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.

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.

ProQuest Information and Learning
300 North Zeeb Road, Ann Arbor, MI 48106-1346 USA
800-521-0600

UMI®

A Search for Astrophysical Point Sources of Neutrinos with Super-Kamiokande

Andrew Lawrence Stachyra

A dissertation submitted in partial fulfillment
of the requirements for the degree of

Doctor of Philosophy

University of Washington

2002

Program Authorized to Offer Degree: Physics

UMI Number: 3041062

UMI Microform 3041062

ProQuest Information and Learning Company
300 North Zeeb Road
P.O. Box 1346
Ann Arbor, MI 48106-1346

In presenting this dissertation in partial fulfillment of the requirements for the Doctoral degree at the University of Washington, I agree that the Library shall make its copies freely available for inspection. I further agree that extensive copying of this dissertation is allowable only for scholarly purposes, consistent with "fair use" as prescribed in the U.S. Copyright Law. Requests for copying or reproduction of this dissertation may be referred to Bell and Howell Information and Learning, 300 North Zeeb Road, Ann Arbor, MI 48106-1346, to whom the author has granted "the right to reproduce and sell (a) copies of the manuscript in microform and/or (b) printed copies of the manuscript made from microform."

Signature Andrew L. [illegible]

Date 2/8/2002

University of Washington
Graduate School

This is to certify that I have examined this copy of a doctoral dissertation by

Andrew Lawrence Stachyra

and have found that it is complete and satisfactory in all respects,
and that any and all revisions required by the final
examining committee have been made.

Chair of Supervisory Committee:

R. Jeffrey Wilkes

Reading Committee:

R. Jeffrey Wilkes

Gerald A Miller

Gerald A. Miller

Christopher W. Stubbs

Date: 2/08/02

University of Washington

Abstract

A Search for Astrophysical Point Sources of Neutrinos with Super-Kamiokande

Andrew Lawrence Stachyra

Chair of Supervisory Committee:
Professor R. Jeffrey Wilkes
Department of Physics

Using a data sample of 2037 upward-going muons gathered by the Super-Kamiokande detector during 1479.6 days of average effective live time between April 1, 1996, and December 8, 2000, a search for evidence of astrophysical point sources of neutrinos is attempted. The results of two different types of all-sky surveys are presented, as well as 90% upper level confidence limits on the maximum observationally consistent neutrino flux that might be coming from selected source candidates, and the data is also examined for the possibility of a time correlated excess of muons associated with the 1997 extended flaring episode of Markarian 501. No evidence for astrophysical point sources of neutrinos is found. In addition to explanations of the data analysis procedures and results, discussions of the detector design and the upward-going muon data reduction are also included.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

There are many, many people who have helped me to complete the personal journey that this document represents; so many, in fact, and so many debts of gratitude owed to them, that I couldn't possibly imagine writing a passage that would be simultaneously long enough to acknowledge them all and yet still short enough that an average reader would ever finish it. These acknowledgments, therefore, will necessarily leave a great deal unstated, but I would like to make it clear that I am ever mindful of the role that others have played in my graduate education, and the omission of any person or deed from these pages should in no way be construed as a lapse of memory or awareness. Rather, any such omission is almost certainly just a casualty of the struggle for brevity. On that note, I would like to open my remarks with a quite general expression of gratitude especially to all of those many individuals whose names and numerous little contributions I will not manage to include in the paragraphs that follow. Rest assured, those of you, my friends, relatives, and colleagues who may come away from reading this disappointed not to find yourselves mentioned, you are not forgotten. Your help or interest in my progress was noticed and appreciated. Thank you.

Among those whom I do have enough space to recognize explicitly, my advisor, R. Jeffrey Wilkes, stands out prominently, and I would like to thank him for the many positive ways in which he has influenced my development into a fully fledged independent researcher. He leads by example, and much of the transformation that graduate school has helped me to effect within myself has originated in my attempts to emulate him.

I would also like to thank Shige Matsuno for his unfailingly prompt and detailed written responses to the many questions that I e-mailed to him over the last few years about the techniques that he uses to perform his own astrophysical point source analysis, and before that about his initial design of the off-site upward-going muon data reduction program, from which I borrowed heavily in order to write the code for the merged on-site/off-site reduction program described in chapter three. I also appreciate, furthermore, his consistent willingness to shoulder the burden of many of the less appetizing responsibilities of leadership within the upward-going muon group that have seemed so frequently to fall to him.

John G. Learned has been, more than he probably realizes, a terrific teacher and technical resource at a small but important number of critical junctures in the last year or two, during the most active phases of development of my analysis, and I thank him for those few crucial lectures and e-mail essays that he happened fortuitously to deliver to me at just the moment when I needed them most. He certainly knows how to search for point sources.

Besides the people already mentioned, most of the members of the upward-going muon analysis group, and indeed, many of my other collaborators as well, have taught me something new or otherwise helped me in some small but vital way at some point during my time with them, and I feel a sense of debt to almost all of them. Unfortunately, their names (many dozens) and contributions (probably hundreds) are simply too numerous to include here, but I would, however, like to recognize especially from among this group two of my earliest coworkers here at UW, Jeff George and Larry Wai, who together with me made up the original outer detector data acquisition software team. I had to collaborate closely for many months with both of them as members of that team, and the camaraderie they provided during that time, one of the most exciting, yet also most stressful phases of my entire graduate school career, was a tremendous comfort to me. Additionally, as the senior graduate student in my research group at the time when I initially joined it, Jeff George also served as my first UNIX guru, and his memory continues to serve as a role model, now that I am serving as guru for others. I also thank Hans Berns and Eric Zager, both of whom have been present during my entire career with my research group, for their consistently pleasant influence on my workday, and Steve Strausz, for suggesting to me the idea of searching for evidence of a neutrino flux in correlation with the 1997 γ-ray flaring episode of Markarian 501.

Furthermore, I would like to thank for their time all of the members who served on my committee: my advisor again, plus Gerald A. Miller, Oscar E. Vilches, James M. Mayer, Christopher W. Stubbs, and the late Kenneth K. Young. I am grateful, as well, to two fellow students, David Leichtman and Ruth Van de Water, who listened patiently to a practice version of my thesis defense talk, then offered numerous constructive criticisms which ultimately improved the quality of the finished product dramatically.

On a more personal level, most of my aunts, uncles, cousins, siblings, grandparents, friends, and their spouses or other companions have taken an interest in what I do, and just about all of them have at one point or another provided encouragement, pleasant distraction, or succor when I needed it. I know that if I were to try to write down an exhaustive list of all of the people who fall into this category, it would be a very long list indeed, and I'd surely realize with regret much later on that I'd left one or two of them out or misspelled somebody's name, so I won't even try to do it. But I want to emphasize that I have always been conscious of the things that they have provided, and that I am grateful to them.

Although I dare not even attempt to name all of the friends and relatives who have supported me outside of the office during my time in graduate school, picking out just a few of the ones from among that group who have had a particularly direct impact on my work is not too difficult: to my mind, there are three of them who stand out especially as having played either a more consistent role in my day-to-day experience at the university, or else provided a more profound influence than any of the others. The first of these, Phillip Geissbühler, my fellow student, served as my lunch companion several times per week for a substantial fraction of my years in graduate school, and he invariably managed to remind me in those days, always with good humor and bonhomie, that no matter how unfortunate my circumstances seemed to me, there was always somebody else (usually him) who had it worse. His commiseration was invaluable to me. He also proofread, in various states of revision, essentially all of this thesis, cover to cover, and for that especially I thank him.

The other two people, of course, are my parents, Ann M. and Frank E. Stachyra, who have supported me consistently for three decades. Thank you.

Andrew L. Stachyra
Seattle
February, 2002

Chapter 1

Current Status of Neutrino Astronomy

1.1 Introduction

In 1960, Greisen[1][2], Markov[3] and Zheleznykh[4] proposed the idea of using liquid-filled underground Cherenkov detectors to make directionally sensitive flux measurements of neutrino-induced upward-going muons. Among other goals, it was the intent of those authors that such devices would be used to usher in a new kind of astronomy, based on the detection of high energy neutrinos rather than of optical or radio wavelength photons, which were the two primary modes of observation in use at the time. This dissertation is about an experiment and ensuing analysis, just one among an ongoing series by a small community of researchers, which has been performed in response to that initial vision.

At the time of this writing, two astrophysical point sources of neutrinos have so far been observed: the sun[5][6][7][8][9][10] and, for the few seconds immediately following its initial explosion, supernova 1987A[11][12][13]. This thesis, however, will address itself toward neither of those two objects, concentrating instead on a search for evidence of extrasolar, non-cataclysmic sources.

Although the existence of detectably bright sources, unlike that of, for example, the Higgs boson or of cosmologically significant dark matter, is not so strongly expected that discovery would provide the missing piece to a long-standing scientific puzzle, there are nevertheless compelling reasons to suppose that a sufficiently sensitive search might one day yield results. The argument in favor of attempting either high energy neutrino or γ-ray astronomy, as it is typically presented[14][15], begins by pointing out that the cosmic ray energy spectrum varies over at least 13 orders of magnitude, from primary particles of less than 10^7 eV[16] total energy to those with more than 10^{20} eV[17]. They are known to consist primarily of positive ions, and there is some consensus[18][19][20] that supernova remnants are the most likely candidate sources of particles up to about 10^{15} eV, but there is as yet little direct evidence to confirm this hypothesis, and little is known for sure about what is powering the rest of the spectrum, except that a combination of several different mechanisms is probably needed in order to account for the full range of observations over all energies.

It is noteworthy, furthermore, that cosmic rays are extremely abundant, in the sense that they provide a substantial contribution to the energy balance of the galaxy[21][22]: at an energy density of ~ 1 eV/cm^3, they are as significant a presence in interstellar space as starlight, or the galactic magnetic field, or the cosmic microwave background radiation, all of which are estimated to have energy densities in the range 0.1 to 1 eV/cm^3. That there are unknown natural celestial objects out there accelerating cosmologically significant quantities of ions to create the cosmic rays observed on Earth is seen as an important justification for attempting to observe astrophysical neutrinos: it has long been understood that cosmic rays striking the atmosphere (and ions accelerated at man-made facilities as well, for that matter, striking a target at the end of a beam pipe, under suitable circumstances) create copious quantities of neutrinos[23] in air showers through the decay of the secondary mesons that arise from the initial interaction of the primary particle. It seems quite plausible that perhaps some fraction of the cosmic ray ions are never even reaching us, interacting instead, via the same mechanisms, in agglomerations of matter much nearer to the sources, so that the sources produce not just the cosmic ray ions but also a flux of neutrinos as well.

This line of reasoning can be fleshed out further by pausing to consider a few of the meson decays which figure most prominently in the production of cosmic ray showers on the Earth:

$$\pi^\pm \rightarrow \mu^\pm + \nu_\mu(\overline{\nu}_\mu) \tag{1.1}$$

$$\pi^0 \rightarrow \gamma + \gamma \tag{1.2}$$

$$\mu^\pm \rightarrow e^\pm + \nu_e(\overline{\nu}_e) + \overline{\nu}_\mu(\nu_\mu). \tag{1.3}$$

In practice of course, real cosmic rays generate not only pions, but also modest quantities of other mesons as well, such as kaons and so forth, so the model implied by the equations above is a bit of an oversimplification. But qualitatively, the behavior of a more realistic model including all of the relevant physics would not be all that much different: the process of meson decay ultimately results, inevitably, in large numbers, roughly comparable in overall scale, of electrons, neutrinos and γ-rays. The implication of this, therefore, is that any celestial object which emits neutrinos through the production and subsequent decay of mesons as it has been outlined above must also release roughly comparable amounts of energy as γ-rays. Optimism about finding such sources of neutrinos is further bolstered by the fact that several point sources of TeV γ-rays, the energies of which fall well within the sensitivity range of Super-Kamiokande to neutrino-induced upward-going muons, have indeed been discovered[24] with atmospheric Cherenkov telescopes in just the last few years. Of course, it's also possible that other power mechanisms which do not produce a corresponding emission of neutrinos, such as inverse Compton scattering on relativistic electrons[25], might instead be at work in these cases, but nevertheless, that any phenomenon consistent with astrophysical neutrino production has been observed at all is seen as encouraging.

The utility of both high energy neutrinos and photons as astronomical tools lies in their lack of electrical charge. The point of origin of the cosmic ray ions in general cannot be discerned because they do not travel in straight lines: they are bent by the magnetic field of interstellar galactic space, which is estimated to have an average value of around 3 μG[26], corresponding to a gyroradius of about $10^3 \times R$ cm[27], where R is the rigidity in units of volts. The height of the galactic gaseous disk is about 5×10^{20} cm, so ions with rigidities of less than $\sim 10^{17}$ V have a gyroradius smaller than the thickness of the galaxy, and probably loop around through the magnetic field many times before reaching Earth. Neutral particles, on the other hand, travel straight from their source, unaffected by those fields, and thus offer a unique investigative opportunity.

Furthermore, the neutrino's extremely small interaction cross section, the same property which makes detection so technically challenging, instead works to an experimenter's advantage in studying possible extragalactic point sources: the radiation fields between the Earth and the remotest sources are opaque[28][29][30] to γ-rays, because the γs are very likely to pair produce with photons in those fields as they travel over cosmologically vast distances. The cross section[31] for this interaction has a kinematic threshold governed by

$$E_\gamma \epsilon (1 - \cos\theta) \geq 2(m_e c^2)^2, \tag{1.4}$$

and peaks when

$$E_\gamma \epsilon \approx 2(m_e c^2)^2 \approx 0.5\ (\mathrm{MeV})^2, \tag{1.5}$$

where θ is the encounter angle between the two photons, E_γ is the energy of the γ-ray, ϵ is the energy of the soft photon in the radiation field, and m_e is the mass of an electron. Thus, the cosmic microwave background radiation, with $\epsilon \sim 2 \times 10^{-4}$ eV, has its greatest effect on γs of around 10^{15} eV, while ordinary starlight, concentrated primarily around the visible range of $\sim$2 eV, exerts its greatest effect on γs of $\sim 10^{11}$ eV. Experiments like Super-Kamiokande, which are sensitive to upward-going muons induced by neutrinos in this energy domain, have the potential to make a unique contribution, by observing astrophysical neutrinos at energies at which photons may not be able to reach the Earth.

Despite all of their advantages, neutrino detectors do also have some inherent limitations. The neutrino interaction cross section rises without bound as neutrino energy increases, and at neutrino energies of 10s to 100s of TeV, the cross section becomes large enough that travel through the diameter of the Earth starts to cause non-negligible attenuation[32]. Thus, an analysis such as this one, which identifies neutrino-induced muons by their upward-going direction, is mostly limited to neutrino energies of less than perhaps a PeV or so, except as the direction approaches the horizon, where the path length through the Earth becomes much shorter.

1.2 Favored Types of Candidate Source Objects

The current and recent theoretically oriented literature on neutrino astronomy is rich in speculation about what types of celestial objects would present the most favorable conditions for generating significant numbers of neutrinos. In the loosest view, one can justify placing on the initial, first guess list of neutrino source candidates any object for which evidence of significant non-thermal emission has already been observed, simply by appealing to the many uncertainties in the taxonomy of the astronomical zoo: the internal dynamics of most such objects are understood only conjecturally at best, and that by itself leaves a fair amount of room for the possibility of neutrino production. The goal of this section, however, will be to pare down the long initial list into a much shorter one, containing only those types of objects for which there is broad consensus among the experts that special attention is merited.

Two of the objects which will be discussed below, pulsars and the expanding shells of young supernova remnants, exist within our galaxy, and the third, active galactic nuclei (AGN), are found only in other galaxies. When considering objects within our own galaxy, it's worth noting that the galactic center, simply by virtue of its large contribution to the overall galactic mass, is particularly likely to contain neutrino emitting objects if any exist, so in addition to the types of objects listed below, it should also be considered a promising source candidate.

1.2.1 Pulsars

Pulsars are assumed to be rapidly spinning neutron stars formed from the collapsed core material remaining after a supernova explosion. They usually have strong magnetic fields, typically $\sim 10^{12}$ G at the surface at the time of formation; small radii, typically $\sim 10^6$ cm; and rotation periods that vary from ms to s, becoming progressively longer as the pulsar ages[33]. Pulsars derive their name from their characteristic emission of short, pulsed bursts of radio waves, by which they were first discovered. They are dynamically complicated objects, consisting of rotating magnetic dipoles, relativistically rotating plasmas, and sometimes in binary systems, disks of matter accreting gradually from the partner to the neutron star. A small number of different models are proposed, each of which predicts a different kind of mechanism capable of accelerating the ions in the magnetosphere into some target medium, thereby producing charged pions, the decays of which eventually result in a flux of neutrinos, as well as neutral pions, which produce a commensurate flux of γs.

In one scenario[34], the pulsar's angular velocity and magnetic moment are parallel, and the dynamics of the magnetosphere result in an electrostatic potential capable of accelerating a positively charge ion up to

$$E_{max} = 3 \times 10^{18} \frac{ZR^3B}{P} \mathrm{eV},$$

where Z is the ion's charge, R is the radius in units of 10^6 cm, B is the magnetic field at the polar surface of the star in units of 10^{12} G, and P is the period in units of ms. For a typical case like the Crab Pulsar, E_{max} is estimated to be about 2×10^{16} eV. If the pulsar's angular velocity is instead more perpendicular to the magnetic moment[35], the rotating magnetic dipole will give rise to intense, yet slowly varying induction fields which may be capable of accelerating ions to significant energies over time scales which are short compared to the characteristic ~ms period of rotation. The maximum possible energies for particles accelerated by the Crab in this scenario are estimated to be about 10^{15} to 10^{16} eV. In either of these cases, the likely target material against which the accelerated ions would collide would be the supernova nebula surrounding the pulsar, especially during the period of relatively high density immediately following the explosion[36][37]. It is worth noting that the γ-ray astronomy results in the energy range of this neutrino analysis seem to support at least the general notion that a unary (i.e., not part of a binary system) pulsar might be capable of such acceleration: three of the four known galactic point sources of TeV γ-rays[24] (Crab, Vela, and PSR1706-44) appear to be plerions, which are a specific subclass of supernova remnant that contain a pulsar at the center.

X-ray binary systems form the other major class of pulsar-powered neutrino production models. These are systems in which the neutron star exerts some kind of strong, profound influence on the atmosphere of a non-compact partner star orbiting nearby. The extreme conditions in the plasma flowing off of the companion are ideal for the production of shock fronts, which may give rise to shock acceleration[18]. In one scenario, matter accreting from the companion onto the compact object[38] (which technically, in some models[39] at least, need not be a neutron star; in this one case, a black hole will also work) could provide the driving mechanism. In another, the collision of a relativistic magnetohydrodynamic wind[40] of electron and positron plasma originating in the pulsar's magnetosphere, impinging upon the solar wind of the companion[41], may create the shock front necessary for acceleration. As with the unary systems discussed above, there is some experimental evidence supporting the hypothesis of very high energy γ-ray production in x-ray binary systems. Several experimenters[42][43][44][45] working independently in the late 1970s and early 1980s reported observation of a statistically significant excess of air showers apparently coming from the direction of Cygnus X-3. Subsequent observations by other researchers, however, failed to corroborate those results, so if the earlier results are to be given any credence at all, one must assume that such emission is only sporadic at best.

1.2.2 Young Supernova Remnants

Another object considered a likely candidate site for the production of significant neutrino flux is the expanding envelope of gaseous debris blown off by a supernova, during the short-term aftermath of the explosion. In a manner analogous to that of the binary system discussed above, it may be possible for a pulsar wind colliding with the inner envelope of the young expanding supernova shell to create a shock front capable of accelerating ions to significantly high energy, which might then interact within the volume of the shell to produce neutrinos[46]. In this scenario, the window of opportunity for point source detection would probably lie somewhere between a few and a few 10s of years, during which time the majority of neutrinos would be produced.

Given the current consensus in the field that expanding supernova nebulae are the likely production site of most of the cosmic ray ions that we see on the Earth below about 10^{15} eV, it is natural to suppose that such environments might make good candidate neutrino sources as well. Indeed, such a model[47] has been developed, but it predicts that most of the neutrino flux would be produced on a characteristic time scale of perhaps 10 to 100 days after the initial explosion, so it is not a favored production mechanism for purposes of this analysis, since no supernovae have been documented near the Earth during the time period covered by or immediately preceding the data set. A notable feature of this model is that the time scale for significant neutrino production is small compared to the time scale for significant cosmic ray acceleration: the accelerative capabilities of supernova blast waves are estimated to dissipate on a comparatively much longer time scale of roughly a few hundred[18] to a thousand[48] years or so. Intriguingly, two of the four known galactic sources of TeV γ-rays[24], Crab and SN1006, are known to be associated with supernova remnants that fall within this longer age range. On the other hand, however, certain detailed features of the observations appear to be most readily explainable by appealing to production models based on inverse Compton scattering[25][24], rather than the decay of hadronically produced pions; this undermines support for the possibility of a significant flux of neutrinos since only the latter process can produce them.

1.2.3 Active Galactic Nuclei

An active galaxy is defined as any galaxy which apparently emits a substantial fraction of its total energy via mechanisms other than the normal processes of stellar and interstellar evolution. The luminosities of many of these objects are observed to vary significantly on time scales ranging from hours to weeks to years[49], and the shortest observed variability time scale for a given object, t_v, places an upper limit r on the physical size of the energy production region, because the luminosity due to one zone could not be causally related to the luminosity from another zone unless they were physically separated by distances such that

$$r \leq c \cdot t_v,$$

i.e., r must be no greater than the distance which could be covered by moving between the two zones at the speed of light in the time t_v. This line of reasoning, combined with the other astronomical observations of these objects, leads researchers to conclude that some of them are radiating energy at rates as high as 10^4 times that of the typical ~kpc-sized galaxy in volumes (presumably the galactic nucleus) measuring much less than 1 pc^3[50].

The most widely favored explanation to account for the incredibly high luminosity in such a compact area, as well as other attributes, is that the active galaxies are powered by accretion onto supermassive black holes of masses ranging from about 10^6 to $10^{9.5}$ $M_\odot$[51][52]. As with the smaller scale accretion-powered models already discussed as possible galactic sites of neutrino generation, the flow of accreting matter in an active galaxy is also considered a likely site of standing shock fronts[53], across which hadrons might be accelerated. Building on this idea, several workers predict that a significant flux of neutrinos[54][55, 56] might be generated in the central engine of an AGN; a notable feature of the models is that pion photoproduction (e.g., $p + \gamma \rightarrow n + \pi^+$) on the thermal radiation field of the core region would contribute significantly to the overall generated flux.

In addition to the astonishing luminosities generated by AGN, another prominent feature is also evidently associated with many of them: the emission of radio waves from long, narrow lobes which are commonly understood to consist of jets of relativistically moving plasma[57] being expelled from the core region. Like the matter swirling about in the accretion disk, the relativistically flowing plasma in the jets is also considered a prime candidate environment for the pion production needed to produce neutrinos, and a model exists[58] that describes the processes involved.

The current picture in ~TeV γ-ray astronomy is that all known extragalactic sources[24] appear to be BL Lacertae objects, which are believed to be simply those AGN which happen by chance to be oriented with their relativistic jets pointed almost directly toward the Earth[59]. Thus, if the γ-ray observations support either of the two neutrino production hypotheses, (and one must bear in mind that the hadronic origin of these γs is not yet certain) one supposes that they are more favorable to the jet-based models, since these seem to include directionality more naturally as a feature than do the diffusive central engine models.

1.3 Other Experiments

A number of other experiments have either published results of searches for astrophysical neutrino sources in recent years, or else are in stages of active construction and development, and should be reporting results within the next few years. Those of greatest current interest are acknowledged here. Their most basic specifications are summarized in table 1.1.

1.3.1 KGF

The KGF experiment was performed in three distinct phases at vertical depths of 7000, and later, 6045 m.w.e. (meters of water equivalent), in the Kolar Gold Fields mine in India. Originally designed as a proton decay experiment, it was constructed of iron absorber layers

Table 1.1 These are the basic specifications for most of the neutrino telescope experiments of current or recent interest to the field. Listings are limited to experimental groups which have either published a neutrino astronomy analysis considered comparable to this one, or appear to be imminently about to do so, based on recent conference proceedings or theses. There are a few collaborations, such as SNO, which probably have the capacity to produce analyses similar to this one, but have not yet shown any evidence of significant effort on this front, presumably because their experiments were built primarily to address other physics topics; cases such as these have been omitted from the list. Dates of operation refer to total known run time, while the number of ν_μs refers to the total size of the data sample on which the published results were based; in some cases data acquisition continued after initial results were reported, but an updated official astrophysics analysis was not necessarily presented.

Detector	Location	Interval of Operation	ν_μ Effective Area (m^2)	Thresh. (GeV)	Depth (mwe)	Number of ν_μs
KGF	India	1980-1993	30	1	6500	243
Baksan	Caucasus	1980-	200	1	850	700
Fréjus	France	1984-1988	100	1.5	4900	49
IMB	Ohio	1981-1990	300	2	1600	172
Kamiokande	Japan	1983-1995	150	1.6	2700	146
MACRO	Italy	1989-1999	800	1	3700	1100
Super-Kam.	Japan	1996-2001	1000	1.6	2700	2037
Baikal NT-96	Siberia	1996-	200	10	1100	~10s
AMANDA B10	Antarctica	1997-	5000	10	2000	~100s

in phases I and II, and no absorber layers in phase III, interleaved with proportional counter tubes for particle tracking. Due to the slow timing response of the proportional tubes, reconstructed tracks suffered from antipodal degeneracy: track direction vectors could only be determined up to within an overall positive or negative sign. Because of the unusually deep location of the experiment, a very wide swath of solid angle, extending up to 30° above the horizontal in phase I (7000 m.w.e.), and 25° in phases II and III (6045 m.w.e.), was sufficiently shielded by earth so that all atmospheric muons were screened out, and any muons coming from those directions could be identified with certainty as having been produced by a neutrino interaction in the surrounding rock. Thus, even though the antipodal degeneracy meant that the experiment had no ability to distinguish between upward-going and downward-going particles, events fitted near the horizon could still be identified with certainty as neutrino-induced, since both possible track solutions lay in an angular region which could only be penetrated by neutrinos. A point source search was attempted[60], taking into account that each event could have come from either of two directions; flux limits were associated with various celestial objects, but no significant evidence of sources was discovered.

1.3.2 Baksan

The Baksan scintillation telescope is located inside a mountain in the Caucasus region of southern Russia and is constructed of 3132 tanks of liquid scintillator, each viewed by a single photomultiplier tube (PMT). The modular tanks are arranged adjacent to one another like bricks to form the six faces of an approximately cube-shaped structure, with two additional internal horizontal planes intersecting the four walls at intervals of one and two thirds of the height of the structure. Tracks are reconstructed from small numbers of collinear tanks firing in unison, and the timing response of the individual tanks is fast enough to allow the detector to distinguish between upward-going and downward-going particles. Neutrino flux limits due to various source candidates have been calculated[61].

1.3.3 Fréjus

The Fréjus detector, a fine grain tracking calorimeter that was located in a side cavern off a highway tunnel through the French alps, was one of several experiments built originally to search for evidence of proton decay. It was constructed of sets of long, narrow flash chambers which first were grouped together in parallel to form planes, and then, with plates of iron absorber interleaved periodically, the planes were sandwiched together in alternating layers, rotated 90° with respect to each other about the common normal, so that the axes of the flash tubes would be oriented orthogonally. The post-trigger recovery time of the flash chambers was very slow, on the order of seconds, so, like KGF, the reconstructed track vectors of particles crossing the entire detector were antipodally degenerate, but unfortunately the slant depth near the horizontal was not great enough to employ the KGF solution in this case, and Fréjus was therefore unable to make use of such particles in its analysis. Fully contained events, those with interaction vertices within the volume of the detector, could, however, still be identified as neutrino-induced, and the Fréjus collaboration made use of this data sample to place limits on the neutrino flux[62] from various sources.

1.3.4 IMB

IMB, the Irvine-Michigan-Brookhaven detector, built in the Morton-Thiokol salt mine in Ohio, was another experiment originally designed to look for proton decay. It was an immediate predecessor, in multiple ways, of the instrument used for this analysis: not only was it an underground water Cherenkov detector, but also, moreover, the 2048 PMTs which originally lined its walls were reused to make the outer veto shield now in operation in Super-Kamiokande. Confidence limits on point sources of extraterrestrial neutrinos have been reported[63].

1.3.5 Kamiokande

Kamiokande, the Kamioka Nucleon(Neutrino) Decay(Detection) Experiment, another underground water Cherenkov detector, was the other immediate predecessor, besides IMB, to the apparatus used in this thesis, and the two Kamioka experiments were located only a few hundred meters away from each other in the same mine. In most respects, Super-Kamiokande was simply a scaled up version of this original detector: the inner detector PMTs were of nearly the same type, made by the same company, most of the inner detector data electronics were the same, and the basic geometric features of the current detector, like its cylindrical shape and its segmentation into inner and outer subdetectors, were the same in the prototype experiment. A published neutrino astronomy analysis is available[64].

1.3.6 MACRO

MACRO, the Monopole Astrophysics and Cosmic Ray Observatory, located in the Gran Sasso underground physics laboratory in Italy, was a multipurpose detector intended to help address a number of different physics topics. The detector consisted, for the most part, of vertical and horizontal arrays of streamer tubes, interspersed with passive absorber layers of crushed rock and iron, surrounded by tanks of liquid scintillator. The streamer tubes provided fine grain particle tracking, while the scintillation detectors provided the fast timing which was necessary in order to be able to distinguish between upward-going and downward-going muons. The MACRO collaboration enjoys the distinction of having published the most thorough astrophysics analysis available from any neutrino detection experiment to date[65].

1.3.7 Baikal

The Baikal project[66] is a large water Cherenkov detector at the bottom of Lake Baikal, in Siberia, which has been undergoing sporadic development and construction since the 1980s. PMTs are attached to vertical strings of cables anchored to the lake bed and accessed for maintenance and hardware upgrades during a 5 to 7 week period in winter when the thick ice sheet that forms on the surface can be used as a staging platform. Physics results have been reported[67] for a modest sample of atmospheric neutrino-induced upward-going muon events which has been culled from 70.3 live days of raw data acquired with configuration version NT-96, a set of 96 PMTs mounted on four strings. NT-200, a set of 192 PMTs mounted on 8 strings, has been operational since 1998.

1.3.8 AMANDA

AMANDA[68] is another large Cherenkov detector, but one which uses the Antarctic polar ice cap as the detection medium rather than water, the more usual choice. This experiment has been undergoing development since the early 1990s and has been built in three stages, AMANDA-A, AMANDA-B, and AMANDA-II. The 130 live day sample of data available from AMANDA-B, a set of 289 PMTs mounted on ten strings, is the subject of current analysis efforts[69]; several hundred neutrino-induced upward-going muon events have been identified, and official, published results are expected soon.

Chapter 2

The Detector

This chapter will address the design and operation of the experimental apparatus which, over a five year period of operation between 1996 and 2001, provided the data for this analysis. During the summer of 2001, in order to perform a scheduled hardware overhaul, the 50 kton water tank, which held the interaction medium for the detector as well as all of its light-sensing photomultiplier tubes (PMTs), was drained for the first time since the initial commissioning of the experiment. In autumn, following the repair work, as the detector was being refilled with water, a sudden, violent, and spontaneous chain reaction resulted in the successive implosion of about 7800 of the PMTs, destroying somewhat more than half of the experiment's photosensitive surface area. Although there are already plans to redesign and rebuild, the unfortunate event is still so recent that the exact configuration in which a more robust future version of the project will continue to run is not yet certain. The descriptions contained in this chapter of the major detector components and how they operated together should be understood, therefore, as limited in their validity to the time period when the original version of the instrument was fully intact. In the interests of succinctness, other references will be cited for technical details wherever possible.

2.1 Cherenkov Radiation

The Super-Kamiokande experiment did not "see" neutrinos directly; rather, it did so by observing the Cherenkov radiation emitted by charged particles created by or energized through the interaction of the neutrinos in the matter in and around the tank. Specifically, the reactions of interest in this analysis were of the type

$$\nu_\mu(\overline{\nu}_\mu) + X \to \mu^-(\mu^+) + \text{hadrons},$$

where X represents, potentially, any nucleon in the rock below or adjacent to the detector, and it should be understood, furthermore, that the hadrons created in the rock did not in general penetrate appreciably into the detector; the outgoing muon was the only measurable product of the interaction. At the neutrino energies of interest to this analysis, the outgoing muons were generally very forward scattered, and were for the most part aligned with the parent neutrino arrival directions to within a few degrees. It was this tendency toward forward scattering that formed the basis of the method for distinguishing muons induced by neutrinos in the rock surrounding the detector from the principle background events in this analysis, namely, atmospheric muons, so called because of their origin in the decay of mesons induced by cosmic rays in the upper atmosphere. To be able to make this distinction, and thus screen out non-neutrino events that might otherwise have contaminated the data sample, the directions of the outgoing muons (and thus, by extension, the directions of the incoming neutrinos) in reactions like 2.1 had to be limited to those that would have required the neutrinos to travel through enough earth in order to reach the detector so that a neutrino would have been the only particle which could have penetrated that deeply. In practice, this ended up meaning that upward-going muons were identified almost exclusively as neutrino-induced, while downward-going muons were assumed to be almost exclusively atmospheric, and were presumed to be just simply those muons which happened to be energetic enough that they were able to penetrate the relatively thin (as compared to the diameter of the Earth) $\sim$1 km or so of mountain that shielded the detector from above.

Thus it was in this restricted sense that Super-Kamiokande "saw" the neutrinos associated with the upward-going muons: it recorded the intensity and timing of Cherenkov light emanating from the muons as the light splashed upon the wall-mounted PMTs, then the recorded information was used to estimate the set of muon track parameters which would most likely have generated the PMT data that were seen, and finally, the inferred muon track direction was accepted as approximately representative of the neutrino arrival direction, to within some modest uncertainty, due to the aforementioned close alignment between them. From this description it can be seen that the generation of the Cherenkov radiation is fundamental to the detection process, and because of this fundamental role, the Cherenkov mechanism itself merits some discussion, which is developed below.

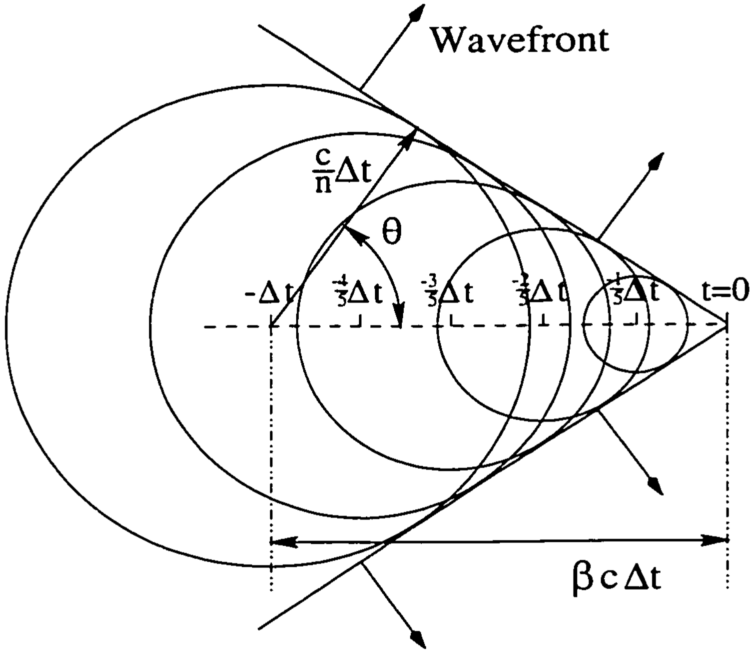

Figure 2.1 The circles in this diagram represent spherical "wavelets" of field information propagating through the medium at speed c/n, each one corresponding to a successively earlier labeled moment in time. At particle speeds above $\beta = 1/n$, the particle begins to move faster than the wavelets, and thus they pile up on each other to form a kind of shock front, which is observed experimentally as Cherenkov radiation. The relationship between θ, β, and n in equation 2.1 follows as a simple consequence of geometry.

Cherenkov radiation can only be emitted if charged particles are traveling through a transparent medium at speeds greater than that of light in the medium: c/n. The radiation is emitted symmetrically about the axis of travel of the emitting particle, in a cone-shaped pattern with an opening angle such that

$$\cos(\theta) = 1/\beta n. \tag{2.1}$$

The magnitude of the opening angle can be derived from first principles[70] by calculating the electromagnetic fields at far away points in the dielectric medium as the particle travels through it, but such a discussion is intensely mathematical and beyond the scope of this document. As an alternative to this, figure 2.1 offers a qualitative comparison between Cherenkov radiation and the shock wave created by the bow of a fast moving boat in water or a supersonic airplane. In pure water, for which $n \approx 1.33$ at the wavelengths and temperatures relevant to this experiment, a muon which penetrated any significant distance into the detector traveled all but the last meter or two with $\beta \approx 1$, and thus θ was about 41.4°, along most of its path.

Equation 2.1 ceases to define a real physical angle if β drops below $1/n$, for then $\cos(\theta) > 1$. Thus, Cherenkov radiation can only be emitted when β is above this threshold, and as particles lose energy, near the end of their trajectories they will cease to emit light, but continue to travel a little farther before stopping completely. In water, this threshold corresponds to a relativistic γ factor of 1.52; a muon of this energy will typically travel a few 10s of cm before stopping.

Besides its transmission angle, Cherenkov radiation also exhibits a few other noteworthy characteristics as well. The light is coherent, and it is also linearly polarized, along the plane containing the observation point and the particle trajectory. It emits energy per unit length

$$\frac{dE}{dx} = \frac{z^2 e^2}{c^2} \int_{n(\omega)>\beta^{-1}} \left(1 - \frac{1}{\beta^2 n^2(\omega)}\right) \omega \, d\omega. \tag{2.2}$$

For muons, or for any particle with $z = 1$ traveling through water in the $\beta \approx 1$ limit, this expression gives about 500 eV/cm in the visible range (400 to 700 nm) of light, which corresponds to roughly 200 photons/cm. To give a sense of scale to these numbers, a minimum ionizing muon in water loses about 2 MeV/cm due to ionization, so the Cherenkov light used by the detector represents only a very tiny fraction of all of the energy lost.

2.2 Physical Description

Figure 2.2 shows a cut away view of the detector that summarizes its physical layout. The basic components of the experiment were the water tank, the PMTs, their attendant electronics, the data acquisition computers and software, and the clean air and water systems. All of these will be described briefly in the subsections which follow.

The detector was located inside of Mt. Ikenoyama in Gifu Prefecture, Japan, at about $137°18^{\mathrm{m}}37^{\mathrm{s}}$ east (WGS 84) longitude and $36°25^{\mathrm{m}}33^{\mathrm{s}}$ north latitude, in a mine owned by the Kamioka Mining and Smelting Company. The underground location offered a fairly uniform shield of about 1 km of rock in all directions above the horizon, which was equivalent to about 2700 m of water. The flux of cosmic ray particles at sea level[71] from all 2π steradians

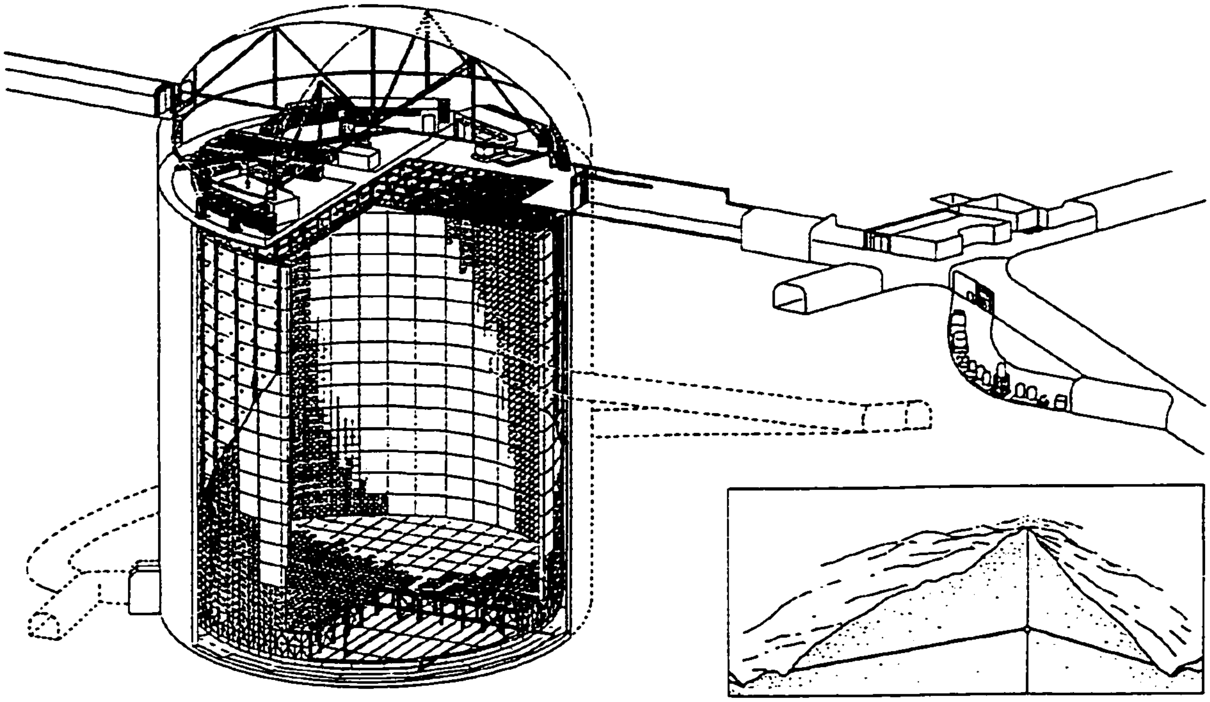

Figure 2.2 The basic features of the detector location and architecture are shown here: it was under a mountain, the water tank and photomultiplier tubes were arranged in a concentric cylindrical geometry, a large dome was hollowed out of the rock in the space above the detector and held most of the calibration equipment and fast readout electronics, and a side cavern housed a control room.

was about 2×10^2 m^{-2} s^{-1}, and the downward-going muon rate in the detector was about 2 Hz over a directionally averaged effective area of about 10^3 m^2, thus the rock shielding afforded a background reduction factor of about 10^5, without which the experiment would not have been technologically feasible.

