

STACEE OBSERVATIONS OF BL LAC OBJECTS

Thomas Lindner
Department of Physics
McGill University, Montreal
August, 2006

A thesis submitted to McGill University
in partial fulfillment of the requirements
of the degree of Doctor of Philosophy

©Thomas Lindner, 2006



Library and
Archives Canada

Bibliothèque et
Archives Canada

Published Heritage
Branch

Direction du
Patrimoine de l'édition

395 Wellington Street
Ottawa ON K1A 0N4
Canada

395, rue Wellington
Ottawa ON K1A 0N4
Canada

Your file Votre référence

ISBN: 978-0-494-32209-3

Our file Notre référence

ISBN: 978-0-494-32209-3

NOTICE:

The author has granted a non-exclusive license allowing Library and Archives Canada to reproduce, publish, archive, preserve, conserve, communicate to the public by telecommunication or on the Internet, loan, distribute and sell theses worldwide, for commercial or non-commercial purposes, in microform, paper, electronic and/or any other formats.

The author retains copyright ownership and moral rights in this thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without the author's permission.

AVIS:

L'auteur a accordé une licence non exclusive permettant à la Bibliothèque et Archives Canada de reproduire, publier, archiver, sauvegarder, conserver, transmettre au public par télécommunication ou par l'Internet, prêter, distribuer et vendre des thèses partout dans le monde, à des fins commerciales ou autres, sur support microforme, papier, électronique et/ou autres formats.

L'auteur conserve la propriété du droit d'auteur et des droits moraux qui protègent cette thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

In compliance with the Canadian Privacy Act some supporting forms may have been removed from this thesis.

Conformément à la loi canadienne sur la protection de la vie privée, quelques formulaires secondaires ont été enlevés de cette thèse.

While these forms may be included in the document page count, their removal does not represent any loss of content from the thesis.

Bien que ces formulaires aient inclus dans la pagination, il n'y aura aucun contenu manquant.


Canada

For Paula, John, Mum and Dad.

It is only too likely that neither of us has any knowledge to boast of: but he thinks he knows something which he does not know, whereas I am quite conscious of my ignorance. At any rate it seems that I am wiser than he is to this small extent, that I do not think that I know what I do not know.

— Socrates

Abstract

STACEE is an atmospheric Cherenkov telescope that was built to detect γ -rays of cosmic origin. By using a large primary mirror area we have reduced the detector's energy threshold to below 200 GeV. The present work continues STACEE's investigation of BL Lac objects, focusing on Markarian 421, 3C 66A and OJ 287. An important part of this work has been the development of software techniques for the correction of systematic biases and the improvement of background rejection. As a result of these improvements, STACEE made a convincing detection of Markarian 421 during the 2002-2004 seasons. The statistical excess was 3.9σ for the 2002-2003 season and 10.9σ for the 2003-2004 season; the latter represents the strongest single season detection ever achieved by STACEE. We compare our measured Mrk 421 rates to those of other instruments and present spectral measurements during various flaring states in 2003-2004. The LBLs 3C 66A and OJ 287 were not detected. We discuss the significance of our flux upper limits in the context of emission models for these two sources.

Résumé

STACEE est un télescope à effet Cherenkov atmosphérique construit pour détecter des rayons γ d'origine cosmique. Nous avons abaissé le seuil énergétique de l'expérience en dessous de 200 GeV en utilisant un miroir primaire de grande surface. Le travail présenté ici approfondit les recherches de STACEE concernant les objets BL Lac, principalement Markarian 421, 3C 66A et OJ 287. Une partie importante de ce travail a porté sur le développement de techniques logicielles pour la correction de problèmes systématiques ainsi que l'amélioration de la rejection du bruit de fond. Grâce à ces améliorations, STACEE a détecté de façon convaincante Markarian 421 durant les saisons 2002-2004. Nous avons obtenu un excès statistique de 3.9σ pour la saison 2002-2003 et un excès de 10.9σ pour la saison 2003-2004; ce dernier relevé représente la détection saisonnière la plus significative que STACEE ait jamais réalisée. Nous comparons ici nos mesures concernant les taux de Mrk 421 avec ceux provenant d'autres instruments et nous présentons nos mesures spectrales durant différentes périodes de haute activité de la saison 2003-2004. Les objets LBLs 3C 66A et OJ 287, eux, n'ont pas été détectés. Nous discutons ici de la signification des limites supérieures de nos flux dans le contexte de modèles d'émission provenant de deux sources différentes.

Acknowledgements

This thesis would not have been possible without a lot of help from others. I'll start by thanking my supervisor, David Hanna. David taught me a great deal about the practical side of being an experimental physicist, as well as about the difference between *will* and *shall*. He provided me with a tremendous amount of freedom to decide how to organize myself; this was the perfect environment for me to learn and work. This thesis would not have occurred without David's encouragement. Much the same is true of my unofficial supervisor, Ken Ragan, who was always available for advice of any kind.

I am also thankful for the help of my other STACEE colleagues at McGill. John Kildea was an essential part of many aspects of this thesis, as well as being a good friend. Pascal Fortin provided me with regular doses of skepticism and Carsten Mueller with regular doses of humour. It was from them that I learned most of the details about how STACEE really worked.

The STACEE collaboration is a small group and it has been a pleasure to work as part of it. Every member of the group has happily responded to my pestering emails over the years. I'd particularly like to thank Reshmi Mukherjee for the week of collaboration in New York, as well as Jesse Ball and Alex Jarvis for putting up with me during many shifts. Alex also taught us what a 'moo point' is (it's a point made by a cow: it doesn't matter).

I want to thank Wei Cui for providing many of the Whipple results that are used in this thesis.

I also need to thank all the many people at the NSTTF for making STACEE possible. In addition, Paul Mercure, our system administrator at McGill, spent an inordinate amount of time over the last few years fixing my computer problems (many of which were my fault), a fact that I am very grateful for.

During my PhD I received financial support from the Carl Reinhardt McGill Major fellowship and our group received funding from NSERC and FQRNT. The support from these groups is greatly appreciated.

My friends at McGill made day-to-day work much more enjoyable and helped me

explore Montreal. They also often set me straight on various basic points of astronomy and physics. I also want to thank the Chiefs for being the only hockey team in the world that would take a player like me.

My family has been a great support in the writing of my thesis. In addition to the tedious correction of my spelling mistakes, my mother, father and brother have also provided encouragement whenever I felt that the task might be impossible. It'd also like to thank the Navratil family who fed and entertained me during the long months of writing.

Finally, and most importantly, this thesis would not have been possible without the love and support of Paula, who rarely complained about the fact that I was spending so much time with someone called 'Stacey'.

Statement of original work

I first joined the STACEE collaboration in September 2002. STACEE was a well established experiment, whose construction was essentially complete. The majority of my effort has therefore involved making observations, calibrating the detector and developing software analysis tools. I spent a significant amount of time on shift in Albuquerque, New Mexico. Shift time was spent making observations of astronomical sources and calibrating, repairing and replacing various elements of the STACEE detector. One particular task that I concentrated on was the calibration of the PMT gains. I implemented a system for the continuous monitoring of the relative gains; I also took over the task of periodically increasing the PMT high voltage in order to maintain a constant absolute gain.

Another important task has been the development of a more organized process for analyzing STACEE data. There are a number of different programs involved in STACEE data analysis, each written by different people. To ensure the proper functioning of these programs I implemented a tagging scheme using the Concurrent Versions System (CVS); this allowed for the creation of official releases of the software package. I also created and improved a large body of text- and web-based documentation of the software analysis programs. The goal of this work was to simplify STACEE data analysis. We were largely successful with this goal and I was able to provide basic analysis of a large number of STACEE data sets; far more data sets, in fact, than are presented in this work.

The core of my work, however, has been developing and testing different analysis techniques, in particular those associated with our Flash Analogue-to-Digital Converters. I largely developed, wrote and tested the Library Padding scheme, which we use to correct a systematic bias of our γ -ray measurement. I also performed a large amount of work on testing and improving our techniques for cosmic-ray rejection and energy reconstruction, though the initial implementation of both these techniques was the work of other STACEE collaborators. A large part of my work involved ensuring that our real results matched our simulated predictions. The result of these developments is that STACEE results are now both more precise and more accurate than before.

Contents

Abstract	iii
Résumé	v
Acknowledgements	vii
Statement of original work	ix
1 Introduction	1
1.1 γ -ray Astronomy	2
1.1.1 Space-based γ -ray Astronomy	2
1.1.2 Ground-based γ -ray Astronomy	5
1.2 STACEE and the High-Energy Gap	8
2 Active Galactic Nuclei	11
2.1 The AGN paradigm	13
2.2 BL Lacertae Objects	16
2.2.1 Emission Mechanisms of BL Lacs	17
2.2.2 EBL Absorption	21
2.3 Markarian 421	23
2.3.1 Variability	24
2.3.2 Spectral Evolution	28
2.4 Searches for new VHE BL Lacs	30
2.4.1 3C 66A	31
2.4.2 OJ 287	32
3 Very High Energy γ-rays and Extensive Air Showers	37
3.1 γ -ray Extensive Air Showers	37
3.1.1 Cherenkov Light Production in Extensive Air Showers	39
3.2 Cosmic-rays: the Background	43
4 STACEE detector	47
4.1 The Solar Farm Concept	47
4.2 The STACEE Detector	49
4.3 STACEE Optical System	50
4.3.1 Heliostats	51

