiments seem to require that the ionizing effect of such a neutron cannot be larger than that of a gamma-ray, therefore μ can seemingly not be larger than $e \cdot (10^{-13}\text{cm})$.

For the time being I am not yet confident to publish something about this idea, and first turn to you, dear Radioactives, with the question of what the prospect of the experimental demonstration of such a neutron is, if it has the same or 10 times larger penetration power than a gamma-ray.

I admit that my way out seems to be little probable from the onset, because the neutrons, if they exist, should long have been seen. But nothing ventured, nothing gained, and the seriousness of the situation with the continuous beta-spectrum is illuminated by a remark of the predecessor in my office, Mr. Debye, when we met recently in Brussels: "Oh, about that one should not even think, and same for the new taxes." Therefore one should seriously discuss any way to rescue. — Therefore, dear Radioactives, please

scrutinize and judge. — Unfortunately I cannot appear personally in Tübingen, since I am not available due to a ball which will be held here in the night from the 6. to the 7. of December. — With many regards to you, and also to Mr. Baek, your most humble servant

signed W. Pauli

Bibliography

- [1] J. N. Bahcall (with an appendix jointly with R. Davis Jr.), *Neutrino Astrophysics*, Cambridge University Press (1989)
- [2] E. Batschelet, *Circular Statistics in Biology*, Academic Press (1981)
- [3] G. Battistoni et al., *Phys. Lett.*, **155B**, 465 (1985)
G. Battistoni et al., *Proc. 19 ICRC*, **1**, 62, La Jolla (1985)
- [4] V. S. Berezinsky, S. I. Grigoreva, G. T. Zatsepin, *Astrophys. and Space Science*, **36**, 17 (1975)
- [5] V. S. Berezinsky, V. L. Ginzburg, *Mon. Not. Royal Astr. Soc.*, **194**, 3 (1981)
- [6] V. S. Berezinsky, *Astrophysics of Cosmic Rays*, Elsevier (1991)
- [7] Ch. Berger et al., *Z. Phys. C*, **48**, 221 (1990)
- [8] L. B. Bezrukov, E. V. Bugaev, *Sov. Jou. Nucl. Phys.*, **33**, 635 (1981)
- [9] R. Bionta et al., *Phys. Rev. Lett.*, **58**, 1494 (1987)
- [10] R. Blandford, D. Eichler, *Phys. Reports*, **154**, 1 (1987)
- [11] K. T. S. Brazier et al., *Proc. 21 ICRC*, **2**, 292, Adelaide (1990)
- [12] K. T. S. Brazier et al., *Proc. 21 ICRC*, **2**, 296, Adelaide (1990)
- [13] K. T. S. Brazier et al., *Proc. 21 ICRC*, **2**, 300, Adelaide (1990)
- [14] K. T. S. Brazier et al., *Proc. 21 ICRC*, **2**, 379, Adelaide (1990)
- [15] C. Brink et al., *Proc. 21 ICRC*, **2**, 283, Adelaide (1990)
- [16] David William Casper: *Experimental Neutrino Physics and Astrophysics with the IMB-3 Detector*, PhD Thesis, University of Michigan (1990)
- [17] G. L. Cassiday et al., *Nucl. Phys.*, **B14A**, 291 (1990)
- [18] C. J. Cesarsky, T. Montmerle, *Space Science Rev.*, **36**, 173 (1983)
- [19] P. M. Chadwick, *Astron. Astrophys.*, **151**, L1 (1985)
- [20] M. L. Cherry, *Nucl. Phys.*, **B14A**, 38 (1990)
- [21] G. V. Chibisov, *Sov. Phys. Usp.*, **19**, 624 (1976)
- [22] A. E. Chudakov, O. G. Ryazhskaya, *Proc. Neutrino 77*, **1**, 155 (1978)
A. E. Chudakov et al., *Proc. 16 ICRC*, **10**, 287, Kyoto (1979)
M. M. Boliev et al., *Proc. 18 ICRC*, **7**, 120, Bangalore (1983)
- [23] D. Ciampa, R. W. Clay, *Univ. of Adelaide preprint ADP-88-94*, (1988)
- [24] D. Ciampa, R. W. Clay, P. G. Edwards, *Astrophys. J.*, **346**, 151 (1989)
- [25] C. L. Cowan et al., *Science*, **124**, 103 (1956)
- [26] M. F. Crouch et al., *Phys. Rev. D*, **18**, 2239 (1978)
- [27] G. Danby et al., *Phys. Rev. Lett.*, **9**, 36 (1962)
- [28] R. Davis Jr., *Phys. Rev.*, **97**, 766 (1955)
- [29] R. Davis Jr., D. S. Harmer, K. C. Hoffman, *Phys. Rev. Lett.*, **20**, 1205 (1968)
- [30] *The solar neutrino results of the Homestake experiment are summarized in:*