2.2.1 The Tank and Photomultiplier Tubes

The water tank measured 41.4 m in height by 39.3 m in diameter and held about 50,000 metric tons of water, supplied by an underground spring from within the mine. It was sheathed in welded stainless steel plates, backed by a layer of concrete to smooth out the uneven, excavated rock walls of the surrounding cavern. It was divided by a 50 cm thick layer of scaffolding into an inner detector, or ID, which measured 38.2 m in height by 35.8 m in diameter, and an outer detector, or OD, which formed a cylindrical shell around the scaffolding of 2.1 to 2.2 m in thickness.

2.2.1.1 Inner Detector

There were 11,146 ID PMTs (Hamamatsu model number R3600-02), each measuring 50.8 cm in diameter, mounted on this scaffolding mostly in racks of twelve, as shown in figure 2.3, providing the ID with 40% photocathode coverage. The areas between the PMT faces were covered with black polyethylene sheets to minimize reflection of any light which was not absorbed on a photocathode, and to help create an optical barrier between the ID and OD. Besides their unprecedented diameter, the PMTs themselves, which were an improved[72] version of a model originally developed for the Kamiokande[73] experiment, had two other noteworthy characteristics that distinguished them from ordinary, generic tubes: a special bialkali (Sb-K-Cs) photocathode chosen especially for its high sensitivity to the Cherenkov spectrum, and carefully tuned, high-gain initial amplification stages. Paying special attention to the design of the initial amplification stages was important for good performance at the lowest light levels, because the amplification mechanism of a PMT involves accelerating a cascade of electrons through a series of dynodes held at ever-higher voltage levels, and the statistical variability of the cascade at the earliest stages, where the numbers of electrons are smallest, profoundly influences the charge resolution at the anode, and thereby, the energy resolution of the experiment. This, however, was a point mostly relevant to solar neutrino analyses, and not to this one, since the upward-going muons generally deposit large amounts of energy in the detector, and only the minimum threshold energy of the sample is known anyway, rather than the total energy on an event-by-event basis. A summary of ID PMT characteristics is available in reference [74]. Just outside of the ID, the PMT support scaffolding was surrounded by Helmholtz coils, which canceled enough of the Earth's 450 mG geomagnetic field to bring the ambient fields in the ID PMTs within the operating specifications of 100 mG.

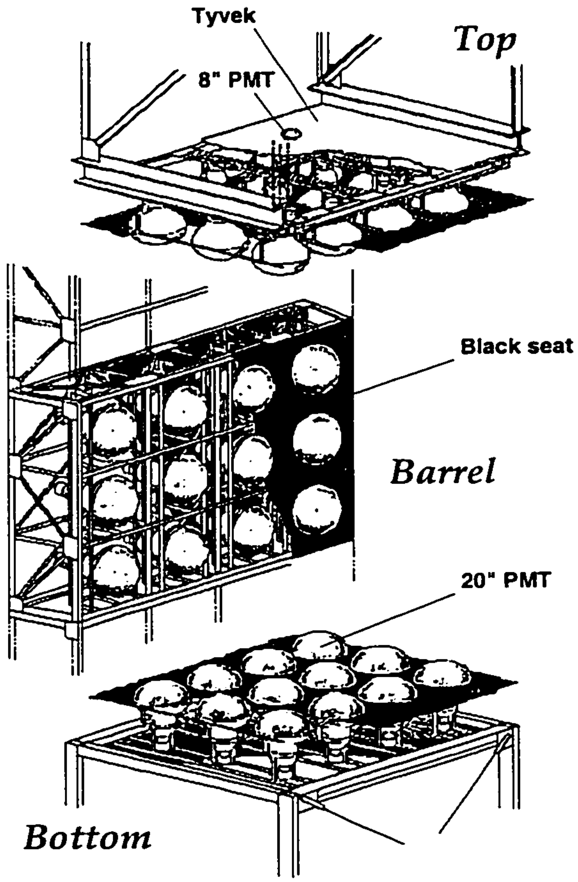

Figure 2.3 To construct the detector, the PMTs were mounted on modular stainless steel frames, ID tubes in groups of twelve on one side and OD tubes in groups of two on the other. Then the modules were, in turn, bolted together to form a continuous support structure.

2.2.1.2 Outer Detector

On the opposite side of each rack of twelve ID PMTs were two outward-facing OD PMTs, each measuring 20 cm in diameter, mounted on a line along the diagonal. Since there were in total only 1885 of these, and they were much smaller than the ID PMTs, they provided only about 1% photocathode coverage of the OD inner wall (i.e., the outer surface of the support scaffolding), which was itself less than half of the surface area of the OD, since there was an uninstrumented outer wall as well. The OD PMTs, which were recycled from

IMB, a previous generation water Cherenkov detection experiment, were much older than the ID PMTs and were therefore expected even at the time the experiment was designed to suffer a much higher failure rate. Since the OD was crucial to some of the analyses as a veto shield, it was important to ensure as much as possible that the clusters of dead PMTs which were expected to develop randomly did not lead in practice to holes through which particles could slip from outside of the detector; this concern, coupled with the comparatively sparse coverage, made it necessary to resort to additional measures to increase the light collection efficiency of the OD, so that the remaining neighboring PMTs would be able to compensate, as much possible, for the dead ones. One of the strategies adopted was to cover all of the walls of the OD with white Tyvek®, a kind of thin fibrous polyethylene sheeting with about 80% reflectivity to photons in the Cherenkov wavelength range; this strategy degraded position resolution of course, but at the same time also ensured that at no position in the OD did particle detection depend on the reliable operation of only one or two PMTs. The Tyvek® served, additionally, to help separate the ID and OD optically from one another. Besides the use of reflective Tyvek®, the other light collection strategy adopted was to augment the collection capability of the OD PMTs themselves by optically coupling them to 60 cm × 60 cm wavelength shifting plates[75], which were transparent plastic panels doped with a compound capable of absorbing Cherenkov light and re-emitting it isotropically at a slightly longer wavelength. An appreciable fraction of the light eventually made its way to the photocathode via total internal reflection, effectively doubling the OD photocathode coverage, at a small expense in timing resolution.

2.2.1.3 On Top of the Tank

On top of the tank, in a 15 m tall domed cavern, there were five electronics huts. Four of them, the quadrant huts, held power supplies and readout electronics for all of the PMTs, and the fifth, the central hut, held mainly trigger electronics plus some additional OD-related readout hardware. The top of the tank was also used as a staging area for myriad calibration exercises, and served as a storage area for all of the attendant equipment. The walls and ceiling of the cavern were sealed with a thick polyurethane resin which was impervious to the radon that emanated constantly from the surrounding rock, and would have compromised the solar neutrino analysis if it had been allowed to contaminate the experiment.

2.2.2 Data Acquisition Electronics

Figures 2.4 and 2.5 show block diagrams of the data flow through the ID and OD electronics, respectively, and the first few stages of real-time computer processing. The next few subsections will explain the function of each of the components in the diagrams.

2.2.2.1 Inner Detector

When a PMT fired, it generated a transient signal of mV magnitude lasting for nanosecond time scales, which traveled down 70 m of cable to arrive at a custom-designed circuit board called an analog timing module (ATM)[76]. The ATM had several distinct parts: a 0.2 pC

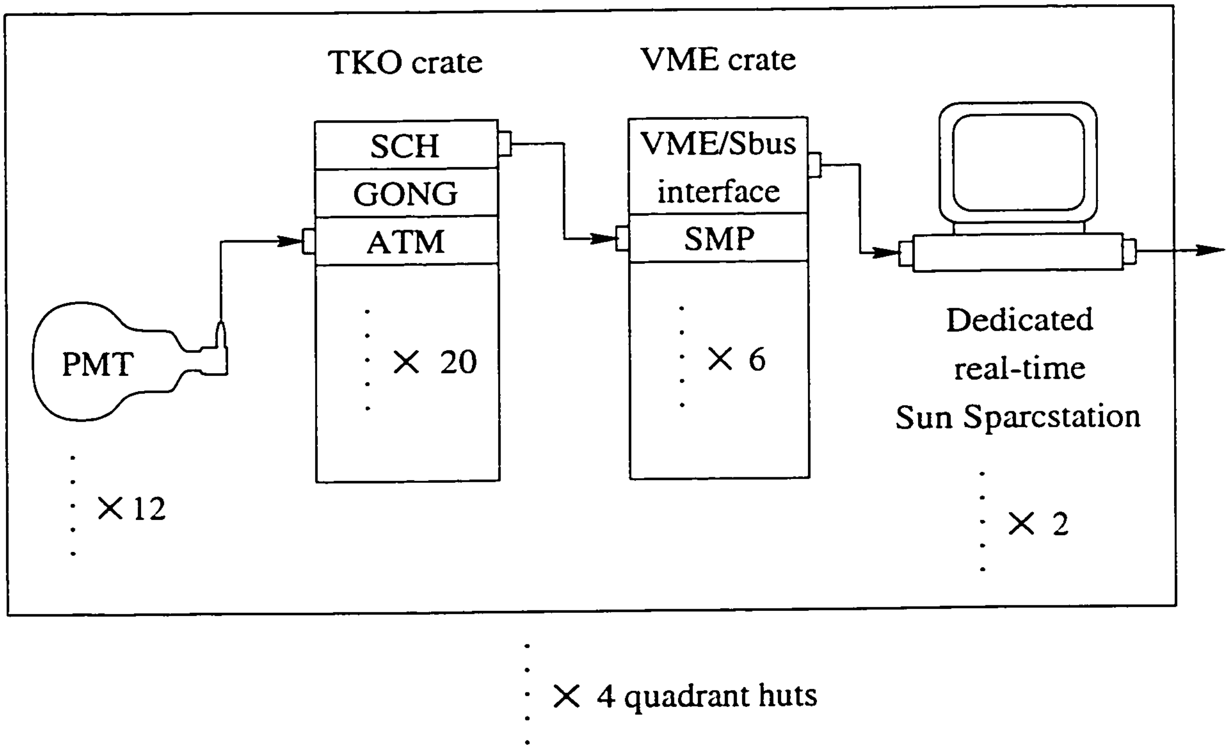

Figure 2.4 The major components of the inner detector electronics, with arrows showing path of flow of the data, are shown here.

resolution charge-to-analog converter (QAC), which integrated the current of the PMT output pulse; a 0.3 ns resolution time-to-analog converter (TAC), which registered the relative arrival time of the pulse against that of the detector trigger signal; and a discriminator, which generated a 200 ns logic pulse that was summed with those of all the other 11,145 ID PMTs to monitor their overall activity level, and trigger the detector when necessary. The ATM had one other noteworthy feature: there were two QACs and TACs for each input channel, in order to be able to register two events occurring in rapid succession, (for example, a muon which stopped in the detector and generated a decay electron) as compared to the μs time scales on which the digital readout occurred.

The ATMs were housed physically in Tristan/KEK online (TKO)[77] standard electronics crates, 20 modules per crate, with 12 channels per module, and there were two other modules included in each crate which aided in the orderly transfer of data to the next stage of processing. One of them, the go/no-go (GONG) module, accepted detector trigger signals, generated by a suite of electronics modules in the central hut, and distributed them to the ATM modules via the TKO crate's backplane. This module also latched the unique event number associated with each trigger, with which the data from each event had to be tagged in order to permit the data from all the different parts of the detector to be assembled into a single unit further down the chain of flow. The other module, the super control header (SCH), was used to read out the events stored in ATM memory, and send them to the next stage of processing.

From each SCH, data were transferred to a super memory partner (SMP) module, which plugged into a VME crate. There were 6 SMP modules per VME crate, and 2 VME crates per quadrant hut. The SMPs in each crate were read out by an Sbus/VME interface board, which transferred the data to a real-time data acquisition computer. There was one such computer, a Sun Sparcstation, specially devoted to each VME crate.

2.2.2.2 Outer Detector

The OD data flow began at the paddle cards, which were both power distribution and signal extraction modules for the PMTs. Each paddle card powered 12 PMTs at a single high

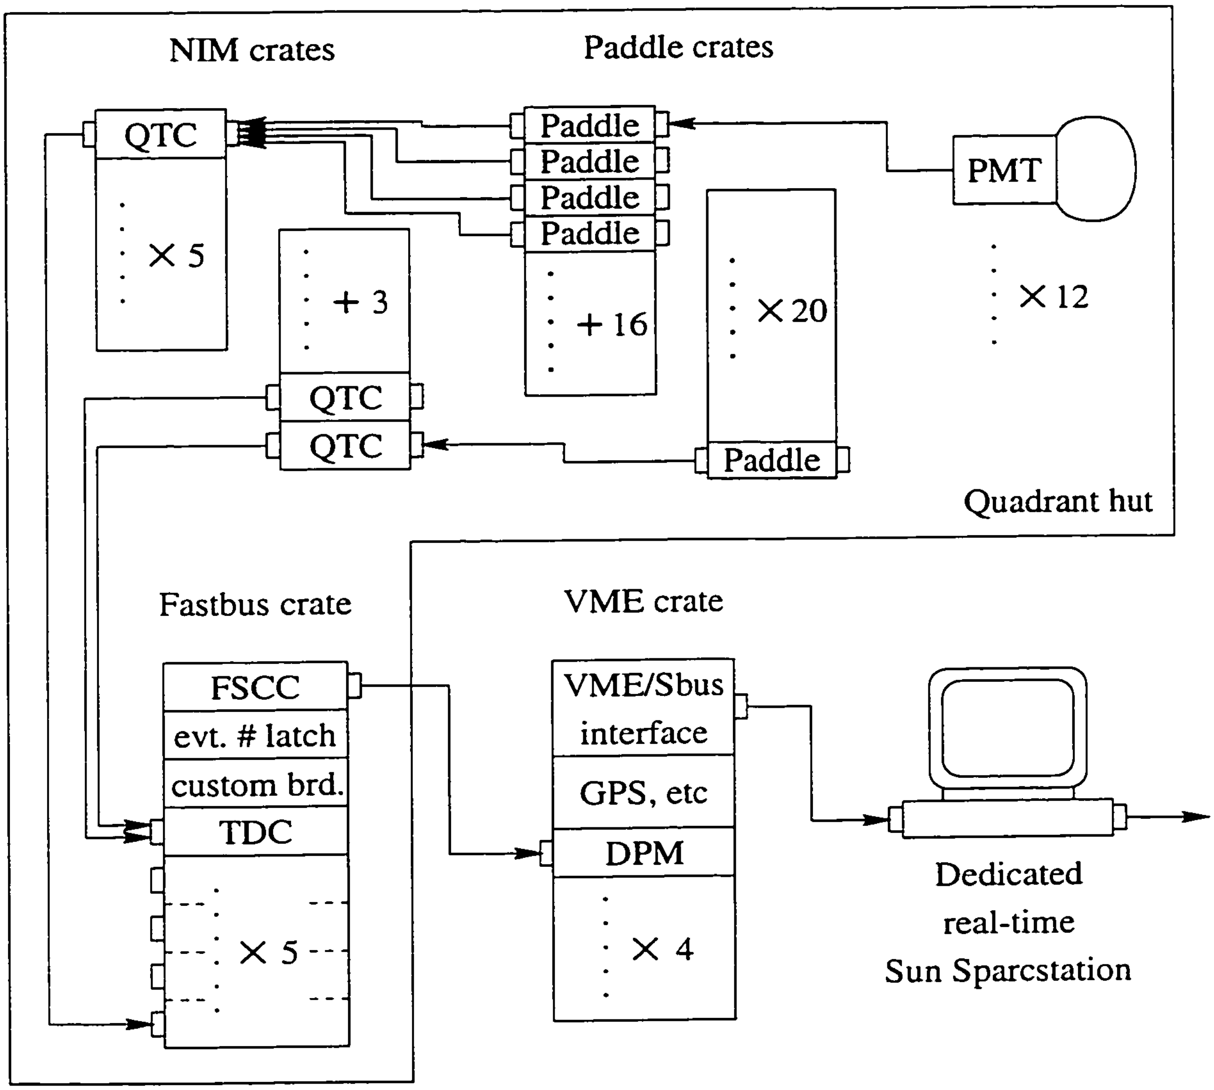

Figure 2.5 The major components of the outer detector electronics, with arrows showing path of flow of the data, are shown here.

voltage setting, and the transient electrical pulse that was emitted by the anode of the PMT when it fired traveled back along the power supply line for that PMT to the card, through a capacitor, and on to the custom-built charge-to-time converter (QTC)[78] modules. The job of the QTC modules was to transform the raw PMT signals into differential ECL signals, so that they were readable by the time-to-digital converters (TDCs), while preserving both the relative timing and brightness information that the PMT measured. The QTC modules did this by producing the leading edge of the output pulse at a fixed time after the PMT input was received, so that the relative timing of all the channels was preserved, and producing the falling edge after some delay which was proportional to the integrated current (and thus the total collected charge at the PMT anode) of the signal. The QTCs also produced a 200 ns logic pulse, which was summed over all channels in the QTC, and eventually over all 1885 channels in the OD, to be used for a trigger.

The TDCs were fastbus modules, and operated within a fastbus crate. They had 96 channels each, and five TDCs were required per quadrant hut. The fastbus crates held three other modules as well: an event number latch, a fastbus smart crate controller (FSCC) which read out the signals recorded by the TDCs, and a custom-built interface board, which was used mainly to relay control signals from the dedicated OD real-time data acquisition computer.

From each of the four FSCCs, the TDC information was transferred, via a two module combination called the DM-115/DC-2, to a pair of VME memory boards in the central hut, which swapped back and forth between accepting data and being read out. The OD VME crate also contained several modules for incorporating the absolute times of the events into the data stream, via 2 km of fiber optic filament, from a GPS receiver located outside of the mine. The absolute event times were important because they were expected to be quantities of considerable interest if the experiment ever detected a supernova.

2.2.2.3 Triggers

There were only two ways that the detector could generate its own trigger, without external prompting from some outside source such as a calibration apparatus: either the number of ID PMTs firing in some time interval had to exceed some threshold, or the number of OD PMTs firing had to do so. As has already been mentioned, discriminators in the ATMs and the QTCs both produced 200 ns logic pulses. The outputs of all of the ID and OD discriminators were joined together into two global sums, which were monitored constantly by more discriminators. The ID trigger logic was slightly complicated: it had three different threshold settings, which corresponded to progressively greater numbers of PMTs firing nearly simultaneously. The triggers they produced were, of course, functionally indistinguishable from one another, as well as from those generated via any other method, as far as the detector hardware was concerned, but the number of thresholds exceeded was recorded in the data as a bit flag in a special variable which encoded what part of the system (ID, OD, calibration apparatus, etc.) initiated the trigger. This special data word in the event header could be used as a quick and simple way to distinguish between higher and lower energy events, so that data reduction programs could make rapid decisions about how to treat each event without wasting a lot of CPU time. Only the highest ID trigger threshold, which corresponded to 31 PMTs firing within 200 ns of each other (i.e., a simultaneous overlap, at some instant in time, of 31 of the 200 ns ATM output logic pulses), was of interest to this analysis; the lower thresholds were used mostly by investigators performing solar neutrino analyses. The OD trigger logic had only one threshold, corresponding to 19 PMTs firing together in the time window.

2.2.3 Real-Time Computing

The real-time computers at the experimental site had two primary tasks: sorting the data which were streaming in asynchronously from the many different detector subcomponents into an ordered series of events, and re-organizing the data from the format used by the online software into the format recognized by much of the analysis software. In addition, a computer outside of the mine, in a nearby office building, performed two auxiliary tasks in real time: translating the raw ADC counts of the data acquisition system into standard units like photoelectrons (p.e.) and ns, and running data reduction programs on this fully calibrated data set.

Before		After
Unit 1:		Event 1:
Event 1		Unit 1
Event 2		Unit 2
⋮		⋮
Event n		Unit m
Unit 2:		Event 2:
Event 1		Unit 1
⋮	$\xRightarrow{\text{sorting}}$	⋮
Event n		Unit m
Unit 3:		Event 3:
⋮		⋮
Unit m:		Event n:
Event 1		Unit 1
⋮		⋮
Event n		Unit m

Figure 2.6 This illustrates the sorting problem that the real-time computers addressed: the data, in small batches of ordered events, streamed in from various detector subcomponents, (ATMs, TDCs, etc) labeled simply as "units", and had to be re-ordered so that all the data from the first event came first, all the data from the second event came second, etc.

The first of these problems, sorting the data into a series of events, is common to all large, multiple-channel physics experiments, and is illustrated in figure 2.6: the data were read into the data acquisition computer, ordered at the primary level by detector unit and sub-ordered within the detector unit by event, and had to be re-ordered by event first, and unit second. The detector unit represented different things in different situations: for the eight computers interfaced to the ID VME crates, it was an ATM, for the OD data acquisition computer, it was a TDC, and since the VME interfaced workstations themselves eventually funneled their pre-sorted data into a main host computer for final sorting and event assembly, they too formed a set of detector units in the above sense.

After all of the data from a batch of events were completely sorted by the main host, it was sent on to another machine to be recast into an alternative internal format, called ZEBRA[79]. The need for the different formats arose because the real-time computing software was mostly C-language based, in order to be compatible with various vendor-supplied software libraries, while most of the off-line data analysis software was Fortran-77 based, because of a desire to be able easily to take advantage of pre-existing libraries of analysis software, developed for use on earlier experiments, and also because a sizable minority of senior collaborators were not adept at writing analysis programs in any other way. The ZEBRA system made it possible to create and manipulate efficiently the complicated dynamic data structures, such as linked lists, which arose naturally in a large multi-channel physics experiment, but which had no analogue in Fortran-77.

From the reformatting computer, the data were sent out of the mine via fiber-optic link to another computer which applied calibration constants from periodically updated tables to translate the raw ADC counts of the ATMs and TDCs into real numbers of photoelectrons and nanoseconds. In addition, many of the analysis working groups had data reduction processes that worked on these data while they resided temporarily on hard disk. After these processes finished, the calibrated times and charges were deleted in order to save on file space, and the raw ADC counts for each event were archived to tape, where they could be retrieved later and expanded again into calibrated events for re-analysis, if desired.

2.2.4 Clean Air and Water

The experimental facility included systems to supply both clean air and water. The purpose of the clean air system was primarily to limit the concentration of radon in the detector, because radon β-decay events were a major source of background in the solar neutrino analysis. Its presence had little impact, however, upon this analysis, and it will not be discussed further here.

The presence of the clean water system, on the other hand, did appreciably affect this analysis, for its purpose was not only to reduce the levels of radon in the tank, but also to maximize the water transparency. The water purification system had nine main components, presented here in order of flow:

1 μm Filter: Removed large dust particles.

Heat Exchanger: Cooled the water to 14° C, in order to suppress bacteria growth. The PMTs and the pumps which circulated the water through the purification system generated enough heat that they could have raised the temperature of the tank significantly over time, if this hadn't been in place.

Ion Exchanger: Removed metal ions.

UV Sterilizer: Killed bacteria.

Vacuum De-gasifier: Removed dissolved gas, including 99% of the oxygen and 96% of the radon.

Cartridge Polisher: Acted as a second stage, high performance ion exchanger.

Ultra Filter: Removed particles down to ~10 nm in size, but in the course of normal operation diverted about 10% of the input water away from the primary output. The remaining 90% was considered clean at this point, and was returned to the tank.

Buffer Tank: Temporarily held the water diverted by the ultra filter.

Reverse Osmosis Membrane: Performed additional filtration on the water in the buffer tank, then sent it back to the heat exchanger.

The system processed about 50 tons/hour, and regularly maintained the mean attenuation length for the relevant light wavelengths at or above 70 m.

2.3 Calibration

In order to relate the raw ADC counts in the ATMs and TDCs to an absolute energy deposition and direction of particle travel in the water, there was a series of three distinct relationships, between a series of four different types of quantities, that had to be parameterized each in terms of the next earlier one up the chain of reactions in the detector. Working backwards up the chain, starting with the raw ADC counts, one first had to relate those values to the actual arrival times and light intensities at the photocathodes; this relationship was called the timing and gain calibration of the PMTs. The next job was to relate the reduced quantity of light that actually reached the PMTs to the quantity of light originally generated by the Cherenkov emitting particle; this was called measuring the light's attenuation length. And finally, there was the task of relating the amount of Cherenkov light emitted to the absolute energy deposition (and thus making a quantitative statement about the absolute energy scale) of the particle; this relationship was called the absolute energy calibration.

The PMT calibrations and the water transparency measurements were equally important to all analyses, since they affected how particle interactions in the detector were reconstructed. It may be argued that the energy calibration, on the other hand, was not as crucial to this analysis as it was to many of the others, because it was not as usable: unlike the energies of most other categories of neutrino-induced events acquired by the detector, the energies of the upward-going muons weren't known on an event-by-event basis, for in general, the full energy of an outgoing lepton exiting an interaction vertex could be estimated only if the entire path of that lepton were contained within the detector. Since upward-going muons generally deposited a large and randomly variable fraction of their energy in the rock outside of the detector, the kinds of event-by-event spectral analyses that were typically applied to the fully contained atmospheric and solar neutrino data sets, which were the main point of energy calibration in the first place, were impossible to perform on the upward-going muons, thus the energy calibration would seem to be of lesser value to this analysis than to the others. Nevertheless, despite its diminished role in this particular case, the energy calibration was not completely without significance: it was intimately connected with the minimum energy threshold requirement, which was used to veto below-threshold candidate events from the final data sample, and which furthermore must of course be well-defined in order to be able to compare the results of this experiment to those of any others. Ultimately, all of the many, many different kinds of calibration exercises that were carried out had some influence on the final results of this analysis, and so a synopsis of those calibration activities is included below.

2.3.1 Photomultiplier Tubes

2.3.1.1 Relative Gain

The amount of amplification, or gain, provided by the PMTs was a function of the voltage applied to the dynodes, but the exact gain provided by a given voltage varied by a few percent from one PMT to the next, due to minute variations in their construction. It was desirable to have the voltages of all the PMTs set to produce about the same gain in each, therefore, even before any water was added to the Super-Kamiokande tanks, one of the first calibration exercises attempted was to shine a light at various positions in the tank, at various PMT voltages, and choose the voltage for each tube which best matched its performance to its neighbors. The light for this type of calibration exercise, which has been repeated regularly, even after the water was added to the tank, came from a Xenon flash lamp, piped through several optical elements to a scintillator ball, which absorbed and re-emitted the light from the lamp at a wavelength of around 450 nm. The system was equipped with a monitor, in order to renormalize the variation in light intensity from one flash to the next. The calibration analysis made appropriate corrections to account for this shot-to-shot variation in intensity, as well as for the distance between the scintillator ball and each PMT, the acceptance of each, and the anisotropy of the light emission from the ball, which was measured externally.

2.3.1.2 Absolute Gain

The absolute gain of the PMTs, defined as the relationship between the number of electrons emitted at the photocathodes and the final amounts of charge integrated in the ATMs, was determined by creating very dim calibration events that didn't produce enough light to fire most of the PMTs most of the time, and then looking at the behavior of the channels that did react, since those reactions were virtually guaranteed to have been produced by single photons. The apparatus used for this type of calibration study was a module containing a ^{252}Cf source surrounded by Ni wire wool. Neutrons emitted by the decaying ^{252}Cf were thermalized in the surrounding water, and some of them drifted back into the module to be absorbed by the Ni nuclei, forming excited states which decayed by emitting 6 to 9 Mev γ-rays. These γ-rays cascaded electromagnetically in the detector, producing electrons which emitted the actual Cherenkov light.

As a part of the gain calibration process, pedestal data, for which the ATMs were disconnected from their respective PMTs and triggered on null inputs, were also taken several times per day. The end result was that the mean charge produced by single p.e. hits on the PMTs was 2.055 pC, corresponding to a gain of around 1.3×10^7.

2.3.1.3 Timing

The leading edge arrival time of a PMT output pulse as it was registered by the ATM was influenced primarily by two things: the length of the signal cable, and the size of the output pulse. Longer signal cables, of course, resulted in longer travel times between the PMTs and ATMs, while the leading edges of the larger output pulses tended to rise above the discriminator thresholds a little earlier than the smaller ones, causing them to register a little earlier in the ATMs. The light for the timing calibration apparatus was provided by a separate system[80], powered by a laser, which could pipe light through several different fiber optic filaments into a network of diffuser balls located throughout the detector. The results of the gain calibration were coupled with the timing calibration analysis to make a time vs. charge, or "TQ" map, which was used on the real data events to correct the relative timing of each PMT according to the size of the output signal received from it.

2.3.2 Attenuation Length

The attenuation length of the light in the water was measured in two different ways: by viewing a light placed in the water at various distances from a CCD camera, and by comparing the relative intensity of the light arriving at PMTs at various distances from its known production point along the well-defined tracks of carefully selected through-going cosmic ray muons. The attenuation length measured by each method varied over time, but generally ranged between 90 and 120 m over much of the lifetime of the experiment, and the results of the two independent methods generally agreed with each other to within a few percent.

The first method used a tunable nitrogen/dye laser system to produce light at several wavelengths from 337 nm to 580 nm, which was piped through fiber optic filaments to a diffusing ball. The relative intensity was normalized from one light burst to the next with a monitoring system, and the diffusing ball was placed at several positions ranging from 5 to 30 m from the CCD camera, which was positioned near the top of the detector. As the ball was placed farther and farther from the camera, the light reaching the camera of course became dimmer and dimmer, and after taking into account the effect of the camera's decreasing solid angle, any remaining reduction in intensity was attributed to attenuation.

The second method[81] proceeded by using the tank entry and exit point fit information for a large sample of through-going cosmic ray muons to define a path of travel along which all of the Cherenkov light that reached the PMTs had to have been produced. Each muon was assumed to produce Cherenkov light at a steady angle and rate, per unit length of travel, along its entire trajectory through the detector, and the observed intensity at each PMT was corrected to account for its solid angle, with respect to the production point, due both to its distance and to its acceptance at the angle of incidence for which the light was assumed to have arrived at the PMT face. After accounting for the solid angle correction, it was again found that the PMTs suffered an additional intensity reduction, due apparently to attenuation because the reduction factor depended exponentially on distance; following standard practice, the attenuation length was taken as the value of the characteristic exponential decay constant.

2.3.3 Energy

The ultimate goal of the final, absolute energy calibration was to be able to estimate, for any charged particle stopping in the detector, solely on the basis of how much Cherenkov light was emitted (including corrections for loss due to attenuation, the quantum efficiency of the PMTs, etc.), the total amount of energy it deposited via all loss mechanisms. All of the calibration methods relied on special classes of events about which, for one reason or another, extra energy-related information was known or could be inferred. In some cases, the actual energy of the individual events could be inferred directly, on an event-by-event basis; in others, only the shape of the overall spectrum was known, and the best fit calibration was chosen as the one which most closely matched the reconstructed energy distribution of the real data to the theoretical prediction. Despite the added complexity it introduced, reliance on such a variety of different techniques was desirable, because in all cases, Monte Carlo data were used as a tool for checking the self-consistency of the entire system, and the availability of so many semi-independent results that had to be correctly reproduced helped to ensure the integrity of the conclusions.

It is worth noting that, for this analysis, the details of the energy calibration really only affected which of the stopping muons made the cut into the data sample. For the through-going muons, it was enough simply to specify a particle track length (in this analysis, 7 m) in the detector, and that corresponded to some well-defined minimum particle energy threshold (in this analysis, about 1.6 GeV). It was only with the stopping muons, for which the fitting algorithms did not produce reliable track lengths, that one had to resort to estimating the total energy deposition of the event from the amount of light gathered, and determine whether or not it exceeded the threshold.

2.3.3.1 Penetration Range of Stopping Muons

The most important energy calibration events in the upward-going muon analysis were the downward-going cosmic ray induced muons that stopped in the detector. Specifically, the useful events were those which, after stopping, emitted a Michel electron as they decayed via the $\mu^+(\mu^-) \rightarrow e^+(e^-) + \nu_e(\overline{\nu}_e) + \overline{\nu}_\mu(\nu_\mu)$ reaction. Many negatively charged stopping muons did not do this, of course, because generally they were captured into the atomic orbitals of the hydrogen and oxygen in the water as they came to a stop, and as a result of this their wave functions overlapped significantly with the nuclei of their host atoms, enabling the $\mu^- + p \rightarrow n + \nu_\mu$ reaction to compete significantly against free decay.

But, for those muons which did decay freely rather than via capture, the production vertices of the decay electrons could be estimated to within about 50 cm, using a specialized fitting algorithm developed just for analyzing such short, relatively dim tracks, and as a consequence, the total range of the parent muons could be estimated, even though they ceased to emit Cherenkov light after their speed dropped below the production threshold. Since the range of penetration into the tank was correlated with the total energy of the muon at the entry point, and since the attenuation corrected total light output was correlated with the range, an empirical relationship could be mapped out between Cherenkov light output and particle energy. It was this relationship that provided the detector calibration: it permitted an inference about particle range, and thus about total energy deposition in the detector, even in cases in which no Michel electron was observed. The calibration was checked for consistency with Monte Carlo events and was estimated to be accurate to within better than 3%.

2.3.3.2 Other Methods

Due to the similarity between the data samples used in the two endeavors, the calibration method described above was the one of primary interest to this analysis. Several additional methods were also used, however, by the other analysis groups, and they are briefly acknowledged here.

The first of these used stopping muons so near the end of their range that they penetrated only from 10s of cm to ~ 1 m into the detector. The velocities of these muons (and hence their initial energies also, since the particle types and masses were not in question) were estimated by measuring the Cherenkov opening angle at the entry vertex, which was somewhat less than the usual saturation value of $\sim 41.4^\circ$. The estimate of initial muon energy at the point of entry that was derived via this technique was compared to the amount of Cherenkov light actually generated, and the statistical distributions of these quantities were further checked against a set of Monte Carlo events to verify understanding of the detector's behavior.

Another calibration method, also based on stopping muons, involved measuring the spectral distribution of the Michel electrons, which is readily calculable from V-A theory [71] for unbound $\mu^{\pm}$, and with some corrections, is also calculable[82] for bound μ^{-}. Unlike the two stopping muon event types discussed above, the penetration ranges and/or opening angles of the decay electrons were difficult to measure, and therefore these events did not offer any separate information that could be used to make an independent estimate of energy on an event-by-event basis. Instead, this was one of the calibration methods that was done on a statistical basis: putative electron momenta were reconstructed from the total Cherenkov light deposited in the detector, and the correct calibration was derived by comparing the shape of the resulting distribution as a whole to that of the theoretical prediction.

A third method made use of the invariant mass of π^0s created via $\nu + N \rightarrow \nu + N + \pi^0$. A π^0 immediately decays to two γ particles, and the recoiling nucleus does not generally get above the Cherenkov threshold, so these events manifested themselves as two rings, one from the electromagnetic cascade of each γ. In the rest frame of the π^0, these two rings are generated back to back with identical energies which sum to 135 MeV, the rest energy of the π^0, which has been very accurately measured by many other experiments[71]. It is this fact that provided the basis for this calibration exercise: when the detector calibration was done correctly, one could Lorentz boost all two ring events to their CMS frame and one found that the remaining energy of the two particle system (i.e., the invariant mass), for an anomalously large number of particles, was clustered near 135 MeV. This spectral peak provided a convenient benchmark against which the Super-Kamiokande calibration could be checked.

The calibration methods described above, taken together, gauged detector performance in energy ranges from 10s of MeV up to a few GeV, which was certainly more than sufficient for this analysis, since the goal was merely to discern whether candidate events for the upward-going stopping muon sample lay above or below the 1.6 GeV energy threshold. Although further description is beyond the scope of this document, it's worth noting that several additional techniques extended the calibration another order of magnitude lower, into the energy range of the solar neutrino analysis. The experimental facility included a linear accelerator[83] capable of injecting single, monoenergetic electrons into the tank at several different positions. The β decay spectrum of ^{16}N, created from ^{16}O in the water interacting with neutrons produced by a commercial generating apparatus, was used[84] in much the same way as the Michel electron spectrum described above. And finally, the spectrum of the Ni/^{252}Cf source, although used primarily for the absolute gain calibration of the PMTs, was also applied to this problem as well.

In summary, the absolute energy calibration of the detector was performed using many different techniques, and the mutual self-consistency of all the results was assured by comparing calibration data from each method to the predictions of a single detector Monte Carlo. Furthermore, that comparison suggested that the systematic uncertainty in the energy calibration, in the energy range of interest to this analysis, was less than 3%.

Chapter 3

Merged Data Reduction

In the early stages of the experiment, in order to catch mistakes, all data analysis on upward-going muons was performed in parallel by two completely independent groups of researchers who convened at intervals of several months to compare results. However, because of the strain that this practice placed upon human resources, the two groups moved to combine efforts as soon as a consensus was reached that they had arrived at substantially the same conclusions about the data. One of the most time-consuming tasks for both groups was data reduction: upward-going muons were observed at Super-Kamiokande at a rate of approximately 1.4 per day but had to be separated from a background of 2×10^5 downward-going muons per day. Computer-automated direction-fitting algorithms were employed with tremendous success by both groups to recognize downward-going events and cut them away by about 3 orders of magnitude, but approximately 200 events per day still remained which had to be examined with visualization software and culled from the data sample using the pattern recognition ability of the human mind. Fortunately, the different computerized fitting algorithms employed by the two groups were characterized by different strengths and weaknesses; each algorithm failed to eliminate a different subset of background events. Therefore it was possible to reduce the number of ambiguous events (and thus the pattern recognition workload of the human researchers) further to about 75 per day by joining the forces of all the fitting algorithms from both analysis groups into one merged data reduction program.

3.1 General Scheme

The general strategy of the merged data reduction program, summarized in figures 3.1, 3.2, and 3.3, was to select events for which Q_{tot}, the total charge collected from the surfaces of the

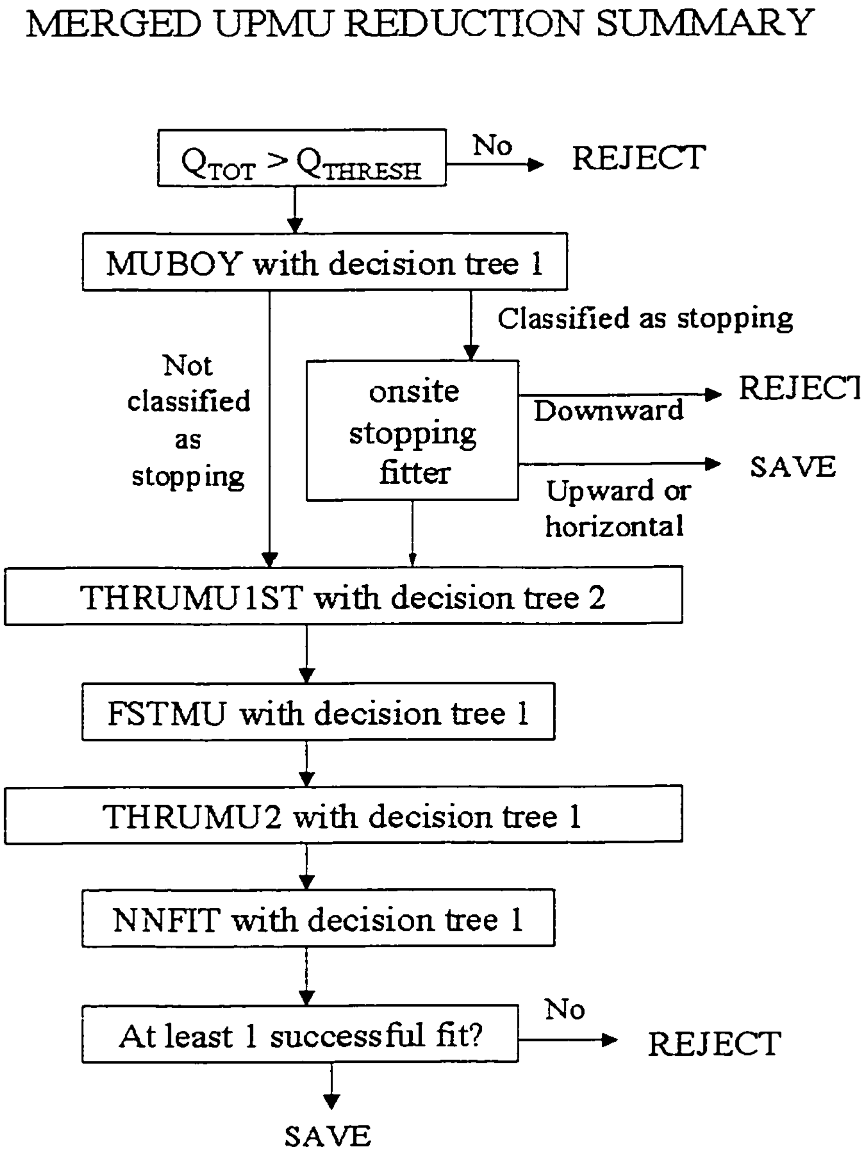

Figure 3.1 This summarizes the general framework of the merged reduction program.