4.3.2	Secondary mirrors	52
4.3.3	PMT Cameras	54
4.3.4	PMT Can Assembly	54
4.4	Optical Calibration	57
4.4.1	Heliostat Pointing	57
4.4.2	Secondary Mirror Alignment	59
4.5	STACEE Electronics System	60
4.5.1	Photomultiplier Tubes and Signal Filtering/Amplification	61
4.5.2	MADDOG: The Trigger System	62
4.5.3	Flash Analog-to-Digital Converters	65
4.6	Electronics Calibration	67
4.6.1	Laser System	67
4.6.2	Gain Calibration	68
4.6.3	Nominal and Effective Discriminator Thresholds	72
4.6.4	Fake Triggers and Radio-Frequency Interference	73
5	STACEE Observations and Basic Data Analysis	75
5.1	STACEE Observations	76
5.1.1	Rate vs Threshold Curves	76
5.1.2	ON-OFF Observation Technique	78
5.2	STACEE Software Analysis Framework	79
5.2.1	Pass0 Program	80
5.2.2	Data Quality Programs	82
5.3	Example Analysis: Crab 2002-2004 Raw γ -ray Rate	86
5.4	AGN Data Sets	88
6	Simulations of the STACEE Detector	91
6.1	Simulation Overview	92
6.1.1	CORSIKA	93
6.1.2	Sandfield	94
6.1.3	Elec	95
6.2	Simulation Details	97
6.2.1	Type and Energy Range of Simulations	97
6.2.2	Pointing of Simulations	99
6.2.3	Computational Requirements of Simulations	100
6.3	Effective Areas and Energy Threshold	101
6.4	Tests of Simulations	104
6.4.1	PMT Rate Comparison	104
6.4.2	Mrk 421 Cosmic-ray Rate Comparison	106
7	Library Padding: Accounting for the Promotion Effect	111
7.1	Promotion and Star Data	113
7.2	Description of the Library Padding Scheme	115
7.2.1	Background Equalization	115
7.2.2	Padding Cut	123

7.3	Real Tests of the Library Padding Scheme	124
7.3.1	Library Padding and the Energy Threshold	126
8	Cosmic-Ray Background Suppression	129
8.1	Description of Grid Alignment Technique	130
8.2	Using the Grid Alignment Technique for Background Suppression	134
8.3	Biases of the Technique	137
8.3.1	Dependence of ξ on Energy	138
8.3.2	Dependence of ξ on Source Position	139
8.3.3	Dependence of ξ on South Camera Fraction	143
8.4	Application of ξ_{cor} to Crab Data	148
8.4.1	Crab Cosmic-ray Comparison	148
8.4.2	Crab γ -ray Comparison	150
9	Energy Reconstruction	155
9.1	Introduction to STACEE Energy Reconstruction	156
9.1.1	Dead-Reckoned Charge Distributions	158
9.1.2	Core Position Determination and STACEE Canting Schemes . . .	160
9.2	Spectral Reconstruction	163
9.2.1	Core/Energy Template Fitting	164
9.2.2	Forward Folding	172
9.3	Crab 2003-2004 Spectral Reconstruction	175
10	AGN Observations and Results	179
10.1	Markarian 421 Observations and Results	179
10.1.1	Multi-wavelength Comparisons	182
10.1.2	Simple Flux Measurement	188
10.1.3	Refined Flux Measurement	193
10.1.4	Discussion	199
10.2	3C 66A Observations and Results	201
10.2.1	3C 66A Rate Upper Limit	202
10.2.2	3C 66A Energy Threshold and Flux Upper Limit	203
10.2.3	3C 66A Discussion	204
10.3	OJ 287 Observations and Results	207
10.3.1	Discussion	209
11	Conclusion	211
A	Correlation Coefficient	215
A.1	Calculation of Error on r	215
A.2	Deficiency of Correlation Coefficient	217
	Glossary	219
	Abbreviations	223

List of Figures

1.1	CGRO Deployment	2
1.2	3EG Skymap	3
1.3	LAT Sketch	4
1.4	Whipple 10 m Telescope	6
1.5	VHE Sky-map	7
2.1	Artist's View of AGN	12
2.2	Orbits around Sgr A*	14
2.3	M87 Jet	15
2.4	3C 279 Spectral Energy Distribution	17
2.5	EBL Absorption	21
2.6	Optical Images of Markarian 421	23
2.7	Mrk 421 Spectral Energy Distribution	25
2.8	Mrk 421 X-ray/ γ -ray Correlation	26
2.9	Mrk 421 X-ray Flares	27
2.10	Spectral Hardening of Mrk 421	28
2.11	Costamante & Ghisellini VHE Candidates	30
2.12	Historical OJ 287 Light Curve	33
2.13	OJ 287 Binary Black Hole Model	34
3.1	γ -ray Shower Cartoon	38
3.2	Longitudinal Development of Extensive Air Shower	39
3.3	Cherenkov Effect	40
3.4	Cherenkov Light Pool from EAS	41
3.5	Timing Profiles of Cherenkov Wavefronts	42
3.6	Cosmic-ray Shower Cartoon	44
4.1	National Solar Thermal Test Facility	49
4.2	STACEE Concept	50
4.3	STACEE Heliostat	51
4.4	Central Tower and Secondaries	53
4.5	Secondary Mirror and PMT Camera	54
4.6	PMT Can Assembly	55
4.7	DTIRC Optical Concentrators	56
4.8	Optical Efficiency	57
4.9	Drift Scan	58

4.10	CCD Pictures of Cameraspots	59
4.11	Electronics Schematic	61
4.12	Cluster Map	63
4.13	Trigger System Diagram	64
4.14	Sample FADC Trace	66
4.15	STACEE Laser Setup	68
4.16	Absolute Gain Measurement	70
4.17	Effective Threshold Calculation	71
4.18	Example of Radio-Frequency Noise	73
5.1	Rate Vs Threshold Curve	76
5.2	Basic Pass0 FADC Analysis	81
5.3	L1 Rates for Stable Pair	84
5.4	L1 Scatter Plots	85
5.5	Hour Angle Distributions	89
6.1	Cherenkov Yield for Different Atmospheric Profiles.	94
6.2	Real and Simulated FADC Traces	97
6.3	Effective Area - Mrk 421	101
6.4	HA-Weighted Effective Area and Response Function - Mrk 421	103
6.5	PMT Rate Comparison - Ratio Distribution	105
6.6	PMT Rates Vs Current	106
6.7	Cosmic-ray Spectral Energy Distribution	107
7.1	Raw Excess Rate for Star Data	114
7.2	Cartoon of Library Padding Procedure	117
7.3	Variance and Current Measurements	119
7.4	Fake FADC Bin Values	120
7.5	Padding Linearity Test	121
7.6	Variance Differences for Iota Corona Borealis	122
7.7	Excess Rate for Iota Corona Borealis	124
7.8	Padded γ -ray Rate for Star Data	125
7.9	Padded Effective Area - Mrk 421	126
8.1	Individual FADC Traces	131
8.2	Summed FADC trace	132
8.3	H/W versus Assumed Core Position	134
8.4	Distributions of H/W for γ -rays and Cosmic-rays	135
8.5	Basic ξ Results	137
8.6	Dependence of ξ on Energy	138
8.7	Off-Alignment Points	139
8.8	Dependence of ξ on Source Position	141
8.9	HA-Weighted ξ Results	142
8.10	Dependence of ξ on Wavefront Direction	143
8.11	Distribution of S_f	144
8.12	Dependence of ξ on S_f	145

8.13	S_f Corrected ξ Distributions	146
8.14	Final Effective Area - Mrk 421	147
8.15	Real Data Cosmic-ray S_f Distribution	148
8.16	Real Data Cosmic-ray ξ_{cor} Distribution	149
8.17	Real Data γ -ray ξ_{cor} Distributions	151
8.18	Final γ -ray Rate after ξ_{cor} Cut	152
8.19	2002-2004 Crab Significance vs Time	153
9.1	Number of Photoelectrons vs Energy	157
9.2	Dead-Reckoned Charge Distributions	158
9.3	Cartoon of Monocanted Heliostat Pointing Geometry	160
9.4	Cartoon of Paracanted Heliostat Pointing Geometry	161
9.5	Choice of 16 Parallel Canted Heliostats	162
9.6	Template Examples	165
9.7	Reconstructed vs Real Energy	167
9.8	Reconstructed Energy Bias and Resolution	168
9.9	Reconstructed Core Resolution	169
9.10	Crab 2003-2004 Cosmic-ray Reconstructed Energy Distribution	170
9.11	Crab 2003-2004 γ -ray Reconstructed Energy Distribution I	171
9.12	Variation of γ -ray Reconstructed Energy Distribution with Spectral Index	175
9.13	χ^2 of Energy Comparison Versus Spectral Index	176
9.14	Crab 2003-2004 γ -ray Reconstructed Energy Distribution II	177
10.1	Mrk 421 ξ_{cor} Distributions	180
10.2	Mrk 421 Cumulative Significance vs Time	181
10.3	STACEE, RXTE and Whipple Light Curve for Mrk 421	183
10.4	Difference in Observation Times	184
10.5	RXTE/STACEE Rate Comparison	185
10.6	Whipple/STACEE Rate Comparison	186
10.7	Coarse Rate Comparison	188
10.8	Mrk 421 2003-2004 Final Effective Area	189
10.9	Mrk 421 2003-2004 Response Curve	191
10.10	Mrk 421 2003-2004 γ -ray Reconstructed Energy Distribution	195
10.11	Mrk 421 2003-2004 Spectral Energy Reconstruction	195
10.12	Mrk 421 Alternate Spectral Energy Reconstruction	196
10.13	Mrk 421 Medium/High State Spectral Energy Reconstruction	198
10.14	Mrk 421 Spectral Energy Distribution	200
10.15	3C 66A ξ_{cor} Distributions	201
10.16	3C 66A Distribution of Pair-wise Significances	202
10.17	3C 66A Final Effective Area and Response Curve	203
10.18	3C 66A Spectral Energy Distribution	205
10.19	OJ 287 ξ_{cor} Distributions	206
10.20	OJ 287 Distribution of Pair-wise Significances	207
10.21	OJ 287 Final Effective Area and Response Curve	208
10.22	OJ 287 Spectral Energy Distribution	209

A.1	Correlation Coefficient Example	216
-----	---	-----

List of Tables

2.1	BL Lac Summary	36
5.1	STACEE Analysis Framework	80
5.2	Analysis Example: Crab Data Quality Cuts	87
5.3	Analysis Example: Crab Raw Results	88
5.4	Summary of AGN Data Sets	88
6.1	Simulation Pointings	100
6.2	Cosmic-ray Rate Comparison	108
7.1	Summary of Star Data	113
9.1	Reconstructed Energy Bins	172
10.1	Summary of Mrk 421 Results for Low, Medium and High Periods	187
10.2	List of Sources of Systematic Error on Energy Scale.	192
10.3	Summary of Mrk 421 2002-2004 Simple Flux Results	194
10.4	Summary of 3C 66A 2003-2004 Final Results	204
10.5	Summary of OJ 287 2003-2004 Final Results	208

Chapter 1

Introduction

The study of γ -rays of cosmic origin is the youngest branch of astronomy. It is also the branch of astronomy that deals with photons of the highest energy and hence, in my opinion, the most interesting type of celestial objects. The environments that produce γ -rays are extreme, involving very dense matter, strong gravity, high magnetic fields or a combination of all three. Examples include black holes, neutron stars and supernova. These types of extreme environments are an exciting laboratory for not only understanding the objects we study, but also for investigating the fundamental nature of the universe.