- R. Davis Jr. et al., *Proc. 2nd Intern. Workshop on Neutrino Telescopes*, 1, Venice (1990)
- [31] B. L. Dingus et al., *Phys. Rev. Lett.*, **61**, 1906 (1988)
- [32] D. Eichler, *Astrophys. J.*, **232**, 106 (1979)
- [33] E. Eichten, I. Hinchliffe, K. D. Lane, C. Quigg, *Rev. Mod. Phys.*, **56**, 579 (1984)
Errata: *Preprint* FERMLAB-PUB-84/17
- [34] S. M. Faber, J. S. Gallagher, *Ann. Rev. Astron. and Astrophys.*, **17**, 135 (1979)
- [35] H. Faissner et al., *Phys. Lett.*, **103B**, 234 (1981)
H. Faissner et al., *Phys. Rev. D*, **28**, 1198 (1983)
- [36] G. Feinberg, *Phys. Rev.*, **110**, 1482 (1958)
- [37] E. Fermi, *Z. Phys.*, **88**, 161 (1934)
E. Fermi, *Nuovo Cim.*, **11**, 1 (1934))
- [38] E. Fernandez, Presented at *Neutrino-90*, Geneva (1990)
- [39] T. K. Gaisser, T. Stanev, *Phys. Rev. D*, **30**, 985 (1984)
- [40] T. K. Gaisser, T. Stanev, *Phys. Rev. D*, **31**, 2770 (1985)
- [41] T. K. Gaisser, A. F. Grillo, *Phys. Rev. D*, **36**, 2752 (1987)
- [42] T. K. Gaisser, *Proc. High Energy Physics in the 1990s*, Snowmass, 2 (1988)
- [43] M. Goldhaber, L. Grodzins, A. W. Sunyar, *Phys. Rev.*, **109**, 1015 (1958)
- [44] K. Greisen, *Proc. Conference on Instrumentation for High-Energy Physics*, Berkeley, 209 (1960)
- [45] K. Greisen, *Ann. Rev. Nucl. Science*, **10**, 63 (1960)
- [46] K. Greisen, *Phys. Rev. Lett.*, **16**, 748 (1965)
- [47] K. S. Hirata et al., *Phys. Rev. Lett.*, **58**, 1490 (1987)
K. S. Hirata et al., *Phys. Rev. D*, **38**, 448 (1988)
- [48] K. S. Hirata et al., *Phys. Rev. Lett.*, **65**, 1297 (1990)
- [49] M. R. Issa, A. W. Wolfendale, *Nature*, **292**, 430 (1981)
M. R. Issa, A. W. Strong, A. W. Wolfendale, *J. Phys. G: Nucl. Phys.*, **7**, 565 (1981)
- [50] J. R. Jokipii, G. Morfill, *Astrophys. J.*, **290**, L1 (1985)
J. R. Jokipii, G. Morfill, *Astrophys. J.*, **312**, 170 (1987)
- [51] L. A. Kelley: *A Search for Very High Energy Gamma Ray Emission from Hercules X-1*, PhD Thesis, University of Hawaii (1991)
- [52] E. W. Kolb, M. S. Turner, T. P. Walker, *Phys. Rev. D*, **32**, 1145 (1985)
- [53] M. R. Krishnaswamy et al., *Proc. Roy. Soc London*, **323**, 489 (1971)
- [54] P. O. Lagage, C. J. Cesarsky, *Astron. Astrophys.*, **125**, 249 (1983)
- [55] R. C. Lamb, *Proc. NATO Workshop VHE gamma-ray astronomy*, Durham (1986)
- [56] R. C. Lamb et al., *Astrophys. J.*, **328**, 113 (1988)
- [57] J. G. Learned, S. Pakvasa, T. Weiler, *Phys. Lett.*, **207B**, 79 (1988)
- [58] H. Lee, Y. Koh, *Chungnam National University Preprint*, CNUPHY-1989-T1 (1989)
- [59] W. Lohmann, R. Kopp, R. Voss, *Preprint* CERN EP/85-03 (1985)