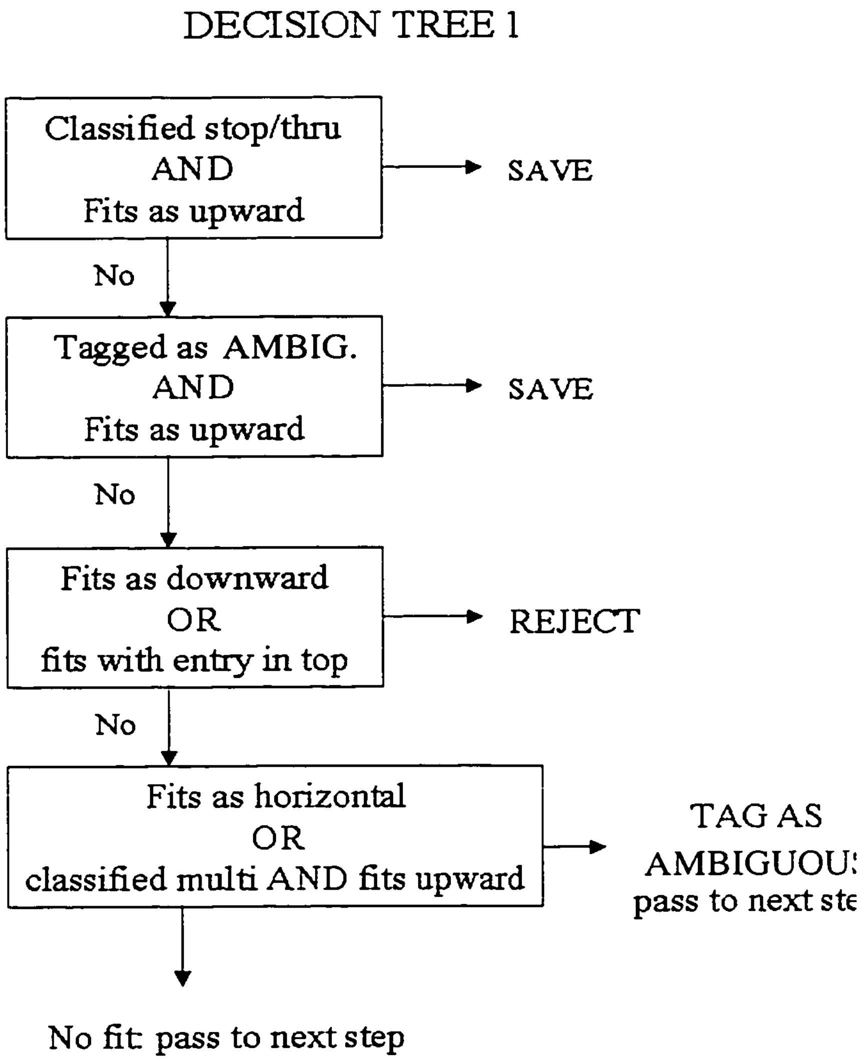

Figure 3.2 This was the decision tree used by 5 of the 6 fitting algorithms.

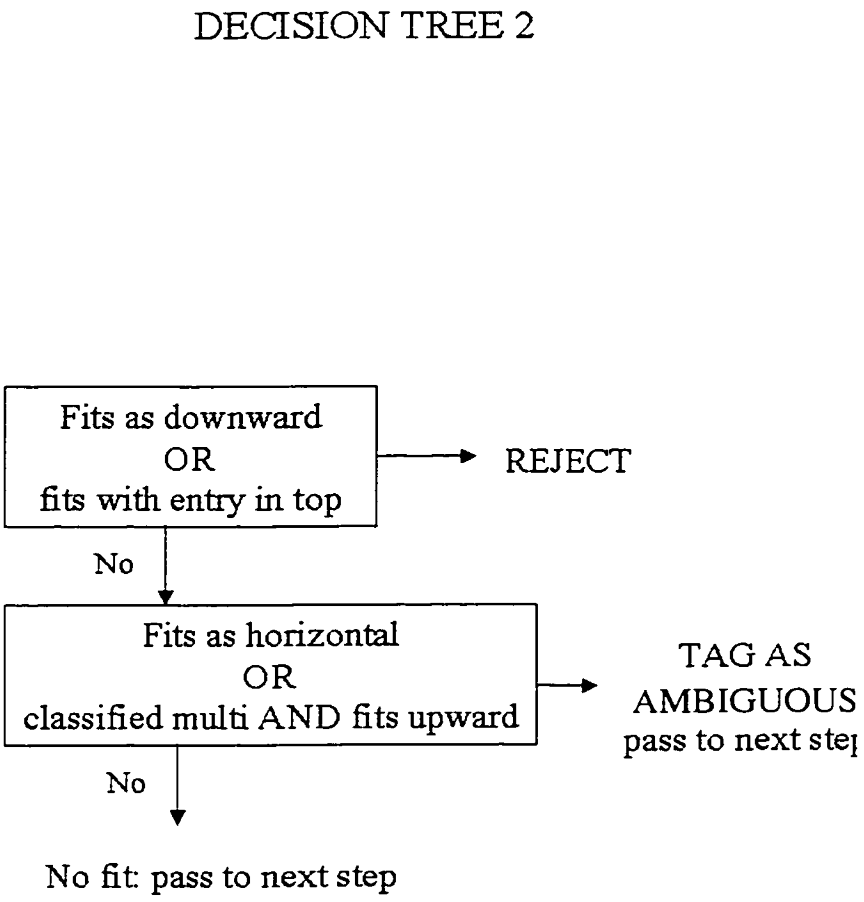

Figure 3.3 This decision tree was used only by thrumulst, which was an algorithm specialized to fitting single through-going muons using minimal CPU time. It had a slightly higher rate than other fitters (but still extremely low) of occasional very inaccurate fits, therefore although it was allowed to reject events as downward-going, confirmation from a second algorithm was required before saving an event as an upward-going candidate, in order to ensure a minimal workload for the human muon track fitters.

ID PMTs, fell in the range 8000 p.e. $< Q_{tot} <$ 1,750,000 p.e., and pass such events through each of the various track fitting subroutines in turn, until an algorithm was found which produced an unambiguous result. Q_{tot} was selected as the discriminant because in practice it really represented more than just an amount of electrical charge: since each photoelectron represented one Cherenkov photon which traveled through the water to interact successfully with a PMT at the edge of the tank, Q_{tot} measured the actual amount of light that the detector saw. Thus, the Q_{tot} cuts really were a way to ensure that the only events allowed to pass to the track fitting algorithms were ones for which the total energy deposited in the detector fell within a certain range.

After passing through the track fitters, events for which an unambiguous fit result could not be found by any algorithm (a few hundred per day out of an input sample of order 2×10^5) were discarded. Of the small percentage of events which were unfittable, a large majority were not single track events–the pattern that one expects from upward-going muons–therefore the effect of this cut on the efficiency for finding upward-going muons was negligible indeed.

3.2 Preliminary Cuts Based on Event Brightness

3.2.1 Maximum Allowable Q_{tot}

By request of one of the Super-Kamiokande collaboration members, events with $Q_{tot} >$ 1,750,000 p.e. were automatically saved, without running the track fitting algorithms on them, for further research studies which were not necessarily related to upward-going muon searches. For simplicity and convenience it was decided also to impose this same value as an upper limit on the maximum allowable light deposition in the upward-going analysis. This cutoff is very conservative because more than 99% of the events in the final upward-going data sample produced below 250,000 p.e. in the detector, and the largest upward-going event ever observed produced 950,000 p.e., therefore if an upward-going muon had ever produced as much as 1,750,000 p.e. it would have been extremely far out in the tail of the distribution. Fortunately, since the events with $Q_{tot} >$ 1,750,000 have already been culled from the raw data sample and placed into compact data summary files, they are readily available for re-analysis should any future investigator ever decide to raise the analysis threshold. It should be emphasized, however, that this is not viewed as a very likely scenario: empirically, the high Q_{tot} events appear to be, with a high degree of purity, essentially all multiple muon events.

3.2.2 Minimum Allowable Q_{tot}

Events with $Q_{tot} < 8000$ p.e. were summarily rejected unless they occurred within 30 μsec of some event with a good fit to both the upward-going and stopping hypotheses. The 30 μsec window was chosen so that it would be several times as long as the 2.2 μsec half-life of a free muon, thus ensuring that essentially all decay electrons remaining after losses due to triggering inefficiency or nuclear capture of the parent muon would be saved. The reason that 8000 p.e., specifically, was chosen as the threshold was to eliminate events with track lengths of less than 7 m without wasting CPU time fitting them, and without inadvertently eliminating any events with track lengths of more than 7 m (see figure 3.4). The 7 m track length threshold was chosen for all events in the final data sample because earlier studies showed[85] that the angular resolution of several of the fitting algorithms was almost independent of track length when the length exceeded 7 m, but varied strongly with lengths shorter than this. It was possible to express the correlation between stopping muon track length and the amount of light gathered at the walls of the detector approximately (see figure 3.5), for track lengths of less than ~30 m, as:

$$Q_{tot} = 21.0 \frac{\text{p.e.}}{\text{cm}} \cdot L_{track} - 1430.0 \text{ p.e.} \tag{3.1}$$

The relationship was linear presumably because the rate of production of Cherenkov light per unit distance saturates as β approaches 1, and muons with sufficient energy to travel more than about 2 meters through the water will do so with $\beta > 0.95$ during most of their trip. The small extrapolated negative offset of 1430 p.e., which interpreted literally would imply the absurd result that muons of sufficiently small L_{track} must have absorbed light from the detector, was presumably an artifact of the non-validity of the linear approximation as L_{track} approached 0, due to the fact that the Cherenkov mechanism begins to turn off in this extreme. Additionally, of course, it's worth remembering that Q_{tot} was really determined

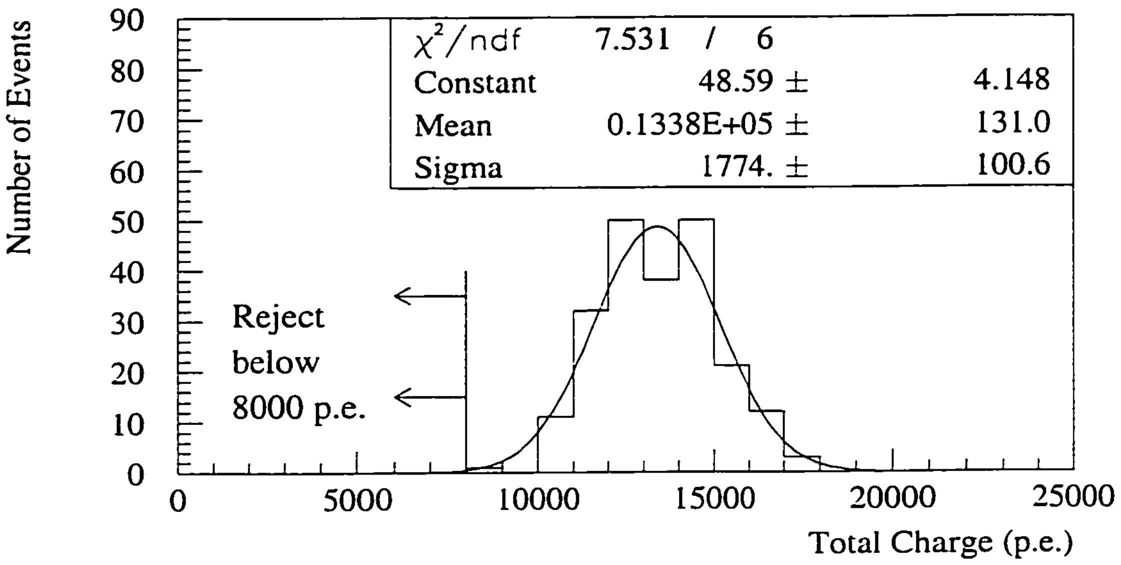

Figure 3.4 This is the distribution of the total charge, Q_{tot}, for stopping muon track lengths between 650 and 750 cm. The track lengths were estimated by selecting only stopping muon events in which a decay electron event followed and in which the position of the electron's short, often convoluted path in the tank could be reliably estimated by a fitting algorithm specialized to point-sources of light. The fitted entry position of the muon at the edge of the tank was defined as the start of the track and the electron's position as the end.

not just by L_{track}, but was also modulated somewhat by the effects of light attenuation in the water, which were ignored here because the goal was just to get a rough estimate of event energy while using minimal CPU time.

As the initial energies of the muons increased to allow deeper penetration into the tank, or to allow the muons to travel all the way through the detector, this approximation became progressively poorer because the typical energies of secondary particles produced by other energy loss mechanisms grew large enough to exceed the Cherenkov threshold, so that they produced additional Cherenkov light of their own. The production of extra light meant that

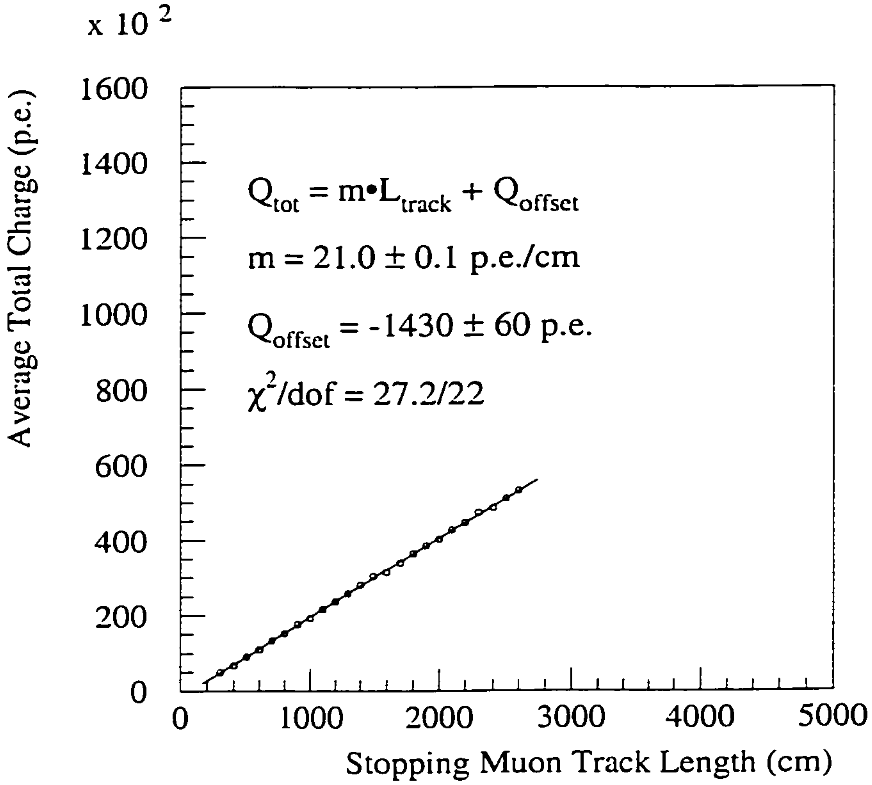

Figure 3.5 The fitted centroids of several Gaussian distributions like the one in figure 3.4 are plotted vs. track length, and the reduced information contained in these points and their fitted error bars is used as the input to a straight-line fit of Q_{tot} vs. L_{track}.

any deviation in Q_{tot} for such events was always systematically higher (see figure 3.6) than the prediction of equation 3.1 for a given track length. There were no known circumstances in which an event of 7 m track length ever produced less than 8000 p.e., therefore the chosen threshold was sufficiently low to guarantee a negligible rate of loss for all events of interest.

3.3 Event Classification and Track Fitting Algorithms

Events with $1,750,000$ p.e. $> Q_{tot} >$ 8000 p.e., of which there were some 2×10^5 per day, constituted a metaphorical ore from which the rare upward-going muons had to be extracted.

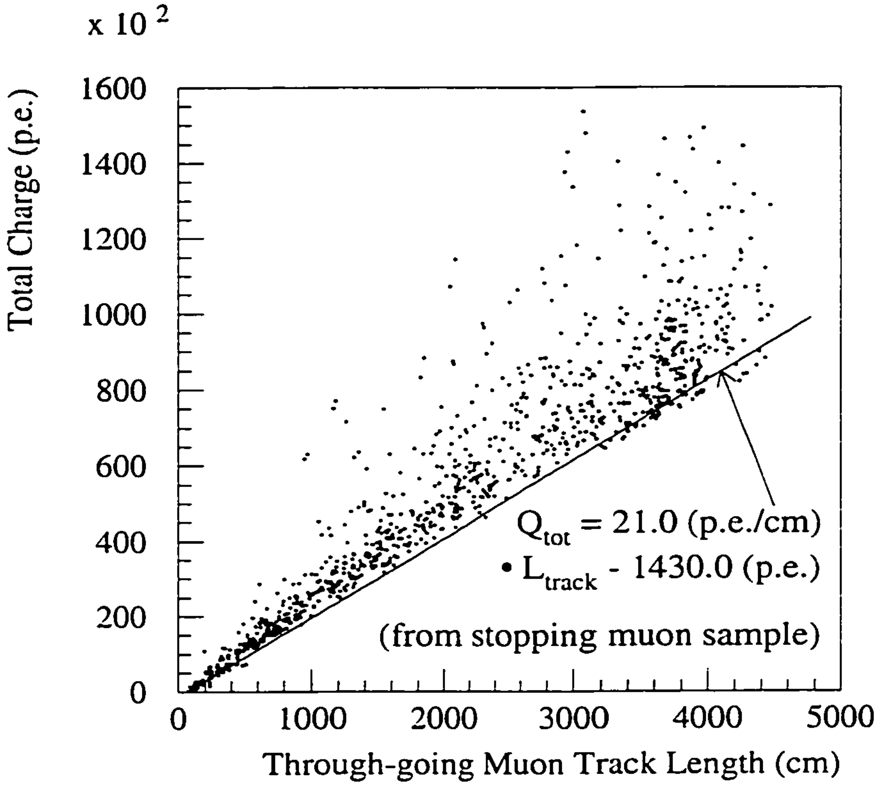

Figure 3.6 This shows measured Q_{tot} plotted against fitted L_{track} for 956 through-going muons. The rate of light production in the detector per unit track length is often much higher than that observed for stopping muons (figures 3.4 and 3.5) because the energies of the events in this sample were higher, and therefore they were able to produce secondary Cherenkov light in a significant fraction of cases.

In order to perform the extraction, the data reduction program had to be able to identify and eliminate events in which the z direction cosine of the distance vector traversed by a particle within the detector was greater than 0. Ideally, such events would be identified by following the canonical parameter estimation procedure of all physical scientists: use the maximum-likelihood method to determine the value of the z direction cosine which would have most probably resulted in the data set that was actually observed. In practice, it was prohibitively difficult to calculate, quickly enough to keep pace with the data stream, for dozens or perhaps even hundreds of trial muon entry and exit positions, the likelihood of the observed arrival times and intensities of the light at all of the PMTs. There were two reasons for this: first, the fact that the detector was instrumented with nearly 13,000 PMTs, each of which yielded two pieces of information, both time and intensity, resulted in a parameter estimation problem with nearly 26,000 degrees of freedom, a prohibitively large number. Second, there were complex physical processes, some of them statistical in nature, which influenced the amount and production angle of the light originating from the parent muon, as well as influencing the light itself as it propagated through the tank and interacted with the PMTs, and therefore a very complicated, detailed Monte Carlo algorithm was required in order to predict reliably the detailed response of the detector to a muon entering and either stopping or exiting at a given pair of vertices. Although such Monte Carlo algorithms existed, were well developed, and indeed, were an important component of much of the data analysis performed by the experimental collaboration, they were not generally fast enough to be usefully incorporated into a track fitting subroutine.

To write track fitting algorithms which recognized muon events and reduced them to a set of fitted track parameters quickly enough to keep up with the data stream, a number of simplifying assumptions had to be made. While this tactic generally greatly increased the speed of the computer code, it also meant that all of the algorithms used in the data reduction failed to fit a few events for which those assumptions were not sufficiently close to true. This provided the basis for the strategy, outlined in section 3.1, of trying different algorithms in succession, until one of them returned a non-null result: since all of the fitters were able to return a useful result the majority of the time, even if one of the fitters failed, the next fitter was usually able to do the job.

3.3.1 Design Principles and Underlying Assumptions

Each of the different fitting algorithms was written by a different programmer, the specific simplifying assumptions that were adopted by each programmer varied from one to the next, and the details of those variations are beyond the scope of this document. Nevertheless, it is still worthwhile to offer a general sense of how the algorithms did their jobs. Many of the fitters took advantage of a peculiarity of the detector's design in order to make an initial guess at an entry point: the backs of the inner detector PMTs were transparent, and light could get into them as the radiating muon traversed the 50 cm wide support structure which housed them, thus very often the earliest triggered PMTs, as well as some of the more brightly illuminated, were the ones very nearest the entry point of the muon into the detector. This was even more likely to be true if the muon entered the detector at a very oblique angle: then part of the Cherenkov cone was guaranteed to hit a PMT near the entry point immediately upon entry. Thus many fitters made initial guesses about track entry points by recognizing clusters of physically proximate, early firing PMTs. For choosing the exit point in the case of through-going muons, a slightly different assumption was often used: the PMTs near the exit were very rarely the latest to fire in the event, but of course most of them usually fired at about the same time, and of the top several dozen most brightly illuminated PMTs, usually a large fraction of them were quite near the exit, so a nearly correct initial guess could often be made by writing an algorithm which was good at recognizing large unusually bright in-time clusters of PMTs.

3.3.2 Events That Challenged the Underlying Assumptions

Several conditions appeared naturally in some of the events which sometimes caused the basic design assumptions to be violated, and often caused especially the simpler fitters to fail.

3.3.2.1 Stopping Muons

The first of these unusual conditions arose occasionally as a muon would stop inside of the detector, and in such cases (see figure 3.8 for an example event display) the pattern of light that splashed upon the walls was a ring surrounding a darkened region that would have been filled in if the muon had continued all the way to the opposite wall. In those situations, the brightest cluster of PMTs was usually on the ring, nowhere even near the projected exit, because the projected exit in such cases always lay near the center of the darkened region. Treating such events as if they were through-going muons thus would have produced egregiously incorrect results for the muon's direction vector, a fatal flaw in a search for upward-going muons since that is the quantity of interest.

3.3.2.2 Multiple Muons

Multiple muons were another confounding natural feature of the data: occasionally, two or three or sometimes even more muons were produced with such narrow opening angles that even after traveling the several km from their production point in the atmosphere to the detector, they were separated laterally by only a few meters, and all of them were recorded in the same event. The patterns of light they produced on the detector walls were readily recognizable to the human eye viewing the event with a display program, but they were also highly variable, and therefore it was difficult to program a computer to do the same. Sometimes, the light from the different particles fell on well separated regions of the detector wall, while at other times, the light from one muon overlapped substantially with that of an adjacent event, so that many of the PMTs received light from several muons, creating bright spots which were not necessarily exit points. Treating such events as though they were only single through-going muons again carried the potential for egregious errors in the event's fitted direction, for the entry point of one of the muons could have gotten mixed up with the exit point of another, or even with a region that was not an exit point at all.

3.3.2.3 Bremsstrahlung Muons

Still a third natural feature of the data that often confounded the fitting algorithms was the existence of events in which a substantial component of the emitted light originated from Cherenkov radiation coming from secondary particles rather than from the parent muon itself. The critical energy for muons, the energy above which radiative energy losses grow larger than ionization losses, is a few hundred GeV; there were certainly some through-going muons in the data sample with energies as high as this and thus the generally accepted explanation for such events was that they were the result of the muon undergoing bremsstrahlung in the tank and producing an electromagnetic cascade containing electrons above the Cherenkov threshold. The radiating cascade products were prone to scattering through macroscopically large angles before they came to rest, with the result that the clean edges of the projected Cherenkov cone upon the walls of the detector were completely washed out by a brilliant flash that illuminated much of the tank.

3.3.2.4 Strategy of Specialization

All of these natural features of the data presented different challenges, and they proved sufficiently difficult that none of the collaboration members who worked on developing fitting algorithms were able to surmount all of them. Many of the fitters therefore were specialized to do a good job on only one or two types of events, and this provided a further explanation of why so many different algorithms were needed in the merged reduction: each brought with it a special skill that most of the other algorithms did not possess. The author of each muon fitting algorithm gave his code a nickname for ease of reference, and these routines plus one event classification algorithm are listed by name in order of their appearance in the merged data reduction program in table 3.1, along with some information about the abilities and efficacies of each.

3.3.3 Goodness of Fit

To estimate whether or not the fit parameters were accurate guesses about the true values of underlying variables, a goodness of fit statistic was needed. Ordinarily, in the context of a more standard parameter estimation formalism, this role would be played by a quantity such as χ^2, which directly represents the value of the likelihood function. In this case, however, to reduce the complexity of the problem, alternative measures had to be used. Unfortunately, these measures varied from one fit algorithm to the next, depending arbitrarily on how each programmer chose to design his code, so there was no easy way to compare the goodness of each fit to that of any other algorithm, because unlike χ^2, these goodness statistics didn't have any universally recognized meaning. The goodness of fit statistics were generally classifiable into two types: those that relied on timing information and those that relied on light intensity. The timing-based goodness of fit parameters typically made use of PMT timing residuals, defined as the observed arrival time at each PMT minus the expected

Table 3.1 This table shows what types of events each muon fitter was primarily used to fit. The numbers in the third column show the percentage of events that were not yet able to be rejected following each stage, for any of the different reasons mentioned in the main text: a null result, an ambiguous result, or an upward-going result. The numbers in the fourth column give the additional reduction factor compared to the immediately previous stage. Reduction factors for all fitters after muboy are low because this table shows fitter performance in situ in the reduction, without accounting for the fact that the most easily fitted events were removed already by the earlier stages. The smf fitter is treated as a special case, since it only received events for which muboy failed if they were classified as possibly stopping, whereas other fitters were allowed to operate on all events for which no decision had yet been made. The numbers in parentheses refer to how it operated on only those events it actually received, as opposed to those not in parentheses, which state the results as if smf had received all events but returned a null result for those not classified as stopping. The numbers at the very bottom show the percentage of events remaining after rejecting events for which none of the six fitters were able get an acceptable fit result.

Name	Specializations	Percent of Initial Events Remaining After Cut	Additional Reduction Factor
stopmulst	Classified events according to type (stopping, through-going, etc.) but did not attempt to fit.	-	-
muboy	Through-going, stopping, multiple, and also classified events according to type	3.8	26.2
smf (stopmu2nd)	Stopping	(87)3.7	(1.15)1.02
thrumulst	Single through-going	2.2	1.7
fstmu	Single through-going	1.5	1.5
hutl (thrumu2nd)	Multiple	0.41	3.5
nnfit	Very bright (bremsstrahlung) muons	0.25	1.6
Final number of events saved after null result cut		0.067	3.8

arrival time for a set of trial fit parameters. For a perfect fit, the answer would be 0 for every PMT, and for real, imperfect fits, the answer was always positive for some PMTs and negative for others. Thus, the width of the residual distribution gave a measure of the goodness of the fit; distributions which were more tightly clustered around 0 were generally better. The light intensity-based methods were predicated on the notion that all of the light in a muon event should lie cleanly within a cone of about 41.4°, the approximate opening angle in water for a singly charged particle with $\beta \approx 1$. The typical tactic was to assign points, for a given set of trial fit parameters, to all of the PMTs which were triggered in a 42° cone about the putative track, and subtract points for all of the PMTs which fired in a conical annulus drawn immediately outside of the borders of the inner one, since that area would be expected to be dark if the trial parameters were indeed the correct ones.

3.3.4 Final Event Categorization and Resultant Action

After being processed by each fitter, an event was typically placed into one of four categories. Most often it was classified as downward-going, in which case it was immediately rejected and processing continued to the next event. Somewhat less often, the fitter returned a null (unfittable or extremely poor fit) result, in which case the program continued onto the next fitter to try to get a result. In a few cases, the reduction returned conflicting, ambiguous, or borderline information. One example of such ambiguity was when events were classified as multiple muons but fit as upward-going; these were unexpected from a theoretical standpoint since all upward-going muons were assumed to come from single neutrino interactions producing only one outgoing charged lepton. Another example would be an event fit as downward-going but too near the horizontal to cut safely. In a very few cases, an event was categorized as a good fit to an upward-going muon; when this happened it was saved to disk for human inspection without further processing, so as to parallel the treatment given to downward-going muons, for which an immediate rejection decision was also made if the fit was good. At the end of the process, any events for which none of the fitters were able to return a non-null result were rejected, on the assumption that if they were simple, single muon events, at least one of the fitters would certainly have recognized it by that point and flagged the event. Events for which an ambiguity arose but could not be resolved were saved to disk for manual inspection using event visualization software.

3.4 Manual Fitting of Events

After events were pre-screened by the data reduction process, typically about 70 or 80 events per live day remained which could not be eliminated by the programmed algorithms. These remaining events had to be further screened using the slow but extremely powerful pattern recognition capabilities of the human mind. Of these events, about 1.4 per day eventually proved to be real upward-going muons. Of the many other saved events which did not ultimately prove to be upward-going, about half turned out to be shallow, nearly horizontal downward-going events which, due simply to the constraints of the fitter resolution, were fit too close to the horizon to be eliminated. Finally, there were also a few dozen events per day for which the data reduction program simply returned wildly incorrect results. These events did not appear to be drawn uniformly from the raw data sample; rather, they could usually be categorized into one of three or four particular types. Typically, there were a few multiple muons, some anomalously bright events, and some corner clipping events, in which a muon entered the tank through the cylindrical wall, traveled a few meters, and exited through the bottom. These types of errors happened, presumably, because some of the fitters suffered from biases which made them prone to occasional failure on a few types of events with very specific patterns of light distribution, and these few rare types of events were simply misfit over and over again. Since these events only numbered a few dozen out of approximately 2×10^5 per day, this is not evidence of any significant systematic bias in the reduction.

A crucial part of the manual event selection process was manually fitting the entry point and direction of the event. This segment of the data reduction was performed by having a human being select trial entry and exit points (or projected exit points, in the case of stopping muons) on the walls of the inner detector, using a mouse and an interactive event display program. (Example screen images depicting two events from this data sample, as rendered by one of the more commonly favored visualization programs, are shown in figures

3.7 and 3.8.) After the human researcher made an initial rough guess about the entry and exit points, the trial points were iteratively refined with the mouse until the fit was deemed acceptable. The primary means of judging fit acceptability was by having the event display program draw the projected edges of the Cherenkov cone upon the wall of the detector. The Cherenkov light patterns which shone upon the wall in the vast majority of upward-going muon events had very clear, distinct edges: the illuminated areas were often tens of meters across and yet dropped off in intensity by an order of magnitude or more in under a meter as one approached the edge. Fitting the events therefore became a problem of edge recognition for the human researchers: the primary means of selecting a fit was to find the pair of entry and exit points for which the resulting projected Cherenkov cone's outline lay most nearly at the edge of the illuminated section of the wall all around its perimeter. This method of manual fitting took only light intensity information into account; an alternative visualization-based method using timing information was also developed and frequently was employed as a backup. To do a timing-based manual fit, the event visualization program displayed a false color image of the timing residuals (defined in section 3.3.3) for a given pair of trial entry and exit points. Since the basis of timing residual dependent fitting methods is the assumption that the best fit will be one in which the residuals for all PMTs have about the same value, the human being chose the best fit by selecting the entry and exit points which resulted in the most uniform color for the false color display.

In order to guarantee that the vagaries of the human attention span would not unduly compromise the efficiency of the process for identifying upward-going muons, a double scan procedure was adopted, and any event that was selected by either of the two human scanners was sent on to a third scanner for final evaluation. Thus, by using both a two-tiered (i.e., involving both initial and final scans) as well as a two-pronged (i.e., involving two completely independent initial scans) approach, the procedure ensured that the inevitable human errors which were bound to occur from time to time would eventually be ferreted out and would ultimately have a negligible impact on the final event sample.

A major drawback to this entire process was that, unlike the automatic fitting algorithms, quantitative measures of goodness of fit were not employed, with the result that, given two different manual fits, there was no systematic way to choose which one was bet-

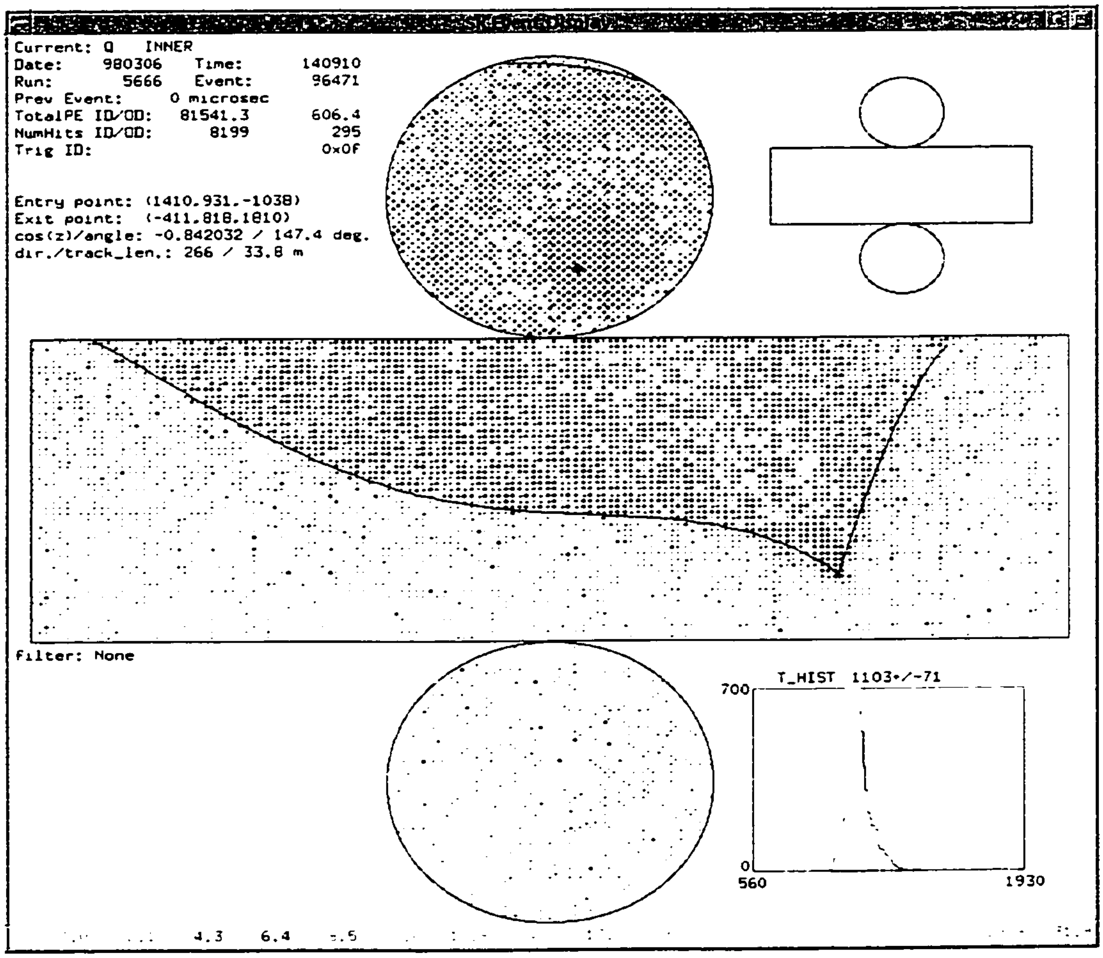

```
Current: Q   INNER
Date:      980306   Time:        140910
Run:         5666   Event:        96471
Prev Event:      0 microsec
TotalPE ID/OD:   81541.3          606.4
NumHits ID/OD:      8199            295
Trig ID:                           0x0F


Entry point: (1410.931.-1038)
Exit point:  (-411.818.1810)
cos(z)/angle: -0.842032 / 147.4 deg.
dir./track_len.: 266 / 33.8 m
```

```
Filter: None
```

Figure 3.7 The main portion of this display depicts the ID as if it had been cut vertically down its cylindrical wall and unrolled into a flat sheet, with the top and bottom end caps folded outward. Each pixel represents a single 50 cm ID PMT. The false colors represent the number of photoelectrons gathered at each PMT; a scale is available at the bottom of the image. The inset at the upper right represents the OD. The event shown here is a through-going muon which entered the detector near the bottom of the cylindrical wall and exited through the ceiling; note the two clusters of PMT hits near the entry and exit points in the OD. The black outline that separates the blue (~ 10 p.e.) from purple (~ 1 p.e.) colored PMTs represents the edge of the Cherenkov light cone for the selected pair of trial entry and exit points; assuming a Cherenkov angle of 41.4° and perfect detector performance, no light should be visible outside of this cone.

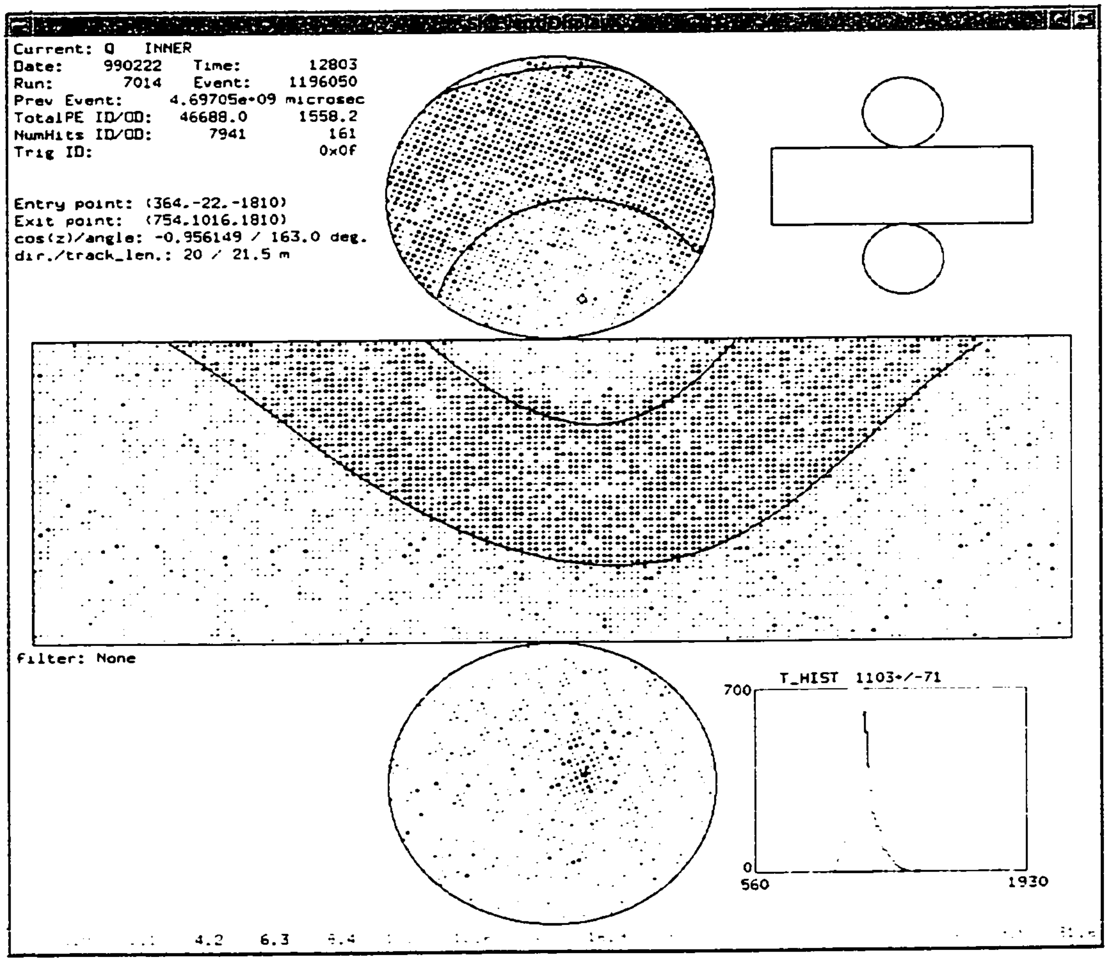

```
Current: Q   INNER
Date:     990222   Time:       12803
Run:        7014   Event:    1196050
Prev Event:      4.69705e+09 microsec
TotalPE ID/OD:   46688.0      1558.2
NumHits ID/OD:      7941         161
Trig ID:                        0x0f

Entry point: (364.-22.-1810)
Exit point:  (754.1016.1810)
cos(z)/angle: -0.956149 / 163.0 deg.
dir./track_len.: 20 / 21.5 m
```

```
Filter: None
```

Figure 3.8 This figure shows a stopping upward-going muon which entered through the bottom end cap and would have exited through the top if it had not come to rest inside of the tank first (for details on how to interpret some of the features of this image, see the caption in figure 3.7). Two characteristics distinguish this event as stopping: first, the Cherenkov pattern in the ID has a large, unilluminated hole in the center, because the muon stopped in the detector before it could reach the ceiling and fill in the rest of the ring; and second, the OD shows only one cluster of hit PMTs, at the bottom of the tank, and no corresponding cluster near the top. The black curve which roughly traces the inner edge of the ring shows the inner edge of the Cherenkov cone, assuming that the muon had traveled in a straight line from the selected entry point, and stopped generating light after 21.5 m. The curve does not lie as perfectly along the border between light (blue) and dark (purple) areas as does the one that circumscribes the pattern on the outside: this is common among stopping muons and indicates that the simplistic assumptions used to calculate the cone's inner edge are violated to a modest degree: stopping muons often scatter around a bit during the last few meters of travel, making thin wispy rings of light that fill in parts of the darkened inner circle slightly.

ter. There were two important advantages, however. One was that the human mind, being much more complicated and intelligent than an automatic fitting algorithm, was less prone to making egregious mistakes, and therefore manual fits tended to have fewer outliers (that is, fewer events in which the final result was off by many times the typical uncertainty for the process) than automatic ones. The other advantage was that manual fits had slightly better resolution than any of the automatic algorithms that collaborators were able to devise: by all accounts[86][85] (besides the references, see also section 4.4 of this work) and measures (i.e., comparison either to Monte Carlo truth data or between pairs of independent human fits), manual fits were typically correct to within 1.5° or better; the best automatic algorithms could only achieve somewhere around 2° to 3°. These benefits were felt to outweigh the drawbacks, therefore manual fits have been used in this and all other upward-going muon analyses for this experiment.