Some of the most interesting of these environments are Active Galactic Nuclei or AGN. AGN are located at the centers of distant galaxies and are characterized by immense luminosities. We now understand AGN to be the result of supermassive black holes that are embedded in the galaxies. These black holes can have a billion times more mass than our own sun and they exert an immense gravitational attraction on the surrounding matter. Particles falling into these black holes gain a substantial amount of energy, much of which is eventually converted into photons. AGN are prodigious sources of all types of radiation, from radio waves to γ -rays. Observation of the high energy photons from AGN allows us to probe the fascinating processes that occur in such an extreme environment, as well as teaching us about the history of galaxies. The particular subject of this work will be *BL Lac objects*, a subset of AGN which are known to be copious emitters of γ -rays.

Detecting γ -rays requires novel techniques, far removed from the traditional telescopes of optical astronomy. The first confirmed successes for ground-based γ -ray astronomy have only occurred within the last two decades. Since then the field has been growing rapidly, with significant developments occurring each year. The STACEE experiment is an example of this progress. STACEE was designed with the goal of lowering the energy threshold of ground-based γ -ray telescopes. This thesis will introduce the STACEE detector and describe our observations of three AGN known as Markarian 421,




Figure 1.1: Deployment of the Compton Gamma-Ray Observatory. Image taken from Space Shuttle Atlantis. Note the Canadarm on right-side of image. Credit: NASA/CGRO

3C 66A and OJ 287. Before jumping into the description of STACEE and AGN, we shall first give a brief historical overview of γ -ray astronomy. An understanding of the field as a whole will better prepare the reader for this study of a particular portion of it.

1.1 γ -ray Astronomy

There are several challenges involved in γ -ray astronomy. The most obvious challenge is the fact that the Earth's atmosphere is opaque to γ -rays; a γ -ray will be absorbed in the atmosphere before reaching the ground. It is therefore not possible to build a detector at ground-level which can directly absorb and measure a cosmic γ -ray. Two different solutions have been found to this problem. The first solution is to place the detector above the interfering atmosphere; the result is space-based γ -ray astronomy. The second solution is to use ground-based telescopes that observe not the γ -ray itself, but rather the radiation produced when the γ -ray interacts in the atmosphere. The result is ground-based γ -ray astronomy, of which the STACEE detector is an example. The following is a brief summary of both branches of γ -ray astronomy.

1.1.1 Space-based γ -ray Astronomy

The first discoveries of cosmic sources of γ -rays did not occur until the 1960s, with the launch of γ -ray detectors onboard satellites. These first satellites led to a number of important milestones. For instance, the NASA OSO-3 satellite, launched in 1967, showed




Figure 1.2: Third EGRET source catalogue, shown in galactic coordinates. The size of the symbol indicates the intensity of the source as seen by EGRET. From [55].

that the galactic disk was a source of γ -rays. This was accomplished despite OSO-3 only having an angular resolution of 25° [67]. The 1970s and 1980s saw the first detection of point sources of γ -rays. The discoveries were made by the SAS-2 satellite which was launched by NASA in 1972 and the COS-B satellite which was launched by the European Space Research Organization in 1975. Observations by these satellites resulted in confirmed detections of several sources, including the Crab and Vela pulsars. In addition, COS-B was the first instrument to unambiguously detect γ -rays from an AGN (called 3C 273) [43, 137].

However, probably the most important milestone in space-based γ -ray astronomy occurred on April 5, 1991, with the launch of the Compton Gamma-Ray Observatory (CGRO). The CGRO was the second of NASA's Great Observatories. An image of the CGRO just after deployment from Space Shuttle Atlantis is shown in Figure 1.1. The CGRO carried on board four different γ -ray experiments. The most important of these was **EGRET**, the *Energetic Gamma-Ray Experiment Telescope*. EGRET was similar in design to the earlier γ -ray detectors, but larger and more sensitive [67, 140]. The bulk of




Figure 1.3: Sketch of the GLAST LAT instrument, showing the 16 individual modules. A cut-away of one of the modules is shown; each module consists of a silicon strip tracker (top) and a CsI calorimeter (bottom). The LAT is surrounded by a plastic scintillator to reject cosmic-rays. From [11].

EGRET's observations occurred between 1991 and 1995. EGRET was used on a more limited basis until 2000, when the CGRO was de-orbited.

The greater sensitivity of EGRET led to an order of magnitude increase in the number of detected γ -ray sources. The Third EGRET Catalogue (3EG) listed a total of 271 individual sources. Sixty-seven of these sources were identified as being AGN. The result was a leap forward in our understanding of AGN, the most important discovery being that some AGN emitted more energy in γ -rays than in the rest of the electromagnetic spectrum combined. In addition to AGN, the 3EG catalogue also contained pulsars, the Large Magellanic Cloud, a Solar Flare, as well as 170 sources whose counterparts at other wavelengths have yet to be determined [55]. Figure 1.2 shows a sky-map of the 3EG sources. All three AGN that are the focus of this work are 3EG sources.

The EGRET detector was a great leap forward for space-based γ -ray astronomy. An equivalent leap forward will be occurring within the next two years with the launch of the GLAST satellite. GLAST, the *Gamma-ray Large Area Space Telescope*, is the successor

to the CGRO. The principal instrument onboard GLAST will be the Large Area Telescope (LAT). The principal advantages of LAT over EGRET will be its greater size and angular acceptance, its improved background rejection capability and its use of silicon strip technology for particle tracking [47, 97].

Figure 1.3 shows a sketch of the LAT instrument. The instrument is composed of 16 identical modules; each module consists of a tracker and a calorimeter. The tracker is composed of a series of tungsten converters sandwiched between silicon strip detectors. An incoming γ -ray will pair-produce in the converter; the tracks of the resulting e^+/e^- will be reconstructed by the silicon strip detectors. The reconstruction will determine the direction of the incoming γ -ray. The e^+/e^- pair will then deposit their energy in the hodoscopic CsI calorimeter. The LAT is surrounded by a plastic scintillator; this allows for the rejection of events that are caused by charged cosmic-rays.

Extrapolating from EGRET results, it is predicted that GLAST will be able to detect approximately 4000 AGN [47]. This wealth of new information will undoubtedly revolutionize AGN research.

1.1.2 Ground-based γ -ray Astronomy

Ground-based γ -ray astronomy has its roots in cosmic-ray research. Cosmic-rays are high energy protons, electrons and ions that continually strike the atmosphere. Cosmic-rays are very energetic, with an energy spectrum extending up to at least 10^{19} eV. Cosmic-ray research started when Victor Hess, of the University of Vienna, did a series of balloon experiments in 1912. The balloon experiments were meant to show that the flux of ‘ionizing radiation’ decreased as one moved away from the Earth’s surface. He found, to his surprise, that the flux of ionizing radiation actually increased with altitude; he surmised therefore that the ionizing radiation was of extraterrestrial origin. This theory and its subsequent confirmation resulted in Hess being awarded the Nobel prize in 1936. It is now known that the ionizing radiation was the product of cascades of particles that are created each time a high energy cosmic-ray interacts with the atmosphere. These cascades of particles are called Extensive Air Showers (EAS) [12].

For our purposes the crucial aspect of EAS is that they can produce a flash of blue/UV Cherenkov light. This optical flash is more penetrating than the cascade of particles and can be more easily detected from sea-level. In 1953, Galbraith and Jelley made the first detection of the Cherenkov flash from cosmic-rays; the *Cherenkov technique* was thereby established [46]. Cosmic-ray research is still an active research field, though it no longer uses the Cherenkov technique. We shall occasionally mention aspects of current cosmic-ray research as they pertain to γ -ray astronomy. However for the purposes of the present




Figure 1.4: Whipple 10 m telescope, on Mt Hopkins, Arizona (circa 2002). Credit: John Kildea [70].

work cosmic-rays will principally play a negative role, since they provide the main background to the ground-based detection of γ -rays. This might seem unexpected, but it occurs because a high energy γ -ray will interact with the atmosphere in much the same way as a high energy cosmic-ray. In particular, a high energy γ -ray will also produce an EAS and a Cherenkov flash. *The Cherenkov technique therefore provides a means with which to detect γ -rays* [65]. Details about the physics involved in EAS and Cherenkov flashes will be described in far more detail in Chapter 3.

The crucial point, however, is that this Cherenkov flash makes possible the ground-based detection of γ -rays. As noted, the principal challenge is the presence of the cosmic-rays. The flux of cosmic-rays greatly exceeds the flux of γ -rays and distinguishing the two types of events is more challenging than one might expect. Because of the large cosmic-ray background a long period passed between the suggestion that Cherenkov-based γ -ray detection was feasible and the first credible detection. This period was marred by a long series of conflicting weak detections and upper limits [107, 149]. The first real detection by a ground-based γ -ray telescope was not until 1989, when the Crab Nebula was convincingly detected by the Whipple telescope [150]. The Whipple telescope, shown in Figure 1.4, consisted of a 10 m diameter primary mirror and a ‘camera’ of 37 Photo-Multiplier Tubes (PMTs). The fast response of the PMT camera meant that the Whipple experiment could capture crude images of the Cherenkov flash. These images proved to be the crucial breakthrough for ground-based γ -ray astronomy. Parameterization of the Cherenkov images allowed the development of powerful techniques for the suppression




Figure 1.5: Catalogue of detected VHE sources, shown in galactic coordinates as of August 2005. From [108].

of cosmic-rays. The resulting increase in sensitivity allowed for the detection of the Crab Nebula. The details of the technique led to the Whipple-type γ -ray instruments being referred to as Imaging Atmospheric Cherenkov Telescopes (IACTs) [149, 150].