- [60] Z. Maki, M. Nakagawa, S. Sakata, *Progr Theor. Phys.*, **28**, 315 (1962)
- [61] K. V. Mardia, *Statistics of Directional Data*, Academic Press (1972)
- [62] M. A. Markov, I. M. Zheleznykh, *Nucl. Phys.*, **27**, 385 (1961)
- [63] M. L. Marshak et al., *Phys. Rev. Lett.*, **54**, 2079 (1985)
- M. L. Marshak et al., *Phys. Rev. Lett.*, **55**, 1965 (1985)
- J. A. Kochocki et al., *Phys. Rev. D*, **42**, 2967 (1990)
- [64] T. Makano, *Proc. 21 ICRC*, **2**, 266, Adelaide (1990)
- [65] S. P. Mikheyev, A. Yu. Smirnov, *Sov. Jou. Nucl. Phys.*, **42**, 913 (1985)
- S. P. Mikheyev, A. Yu. Smirnov, *Sov. Phys. JETP*, **64**, 4 (1986)
- S. P. Mikheyev, A. Yu. Smirnov, *Nuovo Cim.*, **9C**, 17 (1986)
- [66] A. R. North et al., *Proc. 21 ICRC*, **2**, 275, Adelaide (1990)
- [67] Yuichi Oyama, *Experimental Study of Upward-going Muons in Kamiokande*, PhD thesis, University of Tokyo, ICR-Report 193-89-10 (1989)
- [68] W. Pauli, in a letter to physicists at Tübingen, translated in appendix D
W. Pauli, *Proc. VVI Solvay Congress, Brussels*, Paris: Gauthier
Villars), 324 (1933)
- [69] R. D. Peccei, H. R. Quinn, *Phys. Rev. Lett.*, **38**, 1440 (1977)
- R. D. Peccei, H. R. Quinn, *Phys. Rev. D*, **16**, 1791 (1977)
- See reference [98] for further discussion.
- See reference [35] for a claim of the discovery of Axions.
- [70] M. R. Perl et al., *Phys. Rev. Lett.*, **35**, 1489 (1975)
- [71] B. Pontecorvo, *Sov. Phys. JETP*, **6**, 429 (1946)
- [72] W. H. Press, B. P. Flannery, S. A. Teukolsky, W. T. Vetterling, *Numerical Recipes*, Cambridge University Press (1986)
- [73] R. J. Protheroe, D. Kazanas, *Preprint NASA-TM-84974*, NASA Goddard (1983)
- [74] R. J. Protheroe, R. W. Clay, P. R. Gerhardy, *Astrophys. J.*, **280**, L47 (1984)
- [75] R. J. Protheroe, R. W. Clay, *Nature*, **315**, 205 (1985)
- [76] C. Quigg, M. H. Reno, T. P. Walker, *Phys. Rev. Lett.*, **57**, 774 (1986)
- [77] R. C. Rannot et al., *Proc. 21 ICRC*, **2**, 315, Adelaide (1990)
- [78] B. C. Raubenheimer et al., *Proc. 20 ICRC*, **1**, 267, Moscow (1987)
- B. C. Raubenheimer et al., *Proc. 20 ICRC*, **1**, 303, Moscow (1987)
- [79] Lord Rayleigh, *Nature*, **72**, 318 (1905)
- [80] F. Reines, C. L. Cowan, *Phys. Rev.*, **92**, 830 (1953)
- [81] F. Reines, *Proc. Conference on Instrumentation for High-Energy Physics*, Berkeley, 212 (1960)
- [82] F. Reines, *Ann. Rev. Nucl. Science*, **10**, 1 (1960)
- [83] F. Reines et al., *Phys. Rev. D*, **4**, 80 (1971)
- [84] M. H. Reno, C. Quigg, *Phys. Rev. D*, **37**, 657 (1988)
- [85] L. K. Resvanis et al., *Astrophys. J.*, **328**, L9 (1988)
- [86] J. P. Roques et al., *Proc. 19 ICRC*, **1**, 193, La Jolla (1985)
- [87] S. Sakata, T. Inoue, *Progr Theor. Phys.*, **1**, 143 (1946)

- [88] J. D. Scargle, *Astrophys. J.*, **263**, 835 (1982)
- [89] B. V. Sreekantan et al., *Proc. 21 ICRC*, **2**, 340, Adelaide (1990)
- [90] F. W. Stecker et al., *Phys. Rev. Lett.*, **66**, 2697 (1991)
- [91] M. L. Stevenson, *DUMAND Internal Report*, DIR-5-91, unpublished.
- [92] K. Suga et al., *Proc. Workshop on Tech. in UHE gamma-ray astronomy, Proc. 19 ICRC*, **1**, 48, La Jolla (1985)
- [93] Robert C. Svoboda: *A Search for Astrophysical Point Sources of Neutrinos Using a Large Underground Water Čerenkov Detector*, PhD Thesis, University of Hawaii (1985)
- [94] John Updike: "Cosmic Gall", in *Telephone Poles and other poems*, Alfred A. Knopf (1963)
- [95] L. V. Volkova, *Sov. Jou. Nucl. Phys.*, **31**, 784 (1980)
- [96] W. A. Wallis, *Econometrica*, **10**, 229 (1942)
- [97] T. C. Weekes et al., *Astrophys. J.*, **342**, 3789 (1989)
- [98] S. Weinberg, *Phys. Rev. Lett.*, **40**, 223, (1978)
F. Wilczek, *Phys. Rev. Lett.*, **40**, 279, (1978))
- [99] L. Wolfenstein, *Phys. Rev. D*, **17**, 2369 (1978)
L. Wolfenstein, *Phys. Rev. D*, **20**, 2634 (1979)
- [100] F. Zwicky, *Helv. Phys. Acta*, **6**, 110 (1933)