Chapter 4

Preliminary Calculations

Much of the data analysis that will be performed in the chapter following this one will be shaped by, or will require as a basic input, the results of some fundamentally simple calculations which broadly characterize the data or the detector used to gather it. Quantitative statements about how long the detector had to operate in order to acquire the observed number of neutrino-induced events, about the relative sensitivity of the experiment to various regions of the sky, about background events that might obscure a signal, and about detector resolution are developed below.

4.1 The Live Time

In order to make a quantitative statement about the existence or strength of an astrophysical point source, one must take into account the length of time that one has spent observing the sky in any given direction. Since the detector was mounted rigidly in the earth, its field of view of the cosmos was constantly changing as the earth rotated about its axis, and therefore, in principle at least (in practice, it will turn out not to have mattered very much in this case; see figure 5.10), to be able to do the analysis properly, one should have kept track not only of how long the detector was in operation while acquiring a particular data set, but also when it was in operation. To this end, each event which triggered the detector was stamped with a time, as well as various status information which encoded the operational state of the detector when that event was acquired, in order to identify occasional periods when the detector was partly functional but not acquiring data that were useful to the analysis.

The experimental apparatus acquired data with approximately a 90% duty cycle during most of the time period of this analysis, meaning that it was turned off for routine maintenance and for certain types of calibration exercises for roughly 10% of the time, and was left running at all other times. Additionally, while the detector was active, there were occasional interruptions which had to be subtracted from the data-taking time in order to calculate the live time. The most frequent of these was the acquisition of pedestal data, a routine calibration exercise which was more fully automated into the data acquisition scheme than any of the others; it consumed around 2.5% of the data-taking time on average. Other types of dead time such as data loss due to equipment malfunction were also monitored, but did not detract significantly from the live time after the initial start-up phase of the experiment.

The live time for a given interval of data-taking time containing n events was counted by examining each of the $n - 1$ pairs of consecutive events and adding the small time interval between them either to the live or dead time depending on the operational status of the detector encoded in the data for each event. The final results were summarized in intervals of roughly one minute each, which corresponded to many hundreds of events. This time scale was chosen so that the angle of rotation of the Earth about its axis in that time ($\sim 0.25°$) would be small compared to the angular resolution of the experiment, and therefore the granularity in the measured live time would not have contributed significantly to the overall uncertainties in the analysis of, for example, the relative exposure of the experiment to different regions of the sky, had that turned out to be important (section 5.3 includes a discussion of why it did not).

The total live time for upward-going muons in this analysis was different for stopping and through-going events, because the first several hundred days of the data sample were garnered from a different data reduction program than the rest, which were accumulated using the improved algorithm described in chapter 3. The part of the earlier data reduction that was used for identifying stopping muons did not work very well on raw data that was missing information from the outer detector, while the part of the algorithm used for identifying through-going muons was fairly insensitive to this problem, which happened frequently in the weeks and months immediately following the commissioning of the detector, due to occasional failures of the early prototype versions of the data acquisition system that were in use at the time. A decision was made in the initial phases of the experiment to make use of all available upward-going muon data, rather than throw away data in the interests of having equal live times for both the stopping and through-going samples. Although the data acquisition software became more reliable later on, and furthermore the analysis practice was modified so as to impose identical ID and OD availability requirements on all additional events appended to either sample, nevertheless, other collaborators have continued to include the few missing-OD events from the earlier data reduction in their analyses, despite the inconvenience of the live time asymmetry. To expedite comparison with the work of others, the same practice will be adopted here. The live time of the stopping upward-going muon sample used in this analysis was 1458.3 days, and of the through-going sample, 1485.0 days. In some of the studies presented in chapter 5, the through-going and stopping samples will be combined, and some sort of average effective live time for the two will be required. Since the difference between the two results is less than 2%, taking an intermediate value will, of course, not yield a very different result than arbitrarily choosing either of the two extremes, but technically it seems like the most correct way to treat the data, so it will be incorporated into the analysis anyway. In cases where an effective combined live time is needed, it will be estimated by requiring that the total effective combined event rate equal the sum of the separate event rates, i.e.,

$$\frac{n_s + n_t}{T_{tot}} = \frac{n_s}{T_s} + \frac{n_t}{T_t}, \tag{4.1}$$

or

$$T_{tot} = \frac{n_s + n_t}{\frac{n_s}{T_s} + \frac{n_t}{T_t}}, \tag{4.2}$$

where n and T refer to the event numbers and live times, respectively, and the s and t subscripts refer to stopping and through-going muons. This convention is appealing for two reasons: first, because of its natural tendency to conserve the estimated event rate, which is important because the rate of atmospheric neutrinos (and thus the overall event rate in the case of the null hypothesis) must be a well-defined experimental constant, just like the rate of the cosmic ray primaries; and second, because this convention yields the intuitively reasonable result, as T_s approaches T_t, that $T_{tot} = T_s = T_t$. For this analysis, $n_s = 403$ and $n_t = 1634$, thus $T_{tot} = 1479.6$ days.

4.2 The Effective Area

The inner detector was a cylinder measuring 33.8 m across and 36.2 m tall; its most natural mode of use was therefore as a volume detector which was sensitive to events arising from particle interactions within the central region. This analysis, however, uses Super-Kamiokande as a flux detector: all of the events are required to have been triggered by particles which were produced outside of the detector and traversed some distance into it. Analytic expressions for the effective area subtended by Super-Kamiokande, in both the horizon and the equatorial coordinate systems, are derived below.

4.2.1 The Effective Area in Horizon Coordinates

The effective area may be written in terms of the following variables and parameters:

Symbol	Meaning	Nominal Value
h	inner detector height	36.2 m
r	inner detector radius	16.9 m
l	minimum track length in the inner detector	7 m
a	altitude angle	$-90° < a < 90°$

One can calculate the effective area for a given altitude angle a and minimum track length l by imagining all of the different ways that a rigid pole or rod of length l inclined at an angle a could have fit inside of the inner detector, subject to the constraint that one end of the rod, which symbolizes the entry point of a particle into the detector, had to have been touching the detector's edge. For downward-going tracks entering through the top end cap of the inner detector, or for upward-going tracks entering through the bottom end cap, the shape traced out by the most extreme possible locations of the midpoint of the imaginary rod, which of course encloses the set of all possible locations, is a pair of back-to-back lunes (see figure 4.1.a) of width $r - \frac{l\cos a}{2}$ and length $2\sqrt{r^2 - \frac{l^2\cos^2 a}{4}}$. From appendix A.1, the area of the two lunes together must be

$$A_{end} = 2\left[r^2 \arccos\left(\frac{l\cos a}{2r}\right) - \frac{l\cos a}{2}\sqrt{r^2 - \frac{l^2\cos^2 a}{4}}\right].$$

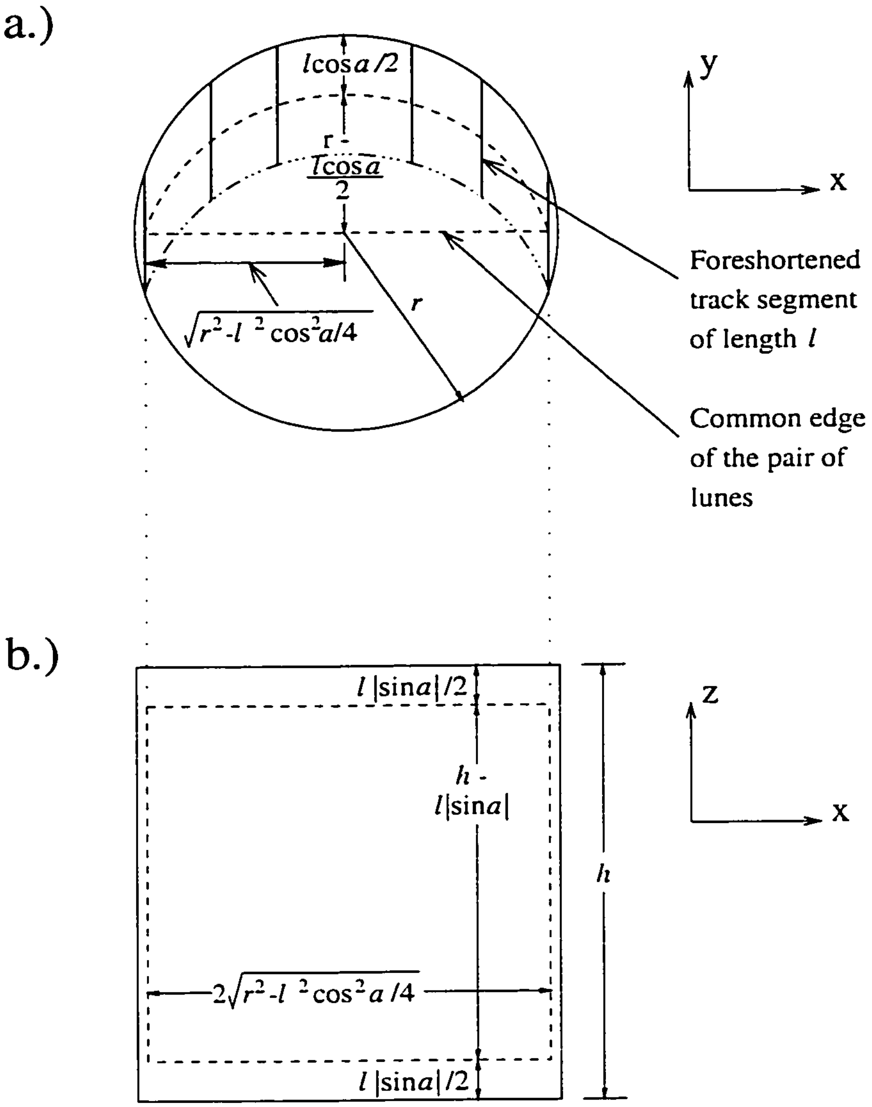

Figure 4.1 a.) This view is along the z-axis. The black lines distributed at intervals along the upper portion of the diagram represent tracks, or alternatively, imaginary rods of length l, inclined at an angle a, entering through an end cap at various locations, at one specific azimuth angle, chosen arbitrarily (they are all equivalent due to the rotational symmetry of the detector) to lie in the yz plane. The midpoints of the tracks trace out the curved edge (the dashed circle segment which passes through them) of one of two back-to-back lunes which enclose all the possible locations for the midpoint of any track entering through an end cap. For simplicity, only one of the lunes is drawn here; their common edge is defined by the dashed line, and the second lune is just the reflection of the first about this axis. b.) This view is along the y-axis. The dashed rectangle is a projection of a curved surface which encloses all the possible locations for the midpoint of any track entering through the side wall. The foreshortened rod is still oriented in the yz-plane.

For tracks entering through the cylindrical side wall, the shape traced out by the most extreme possible locations of the midpoint is a cylindrical segment which projects (figure 4.1.b) as a rectangle of height $h - l|\sin a|$ and width $2\sqrt{r^2 - \frac{l^2\cos^2 a}{4}}$. The effective area due to particles entering through the side wall is therefore

$$A_{wall} = 2\sqrt{r^2 - \frac{l^2\cos^2 a}{4}}(h - l|sina|).$$

The total effective area is simply the sum of A_{end} and A_{wall} together, with each multiplied by an appropriate trigonometric factor in order to account for the foreshortening of those areas from the point of view of the entering particle.

$$\begin{aligned} A &= A_{end}|\sin a| + A_{wall}\cos a \\ &= 2|\sin a|\left[r^2 \arccos\left(\frac{l\cos a}{2r}\right) - \frac{l\cos a}{2}\sqrt{r^2 - \frac{l^2\cos^2 a}{4}}\right] \\ &\quad +2\cos a\sqrt{r^2 - \frac{l^2\cos^2 a}{4}}(h - l|sina|). \end{aligned} \tag{4.3}$$

The expression for the effective area given in equation 4.3 is plotted for $-1 \leq \sin a \leq 1$ in figure 4.2.

4.2.2 The Effective Area in Equatorial Coordinates

Equation 4.3 is readily transformed to equatorial coordinates using the expression[21]:

$$\sin a = \sin\delta\sin\phi + \cos\delta\cos h\cos\phi$$

where δ is the declination, ϕ is the geographic latitude of the experiment (36.4257° N, WGS 84), and h is the hour angle. Since the substitution is direct and the final expression is long it is not written down here, but the result is plotted for upward-going tracks (for all δ, h such that $-1 \leq \sin a \leq 0$) in figure 4.3. Because the detector was left running for almost 90% of the time over the course of several years, and because the Earth swings through a full 360° of hour angle in only one sidereal day, it means that over the time-scales of this analysis there was a quite thorough time-averaging effect over all hour angles for any region of the sky. This average is plotted as a function of declination in figure 4.4.

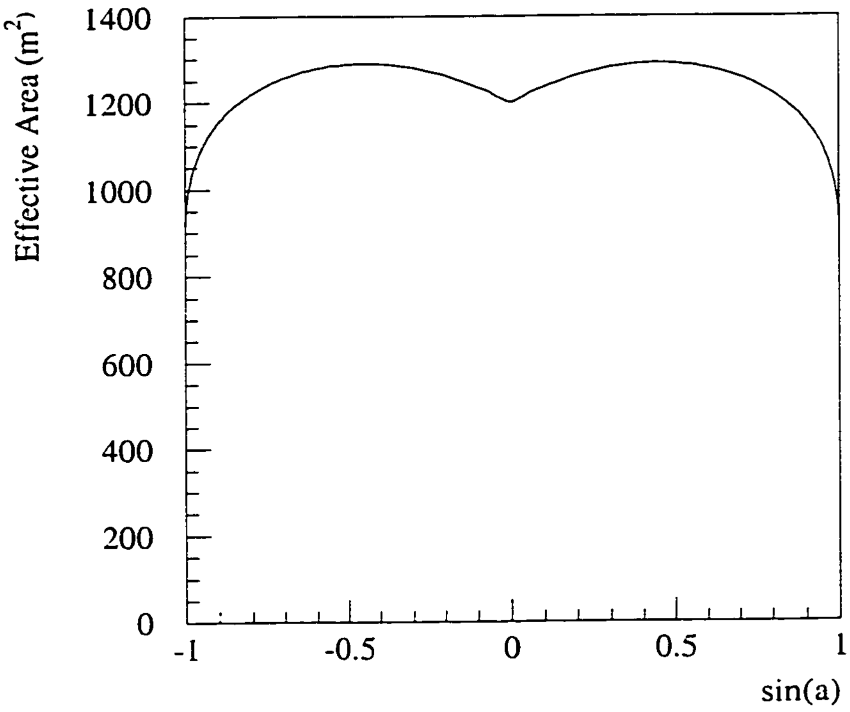

Figure 4.2 The effective area of Super-Kamiokande in m^2 vs. the sine of the altitude angle.

To the extent that the detector acquired its data uniformly in time, with interruptions spaced randomly throughout the sidereal day, the relative exposure of Super-Kamiokande at different declinations will turn out to approach a relative proportionality which scales to the values plotted in figure 4.4, and the exposure at a given declination will approach being identical for all right ascensions. The degree to which the assumption of sidereal uniformity actually approximates the data is examined more closely in section 5.3, and in figure 5.10.

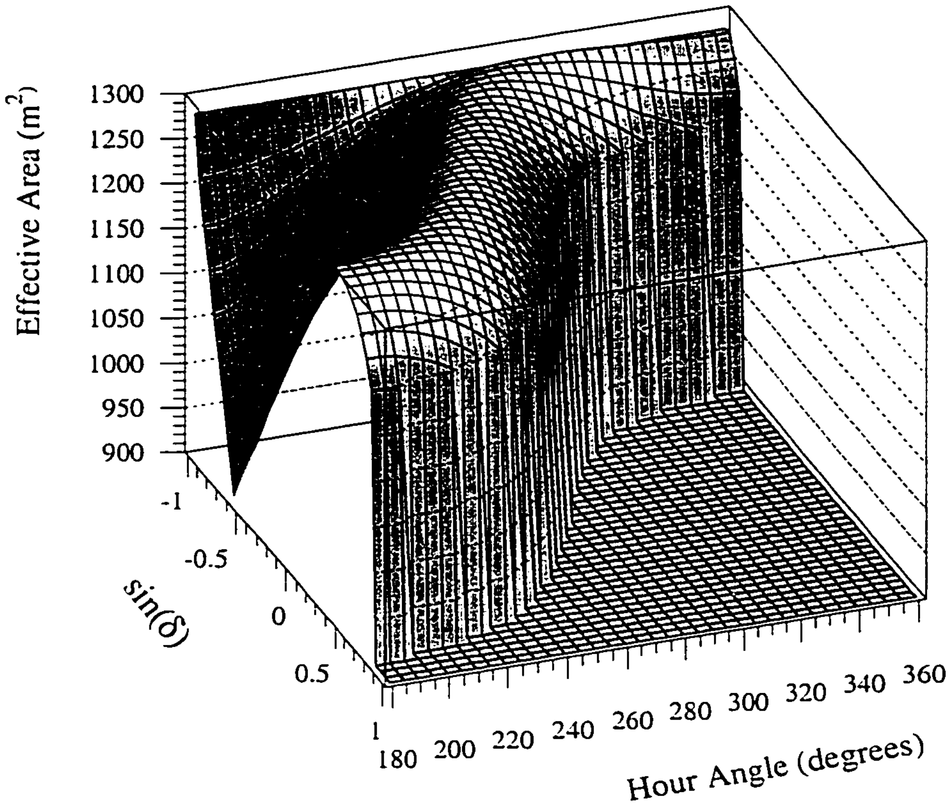

Figure 4.3 A cut-away view of the effective area of Super-Kamiokande for upward-going tracks which are longer than 7 m is plotted as a function of the sine of the declination angle and the hour angle. The function is symmetric about $h = 180°$. The part of the surface which is drawn as a flat plane at 900 m^2 is in a region corresponding to $\sin a > 0$, or downward-going tracks, and thus the effective area for upward-going tracks for δ and h in that region was really 0, which is off the bottom scale of the plot.

4.3 Background Events

4.3.1 The Background in Astrophysical Neutrino Induced Upward-going Muons Due to Atmospheric Neutrinos

Presumably, most of the upward-going muon events acquired at Super-Kamiokande were induced by atmospheric, rather than astrophysical neutrinos; to search for evidence of an

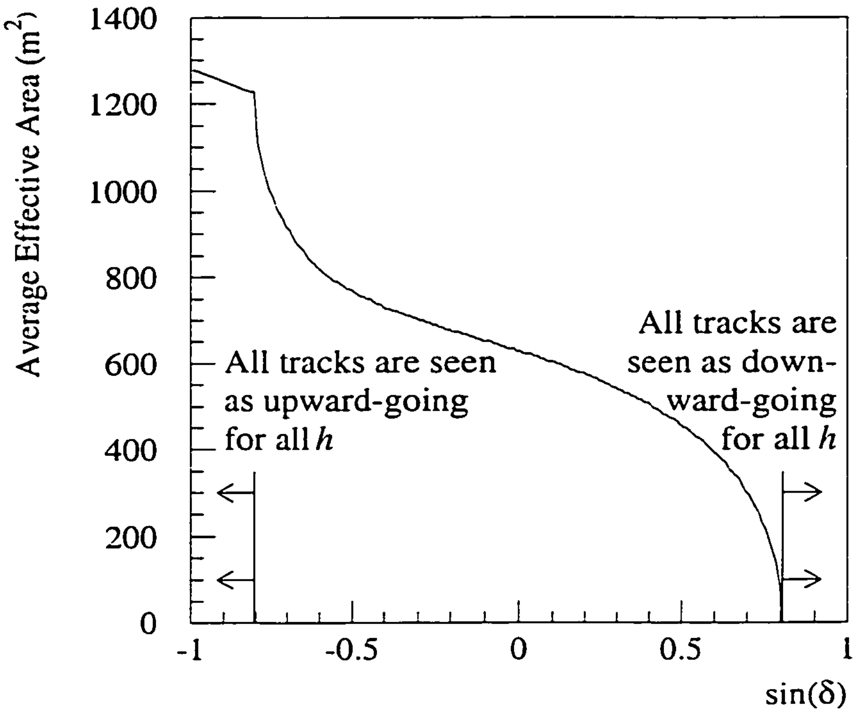

Figure 4.4 This shows the function plotted in figure 4.3, averaged over all hour angles.

astrophysical component means that one must look for a small signal hidden in a large background. Although the flux of the cosmic ray primary particles entering the atmosphere is isotropic, the flux of secondary particles and their decay products, including neutrinos, is not: it varies with altitude angle, and at sufficiently low energies, with azimuth as well. Therefore, to estimate the background present in an on-source region of sky, which is posited to contain both signal and background counts, one cannot simply pick just any region of the sky arbitrarily to serve as an off-source, background-only area. Rather, one must recognize that if the on-source region is imagined as depicted in figure 5.9a, as an area which moves through a band of local horizon coordinates as the Earth rotates, then the off-source region must be chosen symmetrically, as an area or series of areas within the same band, cycling around once per day at the same rate, just ahead of or just behind the on-source region.

4.3.2 The Background in Upward-going Muons Due to Downward-going Muons

The muon events fitted into the upward-going region were themselves not a truly pure sample of neutrino-induced events at the most shallow angles: a few of them were actually slightly downward-going atmospheric muons which were misfit into the upward-going region. This problem occurred because the shape of the mountain above the experiment was irregular, and was thin enough in some azimuthal directions that a fairly large number (relative to the greatly reduced flux of upward-going muons just a few degrees lower in altitude angle) of nearly horizontal downward-going muons were able to penetrate the natural radiation shielding afforded by the rock and trigger the detector. Although the directions of the overwhelming majority of these downward-going muons were fitted correctly into the downward-going region, the disparity in the numbers of events just above and just below the horizon was apparently large enough so that the small fraction of events misfit into the upward-going region was not adequately compensated by upward-going events leaking out in the opposite direction. Studies by other collaborators[86] suggested that the degree of the contamination in the upward-going sample was at the level of a few percent in the narrow slice of angular space between the horizon and about 2° below it. The angular spectrum of this background has not been quantified more precisely, however, than by drawing an estimated boundary broad enough to surround the contaminated region loosely, therefore any algorithm which operates on data from this region will have a small bias from the downward-going muon contamination which can't be corrected. For tests such as the two-point correlation function, described in section 5.1, which don't make any numerical statements about upward-going muon flux, the contamination is simply an additional source of background events and its presence does not compromise the validity of the analysis. But in the case of the flux limit analysis of section 5.3, since it attempts to say something explicit about the possible number of astrophysically induced neutrino events in the data, technically one ought to know, in order really to do the analysis properly, how likely it was that each event analyzed was even a neutrino event in the first place. Therefore, in order to avoid dealing with this extra complication in the flux limit analysis, care is taken to avoid using the small fraction of the data set which comes from the contaminated region.

4.3.3 Estimating the Background Using Statistical Resampling

The observed direction of an event in local horizon coordinates is a function of its arrival time at the Earth and its final direction of travel, in equatorial coordinates, at that time of arrival, regardless of how much it may have wandered around or changed its direction during its journey from the point of production. There isn't really much underlying astronomical significance for the atmospheric neutrino events in that final arrival direction, however, because the cosmic ray primaries that produce those neutrinos wander around and change their direction so much, and become so scrambled while in transit, that ultimately they bombard the atmosphere uniformly from all directions at all times. Therefore, using the observables of those events (their times and local horizon coordinates) to calculate their arrival directions in equatorial coordinates can't possibly yield any astronomical insight, because such events can never point back to a source. Neutrinos from modest numbers of strong astrophysical point sources, on the other hand, (assuming for the moment that they exist) would tend, after translation from local horizon back into equatorial coordinates, to be reconstructed into clusters around the directions of the source objects, rather than the isotropic distribution of the atmospheric neutrinos.

This difference can be exploited by using a common statistical device known as a bootstrap[87] in order to help characterize how closely the data set resembles a pure sample of atmospheric neutrinos, as opposed to a sample with a small number of astrophysical neutrinos from a few point sources mixed in. This technique, which is borrowed from the more general formalism of statistical resampling[88], is implemented by choosing event times and local horizon directions, separately and at random, from the original data set, and pairing them to create a fake data set, then calculating the equatorial coordinates of the fake events just as if they were real. This exercise produces a data set whose experimental observables are distributed, by construction, in the same way as the real one, but whose derived variables, the equatorial coordinates, are scrambled. Since the equatorial coordinates of atmospheric neutrinos must be broadly distributed throughout the sky to begin with, shuffling them does little to change the overall qualitative appearance of that component of the data set. If there are any astrophysical neutrinos mixed in with the data, however, their original celestial directions must be closely distributed about the direction of their source, and redistributing them randomly about the sky will completely defocus the point source induced events to random right ascensions, within the band of declination subtended by the source. (See figure 4.5 for a diagram that illustrates the difference.) In the limit that astrophysical neutrinos are a small component of the data set, as they must be here if they are even present at all, the equatorial coordinates of the resampled data sets provide an estimate of the relatively dominant atmospheric neutrino background, against which the point source induced upward-going muons from the original, unresampled data set will appear as anomalously tight clusters of events.

4.4 Angular Resolution

The correct angular scale at which to search for a point source signal is determined by the overall angular resolution for the entire chain of events leading up to the final measurement of a set of track parameters in the detector. That chain, for purposes of this analysis, would involve an astrophysical neutrino traveling in a straight path from the source, interacting in the rock below the detector to produce a muon inclined at some angle to the parent neutrino, the trajectory of which would suffer further angular degradation with respect to the neutrino as it propagated through the rock to the detector, and finally would be measured only as precisely as the intrinsic resolution of the manual track reconstruction process. As a cross-check of other collaborators' earlier work, the portion of the overall angular uncertainty due to manual track fitting has been estimated independently in this document by measuring the distribution, plotted in figure 4.6, of the space angle difference between pairs of independent manual fits. The result is about 0.9° for the plane-projected angular resolution of the manual fit process. The typical space angular resolution, the quantity of real interest, is best estimated by adding this value to itself in quadrature,

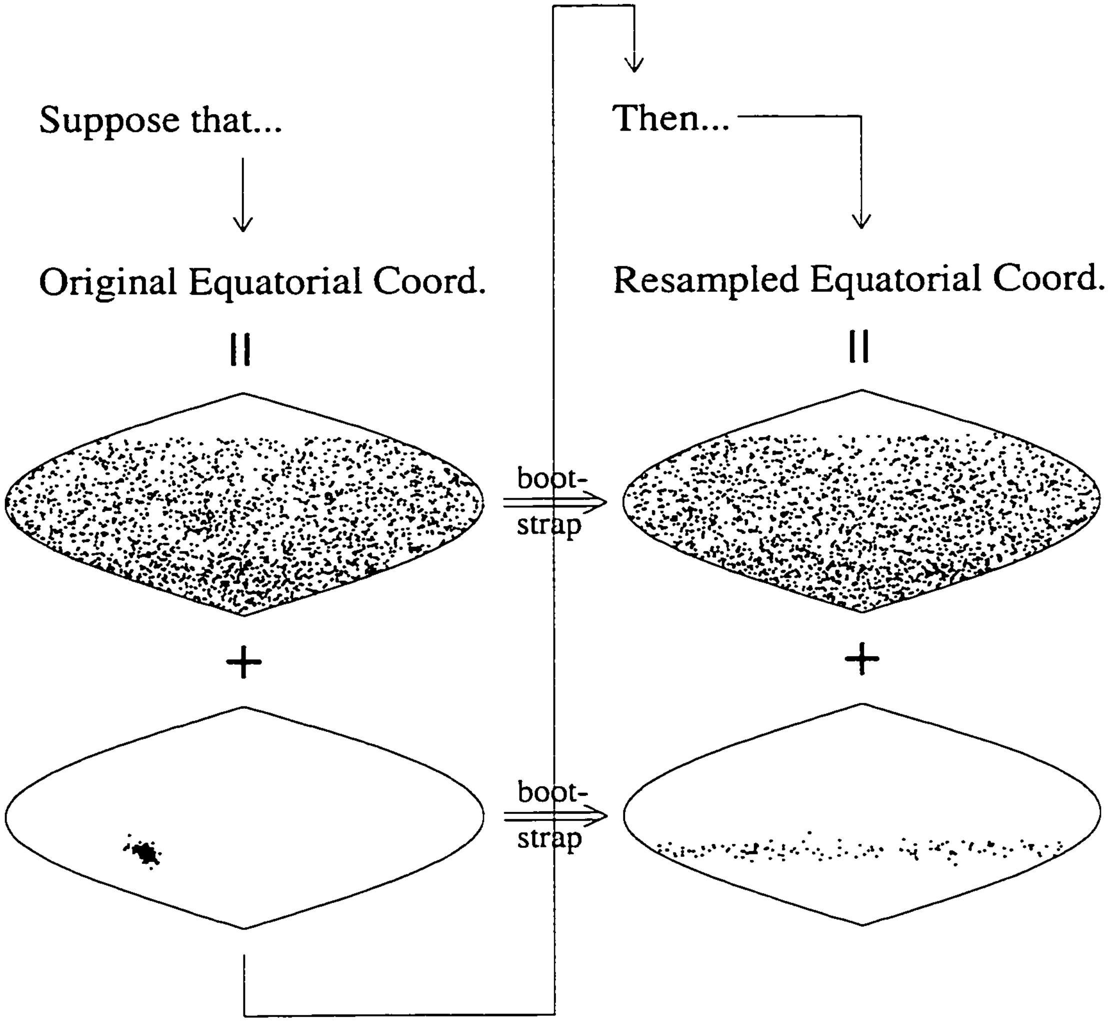

Figure 4.5 The points plotted at the upper left side of this diagram are the reconstructed equatorial coordinates of the real data set of 2037 upward-going muons. Below it are plotted 127 events generated artificially from the direction of the galactic center, with sidereal times chosen to be uniformly random during the period when the artificial source was below the horizon. Times associated with both the real data events and the fake source events were concatenated together to form a single list, and then individual times were paired at random with the local horizon coordinates (not shown here) of those events to generate the two sets of artificial equatorial coordinates shown at right. Some of the analysis exercises of chapter 5 will end up, essentially, comparing a diagram like the one on the upper left hand side to one like that on the upper right, with the operational idea being that, if the apparently isotropic real data actually contains a small admixture of anomalously clustered events like those on the lower left, they will tend to stand out in relief when compared to a data set comprising an isotropic background plus a small admixture of longitudinally defocused events like those on the lower right.

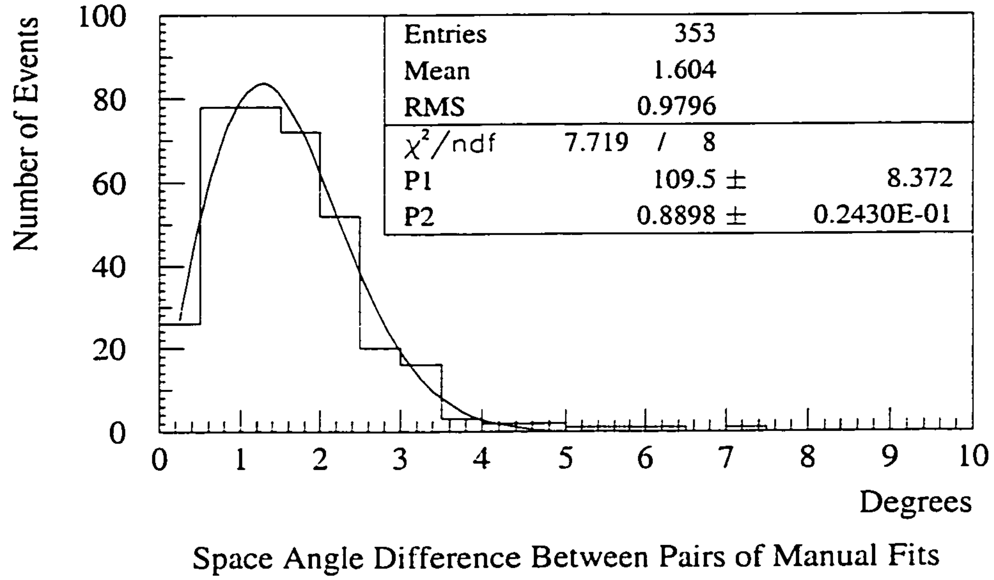

Figure 4.6 The distribution of the space angle between pairs of manually fit upward-going muon track directions for a 353 event subset of the data is plotted here, with a fit to the functional form of equation A.7. The plane-projected RMS angular deviation of a single manual fit from its true value, represented by σ in the notation of appendix A.2, is represented by P2 in this plot, and has a fitted value of 0.89° ± 0.02°. P1 is just a relative normalization parameter between the event distribution and the fit function; its final value is influenced by such details as the number of events in the data sample and the size of the bins into which it is divided, and it is therefore of little fundamental significance. Only track fits separated by less than 5° were used in the fit shown; the small number of outlier events near 6° and 7° were ignored.

since there are two orthogonal planes in which the fitted track direction is free to wander independently. That result, 1.3°, is within the range of the other estimates (0.9°[85] and 1.5°[86]) offered by collaborators for muons which traverse more than 7 m in the detector. The additional contributions to the overall angular resolution, from the muon production angle and from the scattering as the muon propagated through the rock, were not verified independently for the studies presented here. For purposes of this analysis, the value of 4.1° quoted in reference [86] will be accepted as the agreed upon result.

Technically, because the contributions to the resolution from muon production and propagation are energy dependent, while the energies of the individual upward-going muons (and thus, their parent neutrinos) aren't known on an event-by-event basis, the angular resolution for a particular source of neutrinos, whether it be astrophysical or atmospheric, must be estimated by averaging over an assumed energy spectrum, and thus the estimate is somewhat model dependent. Most discussions of neutrino astronomy usually assume that the astrophysical neutrinos, if they exist at all, will be more likely to have a harder spectrum than atmospheric neutrinos, meaning that astrophysical neutrinos will become relatively more prominent at higher energies. If this is so, then 4.1° will be an upper bound on the angular resolution of the astrophysical neutrino events, since the muon production angle and the angular scale of the multiple scattering in the rock tend to decrease as the energy of the parent neutrino increases. Of course, the opposite could hold true: astrophysical neutrinos might have a softer spectrum than atmospheric neutrinos and therefore tend to predominate at lower energies, in which case their overall angular resolution would be somewhat broader than $\sim 4^\circ$, but if that were the case then this analysis, focused as it is on the highest energy sample of neutrino events from the Super-Kamiokande detector, would be examining the wrong data set anyway, and the possibility is not considered further for purposes of this work. While the overall angular resolution for astrophysical neutrinos is, by the above reasoning, unlikely to be any worse than 4.1°, on the other hand, no matter how parallel the trajectory of the initial neutrino and its resulting muon may be, the directional uncertainty can never be lower than the intrinsic resolution of the detector, therefore 1.3° will be taken as a lower bound. The analysis which follows will respect these bounds, and generally will assume that the typical angular resolution of astrophysical neutrinos, if they are present, is likely to be about the same as that of the atmospheric neutrinos: 4°.

Chapter 5

Data Analysis

Figure 5.1 shows the apparent direction of origin, in equatorial coordinates, of the entire 2037 events in the upward-going muon sample. Other investigators in the Super-Kamiokande collaboration have shown[89] that the observed angular distribution may be adequately predicted (in the sense that one can find some set of parameters that lead to a reduced χ^2 of approximately one) by a model in which neutrinos are created by the decay of pions and kaons produced by cosmic rays impinging upon the upper atmosphere of the Earth, and in which the ν_μs and $\overline{\nu}_\mu$s oscillate periodically to some other flavor state which is not as observable as the muon-type state because it does not produce an upward-going muon when it interacts in the rock below the experiment. That analysis shows clearly that atmospherically generated neutrinos constitute the majority of the parent neutrinos of the upward-going muons. But what evidence may be extracted from the data to support or refute the notion of a small astrophysically generated component of the data sample, in addition to the dominant atmospheric contribution? The existence of so many other experiments, described briefly in chapter 1, which have published or hope soon to publish an attempt to answer this question, is testament to the current level of interest by the neutrino community in the answer. This chapter describes some searches for evidence of astrophysical point sources of neutrinos in the upward-going muon data sample from Super-Kamiokande, with results compared whenever possible to those of other experiments.

5.1 The Two-Point Correlation Function

Astronomers need, from time-to-time, to distinguish between cases in which the angular positions of an entire catalog of astrophysical objects are randomly distributed throughout the sky and cases in which most of the objects are randomly distributed throughout the

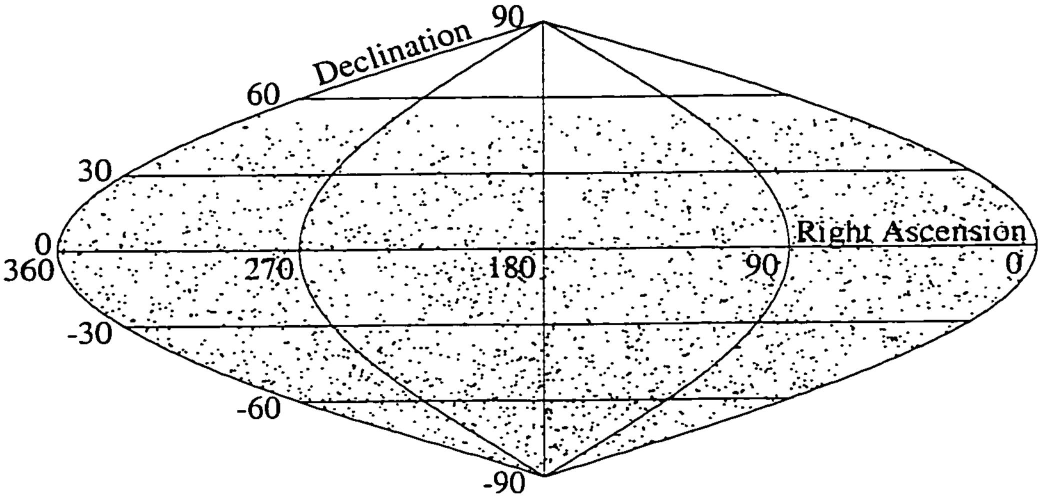

Figure 5.1 These points represent the right ascensions and declinations of 1634 through-going and 403 stopping upward-going muons. The angular distribution of this data set can be adequately explained by assuming that the events all come from neutrinos produced by interactions of cosmic rays in the Earth's atmosphere. The goal of this analysis is to search for evidence of a small additional astrophysical component of the neutrino flux, assumed to be too small to be identified by any but the most sensitive and carefully designed search algorithms, and report that evidence or lack thereof in the most quantitative and succinct terms possible.

sky but a few of the objects are clustered together. Such information can be extracted from the angular positions of a collection of Super-Kamiokande events by borrowing one of the standard analysis tools of conventional photon-based astronomy: the two-point correlation function[90].

This function, $w(\theta)$, is defined by counting the number of pairs of events in the data sample with space angle between $\theta - \delta\theta/2$ and $\theta + \delta\theta/2$, and then quantifying the degree to which that number departs from the value that one would expect if all of the events were drawn randomly from a particular background distribution. This background, $b(\alpha, \delta)$, would be proportional in a conventional astronomy analysis to the solid angle and observation time of an area of sky, but in this case comes from atmospheric neutrinos, which do not arrive at the detector in strict proportion to the exposure in a given celestial direction, due to various physical effects which exert their influence after a cosmic ray primary interacts in the atmosphere.