The emergence of credible ground-based γ -ray astronomy was important because the results were complementary to those provided by space-based γ -ray astronomy. This is because the division between space-based and ground-based γ -ray detectors is also a difference in energy range. The space-based detectors observe lower energy γ -rays. The energy range of EGRET was from 30 MeV to 10 GeV; this is referred to as the High Energy (HE) γ -ray band. The IACTs, on the other hand, observe higher energy γ -rays, in what is referred to as the Very High Energy (VHE) band.¹ The first generation of ground-based γ -ray detectors typically operated in the energy range of approximately 500 GeV to 10 TeV. We shall therefore be using the terms HE and VHE to distinguish between the different energy bands, as well as between the different techniques used in each band.

Since the first Whipple detection, ground-based γ -ray astronomy has gone from strength

¹There are also Low Energy and Medium Energy bands in space-based γ -ray astronomy. But they are far below the range of ground-based γ -ray telescopes and will not concern us further.

to strength. Three years after the first detection, the Whipple group reported the detection of a second VHE γ -ray source [118]. This was the AGN Markarian 421, the first extragalactic source to be detected at VHE energies and one of the subjects of the present work. Since then many more sources of VHE γ -rays have been detected. In addition to Whipple, several other IACTs such as HEGRA, CAT and CANGAROO contributed to this work during the first decade after the Crab Nebula detection. In the past few years, rapid progress has been made by the so-called ‘second generation’ of imaging Cherenkov telescopes, such as HESS, MAGIC and VERITAS. The sensitivity of the second generation telescopes is greatly increased both by the use of multiple telescopes and by increases in the primary mirror area. The most recent catalogue of VHE sources is shown in Figure 1.5; the total number of detected sources is above 30 and is growing quickly. Like the 3EG catalogue, many of these sources are AGN. As we shall detail in Chapter 2, VHE observations of AGN have also led to important advances.

1.2 STACEE and the High-Energy Gap

There is one important detail regarding EGRET and the first generation of ground-based γ -ray detectors: there is a gap between the spectral coverage of the two types of instruments. The highest energy EGRET photons were approximately 10 GeV. This limit was set by the very small fluxes of particles above that energy. In nearly 10 years of observations, EGRET only detected about 1500 γ -rays with $E > 10$ GeV and the majority of them were not associated with any point source [139]. The original imaging Cherenkov telescopes, by contrast, had energy thresholds of ~ 500 GeV. As we shall see, this threshold was set by the requirement that these experiments not trigger on noise from ambient star light.

The result was a gap in our coverage of the electromagnetic spectrum between 10 GeV and 500 GeV. This was unfortunate since there appear to be significant changes in the spectral energy distribution of many sources in this range. For instance, simple extrapolations of the EGRET spectrum for many AGN would lead us to expect these sources to produce copious VHE γ -rays. In reality, the first generation Cherenkov telescopes detected fewer AGN than expected. Clearly, major changes were occurring in the 10 - 500 GeV range. In order to understand these changes, we need to close the gap in our γ -ray coverage. The gap will be partly filled by GLAST when it is launched in 2007; the high energy limit of the LAT detector is expected to be up to 100 GeV for strong sources.

However, it is also possible to lower the energy threshold of ground-based γ -ray detectors. STACEE, the *Solar Tower Atmospheric Cherenkov Effect Experiment*, was built

to do just that. The easiest way of lowering the energy threshold of a ground-based γ -ray detector is to increase its mirror area. STACEE does this by using an array of mirrors at a solar power research facility. Extensive details about the STACEE detector will be given later, but the crucial point is that the STACEE experiment has an energy threshold of approximately 150 GeV.

The STACEE experiment has had several notable successes over the years. STACEE-32, a version of the experiment with 32 primary mirrors, successfully detected the Crab Nebula in 1998/1999 [110]. An upgraded detector, STACEE-48, detected Markarian 421 during a giant flaring episode in 2001 [23]. STACEE-64, the final version of the experiment, has been used to place upper limits on several BL Lac objects [125,28]. The current work continues this tradition with a further investigation of several BL Lac objects. Our work differs from previous studies principally in that it utilizes more fully the potential of the full STACEE-64 detector, in particular our signal digitizing *Flash Analog to Digital Converters* (FADCs). The FADCs have allowed us to develop sophisticated techniques for the correction of systematic errors, the improvement of the detector's sensitivity and the reconstruction of the incident γ -ray energy. The improvement of STACEE's sensitivity is particularly important, since, as the history of the first Whipple discovery shows, improving cosmic-ray rejection is the key to advances in ground-based γ -ray astronomy.

The outline of this work is as follows. In Chapter 2 we will start by giving an explanation of Active Galactic Nuclei, as well as introducing the three particular sources that will be the subject of this work: Markarian 421, 3C 66A and OJ 287. Chapter 3 will then describe extensive air showers and the Cherenkov flashes that VHE γ -rays produce when they enter the atmosphere. Chapter 4 will introduce the STACEE detector and explain how it is used to detect Cherenkov flashes. Chapter 5 will explain the basic elements of STACEE data analysis. Chapter 6 will cover our simulation package, an important subject for a complicated experiment like STACEE. Chapters 7, 8 and 9 will then explain the sophisticated software analysis techniques that have been developed using the information available from our FADCs. Only once all these subjects have been covered will Chapter 10 explain the results and significance of STACEE observations of Markarian 421, 3C 66A and OJ 287. As this outline makes clear, the major focus of this work will be concentrated on describing the STACEE detectors and investigating analysis techniques for improving the quality of our results; the explanation of the results themselves will be relatively brief. An emphasis on detector description and analysis techniques is a natural consequence of the relative infancy of VHE γ -ray astronomy in general and the STACEE experiment in particular.

Chapter 2

Active Galactic Nuclei

The term Active Galactic Nuclei or **AGN** refers to objects whose high luminosity and other peculiar properties can be explained by the presence of supermassive black holes. The standard picture of an AGN is shown in Figure 2.1. At the center of a galaxy sits a black hole, which can have a mass of $10^7 M_{\odot}$ or greater.¹ Surrounding the black hole is an accretion disk of gas and dust that is spiralling into the black hole. The black hole creates an immense gravitational pull and gas that falls into the black hole gains a large amount of energy. It is this large reservoir of energy that powers an AGN. In addition to the black hole and accretion disk, some fraction of AGN possess a relativistic jet. These relativistic jets eject material away from the poles of the black hole. The jet will concern us greatly, as the emission of BL Lac objects is believed to be produced by them.

A standard way of introducing AGN starts by itemizing a large class of extragalactic objects which confused astronomers for many years. These objects, though located far outside our own galaxy, did not match the standard types of galaxies with which astronomers were familiar. These objects were discovered in a variety of different ways, with a variety of different types of instruments. Consequently a bewildering catalogue of objects became recognized: Seyfert galaxies, Fanaroff-Riley galaxies, Optically Violent Variables, Flat or Steep Spectrum Radio Quasars, BL Lac objects and so on. A consensus has gradually emerged over the last 40 years that all these objects can be explained within the AGN framework, namely the concept that the emission is the result of accretion onto a supermassive black hole. The earliest suggestions of this idea came from Salpeter (1964), Zeldovich (1964) and Lynden-Bell (1969) [86, 122, 154]. With this basic picture it can then be shown that the many different objects listed above are explicable by small changes in the model, such as the angle at which we are viewing the object, the rate of

¹A ‘typical’ black hole formed by the collapse of a star will have a mass of $\sim 4 M_{\odot}$; hence the black holes in AGN truly are *supermassive*.




Figure 2.1: Artist's view of an AGN. The cartoon shows the three elements of AGN that will concern us: the central supermassive black hole, the accretion disk and the relativistic jet. The cartoon also shows some of the models for γ -ray production in the relativistic jet, which will be explained in more detail in the text. From [29].

accretion onto the black hole and the black hole mass [143]. The basic AGN paradigm, while perhaps not perfect, is a well established theory. We shall not recount all the triumphs of AGN research, since this would be more distracting than educational. Seyfert galaxies, for instance, are fascinating in their own right, but are tangential to the thrust of this work. It will suffice to note the AGN paradigm has proven to be very successful in explaining the basic properties of all the itemized objects.

In this work we will therefore concentrate on introducing the elements of the AGN model that directly concern us. Section 2.1 will describe the immense reservoir of energy available from accretion onto a supermassive black hole, as well as explaining our understanding of the relativistic jets. Some of the best current evidence for supermassive black holes will also be presented. In section 2.2 we narrow our focus to the subset of AGN that will interest us in this work: BL Lac objects. We will describe the observed properties of these objects and then show how they are explained by the AGN model. In particular, we will describe the models for BL Lac γ -ray emission, which we believe to originate in

relativistic jets that are aligned with our line-of-sight. Finally, in sections 2.3 and 2.4 we will introduce the three BL Lac objects that will be studied in this work: Markarian 421, 3C 66A and OJ 287.

2.1 The AGN paradigm

As already noted, the key features of the AGN paradigm are the supermassive black hole and the disk of accreting gas that surrounds it. The existence of the accretion disk is crucial, since it allows a mechanism for the material to fall into the black holes and hence liberate their potential gravitational energy. Consider an alternate scenario, where the circling material was in the form of isolated stars. The stars would have fixed orbits and have no way of shedding their angular momentum; an isolated star could not, therefore, spiral into the black hole. However, in the case of a disk of accreting gas, angular momentum can be transported outwards through the process of turbulence and the gas can thereby move inwards, gaining kinetic energy as it goes. A large amount of energy can be gained by inspiralling gas because of the depth of the gravity well surrounding the black hole. For a non-rotating black hole, a gas particle in the accretion disk that crosses the event horizon can release an energy equal to 5.7% of its rest mass energy. If the black hole is rotating, the possible released energy is up to 42% of the rest mass energy [39]. Given a large mass accretion rate, the amount of energy that this process makes available for luminosity is immense. It is the massive energy reservoir of supermassive black holes that first led them to be proposed as the explanation for AGN [122]. An AGN can be up to 100 times more luminous than the entire Milky Way, meaning that the AGN's energy source must be prodigious [68].

There have been a number of lines of evidence that suggest that the general picture of supermassive black holes and accretion disks is correct. Among the strongest and most striking are radio observations of material orbiting massive objects in nearby AGN. These observations became possible with the development of Very-Long Baseline Interferometry (VLBI) radio observations, which allow for the resolution of angular structure at the sub-milliarcsecond level. With this tremendously fine angular resolution VLBI observations have allowed for the detection of water maser emission in the accretion disk of the AGN NGC 4258. In this particular case the accretion disk is edge-on to our line-of-sight and observations of the Doppler shifts of individual water maser spots allow us to trace out the velocity profile of the accretion disk gas in the central parsec of this AGN. The velocity structure of the accretion disk is well described by Keplerian motion and allow us to calculate that the gas is orbiting a central object with a mass of greater than 3×10^7




Figure 2.2: Orbits of stars around Sgr A*, the putative supermassive black hole at the center of our galaxy. The plot shows measured positions of seven stars; the positions were measured using data taken during a series of infrared observations during 1995-2003 at the Keck telescopes. A simultaneous fit (dashed lines) for the motion of all seven stars constrains the mass and position of Sgr A*. From [48].

M_{\odot} and a diameter of less than 0.1 pc. The only plausible identification for this central mass is a supermassive black hole [99].

Indeed, there is growing evidence that many or even most galaxies, including our own, have supermassive black holes at their center.² Infrared observations have found a population of stars that are orbiting a massive object at the center of our own galaxy. A diagram of the motion of these stars is shown in Figure 2.2; the stars can be seen to be in orbits with remarkably short periods. Again, the only credible interpretation of the motion of these stars is that they are orbiting a supermassive black hole with a mass $3.7 \times 10^6 M_{\odot}$. In fact, given that one star approaches within 55 AU of this object, the system at

²This does not mean that most galaxies are AGN. The percentage of galaxies containing AGN is only about 1% [68], whereas the percentage containing supermassive black holes is much larger.




Figure 2.3: The relativistic jet in M87, as seen by the Hubble Space Telescope. The central nucleus is the bright spot in the lower-left with the jet extending from it. Credit: NASA/HST/John Biretta (STScI/JHU).

the center of our own galaxy is actually the most constraining proof of the existence of supermassive black holes [48, 128].

The existence of supermassive black holes is therefore well proven. We can turn to the other crucial aspect of AGN: the relativistic jets. Relativistic jets were first observed from an AGN called 3C 279. Early VLBI measurements of this source found two distinct sources, separated by 1.55 ± 0.03 milli-arcsecond; observations four months later found that the angular separation had increased to 1.69 ± 0.02 milli-arcsecond. Given the redshift of 3C 279 ($z = 0.538$), the two sources were separating at an apparent linear velocity ten times the speed of light [151]. It soon became clear that one of the 3C 279 sources was the central supermassive black hole; the other was a blob of material ejected at high velocity from the nuclei. The apparent superluminal motion is the result of the highly relativistic jet being pointed at a small angle from our line of sight; the blobs were not, of course, actually travelling faster than the speed of light. These jets can be seen to extend from sub-parsec scale near the center of AGN to kpc away. A stunning example is the jet of the galaxy M87, shown in Figure 2.3.

These jets are now understood to originate at the supermassive black holes. The process by which these jets is produced is still somewhat uncertain. It is counter-intuitive to imagine that accretion onto a supermassive black hole results in a stream of highly relativistic matter moving away from the nuclei. Nevertheless, there are a number of theories to explain this process. There have been proposals that the jets are the result of either hydrodynamic or radiative acceleration [31]. The dominant proposal, however, is that the

jets are produced by hydromagnetic acceleration. This proposal was made by Blandford and a series of collaborators in the late 1970s [20, 19, 18]. In this model, magnetic fields provide the disk-jet link and can account for the launching, confinement and collimation of the jets. The jet power can be extracted from either the accretion disk or from the angular momentum of the black hole itself [31]. This can be an efficient mechanism for transferring gravitational energy into the jets and hence into observable radiation. Despite the popularity of this theory there is, as yet, no consensus about the details and debate continues [17, 85].

2.2 BL Lacertae Objects

We can now turn to BL Lac objects, which are the particular subject of this work. The term *BL Lac* comes from an object called BL Lacertae. The optical emission of this AGN was found to change considerably between observations. The object was consequently first thought to be a variable star and was named using the convention of variable stars. Several other features of this object would also come to define the class of BL Lac objects; they include broad continuum emission, weak emission lines and polarization [136]. It is now recognized that BL Lacertae itself and BL Lac objects in general are a type of AGN. In particular, they are AGN with a strong relativistic jet that is oriented towards our line-of-sight; the radiation that we observe is produced within the jet. There are a number of lines of evidence that prove that this model is correct; we will present some of the evidence when we come to discuss the emission mechanisms of BL Lac objects.

BL Lac objects have proved to be very important for γ -ray astronomy. The third EGRET Catalogue contained 271 sources, of which 74 were firmly identified with previously known sources. Of these, fully 66 sources were identified as *blazars* [55]; indeed, the observations showed that blazars emit the majority of their power in the form of γ -rays. BL Lac objects and Flat Spectrum Radio Quasars (**FSRQ**) are together described as *blazars* [8, 78]. The emission in FSRQs, like BL Lacs, is believed to originate principally in a jet oriented towards our line-of-sight; the two classes therefore share many characteristics. BL Lacs differ from other blazars by having weak or absent emission lines.

In the VHE γ -ray regime, the list of detected BL Lac objects is smaller. Ten years ago, the list contained only two entries, Markarian 421 and Markarian 501 [118, 119]. Since then, deeper searches and improved instruments have increased this number significantly. The number of VHE-detected BL Lac objects now stands at approximately a dozen, although that number is changing rapidly [108]. BL Lac objects represent a significant




Figure 2.4: Spectral energy distribution of 3C 279. The results range from radio ($\text{Log } \nu = 9.5$) to infrared ($\text{Log } \nu = 13$) to optical ($\text{Log } \nu = 14.7$) to X-ray ($\text{Log } \nu = 18$) to HE γ -ray ($\text{Log } \nu = 23$). Triangles and squares denote data taken at different times. From [92].

fraction of all objects detected at VHE energies, much like they did for EGRET. BL Lac objects are also the only VHE objects that have been detected outside our galaxy, with the possible exception of M87 [14]. With more detections, the sample of VHE-detected BL Lacs continues to push out further in redshift; the most distant detected BL Lac is 1ES 1101-232 at a redshift of $z = 0.186$ [5]. A sample of BL Lacs at different redshifts is important, as will be seen when we discuss EBL absorption of VHE γ -rays.

2.2.1 Emission Mechanisms of BL Lacs

A principal goal of the ongoing study of BL Lac objects is to understand in more detail their emission mechanisms. There are several different observed characteristics that the emission model must explain. The most important properties of BL Lacs, for our purposes, are as follows:

1. The Spectral Energy Distribution (SED) shows a characteristic ‘double hump’ structure. An example of this structure is the SED of the blazar 3C 279, which is shown in Figure 2.4. The highest energy data in this plot is from EGRET.
2. The lower energy hump is peaked near the optical/X-ray range and the higher energy hump is peaked in the γ -ray range.

3. A large fraction (sometimes a majority) of the emitted energy is in the form of γ -rays.
4. The flux from BL Lac objects is highly variable. This variability has been detected at all wavelengths, though it is more pronounced at higher energies [21].
5. X-ray and γ -ray fluxes are generally correlated.

The first conclusion drawn from this evidence is related to the strength of the γ -ray emission. The fact that so much of the emission occurs in the γ -ray range leads to a strong constraint on the environment where the γ -rays are created; namely that it be transparent to the γ -rays. If emitted from a stationary environment, the γ -ray flux would be strongly attenuated by the interaction with ambient, low energy photons. This problem is avoided, however, if the γ -rays are produced in a material that is travelling towards us with a large bulk Lorentz factor. For instance, Maraschi, Ghisellini and Celotti (1992) suggested that the γ -ray emission of 3C 279 occurs in a material with a bulk Lorentz factor of up to 25 [92]. The conclusion of this transparency constraint was that the γ -rays were produced in the relativistic jets that had already been identified in many BL Lacs.

Bulk relativistic motion is therefore a precondition for the emission of γ -rays. The next question is what physical process causes the γ -ray emission, as well as the emission at other wavelengths. The critical observational evidence is the two peaks in the SED of BL Lacs. There is widespread agreement that the low energy peak is the result of synchrotron emission from relativistic electrons. Synchrotron emission occurs when electrons spiral around magnetic field lines. It has been found that the distribution shown in Figure 2.4 can be explained by synchrotron emission from a population of relativistic electrons with a power law distribution of energies. For most models, the power law distribution of electron energies is simply assumed (though not the particular spectral index), rather than being calculated from some model for electron acceleration. We will return to the question of how the electrons are accelerated later. In addition to explaining the shape of the low energy peak of the SED, a synchrotron origin also explains the observed polarization of the radio and optical observations of BL Lacs.

The more contentious question concerns the origin of the emission in the high energy peak. The dominant model is that the high energy emission of BL Lacs is the result of inverse Compton reactions between the relativistic electrons and the synchrotron photons. In an inverse Compton interaction the relativistic electron scatters off the synchrotron photons, bumping the photon energy up to the GeV/TeV range. Because the same photon populations are involved in the low and high energy peaks this scheme is referred to as the Synchrotron Self-Compton (SSC) model. This model has been reasonably successful in

modelling the X-ray and γ -ray emission in a number of blazars [73, 75, 92]; the fits shown in Figure 2.4 employ an SSC model.

An alternative emission model is called the Synchrotron External Compton (SEC) model. In this model the γ -rays are produced by low energy photons being inverse Compton scattered by the relativistic electrons. These are still the same relativistic electrons that produce the synchrotron emission. The difference is that the low energy target photon population is produced external to the jet, rather than being synchrotron-produced photons. For instance, the target photons could be originally from the hot accretion disk. There is some evidence that the FSRQs are better explained by SEC models, whereas BL Lac objects are better explained by SSC models [37, 56, 75].