5.1.1 Derivation of the Function and of its Uncertainty

To make the definition mathematically explicit, suppose that there are N events in the data sample. Imagine picking one event, the ith event, and drawing a half ring around its right ascension and declination, as shown in figure 5.2, which is bounded by an inner radius of $\theta - \delta\theta/2$ and an outer radius of $\theta + \delta\theta/2$, and is oriented in a specific direction: for example, imagine that it lies entirely to one side of the great circle defined by the north celestial pole and the event's right ascension. The background exposure of each of the two half rings will be notated as $db_{1,i}(\theta)$ and $db_{2,i}(\theta)$, their sum as $db_i(\theta)$, and the total integrated background exposure, over all regions of the sky and over all times, will be $B = \int_T \iint_{4\pi} b(\alpha, \delta, T)\, d\alpha\, d\delta\, dT$. For brevity, the θ dependence which appears in some of the above quantities and in others soon to be introduced will be left implicit in most of the discussion that follows below.

Any single event which is drawn at random from the background distribution has a probability $db_{1,i}/B$ of ending up inside of the first of the two half rings and a probability of $1 - db_{1,i}/B$ of ending up outside of it; the probability for observing $n_{1,i}$ events in one of the half rings around the ith event is therefore governed by a binomial distribution:

$$P(n_{1,i}; N-1, db_{1,i}) = \frac{(N-1)!}{n_{1,i}![(N-1) - n_{1,i}]!} \left(\frac{db_{1,i}}{B}\right)^{n_{1,i}} \left(1 - \frac{db_{1,i}}{B}\right)^{[(N-1)-n_{1,i}]} . \tag{5.1}$$

In the limit that $db_{1,i}/B \to 0$, that $N - 1 \approx N$, and that $n_{1,i} \ll N$, this binomial distribution is well approximated by a Poisson distribution,

$$P(n_{1,i}; N-1, db_{1,i}) = \frac{m_{1,i}^{n_{1,i}}}{n_{1,i}!} e^{-m_{1,i}}, \tag{5.2}$$

with the mean expected number of events measured in the ith half ring defined by the newly introduced quantity $m_{1,i} = N(db_{1,i}/B)$. (In practice, in this analysis, it will turn out to be the $m_{1,i}$ and $m_{2,i}$ that are of real interest, since they are used to gauge whether statistically anomalous values of the $n_{1,i}$ and $n_{2,i}$ have been observed, and the m_is will be estimated by statistical resampling, rather than bothering with trying to determine the actual values of

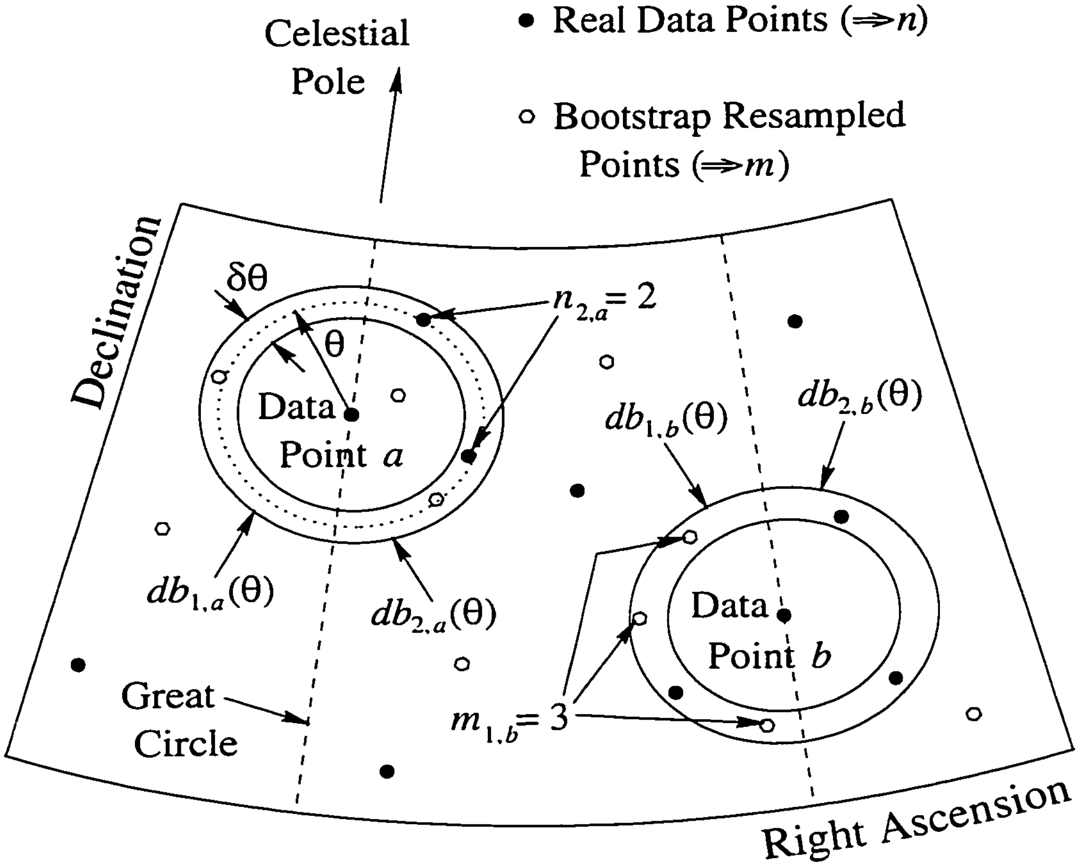

Figure 5.2 This shows a projection of a small portion of the celestial sphere, and each real data point represents the arrival direction of one upward-going muon. To carry out the analysis, rings of width $\delta\theta$ are drawn about all data points over the whole range of θ, although only rings for two data points, at one value of θ, are actually shown here. As explained in the text, $m_i = N(db_i/B)$, which one needs to know in order to gauge whether or not statistically anomalous values of n_i have been observed, is estimated using statistical resampling.

$b(\alpha, \delta, T)$, db_i, and B: these last three quantities are therefore of little further interest to the discussion.) Equation 5.2 can be simplified further by recognizing that the convolution of two Poisson distributions is also a Poisson distribution (see appendix A.3), and thus the probability distribution for finding a total of n_1 events in all of the half rings of radius θ drawn about all N of the events is also Poissonian,

$$P(n_1; N, db_{1,1}, db_{1,2}, \ldots, db_{1,N}) = \frac{m_1^n}{n_1!}e^{-m_1} \tag{5.3}$$

with the global mean for the entire data set given by

$$m_1 = \sum_{i=1}^{N} m_{1,i} = \sum_{i=1}^{N} N\frac{db_{1,i}}{B}. \tag{5.4}$$

The quantity n_1 is just a count of the number of pairs of events separated by an angle between $\theta - \delta\theta/2$ and $\theta + \delta\theta/2$, but because the problem was formulated using half rings, one conveniently avoids the double counting which would have occurred in the sum over all events if full rings had been used. The pairs could just as well have been counted using the second, oppositely oriented set of half-rings, but because the final answer must come out the same, one may drop the subscript and simply write the number of pairs as n, where $n = n_1 = n_2 = (n_1 + n_2)/2$. Such logic does not apply to m_1 and m_2; since these quantities essentially count exposure, which may vary across the sky, it means that in general $m_1 \neq m_2$. Alternatively, thinking in terms of what the bootstrap is doing to the data, this is equivalent to saying that the scrambling does not preserve the double-counting symmetry of the real data, since the only reason that the symmetry exists in the first place is that the rings are all drawn with the real data points at the center, rather than being centered on the bootstrap points, or on none at all. That m_1 and m_2 are not constrained to be identical implies that if one defines

$$m = \frac{m_1 + m_2}{2}, \tag{5.5}$$

it may be used as a more precise estimator of $\langle n \rangle$ than either m_1 or m_2 alone.

At this point, $w(\theta)$, the two-point correlation function, may now at last be defined, using the equation

$$n = m(1 + w),$$

or alternatively,

$$w(\theta) = \frac{n(\theta)}{m(\theta)} - 1. \tag{5.6}$$

The value of $w(\theta)$ which one calculates from the data using counted values of $n(\theta)$ and $m(\theta)$ measures the degree to which the data set seems really to have been drawn from the assumed background distribution in the sense that if all of the events in the data sample are indeed drawn from that background, without any anomalous clustering, then $w(\theta)$ should be measured close to 0, within errors, over all θ, because the $n \approx m$ for all θ.

In a conventional astronomy analysis, the uncertainty in w comes from the statistical uncertainty in n, plus whatever uncertainty is introduced into m by $b(\alpha, \delta, T)$, which would usually be measured by keeping track, in terms of solid angle and time, of how much the telescope had been pointed at various regions of the sky. In this analysis, however, the calculation of $b(\alpha, \delta, T)$ will be ignored, essentially, and $m(\theta)$ will instead be equated with the total number of bootstrap events that fall between $\theta - \delta\theta/2$ and $\theta + \delta\theta/2$ of a real event. So, forgetting about the precise values of db_i and B, and using the definition of m from equation 5.5, one finds that, because the sum $m_1 + m_2$ is a simple, counted, integer quantity, m also contributes an additional statistical uncertainty,

$$\frac{\sigma_m}{m} = \frac{\left(\sqrt{m_1 + m_2}\right) / 2}{\left(m_1 + m_2\right) / 2} = \frac{1}{\sqrt{m_1 + m_2}} = \frac{1}{\sqrt{2m}}, \tag{5.7}$$

to $w(\theta)$ as well. It should be emphasized that, strictly speaking, for purposes of this analysis, the only case that will be considered is the one in which resampling is performed exactly once, i.e., just a single alternative fake data set, of the same size as the original, is generated in order to estimate the background. In principle, there is a small advantage to be gained by iterating the resampling process several times, however, including this possibility complicates the presentation somewhat, while failing to deliver very much added benefit to the analysis, since the total uncertainty on $w(\theta)$ is fundamentally limited by the uncertainty in n, (calculated using equation 5.9 below) which contributes most of the overall uncertainty anyway. Therefore, for the two-point correlation analysis at least, the possibility of iterative resampling will not be considered further in this work.

Because n is drawn, assuming that the null hypothesis of no extra clustering is true, from a Poisson distribution of mean m, its uncertainty is just the width of the distribution:

$$\sigma_n = \sqrt{m} \tag{5.8}$$

In this analysis, m was generally large enough so that one would expect the distribution of n to be narrowly peaked about m, i.e., $n(\theta)$ should be within a few percent of $m(\theta)$ at all θ, assuming the null hypothesis is true. Therefore, approximating $n \approx m$, one may justifiably write, for the null hypothesis, that

$$\frac{\sigma_n}{n} \approx \frac{\sigma_n}{m} = \frac{1}{\sqrt{m}}. \tag{5.9}$$

Recognizing that the uncertainty in the two-point correlation function, $\sigma_{w(\theta)}$, is just the uncertainty in n/m, and also that n/m itself should always be approximately 1 in the null hypothesis, the relative errors of equations 5.7 and 5.9 can be added in quadrature to find that

$$\sigma_{w(\theta)} = \sigma_{(n/m)} = \left(\frac{n}{m}\right)\sqrt{\left(\frac{\sigma_n}{n}\right)^2 + \left(\frac{\sigma_m}{m}\right)^2} \approx (1) \cdot \sqrt{\frac{3}{2m}}. \tag{5.10}$$

The error derived in equation 5.10 gives the magnitude of the fluctuations that one would expect in $w(\theta)$ if the experiment were repeated a large number of times, and therefore one may judge the significance of a non-zero $w(\theta)$ by calculating $w(\theta)/\sigma_{w(\theta)}$.

5.1.2 Results and Interpretation

Although the two-point correlation function is defined over the entire interval $0° < \theta < 180°$, a significant non-zero result for $w(\theta)$ would be expected only at very small angles, at least for the special case of non-random clustering centered, specifically, about a small number of points. The angular scale at which one would expect to see an effect, if any were present, is determined by the angular resolution of the entire detection process, estimated in section 4.4. Figure 5.3 shows $w(\theta)$ plotted in increments of 2°, with error bars calculated according to equation 5.10, for the real data set of 2037 neutrino induced through-going muons, as well as for a number of different partially fake data sets, with point sources artificially superimposed on top of the 2037 real events, in order to illustrate graphically the efficacy of this algorithm for identifying cases of anomalous clustering. Additionally, in order to help give the reader some physical intuition about how the algorithm's capacity to identify anomalously clustered data sets compares to that of a human, the equatorial coordinates of the events from one of those fake data sets are plotted in figure 5.4.

Figure 5.5 shows the distribution of $w(\theta)$ over the entire 180° range of θ, renormalized to units of standard deviation for each annular bin, for the real data set of 2037 events. If the error bars for each annular bin have been estimated correctly, then $w(\theta)$, when renormalized against them in this way, should be normally distributed about zero with a characteristic width of one (assuming that the null hypothesis is true). The closeness of the width of the fitted normal distribution in the plot to this theoretically expected value confirms that

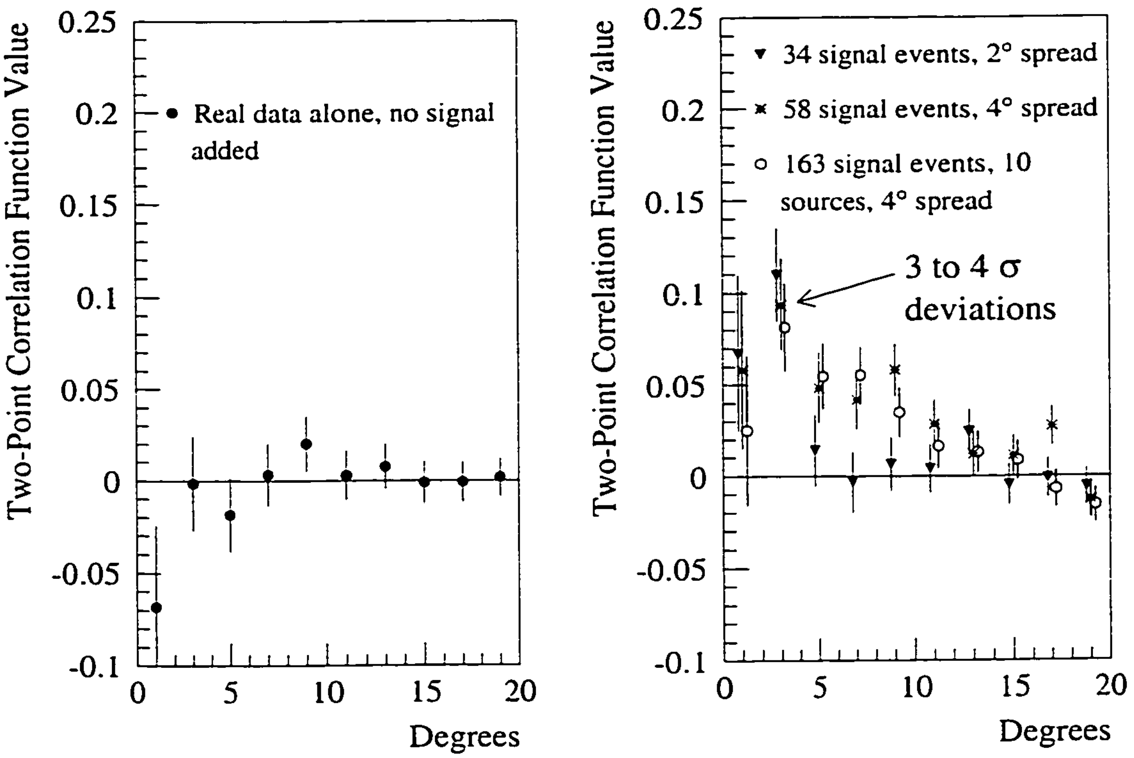

a.) b.)

Figure 5.3 **a.)** The two-point correlation function, $w(\theta)$, plotted for the real data set of 2037 upward-going muons shows no significant upward deviation at small angles, nor at any angle for that matter. **b.)** The correctness of the algorithm which was used to make the plot on the left was checked in order to see if inserting a point source artificially into the data could induce the search algorithm to return a non-null result. Several fake data sets were generated, with signal sizes chosen to be large enough to result in roughly a 3 to 4 sigma deviation in at least one of the annular bins near the angular spread of the source. The first data set consisted of 34 fake signal events with angles distributed normally on a characteristic scale of 2° about the galactic center, superimposed upon the real experimental data. The second data set consisted of 58 signal events from the same general direction, but distributed with a width of 4°, superimposed in the same way. The signal for the third data set (likewise superimposed) was 163 events drawn from 10 normal distributions of width 4° each, but with mean directions chosen at random, with the stipulation that the source position must have been at a low enough declination to be visible by Super-Kamiokande. The equatorial coordinates for the 163 events used in this last case are plotted in figure 5.4. Evidently, the algorithm correctly produces non-null consistent deviations at small angles if clustering is present.

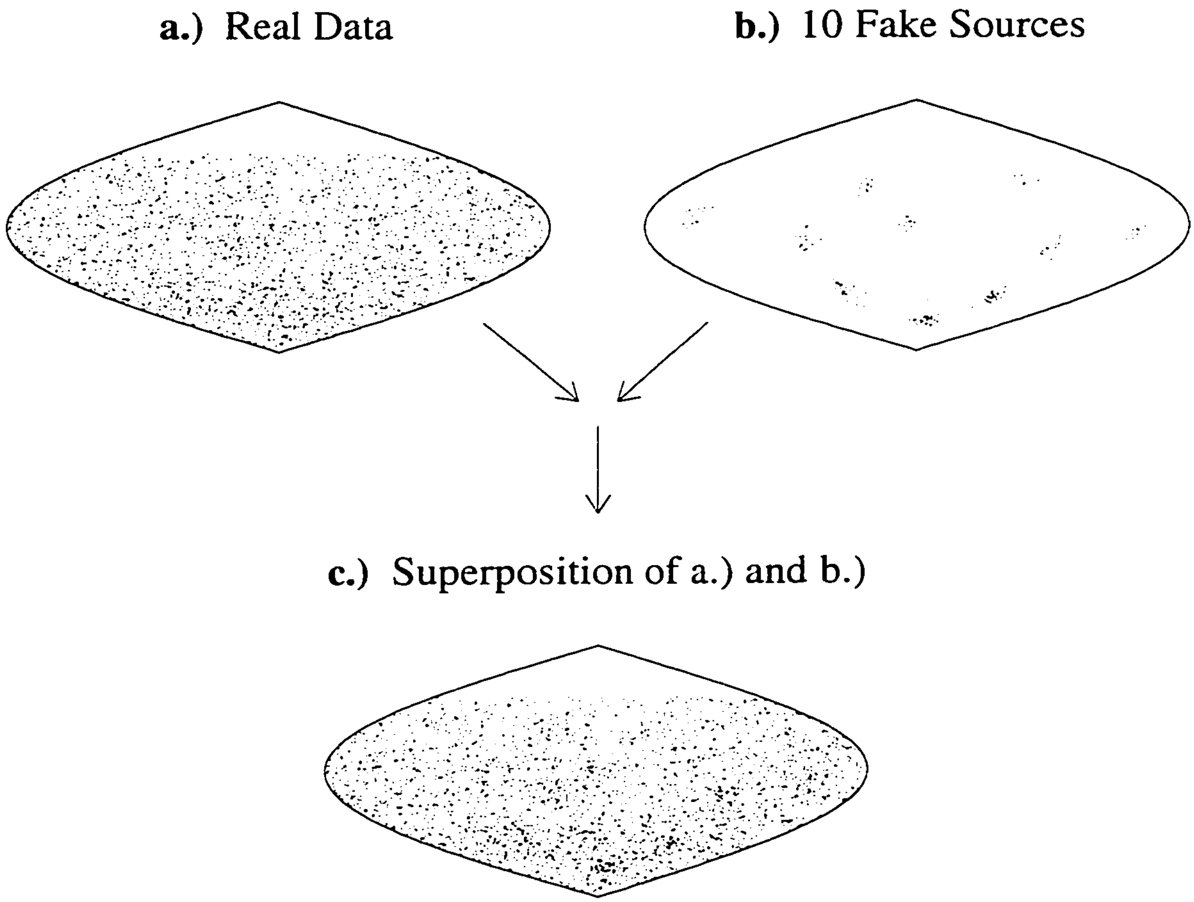

a.) Real Data

b.) 10 Fake Sources

c.) Superposition of a.) and b.)

Figure 5.4 **a.)** This shows the equatorial coordinates of 2037 real data events, by themselves, without any signal added. **b.)** The equatorial coordinates of 163 fake signal events, distributed in 10 smaller clusters, each of which represents a point source, are plotted here. Each fake source is allotted the same number of events initially, with random sidereal times chosen uniformly, but only those events which would have actually been seen as upward-going by the detector are plotted here. Thus, while the sources are chosen to have equal strength, those in the southern sky generally contribute more events to the data set, since the exposure tends to be greatest at those declinations. **c.)** This shows the superposition of the real data set with the artificially generated point source signal events. Figure 5.3b demonstrates that the two-point correlation function formalism can recognize this case correctly as one of anomalous clustering, although it may be difficult for a human being to do so.

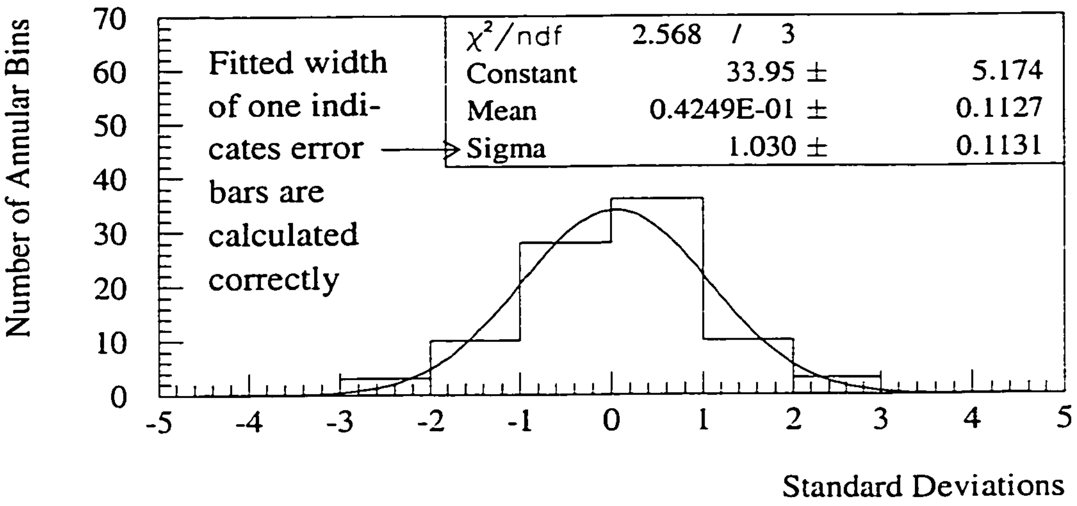

Figure 5.5 This shows the distribution of the two-point correlation function over all angles, for the real data sample of 2037 data events. The excursion of $w(\theta)$ away from zero, on the horizontal axis, is plotted after being renormalized at each annular bin against the uncertainty for that bin. The fitted width of the distribution demonstrates that the error bars used for the renormalization have been calculated correctly.

equation 5.10 does indeed calculate those error bars correctly. As can be seen in figure 5.3.b, any clustering effects due to point sources are overwhelmingly likely to be found in the first few annular bins above 0°. The two-point correlation function for real data, figure 5.3.a, does not deviate substantially from its expected value of zero at any point along its entire range, and in particular it does not do so at small angles, therefore, this test returns a null result: there do not appear to be any astrophysical point sources of neutrinos hidden in the data, at least on the basis of this simple test.

However, a fair treatment of the data requires full disclosure of the inadequacies of any algorithm used to analyze it. In this analysis, the statistical resampling method used to estimate the background has a flaw worth noting: it does a poor job of shuffling the directions of events near the celestial poles, because the process described in section 4.3.3 only randomizes the right ascension of the resampled data set, while the declination always remains the same as in the original. Thus, any signal events near the poles are much more likely to end up resampled back into the signal region than events near the celestial equator.

Since the background estimation method relies on the premise that the resampling process will strongly tend to throw events out of the signal region into a random position in some much larger patch of sky, it will tend to yield artificially inflated background estimates for astrophysical signals near the poles, where the premise does not hold true. Unfortunately, as depicted in figure 4.4, the south celestial pole is the region of the sky with the greatest exposure. Therefore, the two-point correlation function analysis, and indeed, any other search algorithm which relies for estimation of atmospheric background on a resampling technique like the one used here, gives its weakest results in exactly the same place where the experimental apparatus is most sensitive.

5.2 Tiled Solid Angle All-Sky Survey

Apart from considerations about the efficacy of statistical resampling at low southern declinations, the two-point correlation function itself has at least two drawbacks as well: first, it returns only a global measurement of the amount of clustering in the data set, without giving any detail about its variation as a function of direction, and second, the geometric construction of hundreds of overlapping rings of different sizes, upon which the method relies in order to produce its results, does not lend itself well to interpretation by human physical intuition, even though the basis of the test is fundamentally quite simple. Therefore, some other type of test which does not suffer from these deficiencies is still desirable. Such an alternative, inspired by a similar tiled all-sky analysis originally performed on a data set of downward-going muons acquired by the MACRO[91] experiment, but modified somewhat in order to account for the necessarily much smaller statistical size of a neutrino data sample, is described in this section.

5.2.1 Description of the Search Algorithm

The basic structure of the algorithm is simple: to perform the search, the sky is tiled into angular bins, the number of upward-going muons in each bin is counted and compared to an average background estimate, and improbably highly populated bins, if any were found, would be taken as evidence of anomalous clustering, implying neutrino point sources somewhere near the directions of those bins. The optimal bin size for a point source discovery would be one near the angular spread of the signal events coming from the source: selecting a much smaller size would tend to dilute the significance of the bin closest to the source by spreading the signal events thinly over many adjacent bins, while selecting a much larger size would tend to dilute the significance by juxtaposing an unnecessarily large number of background events against the signal. Because more energetic incoming neutrinos will tend to produce more collimated outgoing muons, the angular spread of a point source would depend on its energy spectrum, but for purposes of this analysis, an angular resolution of about 4°, commensurate with that of the atmospheric neutrinos, will be assumed. The bin edges are chosen so that each bin will measure, as nearly as possible, $8^\circ \times 8^\circ$ in space angle, and an integral number of them will fit within the declination limits of -90° and 53.58° which are accessible to upward-going muons at the latitude of the experimental site.

After the boundaries of the angular bins are drawn, and the numbers of upward-going muon track directions which fall into each are counted, the next task is to estimate an expected background rate for each, so that one can evaluate whether or not any of the bins show an improbably high upward fluctuation in their total numbers of events. As in the case of the two-point correlation function, the bootstrap statistical resampling method described in section 4.3.3 is used here, although this time it is iterated 1000 times in order to bring the additional statistical error incurred by the resampling process down to a negligible level, as compared to the inherent statistical error of the data set itself. Now consider the following: whatever the true average long-term neutrino flux through each bin is, whether it is due only to the atmospheric neutrino background or also includes a point source signal, the observed muons are certainly Poisson distributed with respect to it. But, as emphasized in section 4.3.3, estimating the background with a bootstrap only produces valid results if the event times and horizon coordinates are completely uncorrelated with each other, which would not at all be the case for point source events arriving from a specific direction in equatorial coordinates. This fact forms the basis of the search algorithm: if there are indeed point sources in the data, the resampling process will reconstruct their directions to completely random right ascensions, and the background estimate near the direction of the source will be artificially low, causing the observed number of real data events to stand out as an anomalous upward fluctuation.

To find such anomalous fluctuations if they exist, one must calculate something akin to a percentile ranking for each bin and then show that some of the bins have remarkably improbable rankings, given the total number of bins in the data set. It is here, however, that one encounters an additional complicating factor: discrete probability distributions offer two obvious choices of statistic to be used for this purpose: the percentile ranking could either be the probability of observing a greater number of events than was actually observed for a bin of estimated mean background m (i.e., $P(n > n_{obs}; m)$), or it could be the probability of observing a greater or equal number of events (i.e., $P(n \geq n_{obs}; m)$). This distinction is not trivial, especially for m and n_{obs} near zero, where the two different probability choices may differ by 10s of percentage points, and even though solid angle bins with m and n_{obs} near zero actually represent downward rather then upward fluctuations, nevertheless, they are still important too, for ideally, one would like to plot the entire distribution of the percentile rankings (including downward as well as upward fluctuations) over all bins and show that the assumption of Poisson statistics is self-consistent throughout: for example, about 90% of the events ought to fall below the 90th percentile, and 80% below the 80th, etc.

An acceptable method for calculating the exact percentile ranking can be reasoned out by imagining a simpler case in which a large number of solid angle bins all have about the same mean background expectation m, and the values of $n_{obs} = 0, 1, 2, 3, \ldots, \infty$ for the various bins are Poisson distributed, as would be expected. In such a case, for a given value of n_{obs}, if one labels $P(n > n_{obs}; m) = a_{obs}$ and $P(n = n_{obs}; m) = b_{obs}$ for brevity, then one expects that overall a fraction b_{obs} bins will end up with $n = n_{obs}$, a fraction a_{obs} will have $n > n_{obs}$, and a fraction $1 - a_{obs} - b_{obs}$ will have $n < n_{obs}$, for that is the implication of saying that n is Poisson distributed with mean m among the bins. Counting up in percentile space from very improbably high values of n to improbably low values, the bins with $n = n_{obs}$ may be thought of, collectively, as filling the entire percentile space between a_{obs} and $a_{obs} + b_{obs}$, but since the only information available about the bins is the value of n, which is identical for all of them, there is no way to rank them more precisely within that space. Thus, it is perfectly valid to rank them in random order, and that will also be the adopted solution in the context of the larger problem with many bins of different, non-identical m values:

both $P(n > n_{obs}; m))$ and $P(n \geq n_{obs}; m))$ will be calculated for each bin, and a percentile ranking will be chosen at random, using a pseudo-random number generator, in the interval between the two. In this way, the percentile ranking is allowed to assume the smoothly varying character associated with numbers drawn from a continuous distribution, yet the artificial construction of this continuous quality does not come at the cost of having to accept together with it the injection of any artificial systematic bias.

5.2.2 Results and Interpretation

Figure 5.6 shows the end result of this process. The histogram in the lower half of the plot shows the number of bins versus the log of the probability associated with the percentile ranking described above. The least probable angular bin from the map at the top, shown farthest to the left in the histogram, which holds 8 upward-going muons with an expected background of about 2.8, has a calculated probability ranking of about 4.4×10^{-3}, meaning that slightly fewer than 1 out of every 200 bins ought to show an upward fluctuation at least as prominent. Since there were 598 bins in the map, such a fluctuation is entirely unremarkable, and is to be taken as evidence of a null result; this test apparently finds no astrophysical neutrino point sources. It is gratifying, furthermore, that the percentile ranking (plotted as $\log(P)$ along the horizontal axis of the histogram) seems to track smoothly with the numbers of angular bins actually given that ranking: the two quantities can be fitted to a line of slope one. Such a good fit is to be taken as evidence of the correctness of the assumption that the numbers of events are Poisson distributed with respect to the background estimates as calculated by the resampling process.

With any statistical investigation that returns a null result, it's tempting to wonder whether that result is genuine or whether it is due simply to some fundamental shortcoming of the test, such that it would be incapable of ever returning a non-null result no matter what the data looked like. Figures 5.7 and 5.8 dispel this concern by showing the output of the algorithm when 16 and 219 fake events, respectively, clustered near the galactic center, are added to the real data sample. The figures demonstrate unambiguously that a sufficiently highly concentrated clustering of upward-going muon directions in some region of the sky

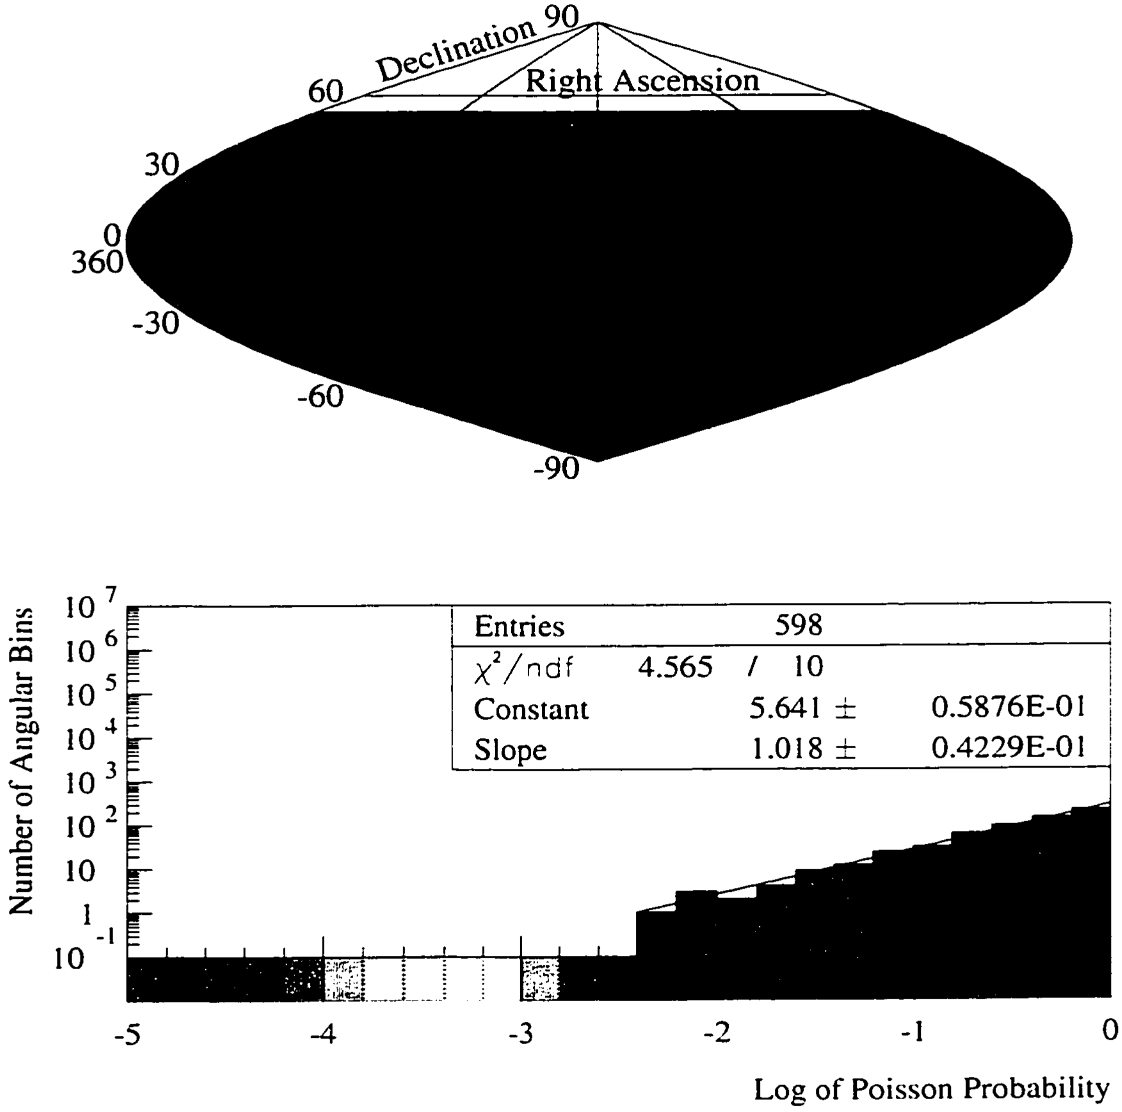

Figure 5.6 The map in the upper half of this figure, in equatorial coordinates, shows 598 bins measuring roughly 8° x 8° in space angle, in the declination range $-90° < \delta < 53.58°$. The colors indicate the integrated likelihood, assuming the observed events in each bin are Poisson distributed with a mean value as calculated via statistical resampling, of seeing a more upward fluctuation than actually observed. As indicated by the fitted line's slope of exactly one, the numbers of angular bins actually observed falling into each cumulative probability bin of the histogram do indeed seem to be distributed with the frequencies that were predicted, and furthermore there are no particularly unlikely upward fluctuations. Thus the assumption of Poisson statistics apparently holds true, and the lack of outliers indicates that there don't seem to be any point sources hidden in the data.

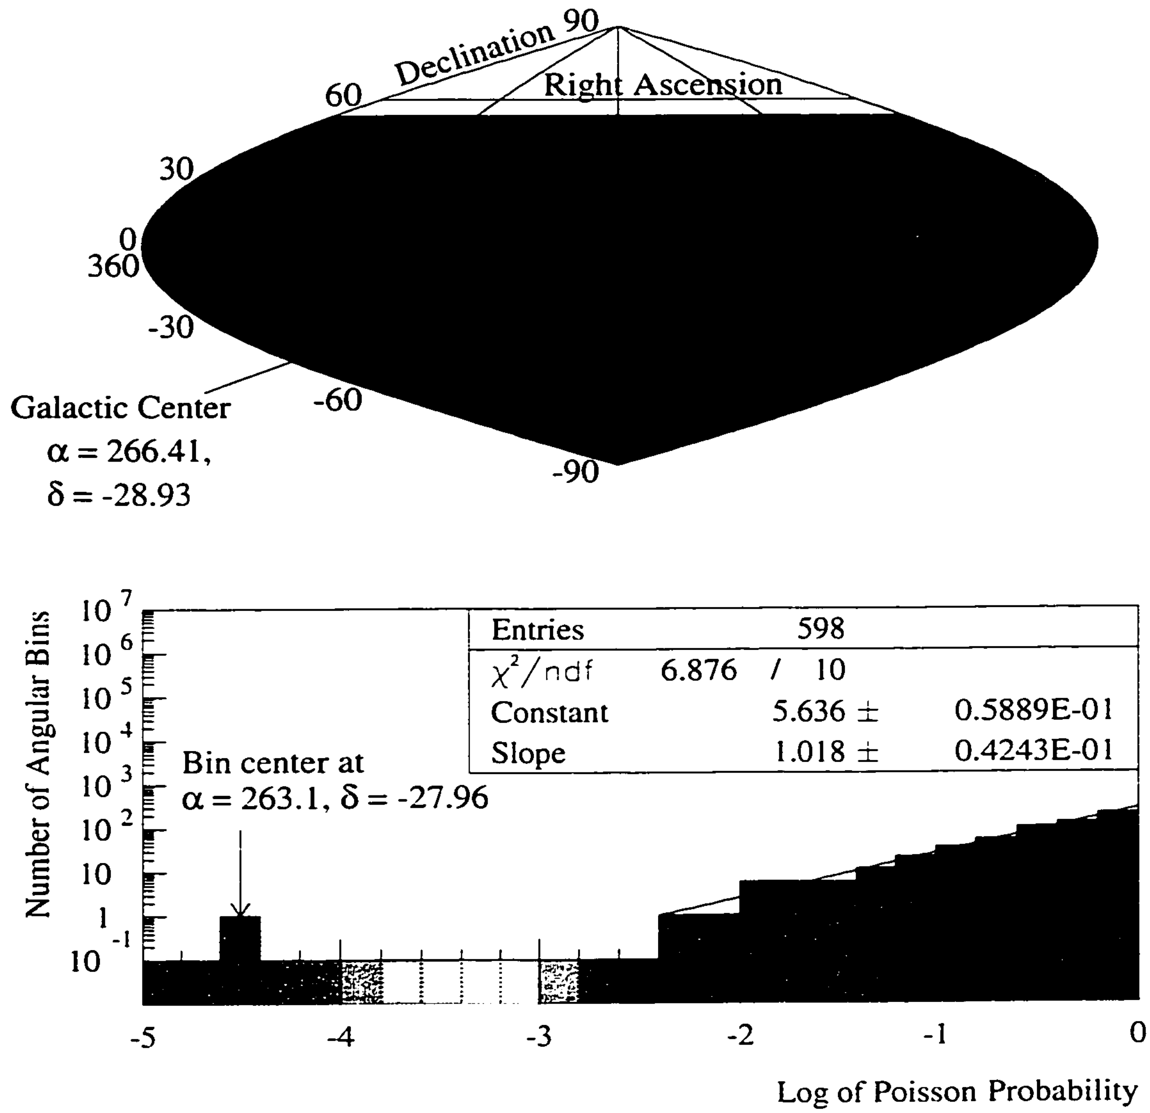

Figure 5.7 This figure shows the same data set as in figure 5.6, treated in the same way, but with an extra 16 fake signal events, Gaussian distributed with an RMS of 4° about the galactic center, added in. The probability of having observed, simply by chance, at least as many events as were counted, including all the fake ones that were added, in the bin superimposed on the galactic center, is only about one in 10,000, still relatively unlikely even taking into account the trials factor of 598 angular bins. This plot is typical of the type of response that it was hoped the data might have given to this type of search method. That no such response was seen in the real data alone, in figure 5.6, reinforces the notion that there are no point sources visible in the data.

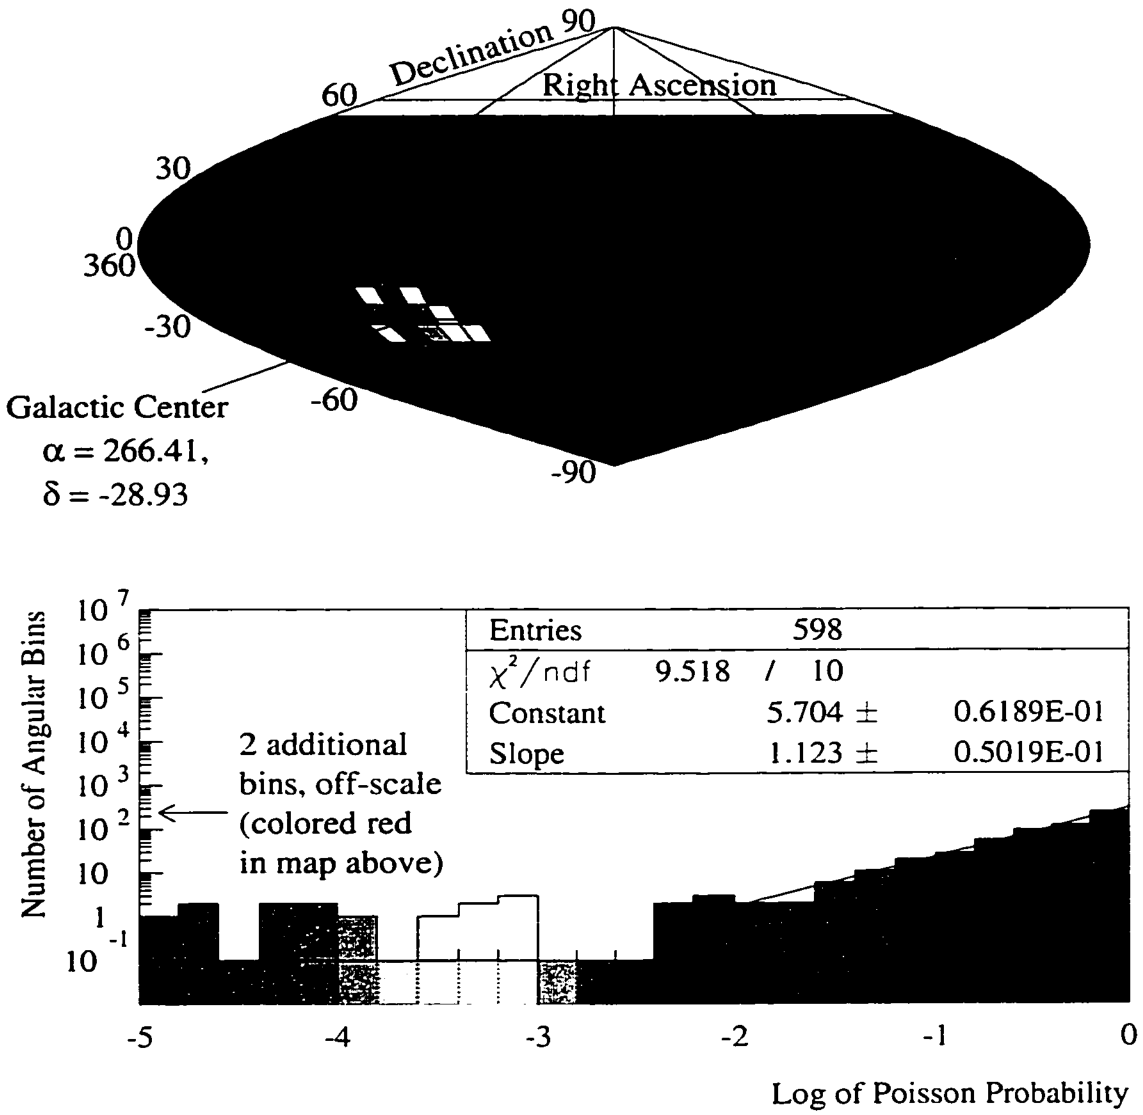

Figure 5.8 This figure is much like figure 5.7, but with 219 fake signal events added to the data set, and distributed evenly about the galactic center over a circle of 20° radius rather than in a Gaussian pattern. This slight modification in distribution helps to make the point of the figure (that the search algorithm gives a reassuringly dramatic response to extremely powerful signals) more obvious to the eye by allowing several bins to fluctuate upward by a modest amount, rather than allowing only a few of the bins, nearest the source, to receive most of the events and fluctuate very far off the probability scale along the horizontal axis of the histogram.

will cause the percentile, or probability values, of the angular bins near the centroid of the clustering to register as improbable upward fluctuations. The implication of this is reassuring: with a suitable choice of threshold of bin improbability, the search algorithm would indeed be able to find astrophysical neutrino point sources that were inducing upward-going muons, as long as they emitted their neutrinos brilliantly enough and for long enough.

5.3 Upper Limits on Point Source Signal Strength

Several other detectors besides Super-Kamiokande have been built which have had the ability to detect upward-going muons. Many of them[60][61][62][63][64][65] have performed analyses to estimate upper limits on the possible strength of neutrino signals which might be coming from the direction of various astrophysical objects of interest. In fact, there are by now enough confidence limit analyses published by other investigators working in this sub-field that they may at this point arguably be viewed as providing a kind of overall performance benchmark for each telescope experiment. In light of this, a thesis on astrophysical neutrino point sources could not be considered complete if it did not include such a study as well. Some algorithms for performing a confidence limit analysis, based on both Classical and Bayesian statistical methods, are developed here.

The analysis in both cases is readily divided into three parts: first, an off-source area is observed in order to estimate the background rate of atmospheric neutrinos; second, the background estimate is combined with on-source data to re-estimate the combined rate of background and signal together (assuming, of course, that such a signal even exists); and third, a probability density describing the combined predictions for both the background and signal is integrated to calculate a coverage probability for the signal from the on-source direction. The analysis is complicated by the fact that the background events are atmospheric neutrinos, therefore their flux varies with direction, and one cannot arbitrarily choose just any area of the sky to serve as an off-source region. The proper way to choose the off-source region is demonstrated in figure 5.9: the on-source region is defined first as an area of solid angle centered on a putative point source of neutrinos, which cycles through a path in local horizon coordinate space as the Earth rotates about its axis, and the off-source

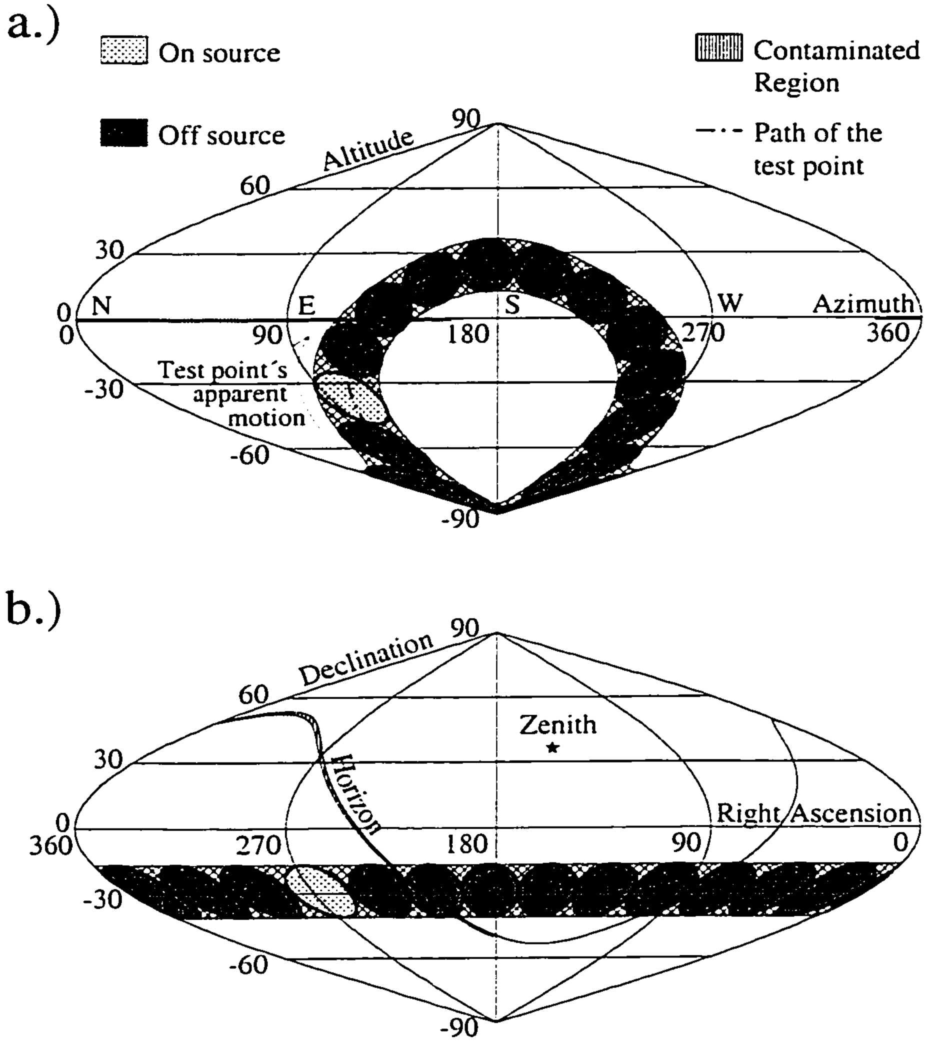

Figure 5.9 **a.)** The on-source region is defined as a patch of sky in the upward-going (altitude angle < 0) region of the plot, excluding the portion contaminated by downward-going muons, extending 12° in all directions from the test point. The test point cycles along the dash-dotted line in local horizon coordinates once per sidereal day. The off-source region is defined as a series of regions of the same size as the on-source region, lined up end-to-end along that same apparent path, but observed, like the on-source region, only at times when they are below the horizon. Thus, the off-source regions pictured above the horizon technically are not off-source at the time to which this diagram corresponds: they are really temporarily null regions, which are labeled off-source because they will become off-source a short time later, as soon as the rotation of the Earth brings them below the horizon again. **b.)** The same on-source and off-source regions are also shown in equatorial coordinates, where they have a simpler form. The test point used in this example is the galactic center.

region is defined afterwards as a series of regions of the same angular size as the on-source region, filling up as much as possible of the remaining angular region below the horizon, along that same path.

Although it is easier to think about the variation of the atmospheric neutrino flux in local horizon coordinates, and therefore also easier in that system to think about how properly to choose off-source regions, in practice it is slightly simpler to do the analysis in equatorial coordinates; figure 5.9.a is therefore intended mostly for use as a pedagogical tool, while figure 5.9.b more accurately represents how the analysis was really done. The added simplicity that one enjoys in equatorial coordinates arises because the on-source and off-source angular cones are all fixed in that system at stationary positions in the same band of declination, and are spaced out only by differing right ascensions, while it is instead the 2π steradians of upward-going solid angle which sweep across the coordinate system once per sidereal day. In practice, it turns out to be much more convenient to convert 2037 pairs of local horizon coordinates and event times to equatorial coordinates, and check to see if each event direction falls within any of the fixed, constant boundaries as they appear in that system, rather than try to redraw the boundaries of all of those circles as they appear in local horizon coordinates at 2037 different event times, and attempt to perform the same check.

In order to obtain an unbiased estimate of the atmospheric neutrino background from the off-source region, the ratio of time spent on-source versus off-source must be the same in all directions in the upward-going region, because the rate of atmospheric neutrino events varies with direction. This implies that the data used in the analysis must be drawn evenly from all sidereal times, or else the on-source and off-source regions will tend to hover preferentially in certain regions of horizon coordinate space. Unfortunately, practical considerations dictated that the detector had to be turned off periodically, therefore the raw data was naturally somewhat unevenly distributed throughout the sidereal day. Figure 5.10 demonstrates that this unevenness was small, only about ±1% over the 537.8 live day subset of the data analyzed in the plot, and of course one must bear in mind that the live time must be even more smoothly distributed about the sidereal day when longer data acquisition periods are included in the average. The statistical error associated with counting on-source and

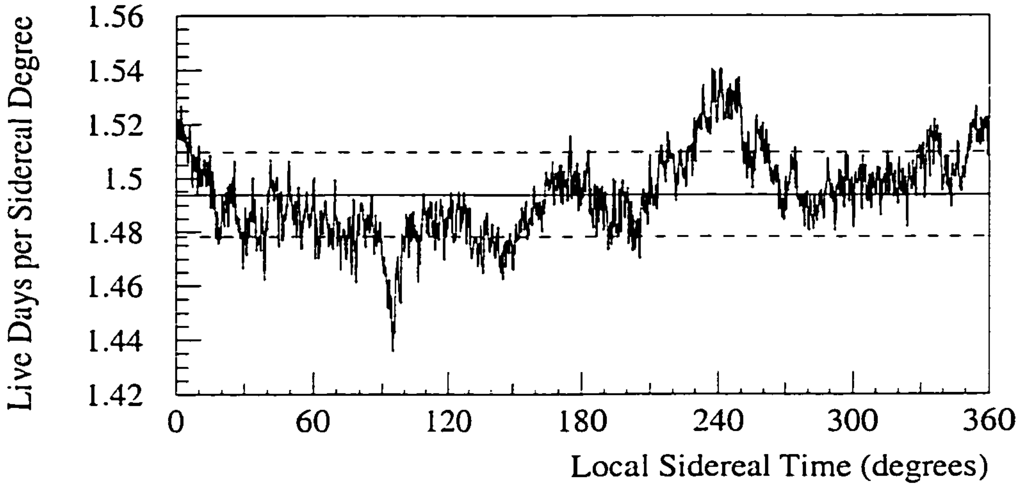

Figure 5.10 The live time was distributed slightly unevenly about the sidereal day because the detector was turned off periodically for scheduled maintenance and calibration, and also because parts of it were occasionally turned off automatically or went dead as a normal part of the data acquisition process. For a 537.8 live solar day subset of the data, the RMS deviation about the mean, shown by the dashed and solid lines, respectively, was about 1%.

off-source events will be appreciably larger than this for all potential point sources tested in this analysis, therefore one may take the unevenness of the live time distribution as a negligibly small systematic error, and not bother with the unnecessary complications of making corrections for it, even though one may know how.

5.3.1 Derivation of the Classical Analysis

To calculate the maximum plausible rate of upward-going muons attributable to a specific region of the sky, one starts by carefully outlining the parameters of the problem: there will be a number N of upward-going muons that have been observed in an off-source area of exposure X, and a number n of upward-going muons that have been observed in an on-source area of exposure x. The differing sizes of X and x could in principle correspond (and indeed, in similar types of analyses of data sets acquired with more traditional kinds of telescopes that have to be aimed, and view only small patches of sky at any one time, often do correspond) simply to different amounts of time spent observing same-sized off-source and on-source areas, but in this case will refer instead to the different angular sizes of the two regions, as described by figure 5.9. To solve the problem, a normalized function $l(s; N, n, X, x)$ will be constructed which estimates the relative likelihood that the rate s of astrophysically induced "signal" upward-going muons takes on any particular value, and the function will be integrated over s from 0 to some value s_f such that

$$\int_0^{s_f} l(s; N, n, X, x)\, ds = f, \qquad 0 < f < 1. \tag{5.11}$$

Then one says that s_f is an upper limit on s at a confidence level of $(100 \times f)\%$. The purpose of this section is to derive the form of $l(s; N, n, X, x)$, using classical statistical methods.

To guide one's thinking about the correct form of l, one may practice by posing and solving a much simpler problem: given only N and X, form an estimator $Q(b; N, X)$ that states the likelihood, based on the observed data N, that the true value for the atmospheric background b in the off-source area takes on any value from 0 to ∞. To determine the form of Q, start by imagining that the true value of b were a known quantity, $\hat{b}$, and ask what would be the probability $P(\hat{N}; \hat{b}, X)$, during an exposure X, for observing $\hat{N}$ upward-going muons, where $\hat{N}$ is now understood to be a variable quantity $0 \leq \hat{N} \leq \infty$, rather than a measured feature of the data, like N. Because there are an enormous number of atmospheric neutrinos streaming through the Earth compared to the number that actually interact, the probability of observing $\hat{N}$ interactions must be described by a Poisson distribution, and therefore

$$P(\hat{N}; \hat{b}, X) = \frac{(\hat{b}X)^{\hat{N}} e^{-\hat{b}X}}{\hat{N}!} \tag{5.12}$$

Reference [92] illustrates, through a lengthy derivation based on general principles, the intuitive result that the form of Q must be the same as that of P, with the possible exception of a different normalization constant due to the fact that N and b do not share the same units, and therefore

$$Q(b; N, X) \propto \frac{(bX)^{N} e^{-bX}}{N!} \tag{5.13}$$

In a similar fashion, one may construct an estimator for the sum $s+b$ in the on-source region:

$$q(s,b;n,x) \propto \frac{[(s+b)x]^n e^{-(s+b)x}}{n!}. \tag{5.14}$$

The functions Q and q may be combined together to form a joint likelihood function,

$$L(s,b;N,n,X,x) = C(N,n,X,x) b^N e^{-bX} (s+b)^n e^{-(s+b)x}, \tag{5.15}$$

which gives a global likelihood that estimates the joint probability that any pair of test values s and b are the true parameters describing the parent distribution from which the observed data set was drawn. The value of the normalization constant, C, may be calculated exactly by writing the $(s+b)^n$ factors in a binomial expansion and applying the definition of the Γ function,

$$\Gamma(m) = \int_0^\infty y^{m-1} e^{-y}\, dy = (m-1)!,$$

as appropriate:

$$\begin{aligned}
\frac{1}{C(N,n,X,x)} &= \int_0^\infty \int_0^\infty b^N e^{-bX} (s+b)^n e^{-(s+b)x}\, db\, ds \\
&= \sum_{j=0}^{n} \frac{n!}{j!(n-j)!} \int_0^\infty b^N b^{n-j} e^{-b(x+X)}\, db \int_0^\infty s^j e^{-sx}\, ds \\
&= \sum_{j=0}^{n} \frac{n!}{j!(n-j)!} \int_0^\infty \frac{[b(x+X)]^N [b(x+X)]^{n-j} e^{-b(x+X)}\, d[b(x+X)]}{(x+X)^{N+n-j+1}} \\
&\quad \times \int_0^\infty \frac{(sx)^j e^{-sx}\, d(sx)}{x^{j+1}} \\
&= \frac{n!}{(x+X)^{N+n+2}} \sum_{j=0}^{n} \frac{(N+n-j)!}{(n-j)!} \left(1+\frac{X}{x}\right)^{j+1}.
\end{aligned} \tag{5.16}$$

In practice, this analysis will not be concerned with making a statement about the estimated background levels associated with test points; only the signal strength is of interest. Therefore, the likelihood function $l(s;N,n,X,x)$ which appears in equation 5.11 will be obtained by integrating $L(s,b;N,n,X,x)$ of equation 5.15 over all possible values of b, again applying the binomial expansion and recognizing the Γ function, as above:

$$\begin{aligned}
l(s;N,n,X,x) &= \int_0^\infty L(s,b;N,n,X,x)\, db \\
&= C \sum_{i=0}^{n} \frac{n!}{i!(n-i)!} \int_0^\infty \frac{[b(x+X)]^N [b(x+X)]^{n-i} e^{-b(x+X)}\, d[b(x+X)]}{(x+X)^{N+n-i+1}}
\end{aligned}$$

$$\begin{aligned}
&\times \frac{x(sx)^i e^{-sx}}{x^{i+1}} \\
&= C\frac{n!}{(x+X)^{N+n+2}} \sum_{i=0}^{n} \frac{(N+n-i)!}{(n-i)!} \left(1+\frac{X}{x}\right)^{i+1} \frac{x(sx)^i e^{-sx}}{i!} \\
&= \sum_{i=0}^{n} W_i \frac{x(sx)^i e^{-sx}}{i!}, \qquad (5.17)
\end{aligned}$$

with W_i defined by

$$\begin{aligned}
W_i &= C(N,n,X,x)\frac{n!}{(x+X)^{N+n+2}} \frac{(N+n-i)!}{(n-i)!} \left(1+\frac{X}{x}\right)^{i+1} \\
&= \frac{(1+\frac{X}{x})^{i+1}\frac{(n+N-i)!}{(n-i)!}}{\sum_{j=0}^{n}(1+\frac{X}{x})^{j+1}\frac{(n+N-j)!}{(n-j)!}}. \qquad (5.18)
\end{aligned}$$

The expressions defined in equations 5.17 and 5.18, substituted into equation 5.11, give an exact prescription for determining a confidence limit on the signal strength s. The off-source and on-source event rates and exposures, N, n, X, and x, are all known quantities that are extracted from the data, therefore the problem of estimating s_f to arbitrary precision is reduced merely to the task of integrating equation 5.11 numerically to ever more precisely specified s_f until the value of the integral becomes arbitrarily close to f.

5.3.2 Derivation of the Bayesian Analysis

Although the classical, or frequentist, statistical approach outlined above is founded on well accepted principles, it is not the only formalism available. Bayesian statistical analysis offers an alternative formulation which may, in principle, arrive at different results from those of the preceding section. As a complement to the above derivation, a Bayesian analysis was also performed.

To estimate, in a Bayesian context, the expected long-term average rate b of atmospherically induced upward-going muons in the on-source region, based only on the data from the off-source region, (i.e., not including any knowledge about the on-source region yet) one begins by writing down Bayes' Theorem for this problem:

$$p(b|NI) = p(b|I)\frac{p(N|bI)}{p(N|I)}, \qquad (5.19)$$

where N is the number of events actually observed during the exposure, X, of the off-source region, and I signifies the assumptions already made about b before analyzing the data. In this case, it shall be assumed that b is the mean of a Poisson distribution, meaning that it is definitely positive, and that it lies between some very small number r and some very large number R, but other than that nothing is known about its scale. The probability density that b will turn out to lie in any given interval db, given complete initial ignorance of its scale, must be:

$$p(b|I) = \begin{cases} \frac{1}{b\ln(R/r)} & \text{if } r < b < R \\ 0 & \text{otherwise} \end{cases} \tag{5.20}$$

because that is the density which, when integrated, gives an even probability that b may be found between any order of magnitude and the next, i.e., for $k > r$, and $Mk < R$, with M representing any arbitrary positive number just as long as it is small enough not to violate the inequality, the value of:

$$\int_k^{Mk} \frac{db}{b\ln(R/r)} = \frac{\ln(M)}{\ln(R/r)}$$

is independent of k. (This form for $p(b|I)$ is called the least informative prior, and this practice is well supported in the literature[93].) Since b is assumed to be the mean of a Poisson distribution, the probability of observing N counts in the off-source region, given b, is just:

$$p(N|bI) = \frac{(bX)^N e^{-bX}}{N!}. \tag{5.21}$$

The value of $p(N|I)$ may be calculated by marginalization, i.e., by carrying out the standard Bayesian procedure in which $p(N|I)$ is expanded in terms of a complete basis of alternative hypotheses about b, and then summed over all of them:

$$\begin{aligned} p(N|I) &= \int_r^R p(Nb|I)\,db \\ &= \int_r^R p(N|bI)p(b|I)\,db \\ &= \int_r^R \frac{(bX)^N e^{-(bX)}}{N!} \cdot \frac{X}{X} \cdot \frac{db}{b\ln(R/r)} \\ &= \frac{1}{\ln(R/r)N!} \int_{b=r}^{b=R} (bX)^{N-1} e^{-bX}\, d(bX) \\ &= \frac{[(N-1)! - \epsilon(r,R)]}{\ln(R/r)N!}, \end{aligned} \tag{5.22}$$

where it is understood that the only detail about the form of $\epsilon(r, R)$ which is of interest is the fact that $\lim_{R\to\infty}^{r\to 0} \epsilon(r, R) = 0$, and where, for the integral form of $\Gamma(N)$, $(N-1)!$

has been substituted, which is allowed since N is assumed to be an integer greater than 0. Combining equations 5.20, 5.21, and 5.22, and taking the limit as $r \to 0$ and $R \to \infty$, results in an explicit expression for equation 5.19:

$$p(b|NI) = \frac{X(bX)^{N-1}e^{-bX}}{(N-1)!}. \tag{5.23}$$

It is easily verified that this probability density is normalized to one, has a mean of N/X and a standard deviation of $\sqrt{N}/X$.

To estimate the combined rate of upward-going muon events in the on-source region due both to point source neutrinos, if they are present, and also the atmospheric neutrinos, by using both the off-source data, as summarized by equation 5.23, and additionally now, folding in the data from the on-source region as well, one again begins by writing down Bayes Theorem for this new problem:

$$p(sb|ni) = p(sb|i)\frac{p(n|sbi)}{p(n|i)} = p(s|bi)p(b|i)\frac{p(n|sbi)}{p(b|i)}. \tag{5.24}$$

Here s is the rate due to the source, n is the number of events observed during the on-source exposure x, and i signifies all of the prior assumptions about either b or s that are already made before analyzing the on-source data, including the results of the off-source analysis described above. Following the example of equation 5.20, we set:

$$p(s|bi) = \begin{cases} \frac{1}{s\ln(R/r)} & \text{if } r < s < R \\ 0 & \text{otherwise} \end{cases} \tag{5.25}$$

again restricting $r < s < R$. It is to be understood, of course, that the r and R of the above equation need not be the same as the ones that appear in equation 5.20: in principle, they should carry separate b and s subscripts, but since the limit has already been taken that $r_b \to 0$ and $R_b \to \infty$, there is no need in this case to maintain the distinction symbolically. The atmospheric neutrino and the astrophysical point source rates are both means of independent Poisson distributions, therefore the likelihood of observing n events given s and b is:

$$p(n|sbi) = \frac{x^n(s+b)^n e^{-(s+b)x}}{n!}. \tag{5.26}$$

Since i in this problem includes the results of the off-source analysis, i.e., symbolically $i = NI$, it means that

$$p(b|i) = p(b|NI) = \frac{X(bX)^{N-1}e^{-bX}}{(N-1)!}, \tag{5.27}$$

as derived in equation 5.23. Following the example of equation 5.22, $p(n|i)$ may also be expressed as a normalization factor, using equations 5.25, 5.26, and 5.27:

$$\begin{aligned} p(n|i) &= \int_0^\infty \int_r^R p(n|sbi)p(s|bi)p(b|i)\, ds\, db \\ &= \frac{x^n X^N}{\ln(R/r)n!(N-1)!} \int_0^\infty \int_r^R (s+b)^n s^{-1} b^{N-1} e^{-(s+b)x} e^{-bX}\, ds\, db. \end{aligned} \tag{5.28}$$

Applying a binomial expansion to $(s+b)^n$, and again recognizing the integral form of the Γ function, the integral quantity may be expressed as:

$$\begin{aligned} &\int_0^\infty \int_r^R (s+b)^n s^{-1} b^{N-1} e^{-(s+b)x} e^{-bX}\, ds\, db \\ &= \sum_{j=0}^{n} \frac{n!}{(j)!(n-j)!} \int_0^\infty \frac{[b(x+X)]^{n+N-j-1} e^{-b(x+X)}\, d[b(x+X)]}{(x+X)^{n+N-j}} \\ &\qquad \times \int_{s=r}^{s=R} \frac{(sx)^{j-1} e^{-sx}\, d(sx)}{x^j} \\ &= \frac{n!}{(x+X)^{n+N}} \sum_{j=0}^{n} \left(1 + \frac{X}{x}\right)^j \frac{(n+N-j-1)!}{(n-j)!} \int_{rx}^{Rx} \frac{(y)^{j-1} e^{-y}\, dy}{j!}, \end{aligned} \tag{5.29}$$

where $y = sx$ has been substituted for brevity. Combining equations 5.25, 5.26, and 5.27 together with the coefficients from equation 5.28 and the simplified integrals of equation 5.29, one obtains the result that:

$$p(sb|ni) = \frac{(x+X)^{n+N}(s+b)^n s^{-1} b^{N-1} e^{-sx} e^{-b(x+X)}}{n! \sum_{j=0}^{n} (1 + \frac{X}{x}) \frac{(n+N-j-1)!}{(n-j)!} \int_{rx}^{Rx} \frac{y^{j-1} e^{y}\, dy}{j!}}. \tag{5.30}$$

This expression gives the joint posterior probability distribution for both s and b, using all of the information from all areas of the sky surveyed. Since only the source strength is of interest, b is a nuisance parameter, and can be eliminated by marginalization, i.e., integrating over all possible values. Taking most of the denominator and labeling it

$$K = \sum_{j=0}^{n} \left(1 + \frac{X}{x}\right)^j \frac{(n+N-j-1)!}{(n-j)!} \int_{rx}^{Rx} \frac{y^{j-1} e^{y}\, dy}{j!} \tag{5.31}$$

for brevity, and again applying the binomial expansion to $(s+b)^n$ and the definition of the Γ function, we find that

$$\begin{aligned}
p(s|ni) &= \int_0^\infty p(sb|ni)\,db \\
&= \frac{(x+X)^{n+N}e^{-sx}}{n!\;K}\sum_{i=0}^{n}\frac{n!}{i!(n-i)!}s^{i-1}\int_0^\infty b^{n-i}b^{N-1}e^{-b(x+X)}\,db \\
&= \frac{e^{-sx}}{K}\sum_{i=0}^{n}\left(\frac{x}{x}\right)^{i}\frac{s^{i-1}(x+X)^{i}}{i!} \\
&\qquad \times\int_0^\infty\frac{[b(x+X)]^{n+N-i-1}e^{-b(x+X)}\,d[b(x+X)]}{(n-i)!} \\
&= \frac{1}{K}\sum_{i=0}^{n}\left(1+\frac{X}{x}\right)^{i}\frac{(n+N-i-1)!}{(n-i)!}\frac{x(sx)^{i-1}e^{-sx}}{i!}
\end{aligned} \tag{5.32}$$

where it is understood that the i in $p(sb|ni)$ is not the same as the i used to index the summation. This expression may be rewritten as:

$$p(s|ni) = \sum_{i=0}^{n} W_i\frac{x(sx)^{i-1}e^{-sx}}{i!} \tag{5.33}$$

with

$$W_i = \frac{\left(1+\frac{X}{x}\right)^{i}\frac{(n+N-i-1)!}{(n-i)!}}{\sum_{j=0}^{n}\left(1+\frac{X}{x}\right)^{j}\frac{(n+N-j-1)!}{(n-j)!}\int_{rx}^{Rx}\frac{y^{j-1}e^{y}\,dy}{j!}}. \tag{5.34}$$

The result contained in equation 5.33 merits several comments. The first of these is that a comparison of equation 5.33 to the general form of result 5.23 suggests that the posterior probability distribution for s may be interpreted roughly as a sort of average, weighted by the W_i, of individual posteriors that correspond to the trial assignment, consecutively, of $0, 1, 2, \ldots, n$ on-source events to the signal. Another point is that although most of the integrals in the denominator,

$$\int_{rx}^{Rx} y^{j-1}e^{y}\,dy,$$

approach finite limits as $r \to 0$ and $R \to \infty$, the $j = 0$ case does not approach any limit as $r \to 0$, which forces the mathematically inelegant application of a lower cutoff on the signal strengths which may be considered in this type of analysis. Such a cutoff is justifiable on philosophical grounds, for there is certainly an upper limit to the size of any future neutrino telescope which will ever be built, and to the length of time for which it will be operated, and that places a lower limit on the astrophysical signal strength that any human scientist could ever hope to observe. Estimating such a limit and applying it will make the meaning of the confidence limits quoted later in this dissertation more concrete: the limits only apply to models in which astrophysical point sources exist with strong enough signals so that real human researchers may ever hope to observe them. Such a tactic has an intuitive appeal, for a confidence limit may be thought of as a kind of statistical conjecture about what truth will be discovered in the event that more data are accumulated in the future, and applying a lower cutoff forces one to distinguish between what would be knowable with a modest amount of additional data and what would be knowable with an infinite amount of additional data. This analysis will apply a cutoff of $rx = 0.01$ to all point sources for which confidence limits are quoted, which corresponds to a signal strength such that an experiment 10 times as big as Super-Kamiokande, operated for 10 times as long, could just be expected to see one signal event during its entire lifetime.

Another thing to point out is that the precise form of the results of equations 5.33 and 5.34 follows in part from a somewhat arbitrary set of assumptions contained in equations 5.20 and 5.25. Despite the justifications that were given, those choices of prior probabilities were not the only reasonable choices available. For example, an alternative choice, called the uniform prior, which appears elsewhere in the literature[94] is that

$$p(b|I) = \begin{cases} \frac{1}{R-r} & \text{if } r < b < R \\ 0 & \text{otherwise} \end{cases} \tag{5.35}$$

and the expression for $p(s|bi)$ is assigned a similar form. Coincidentally, if one goes through an alternate Bayesian derivation analogous to the one above, differing only in that expressions of the type that appear in equation 5.35 are substituted in place of the ones given in equations 5.20 and 5.25, the expressions which replace equations 5.33 and 5.34 are those derived in the classical analysis: equations 5.17 and 5.18.

At least one other variation on this type of analysis, with similar but slightly different results, was also available in the literature[95]. In this third case, the author assumed that

$$p(s|bI) = \frac{1}{s+b}$$

for $0 < s < \infty$ and acknowledged that such a definition is unnormalizable, but argued that as long as the final expression for $p(s|ni)$ was normalized, it would be acceptable to gloss over this inconvenient fact and simply use the final result to make a statistical statement. The results were

$$p(s|ni) = \sum_{i=1}^{n} W_i \frac{x(sx)^{i-1} e^{-sx}}{(i-1)!} \tag{5.36}$$

and

$$W_i = \frac{\left(1 + \frac{X}{x}\right)^i \frac{(n+N-i-1)!}{(n-i)!}}{\sum_{j=1}^{n} \left(1 + \frac{X}{x}\right)^j \frac{(n+N-j-1)!}{(n-j)!}}. \tag{5.37}$$

While this result is more elegant in the sense that there is no need to require any lower cutoff on the possible signal strengths that one may contemplate, it does have a different drawback which becomes apparent if one interprets the expression of equation 5.36, as was done earlier, as a weighted sum of individual posteriors each corresponding to a different number of on-source events attributed to the signal: there is no term which corresponds to 0 of the on-source events coming from the signal! This is, of course, a very likely, indeed even expected possibility, and the absence of such a term is to be viewed as a defect in the analysis. Nevertheless, it is an alternative published result, so for completeness of the discussion it is considered here along with the others.

5.3.3 Results and Discussion

To facilitate comparison between the results, and to show that they are relatively insensitive to the choice of assumptions used in their derivation, confidence limits based on all three expressions, using as numerical input the numbers listed in table 5.1, are presented in table 5.2. Furthermore, the two most extreme cases are plotted in figure 5.11, along with the results of six other comparable analyses from other experiments.

There are several points about figure 5.11 that are important enough to merit explicit discussion. For starters, it is worth noting that the limits coming from this analysis are mostly clustered toward the bottom of the plot, below the majority of the previous limits; since these are upper limits, this is good: lower values in later experiments is the trend that one would prefer to see, and it means that this analysis appears to be outperforming

Table 5.1 The leftmost five columns of this table are self-explanatory, but the other three demand some comment. The off-source to on-source size ratio is the number of additional on-source sized circles which could be laid out end-to-end in the off-source region of the same declination band as a given celestial object. It corresponds to (X/x) in the derivations of the preceding pages. The time fraction below the horizon is the fraction of the sidereal day for which a given object was visible in upward-going muons. This fraction, multiplied by the effective live time as calculated by equation 4.2, and by the average effective area for the object's declination as shown in figure 4.4, yields a number which is nearly equal to the exposure, except for one small correction: strictly speaking, the averages plotted in figure 4.4 only apply to celestial objects which did not pass through the downward-contaminated upward-going region shown in figure 5.9 as they traversed the sky. Since events which pointed back to the downward-contaminated region were not used in the analysis, the effective area should not be averaged over it for those objects which passed through it.

Celestial Object	R.A. (2000) (in °)	Decl. (2000) (in °)	On-Source Events (in 12° cone)	Off-Source Events (in 12° cone)	Off-/On-Source Size Ratio	Time Fraction Below Horizon	Exposure (cm^2 s)
SN1987A	83.87	-69.268	34	150	4	1.000	1.61×10^{15}
LMC X-4	83.21	-66.371	38	151	4	1.000	1.61×10^{15}
Circinus X-1	230.17	-57.17	36	271	7	1.000	1.58×10^{15}
Vela Pulsar	128.84	-45.177	32	293	9	0.758	1.17×10^{15}
Centaurus A	201.37	-43.02	37	276	9	0.735	1.13×10^{15}
Vela X-1	135.53	-40.555	31	276	10	0.711	1.09×10^{15}
Galactic Ctr	266.41	-28.93	20	299	12	0.629	9.66×10^{14}
Scorpius X-1	244.98	-15.640	13	333	13	0.563	8.83×10^{14}
3C273	187.28	2.057	26	278	13	0.488	7.85×10^{14}
SS433	287.96	4.983	23	285	13	0.475	7.68×10^{14}
Geminga	98.40	17.77	16	268	13	0.420	6.84×10^{14}
Crab Nebula	83.63	22.014	19	236	12	0.399	6.50×10^{14}
Hercules X-1	254.46	35.342	8	193	11	0.319	5.13×10^{14}
Markarian 421	166.12	38.213	8	178	10	0.296	4.74×10^{14}
Cygnus X-3	308.11	40.958	26	165	10	0.272	4.33×10^{14}

Table 5.2 These are the upper limits, at the 90% confidence level, on the absolute number of signal neutrino events which could be coming from selected celestial objects. The second column uses the algorithm described by equation 5.17, the third uses equation 5.33, and the fourth uses 5.36. Additionally, in the fifth and final column, 90% confidence level upper limits on muon flux are quoted, based on the classical event limits from the second column, plus the exposure values given in table 5.1.

Celestial Object	Source Events, 90% C.L., Classical (or Bayesian Uniform Prior)	Source Events, 90% C.L., Bayesian Least Inform. Prior	Source Events, 90% C.L., Bayesian Alternative Least Inform. Prior	90% C.L. Classically Derived μ Flux ($cm^{-1}s^{-1}$)
SN1987A	10.1	3.1	9.6	6.3×10^{-15}
LMC X-4	12.5	4.5	12.0	7.8×10^{-15}
Circinus X-1	10.4	3.3	9.9	6.9×10^{-15}
Vela Pulsar	10.8	3.7	10.1	9.7×10^{-15}
Centaurus A	16.2	8.7	15.3	1.4×10^{-14}
Vela X-1	13.0	5.8	12.3	1.3×10^{-14}
Galactic Ctr	6.9	2.0	6.4	7.1×10^{-15}
Scorpius X-1	4.3	1.2	4.0	4.8×10^{-15}
3C273	13.1	6.7	12.2	1.7×10^{-14}
SS433	10.1	3.9	9.4	1.3×10^{-14}
Geminga	6.2	3.5	5.8	9.1×10^{-15}
Crab Nebula	8.4	2.9	7.8	1.3×10^{-14}
Hercules X-1	3.8	1.1	3.6	7.5×10^{-15}
Markarian 421	3.8	1.1	3.6	8.1×10^{-15}
Cygnus X-3	17.5	3.3	16.5	4.0×10^{-14}

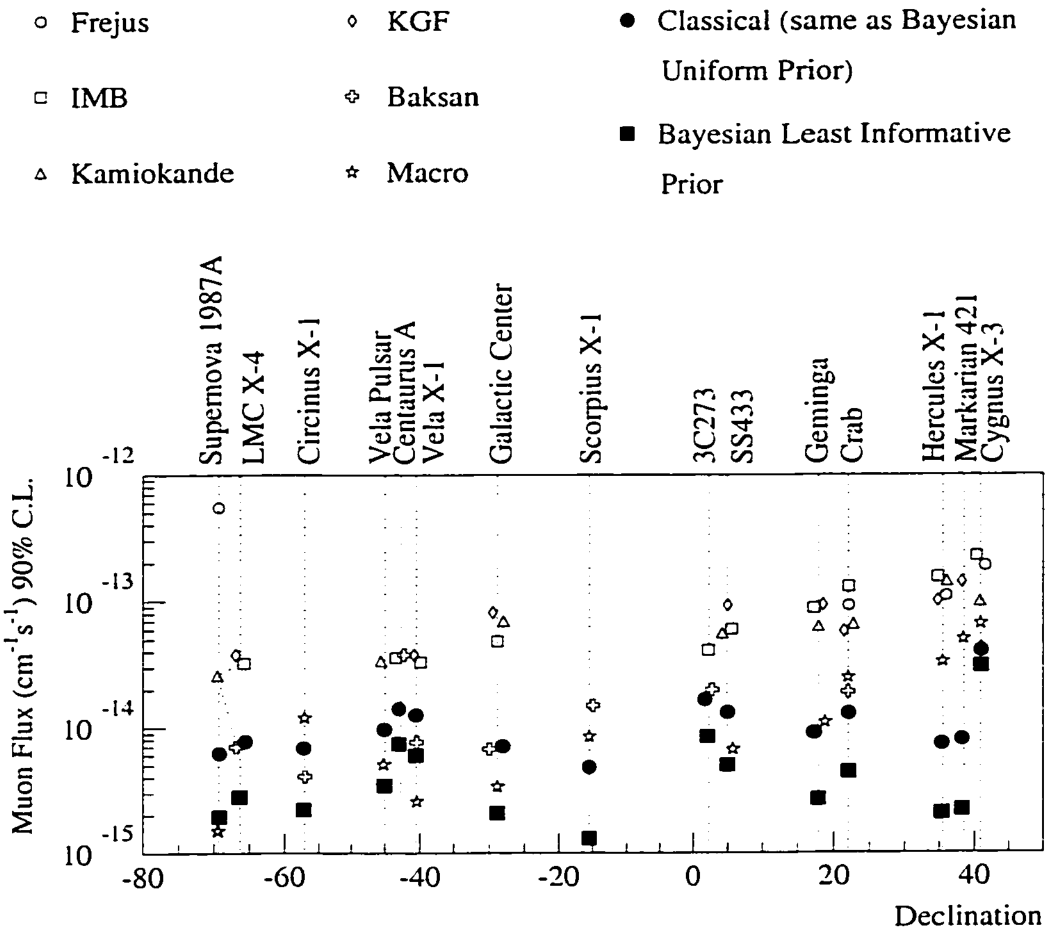

Figure 5.11 These are the 90% confidence level upper limits on anomalous excess upward-going muon flux coming from the directions of various candidate objects which have been favored as likely point sources by other researchers in the neutrino astrophysics community. The values plotted with open markers, which describe the results of six underground experiments which preceded this one, are taken from references [60], [61], [62], [63], [64], and [65]. The values plotted with filled markers are the results of this analysis, as calculated from equation 5.17 in the case of the circles, and 5.33, in the case of the squares. The limits quoted here, being in most cases generally on the low end of the range of values quoted for the various sources, are therefore for the most part as stringent or more stringent than the limits that have been produced by earlier generation experiments.

older analyses that were limited to smaller data sets, as it should. Another important point is that the two experiments which were able to compete most successfully with Super-Kamiokande in setting stringent limits were the two with the next largest data sets (see table 1.1): MACRO, which had 1100 events, and Baksan, which had 700. Thus, not only are the Super-Kamiokande results clustered generally toward the region of the plot where one would expect to see them (the bottom), but furthermore, even in the cases in which the results overlap, that overlap is generally with the confidence limit values quoted by the two experiments from which one would have most strongly expected also to receive similar final numbers.