Both SSC and SEC are examples of *leptonic* emission models. The other class of models for the γ -ray emission are the *hadronic* models [90]. In the leptonic models the jet is assumed to be composed of an e^+/e^- plasma, whereas for hadronic models the jet is composed of protons and electrons [25]. In the hadronic models, synchrotron emission from the relativistic electrons is still the cause of the low energy peak. However, the high energy peak is caused by emission from accelerated protons. The high energy emission can occur in a number of different ways. First, the protons can themselves emit synchrotron radiation. The protons can also interact with ambient photons to produce pion cascades. The pion cascades lead to high energy emission by π^0 decay and by synchrotron emission from the muons produced in π^+/π^- decays. Which of these processes dominates the emission depends on the model and source in question. The proton models certainly require some different assumptions about conditions in the jet (such as higher magnetic fields), but they have also been used with some success to model the observed emission of BL Lacs [25, 101].

One of the original motivations for hadronic emission schemes was that such models also meant that BL Lacs (and AGN in general) could be the source for high energy cosmic-rays [89]. Another important motivation relates to the variability of BL Lac objects that was discussed before. As noted, variability of BL Lac objects has been detected at all wavelengths. In leptonic models for BL Lac emission, the same population of electrons causes both peaks of the SED; we therefore expect that an increase in flux in the low energy peak would be accompanied by an increase in flux in the high energy peak. This is generally what is observed; a *flare* in γ -rays will normally be accompanied by a flare in X-rays. Examples of this type of correlation will be shown when we come to discuss results from Mrk 421. However, there has also been evidence of γ -ray flares that are not accompanied by X-ray flares; these are often referred to as ‘orphan’ flares [76]. If orphan flares prove to be a common phenomenon, then hadronic models may provide a natural

explanation, since the coupling between the electron and proton populations in hadronic models is not as strong as in leptonic models [24].

Figure 2.1 shows examples of all three of these different models occurring in the AGN jet. In reality, we would expect that only one of the processes dominates the production of VHE γ -rays. Because the leptonic models are generally favored over the hadronic models, this work will be generically referring to the high energy peak in the SED as the *Compton Peak*, in the same way as the low energy peak is referred to as the *synchrotron Peak*. This naming does not mean that we exclude the hadronic models; it simply reflects the prevailing consensus.

Suppose we accept that the SSC model is the correct description of how γ -rays are produced in BL Lac objects. There is, however, still another element of the emission process that we need to understand: what causes the acceleration of the electrons in the jet? The SSC models assume a population of relativistic electrons with Lorentz factors of up to $\gamma_{max} \sim 10^{4-6}$. However, the bulk Lorentz factor of the jet as a whole is only $\Gamma \sim 10-50$. In addition, the acceleration and γ -ray production must occur some distance away from the black hole, in order that the γ -ray flux not be suppressed by the high optical photon density from the accretion disk [132]. The distance from the black hole at which the emission occurs varies between 0.003 and 0.03 pc for different models [56, 75]. Clearly, therefore, there must be some mechanism for *in situ* acceleration of electrons within the jet.³

There are different models for how this acceleration occurs. Most BL Lac models assume acceleration is the result of shocks within the jet. Suppose, for instance, that the central engine of the AGN produced blobs of jet material intermittently and with different mass and velocity. The collision of a faster, later blob with a slower, earlier blob would produce an internal shock and allow for the Fermi acceleration of the electrons, which can in turn produce the observed radiation [132]. There are, however, alternative models. For instance, there is the possibility that magnetic reconnection in a magnetically dominated jet could be the cause of particle acceleration [87]. We will return to this possibility when describing the observed variability of Markarian 421.

Finally, there is a subdivision of BL Lac objects that is of importance to us. For many years BL Lac objects had been classified by whether they were discovered in radio or X-ray surveys. Padovani and Giommi (1995) suggested a more physical distinction, based on whether the synchrotron peak occurred at lower (radio-optical) or higher (X-ray) frequencies. The two classes are therefore referred to as *low frequency peaked BL Lacs*

³Note, therefore, that this is a separate question from how the jet as a whole is accelerated. Acceleration of the whole jet is related to models for coupling to the accretion disk and black hole, such as the Blandford mechanism, that we described earlier.




Figure 2.5: EBL absorption curves from the different models in Primack et al. (2001). The attenuation is shown as a function of energy for sources at two different redshifts. The four different curves (solid, dotted, short and long dashed) for each redshift represent different Λ CDM models. From [117].

(**LBL**) and *high frequency peaked BL Lacs* (**HBL**) [112]. Though the exact cause of the HBL/LBL distinction is unclear, the subdivision is important because all BL Lacs that are detected at VHE energies are HBL.

2.2.2 EBL Absorption

As the preceding discussion makes clear, there is still a considerable debate with regards to the intrinsic γ -ray emission mechanisms in BL Lac objects. There is also one additional source of uncertainty in interpreting VHE observations of BL Lacs. That is that the γ -rays from BL Lac objects can be attenuated by interactions with the Extragalactic Background Light (**EBL**) during their journey from the AGN to the earth.

The EBL is composed principally of photons created by the first galaxies; the EBL photons of interest to γ -ray astronomy are in the far IR to near UV range. The EBL photons can interact with VHE γ -rays by the reaction $\gamma\gamma \rightarrow e^+e^-$. Stecker et al. (1992)

noted that this would result in a sharp cutoff of the SED for γ -rays above 100 GeV for sources that were at a large redshift (such as 3C 279 at a redshift of $z = 0.54$) [134]. The attenuation increases with the redshift of the source.

The problem is that the spatial density and spectral shape of the EBL is poorly constrained. Attempts to directly measure the EBL are hampered by strong foreground contamination. Since the EBL density is poorly constrained, its effect on VHE γ -ray attenuation is also poorly understood. Various attempts have been made to model the effects of EBL absorption. An example was the work of Primack et al. (2001) which attempted to predict the EBL density from theoretical expectations of galaxy formation. The resulting predictions for EBL absorption are shown in Figure 2.5. Even if this particular EBL model is incorrect, the figure makes it clear that EBL absorption can be a significant effect, especially for sources at higher redshifts.

EBL absorption therefore presents both an opportunity and a challenge for VHE γ -ray astronomy. It is an opportunity because observations of BL Lac objects may allow us to measure the cutoff in the SED caused by EBL absorption; this in turn may allow us to make a more accurate calculation of the EBL density than is possible by direct measurements. But it is also a challenge, since it is difficult to disentangle the effect of EBL absorption from an SED cutoff that is intrinsic to the BL Lac itself. In order to be able to make meaningful investigations we will have to both improve our understanding of the intrinsic BL Lac observations and detect BL Lacs at a larger variety of different redshifts. Since the redshift increases EBL absorption, a sample of BL Lacs at different redshifts should allow us to disentangle the effects of EBL absorption from the intrinsic SED.

STACEE observations of BL Lac objects hold particular promise for the study of EBL absorption. EBL absorption is less pronounced for lower energy γ -rays. For at least some of the models shown in Figure 2.5 we find that for a source at $z = 0.1$ the EBL attenuation is a factor of 20 at 1 TeV but only a factor of 1.5 at 200 GeV [117]. Since STACEE has a lower energy threshold than most other VHE experiments it will be less affected by EBL absorption. It was therefore hoped that STACEE would be able to detect BL Lac objects out to larger redshift.

STACEE Observations of BL Lac Objects

Despite what is already known, a great deal remains to be discovered regarding the emission of BL Lac objects, as well as their interaction with the EBL. STACEE observations, with their low energy threshold, are well suited to this task. In the following sections we will introduce the three BL Lac objects that are the subject of this work. The first is




Figure 2.6: Images of Mrk 421 in the optical band. Left-side shows the image from the original Markarian survey paper, covering a 16×16 *arcminute* region. Mrk 421 is marked by the two straight lines. Right-side shows the image from recent Hubble observations, covering a 14×14 *arcsecond* region. The lines in this image are the result of Hubble's optical support structure. From [94, 126].

Markarian 421, the AGN most studied at VHE energies. A great deal has already been learned about this source, but there remain many unanswered questions. The other two BL Lacs are 3C 66A and OJ 287. Neither of these sources has been detected at VHE energies, but both are thought to be promising and interesting candidates.

2.3 Markarian 421

Markarian 421 (hereafter Mrk 421) was first discovered in a survey by Dr. Benik Markarian in 1967. The work was conducted at the Byurakan Observatory in what is now Armenia. His survey identified galaxies that had abnormal ultraviolet continuums; he noted in the earliest paper of the survey that this ultraviolet continuum is most likely of a non-stellar (ie non-thermal) origin [93]. The left-side image of Figure 2.6 shows his original image of the source identified as Mrk 421. For comparison, the right-side image shows a more recent Hubble image, which has significantly better angular resolution. Since Mrk 421 has a relatively small redshift of $z = 0.031$, its host galaxy (a giant elliptical) is fully resolved in the Hubble picture.

A couple of years after the initial discovery, it was noted that the optical and radio properties of Mrk 421 were similar to those of BL Lac objects [142]; in particular, the source was variable, polarized and had a broad continuum emission. Observations of

Mrk 421 therefore began to be interpreted in the context of the evolving AGN paradigm; namely that the emission was coming from a relativistic jet oriented towards our line-of-sight. This interpretation was bolstered by observations of Mrk 421 in radio [155] and X-ray bands [103]. Mrk 421 was also detected by EGRET [55]. The detection was not as strong as for some other blazars, but the observations did prove that Mrk 421 had the double peaked SED characteristic of BL Lac objects.

Mrk 421 was the first extra-galactic VHE γ -ray source, detected in 1992 by the Whipple telescope [118]. The detection was confirmed by the HEGRA collaboration in 1996 [114], as well as by many other instruments [81, 115]. This included the detection of the large 2001 flare by the STACEE experiment [23]. Mrk 421 was also only the second source of any type detected in the VHE energy range (the first being the Crab Nebula).⁴ Though the Crab is on average brighter, Mrk 421 is relatively bright all the time and is sometimes much brighter than the Crab. Because of this brightness, Mrk 421 has been observed extensively. We therefore have spectral measurements for Mrk 421 from radio to VHE γ -ray. As an example, Figure 2.7 shows the SED of Mrk 421 from Takahashi et al. (2000) [138]. The VHE results for that plot are a combination of results from the Whipple, HEGRA and CAT experiments.

One goal of the present work is to improve the sensitivity, energy coverage and temporal simultaneity of VHE observations. There are two aspects of Mrk 421 that will be of particular interest for this work: flux variability and spectral variability (correlated with flux variability). STACEE observations of both these characteristics near the Compton peak may hold potential for deciphering the emission mechanism of Mrk 421 in particular and BL Lacs in general. We will describe in more detail these aspects in the following sections.

2.3.1 Variability

Variability of the VHE γ -ray flux of Mrk 421 is a well-observed characteristic. The largest observed flare occurred in 2001, when the VHE flux of Mrk 421 reached 13 times the flux of the Crab; in its quiescent state, the Mrk 421 flux is approximately 40% of the Crab flux [77]. Flux doubling has been seen on timescales of hours [1, 35]. Observations of rapid VHE variability are important, because they allow us to place constraints on the size of the emitting region. For instance, if Mrk 421 varies on the timescale of ΔT_{obs} , causality

⁴Indeed, Weekes (2003) notes that it is fortunate that the Crab Nebula (hereafter simply the ‘Crab’) is on average brighter than Mrk 421. It would have been very difficult to conclusively claim the first VHE detection as being from Mrk 421, with its highly variable emission levels; follow-up confirmation would have been tricky, to say the least [149].




Figure 2.7: Spectral energy distribution of Mrk 421. The plot shows results from Takahashi et al. (2000) for radio through VHE energies. The lines show results of SSC model fits to data. From [138].

arguments imply that the size of the emitting region, R , must be smaller than

$$R \leq \delta c \Delta T_{obs} \quad (2.1)$$

where δ is the Doppler factor of the jet [75]. In addition, observations of VHE variability may allow us to distinguish between different emission models, as well as different *in situ* acceleration models.

However, the variability of Mrk 421 also causes observational challenges. The goal is to construct a more accurate and precise version of the SED shown in Figure 2.7. However, since the source is known to vary we cannot use ten year-old radio, optical or X-ray measurements to construct the SED. We require *multi-wavelength data* that is taken as simultaneously as possible. In particular, observations near both peaks are important in order to be able to properly model the emission mechanism. The synchrotron peak for Mrk 421 (and HBLs in general) occurs in the X-ray range; this can be seen in Figure 2.7. Because of this, current VHE studies of HBL are dependant on X-ray data from the Rossi X-ray Timing Explorer (RXTE). There are two instruments onboard RXTE that are used




Figure 2.8: Correlation of X-ray and γ -rays from Mrk 421, as seen by RXTE and Whipple. These results are from the 2002-2004 seasons. From [21].

to complement VHE observations: the Proportional Counter Array (**PCA**) and the All Sky Monitor (**ASM**). The most important difference between the two instruments is that the PCA is more sensitive, but has a narrower field of view. With a smaller field of view, the PCA monitor a source less often; there is therefore worse temporal coverage of Mrk 421 with the PCA.

Despite the observational challenges, many comparisons have been made between the X-ray and γ -ray flux levels for Mrk 421. An example of this type of correlation is shown in Figure 2.8, which is a plot of the RXTE and Whipple measured rates. In this case the measurements are perfectly simultaneous. The figure shows that there is a fair degree of correlation, which therefore favours the leptonic emission models. Nevertheless, the correlation is not perfect and interpretation of the results is therefore uncertain [21].

In addition to the complementing VHE observations, studies of the X-ray variability of Mrk 421 are interesting in their own right. This is partly because the temporal coverage of RXTE is better than that of ground-based VHE detectors, which can only operate on clear moonless nights. One of the recent interesting lines of evidence concerning variability models for BL Lac objects is based on observations of X-ray flaring activity from Mrk 421 and 501. Cui (2004) [34] and Xue and Cui (2005) did studies of X-ray flaring activity




Figure 2.9: Structure of X-ray flares seen from Mrk 421 by RXTE. The different panels show observations on different timescales; the timescale is shortest for the upper plot. Note that similar type of flaring activity seems to be occurring on all timescales. From [34].

from these objects [34, 153]. They found that there was a scale-invariant hierarchical structure to the flares. The scale-invariant nature of the flaring activity is best shown in Figure 2.9; similar type flares are seen on all timescales. Both Fermi shock and magnetic reconnection models seem to have the potential for explaining this flaring hierarchy. But the magnetic reconnection models are particularly intriguing. The flaring hierarchy seen




Figure 2.10: Whipple observations of the spectral energy distribution of Mrk 421 during different flux states in 2001. The labels ‘I...VIII’ refer to different flux states, with I being the highest flux state and VIII being the lowest. One can clearly see the hardening of the spectrum with increasing overall flux. From [77].

in Mrk 421 is similar to that seen in solar flares and the flaring hierarchy of solar flares is understood to be the result of magnetic reconnection [87]. The observations are therefore suggestive of a link between the two different types of phenomena, though further results and modelling will be required to make any definitive statements.

2.3.2 Spectral Evolution

The second important characteristic of Mrk 421’s VHE emission is its spectral evolution as a function of the flux state. This effect was discovered by both the Whipple and HEGRA collaborations during the large Mrk 421 flare of 2001 [1, 77]. Both groups found that the flux from Mrk 421 was well described as a power law with an exponential cutoff, that is

$$\frac{dN}{dE} \propto E^{-\alpha} e^{-E/E_0} \quad (2.2)$$

where α is the spectral index and E_0 is the exponential cutoff. The crucial discovery was that the spectral index changed with different flux levels. In particular, α varied between 3.0 in the low flux state and 2.0 in the high flux state; the spectrum therefore got *harder* during higher flux states.⁵ The Whipple SED for the different flux states is shown in Figure 2.10. Both groups found that the exponential cutoff was near 4 TeV; unlike the spectral index, they did not find evidence that the cutoff varied with the flux state.

It is possible to explain spectral evolution of Mrk 421 with the SSC model for BL Lac emission. For instance, in the model of Konopelko et al. (2003) the changes in spectral shape between the low and high flux states for the X-ray and γ -ray data can be reproduced by varying a single parameter of the model; namely the maximum energy of the injected electrons, γ_{max} . In that model, increasing γ_{max} results in both the synchrotron and Compton peaks moving to higher energies [73]. Despite this limited success, it is important that the phenomenon of spectral evolution be better understood. As noted earlier, measurements of the hierarchy of X-ray and γ -ray flares may allow us to better understand the *in situ* acceleration mechanism. But our acceleration models will differ if the varying flux levels are the result of changes in the shape of the SED, as opposed to changes in the overall normalization of the SED.

In addition, spectral evolution may provide a method for disentangling effects that are intrinsic to the source from effects of the EBL absorption. Spectral features that vary rapidly with time must be intrinsic to the source; conversely, EBL absorption features should be constant with time. For instance, it is possible that the exponential cutoff noted above is a feature of EBL absorption, since it appears to remain constant during different flux states. Better observations of spectral evolution for Mrk 421, as well as detection of spectral evolution in BL Lacs with higher redshifts, may well hold the key to understanding EBL absorption.

Goal of STACEE Observations of Mrk 421

It is clear from the preceding section that though a great deal is already known about the VHE γ -ray emission of Mrk 421, there are still many questions to answer. The advantage of STACEE observations of Mrk 421 is the lower energy threshold of our experiment. As can be seen in Figure 2.7, the Compton peak for Mrk 421 occurs somewhere between the EGRET and VHE energy ranges. STACEE measurements probe the spectral energy distribution closer to the Compton peak than do observations by Whipple and HEGRA. A measurement of the Compton peak or a better constraint on its position would provide considerable benefits for the modelling of this object.

⁵*Harder* means that the ratio of high energy γ -rays to low energy γ -rays increased.




Figure 2.11: Plot of X-ray versus radio flux from Costamante and Ghisellini (2002). Their VHE candidates are enclosed within the dotted square. The 12 sources that have been detected in VHE γ -rays (as of July 2006) are marked by black dots. The points for Mrk 421, 3C 66A and OJ 287 are marked. In addition, the dividing line between HBL and LBL is shown. Modified from [32].

2.4 Searches for new VHE BL Lacs

In addition to studying objects such as Mrk 421 that have been previously detected, STACEE also searches for new VHE-bright BL Lacs. Detection of other BL Lacs in VHE γ -rays would allow us to determine which characteristics are peculiar to a single source and which are more general. This will aid in our understanding of the emission mechanisms. Also, as noted, it is important to detect BL Lacs at a variety of different redshifts in order to understand the effects of EBL absorption.

STACEE has concentrated its search for new BL Lacs on the VHE candidates listed in Costamante and Ghisellini (2002) (hereafter **C&G**) [32]. C&G list a total of 33 BL Lac objects that they claim are likely to be emitters of VHE γ -rays. The objects have been chosen based on having large X-ray, optical and radio fluxes. The plot of X-ray versus radio flux from C&G is shown in Figure 2.11; the VHE candidates are within the dotted square. There have been seven BL Lac objects detected at VHE energies⁶ since the publication of C&G's paper; all seven objects are on their list of potential emitters. The C&G predictions therefore seem to be a reliable guide.

Figure 2.11 marks the 12 BL Lac objects that have been detected at VHE γ -rays with black dots. It is clear that all 12 sources are HBL. This is interesting, because 3C 66A and OJ 287, the two other sources investigated in this work, are both LBL. Though there are various models explaining the difference between LBL and HBL [45, 112], there is no consensus yet on the subject. STACEE detections of either LBL, or strong upper limits on their VHE emission, would be useful in distinguishing between different models. The following will provide an introduction to 3C 66A and OJ 287.

2.4.1 3C 66A

3C 66A was originally discovered as part of the Third Cambridge (3C) survey of radio sources at 159 MHz [40]. This was a catalogue of 471 radio sources detected in the northern hemisphere. The original 3C source #66 was actually composed of two unrelated objects: the BL Lac object that interests us and a radio galaxy. The BL Lac object became identified as 3C 66A and the radio galaxy as 3C 66B [106]. Further optical and X-ray observations confirmed 3C 66A as part of the BL Lac class [88, 152].

3C 66A is coincident with the EGRET source 3EG 0222+4253 [55]. The EGRET angular resolution was, however, rather poor so there are also other possible identifications of 3EG 0222+4253, such as the pulsar J0218+4232. Kuiper et al. (2000) suggest that both 3C 66A and the pulsar contribute to the EGRET source; in particular, they suggest that the pulsar contributes principally to the lower energy (< 500 MeV) γ -rays and that 3C 66A contributes principally to the higher energy (> 500 GeV) γ -rays [79].