More gratifying still is that even the declination dependence of the results appears to be roughly correct: unlike Super-Kamiokande, both MACRO and Baksan suffered from substantially reduced sensitivity near the horizon, as compared to that near the vertical, and since geometry requires that upward-going events from northern declinations tend to arrive at northern hemisphere detectors most often at shallow angles, just below the horizon, the low horizontal acceptance of these two other detectors tended to diminish their overall sensitivity to potential sources in the northern sky. Location exacerbated this situation even further: MACRO, at about 42.5° north latitude, was about 6.1° further north than Super-Kamiokande, and Baksan, at 43.7°, was 7.3° further, meaning that quite apart from any issues of horizontal vs. vertical sensitivity, in their capacity as upward-going particle detectors, these other two experiments were simply intrinsically more southward-pointing than Super-Kamiokande. Because these two factors tended to augment, rather than cancel out each other's effects, one would naively expect that the confidence limit results from this analysis would tend to compare most favorably to those of MACRO and Baksan at the northernmost declinations, and indeed, the evidence in figure 5.11 appears to show that they do.

Thus, the detailed comparison with previous results shown in figure 5.11 demonstrates at least two important ideas: first, that the point source search capabilities offered by the Super-Kamiokande data set are generally of a higher or at least comparable quality to those of other recent experiments, and second, that this analysis was for the most part probably done correctly, since the results are commensurate with those of the other experiments in the cases where one would suppose that they should be, while differing from the other results in the expected way at high northern declinations.

5.3.4 Extrapolation to Neutrino Luminosity

The analysis of the preceding sections arrives at conclusions about the absolute upward-going muon flux which might possibly be arriving from various sections of the sky, but this is not the quantity of real interest. The quantity of real interest is, of course, the maximum possible neutrino luminosity of the selected celestial objects. To extrapolate to luminosity at the source from the flux at the Earth, three things must be known about the object for which the calculation is being performed: the putative spectrum of the signal neutrinos, in order to be able to convert from flux of muon particles to flux of neutrino energy through the detector; the distance to the object; and the angular dependence of the neutrino emission. The second of these, the distance to the source, is known to within ±10% or better for several of the selected celestial objects, and it is known to nearly that precision for the remainder of them, on the basis of other astronomical analyses. The other two things are unknown, and to make further progress, some assumptions will have to be made, following the examples of other researchers who have attempted similar analyses on data sets from other experiments before this one.

The angular dependence of any neutrinos which might be radiating from putative sources will be assumed isotropic, because that is the most general, least model-dependent way of performing such a calculation. Thus if R is the distance from Earth to some celestial object, the 90% confidence level upper limit on neutrino luminosity at the source is simply $4\pi R^2$ times the 90% confidence limit on the neutrino power per unit area flowing through the detector. Of course, it should be understood that many of the celestial objects which are most likely to emit neutrinos might very well be emitting them in highly collimated beams or cones, but it is not clear at this time what are the most appropriate detailed assumptions to make about the geometric shapes of such things, so in the interest of avoiding excessively speculative results, the relatively simpler assumption of isotropy will be favored here. In any case, a curious reader who wishes to know how the results would differ for the case of an arbitrary beaming model can easily add in the effects of such beaming himself, simply by multiplying by appropriate geometric and duty cycle scaling factors.

A power law distribution will be assumed for the spectrum of the astrophysical neutrinos as observed at the Earth:

$$\frac{dN_\nu}{dE_\nu} = CxE_\nu^{-(\gamma+1)}, \tag{5.38}$$

where x is again the exposure, discussed in the previous sections, the coefficient C sets the overall scale of the power flow per unit area at the Earth's distance from the source, and N_ν represents an absolute number of neutrinos flowing through the entire effective area of the detector over the entire period of the analysis. This is the same type of spectrum that describes the observed energy distribution for cosmic ray primaries. For protons in the GeV to TeV range, $\gamma \approx 1.7$. Astrophysical neutrinos will be assumed, if they exist, to have a spectral index in some nearby range about that value, probably slightly smaller, but not less than $\gamma = 1$. To estimate a confidence limit on the maximum amount of neutrino energy which could have flowed through the detector during the analysis period, one must calculate

$$\int_{E_{thr}}^{\infty} E_\nu \frac{dN_\nu}{dE_\nu}\, dE_\nu = Cx \int_{E_{thr}}^{\infty} E_\nu^{-\gamma}\, dE_\nu, \tag{5.39}$$

where E_{thr} is the minimum energy an upward-going muon is allowed to have deposited in the detector in order to be included in the analysis, and therefore E_{thr} also represents the minimum energy for an observable parent neutrino. The confidence limit on the luminosity, or absolute power output, of a source candidate will then be simply

$$L(E_\nu > E_{thr}) = 4\pi R^2 C \int_{E_{thr}}^{\infty} E_\nu^{-\gamma}\, dE_\nu. \tag{5.40}$$

Evidently, to do this calculation, a value for C must be determined from the confidence limits set on the absolute numbers of upward-going muons which could be coming from the putative source, shown in table 5.2.

To determine the value of C for a neutrino point source candidate, one can exploit the relationship between the observed upward-going muon spectrum at the detector and the parent neutrino spectrum, which may be expressed as

$$\frac{dN_\mu}{dE_\mu} = \int_{E_\nu=E_\mu}^{E_\nu=\infty} \left(\frac{dN_\nu}{dE_\nu}\right) \left(\frac{dP(E_\nu)}{dE_\mu}\right) dE_\nu, \tag{5.41}$$

where $dP(E_\nu)/dE_\mu$ represents the probability, discussed at length in reference [48], that a neutrino which was aimed at the detector interacted in the rock below the detector to produce a muon that traveled all the way through the rock to enter the detector with an initial energy in the interval $[E_\mu, E_\mu + dE_\mu]$. The probability factor is a convolution of both the neutrino interaction cross section and the muon propagation dynamics, and since the the cross section for $\overline{\nu}_\mu$ interactions with nucleons is about half that of ν_μ at the lower end of the energy range of interest, it means that technically, one should specify which case one is attempting to analyze. Substituting the functional form of equation 5.38 for dN_ν/dE_ν, the expression above can be integrated over E_μ to find that

$$\int_{E_{thr}}^{\infty} \frac{dN_\mu}{dE_\mu}\, dE_\mu = Cx \int_{E_\nu=E_{thr}}^{E_\nu=\infty} E_\nu^{-(\gamma+1)} \left[\int_{E_\mu=E_{thr}}^{E_\mu=\infty} \left(\frac{dP(E_\nu)}{dE_\mu} \right) dE_\mu \right] dE_\nu, \tag{5.42}$$

and substituting more compact expressions for two of the integrals yields

$$N_\mu(E_\mu > E_{thr}) = Cx \int_{E_\nu=E_{thr}}^{E_\nu=\infty} E_\nu^{-(\gamma+1)} P(E_\nu; E_\mu > E_{thr})\, dE_\nu. \tag{5.43}$$

By happy coincidence, values of $P(E_\nu; E_\mu > E_{thr})$, averaged over the cross sections for ν_μ and $\overline{\nu}_\mu$, are already tabulated for $E_{thr} = 2$ GeV (the threshold for this analysis, to 1 significant figure) in the literature[48], and therefore the integral in equation 5.43 can be calculated numerically. Those tabulated values are interpolated and plotted in figure 5.12. After doing the integral and substituting values for $N_\mu(E_\mu > E_{thr})$ from table 5.2, and x for each source candidate from table 5.1, equation 5.43 is conveniently reduced to an equation in one unknown: C. The luminosity limits listed for selected sources in table 5.3 and figure 5.13 are calculated by solving for C, using equation 5.43, and substituting the results into equation 5.40. For the sake of simplicity, the convention will be adopted that all values for $N_\mu(E_\mu > E_{thr})$ used in this section are to be those calculated using the classical likelihood-based algorithm derived in section 5.3.1.

There are at least four points that should be emphasized about the results shown in figure 5.13. First, the neutrino luminosity upper limits are, like the upward-going muon flux limits (as they ought to be, since the muon and neutrino results are correlated), relatively competitive compared to those of other experiments: in several cases, for comparable choice of spectral index, the Super-Kamiokande results are an order of magnitude or so lower

Table 5.3 The second column in this table gives the current best estimate of the absolute distance to each object, garnered from various external references. One of the items in this column, the distance to 3C273, merits some special comment, since it is on a much different theoretical footing than all of the other measurements: it was determined via the Hubble law, $D = cz/H_0$. The redshift, $z = 0.158$[96], was very well measured in this case, so most of the uncertainty in the distance determination came from that associated with the Hubble constant, for which the best current value of $H_0 = 71 \pm 3 \pm 7$ km/s/Mpc[97] was assumed here. The third through fifth columns list various corresponding neutrino luminosities, presuming that the signal spectrum is power-law distributed, and assuming various possible spectral indices.

Celestial Object	Distance (kpc)	90% C.L. L_ν, $\gamma = 2.2$ (erg s^{-1})	90% C.L. L_ν, $\gamma = 2.7$ (erg s^{-1})	90% C.L. L_ν, $\gamma = 3.2$ (erg s^{-1})
SN1987A	$5.14 \pm 0.12 \times 10^{1}$ [98]	1.1×10^{40}	1.1×10^{41}	5.5×10^{41}
LMC X-4	$5.2 \pm 0.4 \times 10^{1}$ [98][99]	1.4×10^{40}	1.4×10^{41}	7.0×10^{41}
Vela Pulsar	$3.2\ ^{+0.4}_{-0.3} \times 10^{-1}$[100]	6.7×10^{35}	6.6×10^{36}	3.3×10^{37}
Centaurus A	$3.8 \pm 0.1 \times 10^{3}$ [101]	1.4×10^{44}	1.4×10^{45}	6.8×10^{45}
Vela X-1	$1.75 \pm 0.16 \times 10^{0}$ [102]	2.6×10^{37}	2.6×10^{38}	1.3×10^{39}
Galactic Ctr	$8.0 \pm 0.5 \times 10^{0}$ [103]	3.1×10^{38}	3.0×10^{39}	1.5×10^{40}
Scorpius X-1	$2.8 \pm 0.3 \times 10^{0}$ [104]	2.6×10^{37}	2.5×10^{38}	1.3×10^{39}
3C273	$6.7 \pm 0.7 \times 10^{5}$ [96][97]	5.1×10^{48}	5.0×10^{49}	2.5×10^{50}
SS433	$4.85 \pm 0.2 \times 10^{0}$ [105]	2.1×10^{38}	2.1×10^{39}	1.0×10^{40}
Geminga	$1.6\ ^{+0.6}_{-0.3} \times 10^{-1}$[106][107]	1.6×10^{35}	1.6×10^{36}	7.8×10^{36}
Crab Nebula	$1.93 \pm 0.11 \times 10^{0}$ [108]	3.2×10^{37}	3.1×10^{38}	1.6×10^{39}
Hercules X-1	$5.95 \pm 0.05 \times 10^{0}$ [109]	1.8×10^{38}	1.8×10^{39}	8.8×10^{39}
Cygnus X-3	$1.0 \pm 0.2 \times 10^{1}$ [110][111]	2.7×10^{39}	2.7×10^{40}	1.4×10^{41}

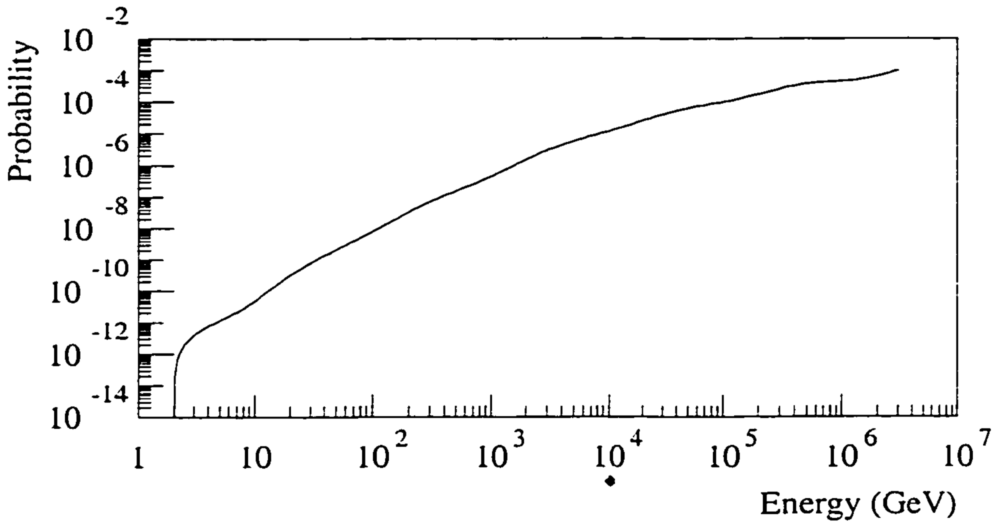

Figure 5.12 This is the convolved probability, vs. initial neutrino energy, that a ν_μ or $\bar{\nu}_\mu$ interacts in the rock below the detector to produce a muon, and then that muon travels through the rock to reach the detector with at least 2 GeV of its energy still remaining.

than those produced by the others (but note that investigators working on MACRO and Baksan apparently did not publish this kind of luminosity extrapolation analysis). The second point is that a modicum of caution is advisable in any attempt to make comparisons between results from different experiments, for the assumed distances to a few of the celestial objects have varied somewhat from one analysis to another, as well as the assumed spectral index for the astrophysical neutrinos. For this presentation, in order to report the results of other experimenters as faithfully as possible, those results have been displayed as originally quoted, rather than being renormalized according to current distance estimates, or to a single choice of spectral index. An undesirable consequence of following this convention is that in some cases, different luminosity limits for a given object are slightly less directly comparable to each other than one might prefer, but fortunately, nevertheless, the trend toward generally more stringent limits is still evident. The third noteworthy point about

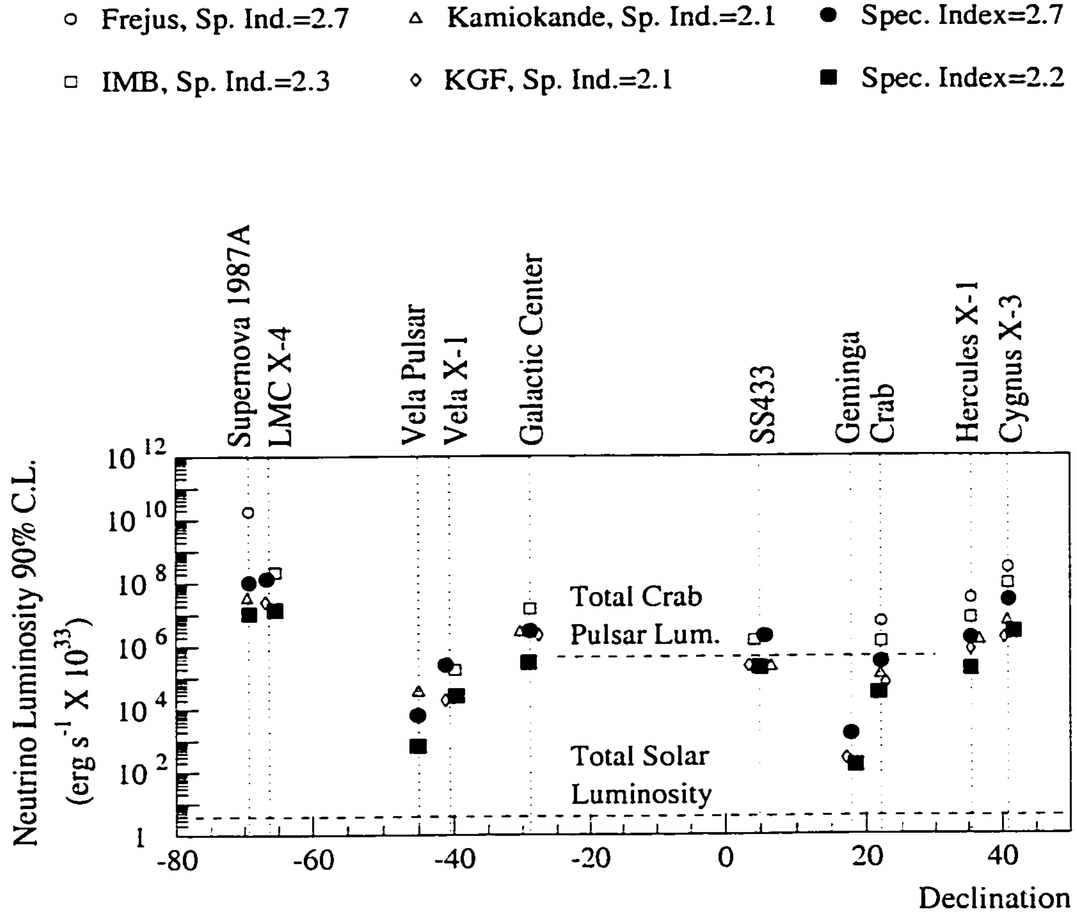

Figure 5.13 These are the 90% upper level confidence limits on the neutrino luminosity of those candidate source objects plotted in figure 5.11 for which absolute distances have been measured by other researchers. As in figure 5.11, filled markers show the results of this analysis under different assumptions (i.e., two different choices of $\gamma + 1$, the power law spectral index mentioned in the text), and open markers show the results of previous experiments, for comparison. Taking care to make comparisons only between results corresponding to nearly equal spectral indices, the general trend is again evident here, just as it also was with the upward-going muon flux confidence limits calculated in the previous section, that this analysis is able at least to match, and in many cases exceed, most of the limits of previous analyses, due to the larger event sample available to it. The total luminosities (i.e., integrated over all energies and particle types) of the sun and of the Crab Pulsar, the standard candle of very high energy γ-ray astronomy, are shown here to give a sense of scale. The comparable sizes of the Crab luminosity and the neutrino luminosity upper limit derived in this analysis suggest that neutrino astronomy is nearing the point in its evolution of significant experimental relevancy, such that next generation detectors may have a realistic chance of observing a signal.

figure 5.13 is that, while an improvement over the previous best results is certainly evident, fortunately, it is not implausibly large; this indicates that the calculations are probably mostly correct.

The fourth point is perhaps the most interesting one: the latest luminosity limit for the Crab Pulsar, the standard candle of γ-ray astronomy, is starting to dip significantly below the estimated total luminosity[112] of that object, which (making an approximate guess about the correct value for the moment of inertia I) can be calculated from the observables Ω and $\dot{\Omega}$ simply by applying the fact that a pulsar emits radiation entirely at the expense of its rotational kinetic energy T_{rot}, which, from elementary rotational dynamics, must decrease as $\dot{T}_{rot} = I\Omega\dot{\Omega}$. This comparability in magnitude is an important point because it means that neutrino astronomy is nearing the stage in its evolution at which confidence limits on potential source strength might someday soon be stringent enough actually to be relevant. It also suggests that a next generation experiment, rather than being limited only to quoting confidence limits, might actually have a chance at seeing a signal.

5.4 The Prolonged 1997 Flaring of Markarian 501

There are currently only eight known astrophysical sources of TeV γ-rays, which are of particular interest to would-be neutrino astronomers because one of the mechanisms that can produce such phenomena, pion decay, can also produce neutrinos in approximately the same energy range to which Super-Kamiokande is sensitive. The earliest of these sources ever to be detected[113], the Crab Nebula, is ordinarily the brightest, and serendipitously, it is also the only one for which the rate of emission does not appear to vary with time, hence it is often used as a natural standard candle against which all other TeV γ-ray sources may be compared. Markarian 501, a BL Lacertae object, has emitted TeV γ-rays through most of the second half of the 1990s at rates varying from 10% to 30% of the Crab Nebula, but between February and October 1997 it underwent an unusually prolonged phase of heightened activity in which its flux in that energy range was consistently about an order of magnitude or more intense than usual. This flaring activity was observed independently, over time scales of several months in each case, by at least 6 different experimental collaborations: Whipple[114], HEGRA[115], CAT[116], and TA[117], which all operate atmospheric Cherenkov telescopes, as well as the Milagro[118] group, which operates a water Cherenkov detector as a wide field of view sky monitor, and the Tibet air shower array[119]. At least three other atmospheric Cherenkov experiments, CAO[120], the University of Durham group[121], and TACTIC[122] also reported a high level of activity during part of this time, but over shorter observation periods of about a month or so each, and unlike the others, the last one of these three did not publish a plot of light intensity vs. time, only evidence for a statistically significant excess when averaged over the entire observation period.

5.4.1 Analysis Scheme and Experimental Observations

The extraordinary length of the 1997 flaring period of Markarian 501, as compared to the hours-long or days-long TeV γ-ray flares that have been observed in the handful of other celestial objects known to produce them, suggests an additional kind of neutrino point source search which one might attempt to carry out on the Super-Kamiokande data set: one that uses the start and stop time of the flare, as measured by photon astronomers, as well as source direction, to define on-source and off-source regions. In practice this plan requires a small modification, because Markarian 501 becomes unobservable to atmospheric Cherenkov telescopes during late autumn and early winter, as most of the time it spends above the horizon at that time of year is only during daylight hours, therefore the precise beginning and ending of the flare were not very well measured. In order to address this practical reality, in addition to the usual requirement that on-source events must lie near Markarian 501 in space angle, a further constraint will be imposed such that they will also be required to fall only within the earliest and latest 1997 observation dates, which were February 10, or Modified Julian Date (MJD) 50489[114], and October 2, or MJD 50723[116], respectively. The time periods from the last recorded measurement of 1996, on approximately August 17 (MJD 50312)[123] until February 10, 1997, and from October 2 until the first recorded measurement of 1998, on February 28 (MJD 50872)[114] will be treated as null periods, neither off-source nor on-source, since there is strong reason to believe flaring may have occurred during some of that time, but little is known for sure. The data-taking periods

Table 5.4 This table shows the precise start and stop times of each timing correlation analysis period, with the live time that corresponds to each. There are several small gaps of missing runs in the data: most notably, runs 1000 to 1444, about 15.8 days of live time, were not analyzed for stopping muons, and runs 5459 to 5499, about 3.9 days of live time, were not analyzed for through-going muons. There were a few other small gaps as well; live time was corrected accordingly in each case.

Start Time	Stop Time	First Run	Last Run	Period Type	Live Time (days)	Fraction of Total Live Time	Real Time (days)
96/03/31 15:00	96/08/17 00:25	1000	2552	Off-src	94.3	0.0637	138.39
96/08/17 00:25	97/02/10 07:03	2553	3523	Null	157.2	0.1062	177.28
97/02/10 07:03	97/10/02 08:58	3524	4867	On-src	199.4	0.1347	234.08
97/10/02 08:58	98/02/28 08:34	4869	5603	Null	124.1	0.0838	148.98
98/02/28 08:34	00/12/08 06:00	5604	9599	Off-src	905.2	0.6115	1013.89

before August 17, 1996, and after February 28, 1998, will be treated as partial off-source regions, with the rest of the off-source events coming from off-source circles as defined by figure 5.9. For convenient management of the raw data files, the Super-Kamiokande data acquisition system was, as a rule, turned off for a minute or two once per day in order to change the data-taking run number; for convenient analysis, the precise boundaries of the on-source, off-source, and null time intervals will be chosen to lie within the nearest such dead period to the photon observation start and stop dates listed above. The data acquisition live times of the various real time periods of interest are summarized in table 5.4.

The timing distribution of all of the upward-going muon events observed within 20° of Markarian 501 during the entire data acquisition period is shown in figure 5.14. Disappointingly, only two events were observed during the known flaring period; even without further exposition one may guess that almost certainly this will not turn out to be a large enough fluctuation above the background to be significant. To make that determination for sure,

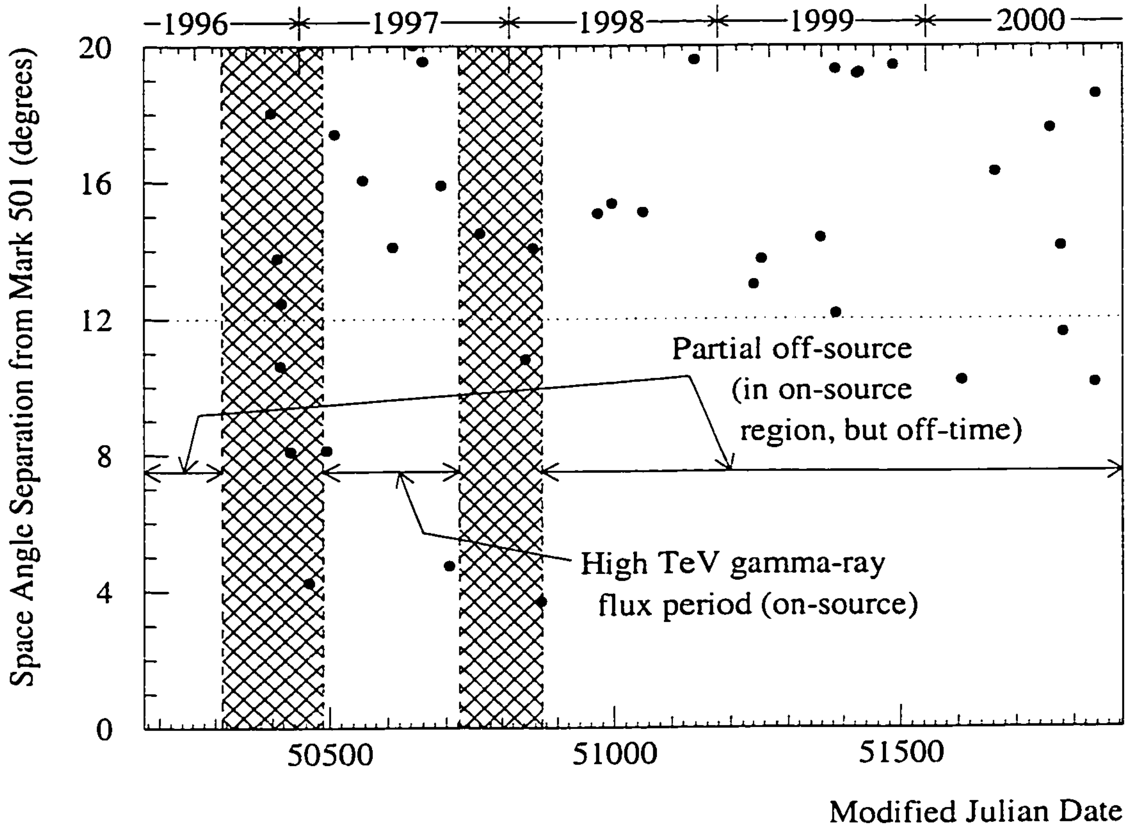

Figure 5.14 This shows the timing distribution of upward-going muons within 20 degrees (roughly 5 σ in estimated angular resolution) of space angle separation from Markarian 501 over the entire acquisition period of the data set. The flare apparently started and ended at times somewhere in the two hatched regions, which mark time periods when the γ-ray emission of the object could not be characterized very precisely because the flux was almost certainly changing quite drastically during some sub-interval of that time, yet no observations were being recorded with atmospheric Cherenkov telescopes. The two events observed near the source during the flare period, 8° away on February 16 and 5° away on September 14, 1997, are consistent with the background expectation.

a background estimate for the on-source time and direction must be provided. Drawing a series of off-source circles of 12° opening angle again, as shown in figure 5.9, 179 directionally off-source events are counted, as well as three off-time events in the Markarian 501 on-source circle. At Markarian 501's 39.755° declination, there is room for 10 such circles. Within the on-source circle, 0.1347 of the live time elapsed while Markarian 501 was flaring, and 0.6752 of the live time elapsed while it was relatively quiet; this latter time period, although it is associated directionally with the on-source region, should nevertheless, because of the time correlation aspect of the study, be counted toward the off-source exposure. Thus the on-source to off-source exposure ratio is $0.1347/10.6752 = 0.01262$. Taking into account the 182 background events stated above, a reasonable first guess for the expected number of events in the on-source circle during the time of the flare would be about 2.30, meaning that the two observed events apparently are actually a slight downward statistical fluctuation. Arguably, that fact alone is sufficient to demonstrate that the Super-Kamiokande experiment did not observe any neutrinos from Markarian 501 during the 1997 flaring period. However, for completeness of the analysis, and also to help contemplate what would have been required in order to claim such a discovery, a brief discussion of statistical significance is appended below.

5.4.2 Statistical Significance

There is a classic paper[124] which describes, for most cases of practical interest, how to estimate the statistical significance of the number of particle counts n observed in a putative on-source region, given an observation of N particles in the off-source region, and an on-source to off-source exposure ratio α (also called x/X in some of the earlier sections of this chapter). The details of that derivation are too complicated to repeat here, but the conclusion is that, if n and N are both drawn from Poisson distributions, and if both are sufficiently large, each greater than about 10 counts or so, and if the null hypothesis of no additional signal counts above background is true (i.e., $\langle n \rangle = \alpha \langle N \rangle$, rather than $\langle n \rangle > \alpha \langle N \rangle$), then in a statistical ensemble of repeated observations, the quantity

$$S = \sqrt{2}\left\{n \ln\left[\frac{1+\alpha}{\alpha}\left(\frac{n}{n+N}\right)\right] + N \ln\left[(1+\alpha)\left(\frac{N}{n+N}\right)\right]\right\}^{1/2} \tag{5.44}$$

will be distributed as the absolute value of a standard normal variable. Thus S is a measure of the statistical significance of the observation: for a given value S_0, the probability of drawing pairs of values n and N such that $S \geq S_0$ is $\int_{S_0}^{\infty} \sqrt{2/\pi} e^{-x^2/2}\, dx$.

In this case, for Markarian 501, $n = 2$ and $N = 182$, so strictly speaking, n is not large enough for S to be an ideal choice for a test statistic. Even so, S is still a useful quantity to calculate, because this analysis is primarily concerned with upward fluctuations of n, and in practice it is mostly for downward fluctuations that S becomes an inaccurate measure of the statistical significance. Figure 5.15 shows the result of two Monte Carlo studies designed to demonstrate this. To perform each study, pairs of numbers n and N were drawn at random from Poisson distributions with averages $\langle n \rangle$ and $\langle N \rangle$ chosen as shown in the figure, and S, the statistical significance, was estimated for each pair using equation 5.44. For purposes of these two studies, even though S was a positive definite quantity, it was allowed to keep its positive value only if it was associated with an upward fluctuation of n as compared to the expectation, i.e., only if $n > \alpha N$, and it was reassigned a negative value if it was associated with a downward fluctuation. To confirm that S is indeed normally distributed in simulated experimental practice (except, of course, as already mentioned, for downward fluctuations in the low $\langle n \rangle$ limit), integral histograms were constructed from the Monte Carlo distributions of S and compared to the analytic curves (i.e., $f(S) = \int_S^{\infty} (1/\sqrt{2\pi}) e^{-x^2/2}\, dx$ for $S > 0$, and $f(S) = \int_{-\infty}^{S} (1/\sqrt{2\pi}) e^{-x^2/2}\, dx$ for $S < 0$) that one would expect the integral histograms approximately to reproduce, assuming that the claim about S being normally distributed were true. Because the histograms represent integral distributions, the height of each bin is calculated by counting the total number of pairs n and N for which their particular S happens to be at least as extreme as (i.e., more positive if the bin abscissa is positive, and more negative if the bin abscissa is negative) the value of the abscissa associated with that bin, and then renormalizing that counted number against the total number of pairs used in the study. Thus, the height, or ordinate, of each bin represents a kind of numerical estimate of the total integrated probability (with the understanding that the integral is carried out to $+\infty$ for positive S and out to $-\infty$ for negative S, in order to be able to display the symmetry of the distribution) for actually obtaining, in a genuinely representative, albeit simulated, situation, a value of S at least as extreme as the one associated with that bin.

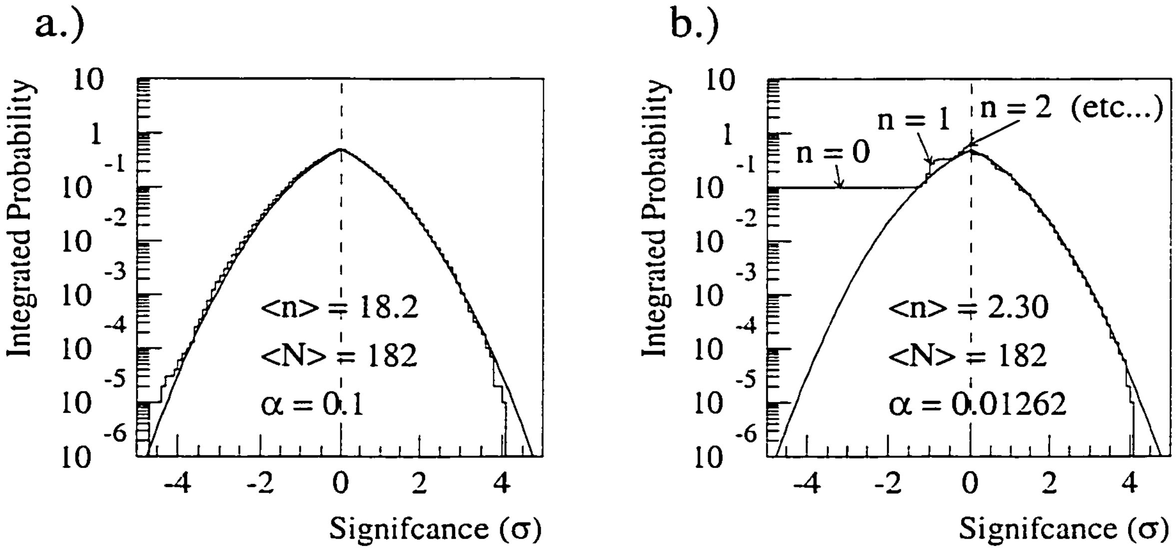

Figure 5.15 a.) The histogram in this plot gives, as described in detail in the main text, the total, integrated empirical Monte Carlo probability for randomly obtaining a value of S at least as extreme as the one associated with the abscissa of each bin, with S calculated using equation 5.44 and with n and N drawn randomly from Poisson distributions of means 18.2 and 182, respectively. The superimposed curve shows the shape that would be expected for this histogram if S were a normally distributed quantity. The close adherence between the histogram and the analytic curve confirms that, indeed, to a good approximation, S is normally distributed. b.) This plot shows the same quantities as the left one, but recalculated using $\langle n \rangle = 2.30$. For low values of $\langle n \rangle$, below about 10 or so, as opposed to the large $\langle n \rangle$ case depicted in the left plot, negative values of S are no longer normally distributed (nor indeed, for the $n = 0$ case, even mathematically defined by equation 5.44); this plot shows an example of such distortion. Fortunately, however, the persistent adherence of the Monte Carlo values to the predicted curve for $S > 0$ even in this case demonstrates empirically the final point of this whole exercise, that the expression of equation 5.44 is still valid for upward fluctuations of S even in the small $\langle n \rangle$ limit.

Ultimately, the expectation is that, if the entire scheme is self-consistent, the shape of the Monte Carlo histograms should closely follow that of the analytic Gaussian integral curves.

The first study, with $\langle n \rangle = 18.2$ and $\langle N \rangle = 182$, demonstrates that S is indeed normally distributed, at least in a typical case where $\langle n \rangle$ and $\langle N \rangle$ are purposefully chosen to be large enough to meet the criterion outlined in reference [124]. Also, the ability of the study to reproduce the expected behavior for S confirms that the basic methodology of the study is apparently correct, and any difference that may be witnessed in the other Monte Carlo study between the histogram and the expectation curve will be attributable solely to the difference in the choice of $\langle n \rangle$ for that study, rather than any kind of human error or misunderstanding that might have crept into the computer code that was used in the implementation. The second study, with $\langle n \rangle = 2.30$, demonstrates the point that initially launched this whole discussion, that values of S corresponding to downward fluctuations of n are, as expected, not normally distributed at all, but those above zero still retain this behavior, thus the S parameter remains a valid way to estimate statistical significance of upward fluctuations even in the low $\langle n \rangle$ limit. In principle, of course, if it were necessary, it would still be possible to calculate the statistical significance of the downward fluctuations correctly, even when both $\langle n \rangle$ and $\langle N \rangle$ were very small, by computing the absolute probabilities, directly from the Poisson distributions, of obtaining various combinations of n and N[125]. But such accuracy is not required for this analysis, and for simplicity's sake, the prescription of equation 5.44 will be adopted as the method of choice for the remainder of this discussion.

Moving on from this long digression about the validity of using S in equation 5.44 to estimate significance, and pushing forward toward a final conclusion, one may notice that several of the event directions nearest Markarian 501 in figure 5.14 fall into the null regions, during which time no photon observations were possible, and this presents a tempting question: suppose a few of those points had been observed, instead, in the on-source region; how many more would have been required in order to lay claim to a discovery of neutrino emission? Figure 5.16 addresses this by plotting S vs. n, for several n greater than the actually observed value of 2, under two slightly different assumptions about how n is varied. From the plot, one can see that, in this case, about eight or nine events in the on-source region (within 12° of Markarian 501) would have been required; even shifting the boundary

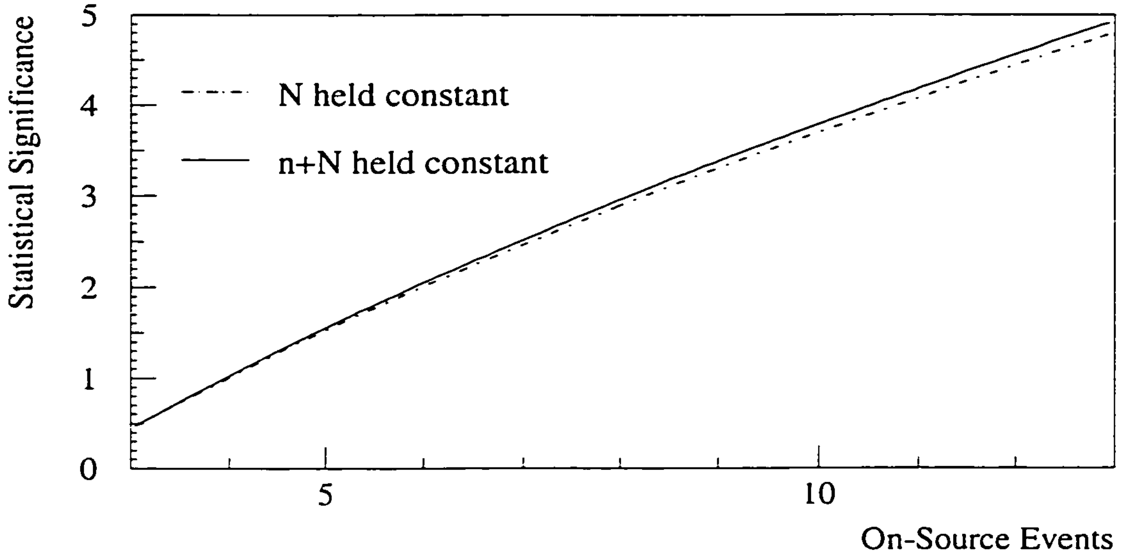

Figure 5.16 The statistical significance S of equation 5.44 is plotted here vs. n under two slightly different schemes: one in which N is held constant at $N = 182$ while n varies, and the other in which $n + N$ is held constant at $n + N = 184$ while both n and N are varied in opposition to each other.

lines of the on-source area by a little bit in order to include an extra event or two would not substantially change the conclusion: there is no evidence of neutrino flux from the 1997 flaring episode.

Chapter 6

Conclusions

6.1 Experimental Inferences

The final result of this analysis is a simple one: Super-Kamiokande finds no significant evidence of astrophysical point sources of neutrinos. This broad, general lack of evidence has been demonstrated by completing three different but all fairly general studies: first and second, by performing two distinct types of all-sky searches for anomalous clustering of neutrino events, and third, by setting upper limits on the maximum numbers of such events that could plausibly be coming from specific favored candidate celestial objects. Additionally, as part of a fourth study which was selected despite the limitations of its scope because of a perceived enhanced likelihood for returning non-null results, a search was attempted for evidence of time correlated neutrino emission from Markarian 501 during the extended γ-ray flare period of 1997, but none was found.

The two-point correlation function analysis, which constituted the first half of the all-sky clustering search, had the advantage of freedom from a priori bias toward any potential source object. The results of that analysis were that, at angles near the angular resolution of the experiment, the two-point correlation function demonstrated completely random correlations between the directions of pairs of neutrino-induced upward-going muon events, meaning that there was no anomalous clustering, and therefore no evidence of neutrino point sources.

A disadvantage of the two-point correlation function is that its behavior, resulting as it does from the interplay of overlapping and concentric rings, is not readily grasped in any simplified form by physical intuition. Also, it says nothing at all about which areas of the sky exhibit the tightest clustering, though such information might be desirable even if that clustering is not statistically significant. To address these shortcomings, as well as to be able to have a second and completely different type of test to support the initial results, a tiled sky map was created, showing the probability for each tile, or angular bin, of observing an upward fluctuation in the total number of events per bin at least as large as that which was actually observed. The results of that analysis were that the estimated observation probabilities for the entire collection of angular bins seemed entirely consistent with chance, with no evidence of anomalous clustering manifest.

The flux limit analysis, while producing less generally applicable results due its limited and direction-specific nature, nevertheless offered a different advantage not available from either of the first two studies: the chance to compare results with those of other experiments. At least six other underground neutrino observatories have published analyses similar to this one, and this analysis gets results comparable to theirs, yet with an increased sensitivity in most cases, consistent with the larger data set and exposure of this experiment. In general, the flux upper limits from most source candidates have been lowered by about an order of magnitude relative to those of four of the six earlier experiments, and without exception they beat the results of all six of the others at the highest northern declinations. Such a comparison demonstrates two things: first, the relative similarity of the results shows that no glaring errors have been made in this analysis, and second, the fact that many of the confidence limits and extrapolated luminosities are lower than those of any other experiment underscores the point that this detector is one of the most sensitive neutrino telescopes currently in operation. This second point is important, because it means that the lack of evidence that Super-Kamiokande sees for neutrino point sources is one of the most definitive statements made on the topic to date.

Last of all, the search for time correlated neutrino events associated with the extended flaring episode of Markarian 501 during 1997 also failed to produce a positive result. This search was of interest because during the months of the flare period, Markarian 501 became the brightest emitter in the sky of TeV γ-rays; it was thought that if this flare were accompanied by a similar upward fluctuation in neutrino flux in the same energy range, such a fortuitous occurrence would represent the best chance of this experiment ever to report a source discovery. Instead, however, it has turned out that the null result here just serves further to confirm those of the earlier sections.

6.2 Future Directions

Judging by the results of this analysis, future workers probably will not be able to lay claim to an unambiguous neutrino point source discovery without a substantial increase of some sort in detection capability. There is not even the barest hint in this data set, the largest of its type, of a statistically marginal possible source, and a modest factor of five or ten in size or observation time is not likely to make a difference in the final conclusions.

A complicating factor in building larger underground detectors of this type is that they are fundamentally limited in how large they can get by the mechanical stability of the surrounding rock. Furthermore, the recent spectacular demise, in a cataclysmic chain reaction of successive implosions, of 6779 of the 11,146 ID PMTs used in this experiment, as well as 1017 of the 1885 OD PMTs, has demonstrated the unexpected result that despite five years of successful operation, the overall mechanical design of Super-Kamiokande, as originally conceived at least, turns out not to be robust in the long term against catastrophic failure. Moreover, besides the fundamental engineering problems that must be addressed in any future experiments, the costs and technical logistics of excavating an arbitrarily large volume of rock scale in proportion to that volume, and impose separate limitations of their own.

Due to these limitations, benthic experiments[28][15] are seen as the wave of the future by many in the neutrino astrophysics community, and the evident need for a large jump in size suggested by this analysis supports the construction of such alternative types of detectors as the best hope for an astrophysical source discovery. The typical design of the deep underwater or ice experiments, the viability of which is currently being demonstrated by Baikal[67] and AMANDA[69], and which will likely soon be further validated by NESTOR[126] and ANTARES[127] as well, uses as its basis several strings of photomultiplier tubes placed at some depth, which are spaced apart from one another laterally to form a lattice. If immense size is the overarching consideration, this type of design is well suited to the problem, because all that is needed in order to increase the size of the detector at any time is merely to add another string to the edge of the sensitive volume.

Benthic neutrino telescopes have an additional advantage also: they can be operated at much higher minimum energy thresholds. In addition to enjoying somewhat looser constraints on their lateral spread than those which are imposed upon underground detectors, the vertical strings can be quite long as well, up to hundreds of meters, and since the threshold is determined by the minimum path length required of a final data sample event through the sensitive volume of the detector, it means that these experiments can have thresholds orders of magnitude greater than that of Super-Kamiokande. Such a capability is seen as a terrific advantage, since there is much speculation that astrophysical neutrinos, if they exist, might have a flatter energy spectrum than the atmospheric neutrinos in the background[28], and therefore even though the overall flux of both types of neutrinos must inevitably be greatly reduced at higher energies, imposing a higher threshold would dramatically improve the signal-to-noise ratio. The results from these future generations of neutrino detectors are anticipated with enthusiasm.

BIBLIOGRAPHY

[1] Kenneth Greisen. Possibility of high-energy neutrino measurements with cosmic rays. In Charlotte E. Mauk, Arthur H. Rosenfeld, and R. K. Wakerling, editors, *Proceedings of an International Conference on Instrumentation for High-Energy Physics*, number 1.h in V, pages 209–211, Berkeley, California, September 1960. Ernest O. Lawrence Radiation Laboratory, Interscience Publishers, Inc.

[2] Kenneth Greisen. Cosmic ray showers. *Annual Review of Nuclear Science*, 10:63–108, 1960.

[3] M. A. Markov. On high energy neutrino physics. In E. C. G. Sudarshen, J. H. Tinlot, and A. C. Melissinos, editors, *Proceedings of the 1960 Annual International Conference on High Energy Physics at Rochester*, S4, pages 578–581, Rochester, New York, August 1960. University of Rochester, Interscience Publishers, Inc.

[4] M. A. Markov and I. M. Zheleznykh. On high energy neutrino physics in cosmic rays. *Nuclear Physics*, 27(3):385–394, October 1961.

[5] Y. Fukuda. Solar neutrino data covering solar cycle 22. *Physical Review Letters*, 77(9):1683–1686, August 1996.

[6] T. Cleveland et al. Measurement of the solar electron neutrino flux with the Homestake chlorine detector. *Astrophysical Journal*, 496(1,1):505–526, March 1998.

[7] J. N. Abdurashitov et al. Measurement of the solar neutrino capture rate with gallium metal. *Physical Review C*, 60(5):055801(1–32), November 1999.

[8] M. Altmann et al. GNO solar neutrino observations: Results for GNO I. *Physics Letters B*, 490(1,2):16–26, September 2000.

[9] S. Fukuda et al. Constraints on neutrino oscillations parameters using 1258 days of Super-Kamiokande solar neutrino data. *Physical Review Letters*, 86(25):5656–5660, June 2001.

[10] Q. R. Ahmad et al. Measurement of the rate of $\nu_e + d \rightarrow p + p + e^-$ interactions produced by ^{8}B solar neutrinos at the Sudbury Neutrino Observatory. *Physical Review Letters*, 87(7):071301(1–6), August 2001.

[11] C. B. Bratton et al. Angular distribution of events frm SN1987A. *Physical Review D*, 37(12):3361–3363, June 1988.

[12] K. S. Hirata et al. Observation in the Kamiokande-II detector of the neutrino burst from supernova SN1987A. *Physical Review D*, 38(2):448–458, July 1988.

[13] E. N. Alexeyev. On possible detection of a neutrino burst on February 23, 1987 at the Baksan underground scintillation telescope. In J. Tran Thanh Van, editor, *The Standard Model, The Supernova 1987A, Proceedings of the Leptonic Session of the XXII Recontre de Moriond*, 1, pages 739–744, Les Arcs, Savoie, France, March 1987. Editiones Frontières.

[14] René A. Ong. Very high energy gamma-ray astronomy. *Physics Reports*, 305(3-4):93–202, November 1998.

[15] Francis Halzen. High-energy neutrino astronomy. *Physics Reports*, 333-334(1-6):349–364, August 2000.

[16] J. A. Simpson. Elemental and isotopic composition of the galactic cosmic rays. *Annual Review of Nuclear and Particle Science*, 33:323–381, 1983.

[17] M. Nagano and A. A. Watson. Observations and implications of the ultrahigh-energy cosmic rays. *Reviews of Modern Physics*, 72(3):689–732, July 2000.

[18] Roger Blandford and David Eichler. Particle acceleration at astrophysical shocks: A theory of cosmic ray origin. *Physics Reports*, 154(1):1–75, October 1987.

[19] W. I. Axford. The origins of high-energy cosmic rays. *Astrophysical Journal Supplement Series*, 90(2):937–944, February 1994.

[20] Peter L. Biermann. The origin of cosmic rays. *Space Science Reviews*, 74(3-4):385–396, November 1995.

[21] Herbert L. Anderson, editor. *A Physicist's Desk Reference.* American Institute of Physics, New York, 1989.

[22] Igor S. Grigoriev and Evgenii Z. Meilikhov, editors. *Handbook of Physical Quantities.* CRC Press, Boca Raton, 1997.

[23] J. G. Wilson. *Cosmic Rays.* Springer-Verlag, New York, 1976.

[24] Michael Catanese and Trevor C. Weekes. Very high energy gamma-ray astronomy. *Publications of the Astronomical Society of the Pacific*, 111(764):1193–1222, October 1999.

[25] C. M. Hoffman, C. Sinnis, P. Fleury, and M. Punch. Gamma-ray astronomy at high energies. *Reviews of Modern Physics*, 71(4):897–936, July 1999.

[26] Malcolm Longair. *High Energy Astrophysics, Volume 2: Stars the Galaxy, and the Interstellar Medium.* Cambridge University Press, Cambridge, 2nd edition, 1994.

[27] Catherine J. Cesarsky. Cosmic ray confinement in the galaxy. *Annual Review of Astronomy and Astrophysics*, 18:289–319, 1980.

[28] John G. Learned and Karl Mannheim. High-energy neutrino astrophysics. *Annual Review of Nuclear and Particle Science*, 50:679–749, 2000.

[29] Robert J. Gould and Gerald Schréder. Opacity of the universe to high-energy photons. *Physical Review Letters*, 16(6):252–254, February 1966.

[30] M. H. Salamon and F. W. Stecker. Absorption of high-energy gamma rays by interactions with extragalactic starlight photons at high redshifts and the high-energy gamma-ray background. *Astrophysical Journal*, 493(2):547–554, February 1998.

[31] J. M. Jauch and F. Rohrlich. *The Theory of Photons and Electrons.* Springer-Verlag, New York, 2nd edition, 1976.

[32] Raj Ghandi, Chris Quigg, Mary Hall Reno, and Ina Sarcevic. Neutrino interactions at ultrahigh energies. *Physical Review D*, 58(9):093009(1–15), November 1998.

[33] A. G. Lyne and F. Graham-Smith. *Pulsar Astronomy.* Cambridge University Press, Cambridge, 1990.

[34] Peter Goldreich and William H. Julian. Pulsar electrodynamics. *Astrophysical Journal*, 157(2):869–880, August 1969.

[35] J. E. Gunn and J. P. Ostriker. Acceleration of high-energy cosmic rays by pulsars. *Physical Review Letters*, 22(14):728–731, April 1969.

[36] V. S. Berezinsky and O. F. Prilutsky. Pulsars and cosmic rays in the dense supernova shells. *Astronomy and Astrophysics*, 66(3):325–334, June 1978.

[37] Humitaka Sato. Pulsars covered by the dense envelopes as high energy neutrino sources. *Progress of Theoretical Physics*, 58(2):549–559, August 1977.

[38] David Eichler and W. Thomas Vestrand. Implications of ultra-high-energy emisson from Hercules X-1. *Nature*, 318(6044):345–347, November 1985.

[39] V. S. Berezinskiĭ, S. V. Bulanov, V. A. Dogiel, and V. S. Ptuskin. *Astrophysics of Cosmic Rays.* North-Holland, Amsterdam, 1990.

[40] M. J. Rees and J. E. Gunn. The origin of the magentic field and relativistic particles in the Crab nebula. *Mothly Notices of the Royal Astronomical Society*, 167(1):1–12, April 1974.

[41] A. K. Harding and T. K. Gaisser. Acceleration by pulsar winds in binary systems. *Astrophysical Journal*, 358(2):561–574, August 1990.

[42] M. Samorski and W. Stamm. Detection of 2×10^{15} to 2×10^{16} eV gamma-rays from Cygnus X-3. *Astrophysical Journal*, 268(1):L17–L21, May 1983.

[43] J. Lloyd-Evans et al. Observation of γ rays $> 10^{15}$ eV from Cygnus X-3. *Nature*, 305(5937):784–787, October 1983.

[44] G. Battistoni et al. Observation of a time modulated muon flux in the direction of Cygnus X-3. *Physics Letters B*, 155(5,6):465–467, June 1985.

[45] M. L. Marshak et al. Time distributions for underground muons from the direction of Cygnus X-3. *Physical Review Letters*, 55(19):1965–1968, November 1985.

[46] T. K. Gaisser, Alice K. Harding, and Todor Stanev. Signatures of particle acceleration at SN 1987A. *Astrophysical Journal*, 345(1):423–433, October 1989.

[47] V. S. Berezinsky and V. S. Ptuskin. Radiation from young SN II shells produced by cosmic rays accelerated in shock waves. *Astronomy and Astrophysics*, 215(2):399–408, May 1989.

[48] Thomas K. Gaisser. *Cosmic Rays and Particle Physics.* Cambridge University Press, Cambridge, 1990.

[49] Marie-Helene Ulrich, Laura Maraschi, and C. Megan Urry. Variability of active galactic nuclei. *Annual Review of Astronomy and Astrophysics*, 35:445–502, 1997.

[50] Julian H. Krolik. *Active Galactic Nuclei: From the Central Black Hole to the Galactic Environment.* Princeton University Press, Princeton, 1999.

[51] Ian Robson. *Active Galactic Nuclei.* John Wiley & Sons, New York, 1996.

[52] John Kormendy and Douglas Richstone. Inward bound – the search for supermassive black holes in galactic nuclei. *Annual Review of Astronomy and Astrophysics*, 33:581–624, 1995.

[53] Demosthenes Kazanas and Donald C. Ellison. The central engine of quasars and active galactic nuclei: Hadronic interactions of shock accelerated relativistic protons. *Astrophysical Journal*, 304(1):178–187, May 1986.

[54] Mitchell C. Begelman, Bronisław Rudak, and Marek Sikora. Consequences of relativistic proton injection in active galactic nuclei. *Astrophysical Journal*, 362(1):38–51, October 1990.

[55] F. W. Stecker, C. Done, M. H. Salamon, and P. Sommers. High-energy neutrinos from active galactic nuclei. *Physical Review Letters*, 66(21):2697–2700, May 1991.

[56] F. W. Stecker, C. Done, M. H. Salamon, and P. Sommers. Errata: High-energy neutrinos from active galactic nuclei. *Physical Review Letters*, 69(18):2738, November 1992.

[57] Mitchell C. Begelman, Roger D. Blandford, and Martin J. Rees. Theory of extragalactic radio sources. *Reviews of Modern Physics*, 56(2):255–351, April 1984.

[58] K. Mannheim. The proton blazar. *Astronomy and Astrophysics*, 269(1-2):67–76, March 1993.

[59] C. Megan Urry and Paolo Padovani. Unified schemes for radio-loud active galactic nuclei. *Publications of the Astronomical Society of the Pacific*, 107(715):803–845, September 1995.

[60] H. Adarkar et al. Search for point sources of neutrinos with the KGF underground muon detectors. *Nuovo Cimento*, 21C(6):661–666, Novembre-Dicembre 1998.

[61] E. N. Alexseev et al. The Baksan underground scintillation telescope. *Physics of Particles and Nuclei*, 29(3):254–256, May-June 1998.

[62] Ch. Berger et al. A search for high energy neutrinos from SN1987A, the Crab, Hercules X-1, and Cynus X-3 with the Fréjus detector. *Zeitschrift für Physik C*, 48(2):221–226, October 1990.

[63] R. Svoboda et al. An upper limit on the flux of extraterrestrial neutrinos. *Astrophysical Journal*, 315(2):420–424, April 1987.

[64] Y. Oyama et al. Experimental study of upward-going muons in Kamiokande. *Physical Review D*, 39(6):1481–1491, March 1989.

[65] M. Ambrosio et al. Neutrino astronomy with the MACRO detector. *Astrophysical Journal*, 546(2):1038–1054, January 2001.

[66] I. A. Belolaptikov. The Baikal underwater neutrino telescope: Design, performance, and first results. *Astroparticle Physics*, 7(3):263–282, August 1997.

[67] V. A. Balkanov. Registration of atmospheric neutrinos with the BAIKAL neutrino telescope NT-96. *Astroparticle Physics*, 12(1-2):75–86, October 1999.

[68] E. Andres et al. The AMANDA neutrino telescope: Principle of operation and first results. *Astroparticle Physics*, 13(1):1–20, March 2000.

[69] Scott Matthew Young. *A Search for Point Sources of High Energy Neutrinos with the AMANDA-B10 Neutrino Telescope*. Ph.D. thesis, University of California, Irvine, Department of Physics, 2001.

[70] J. D. Jackson. *Classical Electrodynamics.* John Wiley and Sons, Inc., New York, 3rd edition, 1999.

[71] D. E. Groom et al. *European Physical Journal*, C15(1), 2000.

[72] A. Suzuki et al. Improvement of 20 in. diameter photomultiplier tubes. *Nuclear Instruments and Methods in Physics Research*, A329(1,2):299–313, May 1993.

[73] H. Kume et al. 20 inch diameter photomultiplier. *Nuclear Instruments and Methods in Physics Research*, 205(3):443–449, January 1986.

[74] Syunsuke Kasuga. *Observation of a Small ν_μ/ν_e Ratio of Atmospheric Neutrinos in Super-Kamiokande by the Method of Particle Identification.* Ph.D. thesis, University of Tokyo, Department of Physics, January 1998.

[75] R. Becker-Szendy et al. IMB-3: A large water Cherenkov detector for nucleon decay and neutrino interactions. *Nuclear Instruments and Methods in Physics Research*, A324(1,2):363–382, January 1993.

[76] H. Ikeda et al. Front-end hybrid circuit for Super-Kamiokande. *Nuclear Instruments and Methods in Physics Research*, A320(1,2):310-316, August 1992.

[77] TKO specification. KEK Report, High Energy Accelerator Research Organization (KEK), Tsukuba, Japan, 1985. Identification Numbers: 198524010, KEK-85-10.

[78] Mark D. Messier. *Evidence for Neutrino Mass from Observations of Atmospheric Neutrinos with Super-Kamiokande.* Ph.D. thesis, Boston University, Department of Physics, 1999.

[79] Jamie Shiers, Helge Meinhard, Otto Schaile, and Michel Goossens. *The ZEBRA System.* CERN, CN Division, Program Library Office, Geneva, Switzerland, February 1995.

[80] John W. Flanagan. *A Study of Atmospheric Neutrinos at Super-Kamiokande.* Ph.D. thesis, University of Hawaii, Department of Physics, December 1997.

[81] Masato Shiozawa. *Search for Proton Decay via* $p \rightarrow e^+\pi^0$ *in a Large Water Cherenkov Detector.* Ph.D. thesis, University of Tokyo, Department of Physics, November 1999.

[82] P. Hänggi, R. D. Viollier, U. Raff, and K. Alder. Muon decay in orbit. *Physics Letters*, 51B(2):119–122, July 1974.

[83] M. Nakahata et al. Calibration of Super-Kamiokande using an electron linac. *Nuclear Instruments and Methods in Physics Research*, A421(1,2):113–129, January 1999.

[84] E. Blaufuss et al. ^{16}N as a calibration source for Super-Kamiokande. *Nuclear Instruments and Methods in Physics Research*, A458(3):638–649, February 2001.

[85] Shuichiro Hatakeyama. *Search for Muon Neutrino Oscillations in Kamiokande and Super-Kamiokande.* Ph.D. thesis, Tohoku University, Department of Physics, February 1998.

[86] Y. Fukuda et al. Measurement of the flux and zenith-angle distribution of upward throughgoing muons by Super-Kamiokande. *Physical Review Letters*, 82(13):2644–2548, March 1999.

[87] Bradley Efron and Robert J. Tibshirani. *An Introduction to the Bootstrap.* Chapman & Hall, New York, 1993.

[88] Phillip I. Good. *Resampling Methods A Practical Guide to Data Analysis.* Birkhäuser, Boston, 1999.

[89] Y. Fukuda et al. Evidence for oscillation of atmospheric neutrinos. *Physical Review Letters*, 81(8):1562–1567, August 1998.

[90] P. J. E. Peebles. *The Large Scale Structure of the Universe.* Priceton University Press, Princeton, N.J., 1980.

[91] S. Ahlen et al. Muon astronomy with the MACRO detector. *Astrophysical Journal*, 412(1):301–311, July 1993.

[92] Glen Cowan. *Statistical Data Analysis.* Clarendon Press, Oxford, 1998.

[93] Harold Jeffreys. *Theory of Probability.* Oxford University Press, London, 2nd edition, 1948.

[94] R. J. Protheroe. Statistical analysis in gamma-ray astronomy: The upper limit. *Astronomy Express*, 1(1):33–38, July 1984.

[95] T. J. Loredo. From Laplace to supernova 1987A: Bayesian inference in astrophysics. In P. F. Fougère, editor, *Maximum Entropy and Bayesian Methods*, pages 81–142. Kluwer Academic Publishers, Dordrecht, Netherlands, 1990.

[96] Thierry J.-L. Courvoisier. The bright quasar 3C 273. *Astronomy and Astrophysics Review*, 9(1-2):1–32, December 1998.

[97] Wendy L. Freedman. The Hubble constant and the expansion age of the universe. *Physics Reports*, 333-334(1-6):13–31, August 2000.

[98] Nino Panagia. New distance determination to the LMC. *Memorie della Società Astronomica Italiana*, 69(1):225–235, 1998.

[99] Martin D. Weinberg and Sergei Nikolaev. Structure of the Large Magellanic Cloud from 2MASS. *Astrophysical Journal*, 548(2):712–726, February 2001.

[100] David Legge. Astrometry of southern pulsars. In M. Kramer, N. Wex, and R. Wielebinski, editors, *Pulsar Astronomy-2000 and Beyond, IAU Colloquium 177,*

Astronomical Society of the Pacific Conference Series, 202, pages 141–142, Bonn, Germany, September 1999. Max Planck Institut für Radioastronomie, Sheridan Books.

[101] H. Jerjen, K. C. Freeman, and B. Binggeli. Testing the surface brightness fluctuations method for dwarf elliptical galaxies in the Centaurus A group. *Astronomical Journal*, 119(1):166–176, January 2000.

[102] B. Cameron Reed. Vela OB1: Probable new members and Hertzsprung-Russell diagram. *Astronomical Journal*, 119(4):1855–1859, April 2000.

[103] Mark J. Reid. The distance to the center of the galaxy. *Annual Review of Astronomy and Astrophysics*, 31:345–372, 1993.

[104] C. F. Bradshaw, E. B. Fomalont, and B. J. Geldzahler. High resolution parallax measurements of Scorpius X-1. *Astrophysical Journal*, 512(2):L121–L124, February 1999.

[105] Wolfgang Kundt. SS433. *Memorie della Società Astronomica Italiana*, 70(3/4):1097–1103, 1999.

[106] Patrizia A. Caraveo, Giovanni F. Bignami, Roberto Mignani, and Laurence G. Taff. Parallax observations with the Hubble Space Telescope yield the distance to Geminga. *Astrophysical Journal*, 461(2):L91–L94, April 1996.

[107] P. A. Caraveo, G. F. Bignami, R. Mignani, and L. G. Taff. Geminga in the space telescope era. *Astronomy and Astrophysics Supplement Series*, 120(4):65–68, December 1996.

[108] Virginia Trimble. The distance to the Crab nebula and NP0532. *Publications of the Astronomical Society of the Pacific*, 85(507):579–585, October 1973.

[109] Ian D. Howarth and Bob Wilson. A study of the low-mass X-ray binary HZ Her/Her X-1 using IUE and optical data. *Monthly Notices of the Royal Astronomical Society*, 202(1):347–366, January 1983.

[110] John M. Dickey. A new distance to Cygnus X-3. *Astrophysical Journal*, 273(2):L71–L73, October 1983.

[111] M. F. Cawley and T. C. Weekes. Gamma ray absorption and the distance to Cygnus X-3. *Astronomy and Astrophysics*, 133(1):80–82, April 1984.

[112] Kenneth R. Lang. *Astrophysical Data*. Springer-Verlag, New York, 1992.

[113] T. C. Weekes et al. Observation of TeV gamma rays from the Crab nebula using the atmospheric Cerenkov imaging technique. *Astrophysical Journal*, 342(1):379–395, July 1989.

[114] J. Quinn et al. The flux variability of Markarian 501 in very high energy gamma rays. *Astrophysical Journal*, 518(2):693–698, June 1999.

[115] F. A. Aharonian et al. The temporal characteristics of the TeV gamma-radiation from Mkn 501 in 1997. *Astronomy and Astrophysics*, 342(1):69–86, February 1999.

[116] A. Djannati-Atai et al. Very high energy gamma-ray spectral properties of Mkn 501 from CAT cerenkov telescope observations in 1997. *Astronomy and Astrophysics*, 350(1):17–24, October 1999.

[117] N. Hayashida et al. Observations of TeV gamma-ray flares from Markarian 501 with the Telescope Array prototype. *Astrophysical Journal*, 504(2):L71–L74, September 1998.

[118] R. Atkins et al. TeV observations of Markarian 501 with the Milagrito water cerenkov detector. *Astrophysical Journal*, 525(1):L25–L28, November 1999.

[119] M. Amenomori et al. Detection of multi-TeV gamma-rays from Markarian 501 during an unforeseen flaring state in 1997 with the Tibet air shower array. *Astrophysical Journal*, 532(1):302–307, March 2000.

[120] A. M. Anokhina et al. The application of the multidimensional analysis to the data of Mk 501 observations. *Journal of Physics G*, 26(1):57–65, Jan 2000.

[121] P. M. Chadwick et al. Observations of TeV gamma rays from Markarian 501 at large zenith angles. *Jounal of Physics G*. 25(8):1749–1753, August 1999.

[122] A. K. Tickoo et al. Drive-control system for the TACTIC gamma-ray telescope. *Experimental Astronomy*, 9(2):81–101, 1999.

[123] S. M. Bradbury et al. Detection of γ-rays above 1.5 TeV from Mkn 501. *Astronomy and Astrophysics*, 320(1):L5–L8, April 1997.

[124] Ti-pei Li and Yu-qian Ma. Analysis methods for results in gamma-ray astronomy. *Astrophysical Journal*, 272(1):317–324, September 1983.

[125] D. E. Alexandreas et al. Point source search techniques in ultra high energy gamma ray astronomy. *Nuclear Instruments and Methods in Physics Research*, A328(3):570–577, May 1993.

[126] L. Trasatti et al. NESTOR: An underwater neutrino observatory in the Ionian sea. In A. Bottino, A. di Credico, and P. Monacelli, editors, *Nuclear Physics B Proceedings Supplements: Taup97: Proceedings of the Fifth International Workshop on Topics in Astroparticle and Underground Physics*, 70, pages 442–444, Italy, September 1997. Laboratori Nazionali del Gran Sasso, North-Holland.

[127] F. Montanet et al. Design and expected performance of the ANTARES neutrino telescope. In J. Dumarchez, M. Froissart, and D. Vignaud, editors, *Nuclear Physics*

B Proceedings Supplements: Taup99: Proceedings of the Sixth International Workshop on Topics in Astroparticle and Underground Physics, 87, pages 436–438, France, September 1999. College de France, North-Holland.

Appendix A

Not Quite Trivial Mathematical Facts

A.1 The Area of a Lune

Figure A.1 depicts a lune of width a and length $2b$, demarcated by points JOLMN. Its area is readily calculated by subtracting the area of the isosceles triangle JKLO from the area of the wedge JKLMN. From the diagram, the area of the triangle is just $A_{JKLO} = b(r-a)$. The area of the wedge is

$$\begin{aligned} A_{JKLMN} &= \frac{2\theta}{2\pi}\pi r^2 \\ &= r^2 \arccos\left(\frac{r-a}{r}\right). \end{aligned} \tag{A.1}$$

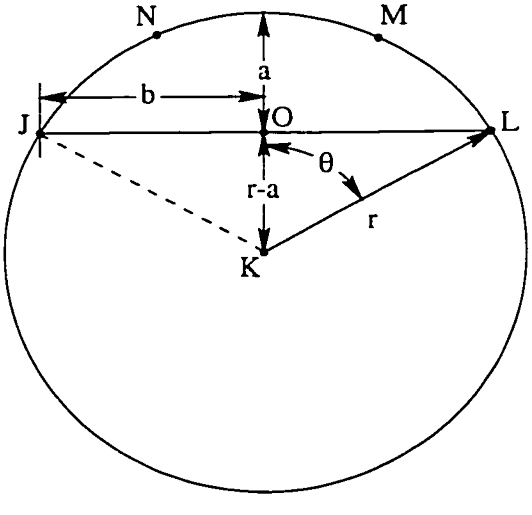

Figure A.1 The lune is the shape defined by JOLMN. To calculate its area, subtract the area of the isosceles triangle JKLO from the area of the wedge JKLMN.

Thus, the area of the lune is just

$$A_{JOLMN} = r^2 \arccos\left(\frac{r-a}{r}\right) - b(r-a). \tag{A.2}$$

This relation is used in calculating the effective area of the detector in section 4.2.

A.2 Estimating Angular Resolution When the True Track Direction Isn't Known

An important factor affecting the final results of this analysis is the whole subtopic of the direction of the upward-going muons tracks: how the arrival angles are measured, what the resolution for that process is, and how that resolution is estimated. Conceptually, the simplest way to estimate the resolution of a fitting process is to fit a bunch of events for which directions are already very precisely known, with much higher accuracy than can be offered by the fitting process being tested, and then simply quantify how close the fitting process came, on average, to guessing the correct value. Such studies can and have been done[86] on Monte Carlo data samples, but for fitting processes like manual reconstruction, which have a randomized element to their functioning that causes them to return a slightly different result each time they are run on an event, a second method is also available: resolution may be estimated by quantifying how much two fit results from the same event differ from each other, without necessarily knowing the true correct answer that should have been returned for that event. The results of this section will be developed with the manual fitting process kept in mind as the eventual application, but in principle they are relevant to any fitting process which does not produce its results deterministically.

The overall scale of the angular resolution will be estimated by measuring, or fitting, a track direction for each event twice, plotting the distribution of the space angle θ between each pair of track directions, and then extracting a characteristic angle of resolution σ as a parameter in some fit function, derived below, which is adjusted to match the shape of the observed space angle distribution. The projected RMS angular deviation for each of the two separate event track fits will be assumed to be normally distributed for each degree of freedom. For example, the first fit will be assumed to vary independently by two small angles x_1 and y_1 which measure the projected angular deviation, in two directions which are mutually orthogonal to each other as well as to the true track direction vector $\vec{v}_t$ of the event, of the fitted direction from the true direction. So, if x_1 is taken as the projection, along the first of the two orthogonal axes of deviation, of the space angle $\theta_1 = \arccos(\frac{\vec{v}_t \cdot \vec{v}_1}{|\vec{v}_t||\vec{v}_1|})$ between $\vec{v}_t$ and the first fit direction vector $\vec{v}_1$, then the distribution of x_1 about $\vec{v}_t$ is assumed to be

$$P(x_1) = \frac{1}{\sigma\sqrt{2\pi}} e^{-x_1^2/2\sigma^2}\, dx_1. \tag{A.3}$$

The distributions of y_1, as well as x_2 and y_2, the projected RMS angular deviations of the second fit, are of course all defined similarly.

Because x_1 and the other projected angles are typically of order 1° or 2°, the space angle difference between the fitted direction of the first fit and its true value is small enough to be approximated simply as $\theta_1 = \sqrt{x_1^2 + y_1^2}$, and the same also holds true for the second fit. To see why this is true, imagine transforming the true fit direction vector $\vec{v}_t$ into a basis in which it lies entirely along the z-axis. In such a coordinate system, the space angle θ_1 is just what is normally identified in spherical coordinates as the polar, or zenith angle θ, without any subscript attached; the subscript will remain in place for the remainder of this discussion, however, because plain unadorned θ, by itself without a subscript, has already been reserved for other use as described above. In this basis, the azimuthal angle will be ϕ_1, and the projections of $\vec{v}_1/|\vec{v}_1|$ onto the x-, y-, and z-axes will be $\sin\theta_1 \cos\phi_1$, $\sin\theta_1 \sin\phi_1$, and $\cos\theta_1$. The angles x_1 and y_1, the projections of θ_1 into the xz and yz planes in this basis, can be found by taking tangents: $\tan x_1 = \sin\theta_1 \cos\phi_1/\cos\theta_1$ and $\tan y_1 = \sin\theta_1 \sin\phi_1/\cos\theta_1$. But, as has been mentioned earlier, x_1, y_1, and θ_1 are all small angles, thus to first order, $\tan x_1 \approx x_1$, $\tan y_1 \approx y_1$, $\cos\theta_1 \approx 1$, and $\sin\theta_1 \approx \theta_1$, and by substituting these approximations for the true expressions, it is seen that $x_1^2 + y_1^2 = \theta_1^2(\cos^2\phi_1 + \sin^2\phi_1)$, or $\theta_1 = \sqrt{x_1^2 + y_1^2}$, as claimed.

Since θ_1 and θ_2 are each sums in quadrature of two independent variables drawn from distributions of characteristic width σ, it is customary in this type of discussion to refer to the typical size of each as $\sqrt{2}\sigma$, since that is the result of adding σ to itself in quadrature. It is this quantity, $\sqrt{2}\sigma$, rather than σ itself, that is of the most interest to the rest of the analysis, because it is the most appropriate quantity to quote as the angular resolution of the track fitting process. In practice, of course, neither θ_1 nor θ_2 are known on an event-by-event basis because $\vec{v}_t$ (that is, $\vec{v}_t$ as expressed in detector coordinates, rather than in whatever system would have made $\vec{v}_t$ parallel to the z-axis) is never known for real, non-Monte Carlo data events; what is actually measured is the space angle θ between the two fits, which can also be expressed in the small angle approximation, as $\theta = \sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2}$. To extract a value of σ from the fitted event directions, the shape of the distribution of θ must be predicted as a function of σ, and the best value of σ will be chosen as the one that most closely matches the observed distribution of the data.

The theoretically expected distribution of θ may be derived, essentially, by working up to it in small steps. One may start by defining an extra variable, $X = (x_1 - x_2)$, and pointing out that, like x_1 and x_2, it is also normally distributed, albeit with a slightly different width. The probability density $P(X)$ may be calculated by integrating over all the probabilities of choosing values for x_1 and x_2 which differ by exactly X. Therefore,

$$\begin{aligned} P(X) &= \frac{1}{2\pi\sigma^2} \int_{-\infty}^{\infty} e^{-x_1^2/2\sigma^2} e^{-(x_1 - X)^2/2\sigma^2} \, dx_1 \\ &= \frac{1}{2\pi\sigma^2} e^{-X^2/4\sigma^2} \int_{-\infty}^{\infty} e^{-(\sqrt{2}x_1 - X/\sqrt{2})^2/2\sigma^2} \, dx_1 \\ &= \frac{1}{2\sqrt{\pi}\sigma} e^{-X^2/4\sigma^2}, \end{aligned} \tag{A.4}$$

and the corresponding probability density for the y variables will be signified by $P(Y)$.

Introducing the variables $R = X^2$ and $S = Y^2$, one may calculate the distributions of R and S with a simple change of variables. Noting that $dR = 2X\,dX$, and thus that $dX = dR/2R$, the substitution can readily be made:

$$\begin{aligned} P(R)\,dR &= 2 \cdot \frac{1}{2\sqrt{\pi}\sigma} e^{-R/4\sigma^2} \frac{dR}{2\sqrt{R}} \\ &= \frac{e^{-R/4\sigma^2}}{2\sigma\sqrt{\pi R}} \, dR, \end{aligned} \tag{A.5}$$

and again there is a corresponding expression for $P(S)\,dS$. The extra factor of 2 written explicitly at the front of the expression on the first line comes from the fact that there are two values of X (one positive and one negative) corresponding to every value of R.

To carry the line of reasoning one step further, one may define a third variable,

$$T = R + S = X^2 + Y^2 = (x_1 - x_2)^2 + (y_1 - y_2)^2,$$

and derive its distribution. In analogy with the probability distribution of X derived above, that of T is just the integral over all the probabilities for choosing values of R and S which together add up to T. Written explicitly, it is

$$\begin{aligned} P(T) &= \frac{1}{4\pi\sigma^2}\int_0^T \frac{e^{-R/4\sigma^2}}{\sqrt{R}}\frac{e^{-(T-R)/4\sigma^2}}{\sqrt{(T-R)}}\,dR \\ &= \frac{e^{-T/4\sigma^2}}{4\pi\sigma^2}\int_0^T \frac{dR}{\sqrt{R(T-R)}} \\ &= \frac{e^{-T/4\sigma^2}}{4\sigma^2}. \end{aligned} \tag{A.6}$$

The final step is simply to recognize that θ, the space angle between the two fit results, defined above, is just $\sqrt{T}$, and make one more variable substitution. Since $dT = 2\theta\, d\theta$, the final substitution gives

$$P(\theta)\,d\theta = \frac{\theta e^{-\theta^2/4\sigma^2}}{2\sigma^2}\,d\theta. \tag{A.7}$$

A fit to this function of the distribution of space angle differences between pairs of manually fitted muon track directions is shown in section 4.4.

A.3 The Convolution of Two Independent Poisson Distributions

Suppose that k_1 and k_2 are independent, Poisson distributed variables, i.e., that the probability of measuring k_1 counts in some kind of counting experiment is

$$P(k_1;\mu_1) = \frac{\mu_1^{k_1}e^{-\mu_1}}{k_1!}, \tag{A.8}$$

and $P(k_2;\mu_2)$ has the identical form. If these two variable are added together to form a third variable, k, it can be shown that k is Poisson distributed as well.

The probability for adding randomly drawn k_1 and k_2 together to get any given result for k can be calculated by adding together all of the individual joint probabilities for choosing k_1 and k_2 such that their sum is k. This may be written as

$$P(k = k_1 + k_2;\mu_1,\mu_2) = \sum_{\substack{k_1=0,\\ k_2=k-k_1}}^{k_1=k} \frac{\mu_1^{k_1}e^{-\mu_1}}{k_1!}\frac{\mu_2^{k_2}e^{-\mu_2}}{k_2!}. \tag{A.9}$$

Imposing the constraint that $k_2 = k - k_1$ explicitly, the expression can be simplified to find that

$$\begin{aligned} P(k; \mu_1, \mu_2) &= \sum_{k_1=0}^{k} \frac{\mu_1^{k_1} e^{-\mu_1}}{k_1!} \frac{\mu_2^{k-k_1} e^{-\mu_2}}{(k-k_1)!} \\ &= \left[\sum_{k_1=0}^{k} \frac{k!}{(k-k_1)!k_1!} \mu_1^{k_1} \mu_2^{k-k_1} \right] \frac{e^{-(\mu_1+\mu_2)}}{k!} \\ &= \frac{(\mu_1+\mu_2)^k e^{-(\mu_1+\mu_2)}}{k!}, \end{aligned} \tag{A.10}$$

where it has been recognized that the quantity in square brackets on the second line is simply the binomial expansion. This fact is applied in section 5.1.1 to justify the simplification from equation 5.2 to 5.3.

VITA

Andrew Lawrence Stachyra was born in Chicago, Illinois on August 31, 1971, to Frank E. and Ann M. Stachyra, and grew up during most of his childhood in the nearby suburb of Oak Park. He received a B.S. in physics from Yale University in 1993 and an M.S. in physics from the University of Washington in 1994. His second proudest achievement during his time as a student at UW was attaining the summit of Mt. Rainier on August 13, 2000. At this point in life, he is contemplating what mountain to climb next.