There is therefore good evidence that 3C 66A has been detected in HE γ -rays, but there have as yet been no convincing detections of 3C 66A at VHE energies. There were reports of detections of 3C 66A by the Crimean GT-48 telescope [135], but the reported flux levels are significantly higher than the flux upper limits measured by Whipple and HEGRA [2, 62]. The observations by the different instruments were not simultaneous and

⁶As of July 2006.

since 3C 66A is known to have a variable flux, the Crimean results are not formally inconsistent with the Whipple and HEGRA results. Nevertheless, the detection is normally treated as being only tentative.

3C 66A is one of the potential emitters of VHE γ -rays listed in C&G. As noted, one of the interesting aspects of 3C 66A is that it is an LBL; a solid detection of the source would therefore be of some importance. Another interesting aspect of 3C 66A is that it has a relatively large redshift of $z = 0.444$; EBL absorption of the VHE γ -rays from a source as far away as 3C 66A should be rather severe. We might therefore not expect much chance of detecting 3C 66A with STACEE, especially since the C&G predictions do not account for EBL absorption.

There are two factors that mitigate these concerns and make 3C 66A worthy of STACEE observations. The first is that there is some question as to the accuracy of the redshift measurement for 3C 66A. Both the original redshift measurement in Miller, French and Hawley (1978) [98] and a follow-up measurement in Lanzetta, Turnshek and Sandoval (1993) [80] are presented with disclaimers regarding their accuracy [27]. If the true redshift of 3C 66A is lower than 0.444, EBL absorption may not be a significant effect.

Second, even if the redshift is correct, the EBL absorption decreases with decreasing γ -ray energy, as shown in Section 2.2.2. While 1 TeV γ -rays from 3C 66A might be strongly absorbed by the EBL, 200 GeV γ -rays will be much less so. STACEE observations, with our low energy threshold, may therefore be more interesting than those of more sensitive instruments (Whipple, HEGRA) at higher energies.

STACEE observed 3C 66A in late 2003 as part of a multi-wavelength observation campaign, which included instruments from radio to VHE γ -rays. Many of the results of this campaign have been summarized in Böttcher et al (2005) [24]. Initial STACEE results from these observations have already been presented in Bramel et al. (2005) [28]. In this work we will present a re-analysis of this data set, using a background suppression technique that significantly improves the sensitivity of the STACEE experiment.

2.4.2 OJ 287

OJ 287 was originally discovered in 1968 as part of the Ohio State University Survey of Radio Sources at 1415 MHz [38]. This was a catalogue of 1200 radio sources detected in a portion of the northern hemisphere. OJ 287 was soon identified with an optical source of magnitude ~ 14.5 [16]. Further radio, infrared and optical study showed that the object belonged to the BL Lac class [136]. The redshift of the source was found and confirmed to be $z = 0.306$ [98, 131].

The most distinctive aspect of OJ 287 is that it appears to have a 12 year periodicity




Figure 2.12: Historical OJ 287 light curves. Plot shows the V-band flux as a function of time for the last hundred years. From [145].

in its optical emission. Figure 2.12 shows the historical light curve for OJ 287; the plot includes more than 100 years of observations. It can be seen that OJ 287 appears to become much more active every 12 years. The last complete outburst occurred in 1994-1995. The most popular theory for explaining the 12-year periodicity of OJ 287 is that the AGN contains *two* supermassive black holes. This theory was first proposed by Sillanpää et al. in 1988 [129]. In this theory there is a primary $\sim 10^{10} M_{\odot}$ black hole at the center of OJ 287, with a standard accretion disk and superluminal jet. However this system is being perturbed by a secondary black hole with a mass of $\sim 10^7 M_{\odot}$, which is in an eccentric orbit around the primary black hole. The period of the orbit was estimated to be 11.65 years. The close approach of the secondary black hole results in a disruption of the accretion disk of the primary black hole and an increase in the rate of inflowing matter. The increased accretion rate results in a more luminous superluminal jet. A cartoon of one particular binary model from Sillanpää et al. (1988) is shown in Figure 2.13.

This theory received some confirmation with successful predictions about the 1994 optical outburst, which was extensively monitored [130], as well as predictions about a 1998 dimming of the source [144]. Nevertheless, the existence of a true periodicity in the OJ 287 light curve, let alone the correctness of the binary supermassive black hole model,




Figure 2.13: Example of a binary black hole model for OJ 287. Cartoon shows the central primary supermassive black hole with its dashed accretion disk, as well as the orbit of the secondary supermassive black hole. The relative sizes of the primary accretion disk and secondary orbit are to scale. From [129].

is still a matter of some debate [69]. Resolution of many of these questions may occur with observations of the expected outburst in 2005-2007.

Since mergers between galaxies should be reasonably common, galaxies with a pair of supermassive black holes are generally expected [121]. Observations of such binary supermassive black holes are of considerable interest, partly because measurements of the dynamics of such a system may provide accurate measurements of the mass of the black holes. In addition, the measurement of the frequency of such systems is important for gravity wave experiments since they can be used to predict the rate of double black hole mergers [145]. Finally, there are theories which suggest that AGN activity is often related to the merger of two galaxies. The proposal is that mergers or tidal disruptions of galaxies allow for more rapid accumulation of gas at the center of the galaxy; consequently there is more fuel for the AGN [61]. There is observational evidence to support this theory; investigators have found that many quasars occur in galaxies that seem to be either disturbed or visibly merging with another galaxy [63]. Merging galaxies might well result in a binary black hole system such as OJ 287. For all these reasons confirmation of the

binary black hole theory for OJ 287 would be very interesting.

OJ 287 was detected by EGRET. The EGRET detection is relatively weak, with an average integral flux of only $10.6 \pm 3.0 \times 10^{-8}$ photons $\text{cm}^{-2} \text{s}^{-1}$. Interestingly most of the EGRET observations occurred in 1994, during the period of the optical outburst. Indeed, the highest detected EGRET flux for OJ 287 for a single viewing period was $15.8 \pm 6.9 \times 10^{-8}$ photons $\text{cm}^{-2} \text{s}^{-1}$; this was for period 400.5 during October 1994, the same time as the peak of the optical outburst [55]. The October EGRET rate is not statistically different from the average rate; it is nevertheless interesting. It suggests that there is benefit in future γ -ray observations of OJ 287 being concentrated near the 2005-2007 outburst.

OJ 287 is also one of the C&G VHE γ -rays candidates [32]. Like 3C 66A, it is an LBL and has a relatively large redshift. This source is not listed in the Whipple, HEGRA or HESS surveys of AGN [2, 4, 62]; STACEE observations will therefore provide the first measurements of the VHE flux of this object. It is hoped that STACEE observations might shed light on the exciting possibility of OJ 287 being a binary black hole system.

Source	z	Class	VHE detected?	RA	Dec
3C 66A	0.444	LBL	no	2.22h	43.0°
OJ 287	0.306	LBL	no	8.54h	20.1°
Mrk 421	0.031	HBL	yes	11.04h	38.1°

Table 2.1: Summary of important characteristic of BL Lac objects studied in this work: redshift, LBL/HBL class, whether the source is detected by VHE instruments, right ascension and declination.

Summary of BL Lac Sources

Table 2.1 summarizes some of the important characteristics of the BL Lac objects we will be studying in this work. In particular, the table lists the astronomical coordinates of each source. The sensitivity and energy threshold of the STACEE experiment depend quite strongly on where in the sky we are observing. It is therefore useful to note the right ascension (**RA**) and declination (**Dec**) of the sources that we are observing. The declination is the more important quantity, since it defines where a source will transit and consequently what the STACEE energy threshold for a given source will be; this dependence will be explained more fully in later chapters. The right ascension only changes what time of the year we can observe a source.

Chapter 3

Very High Energy γ -rays and Extensive Air Showers

As already noted, the atmosphere is opaque to Very High-Energy (VHE) γ -rays. This does not mean, however, that these VHE γ -rays cannot be observed from the ground. Energetic γ -rays create showers of particles in the atmosphere and the different products of these showers can be detected. It is the existence of these showers that has allowed the development of ground-based γ -ray detectors such as STACEE.

The following sections discuss in detail the processes that occur when a VHE γ -ray enters our atmosphere. Descriptions will be given of the showers of particles created in the atmosphere. Emphasis will be placed on the Cherenkov light produced by the shower, since it is the Cherenkov light that STACEE detects. In addition, a description will also be given of the showers and Cherenkov light produced by VHE cosmic-rays. VHE cosmic-rays constitute the principal background for VHE γ -ray astronomy. Emphasis will therefore be placed on differences in the Cherenkov flash between the γ -rays and cosmic-rays, with a view to background suppression.

3.1 γ -ray Extensive Air Showers

When a VHE γ -ray enters the atmosphere it will pair-produce, creating a high energy e^+/e^- pair. The mean free path for this process is $\sim 47 \text{ g cm}^{-2}$. The atmosphere, by contrast, has a total depth at sea-level of $\sim 1000 \text{ g cm}^{-2}$. This indicates that the first interaction will occur high in the atmosphere. The newly created electrons and positrons will, in turn, undergo bremsstrahlung interactions with other atmospheric nuclei, creating secondary photons that take away some of the particle's energy. The process then repeats, with the result being a geometrically increasing shower of e^+ , e^- and photons. A diagram




Figure 3.1: Cartoon showing the growth of a γ -ray induced Extensive Air Shower. The number of particles will continue to grow until the average particle energy falls below the critical energy.

of this process is shown in Figure 3.1. This cascade of particles is what is referred to as an Extensive Air Shower (EAS) [107].

As the number of particles increases, the energy of each individual particle decreases. The shower will continue to grow while the energy of the electrons and positrons is above the critical energy (83 MeV in air). Once the energy of the electrons and positrons falls below the critical energy, ionization will start to dominate over bremsstrahlung as the dominant energy loss mechanism. At the same time, Compton scattering will start to dominate over pair-production as the energy loss mechanism for the photons. As few new particles will be created after this point, the number of particles (e^+/e^- and photons) will have reached a maximum. The location where this occurs is referred to as *shower maximum*.

Figure 3.2 shows the development of the number of particles in a γ -ray EAS as a function of depth into the atmosphere. The figure shows that shower maximum occurs deeper into the atmosphere for higher energy γ -rays. This is no surprise, since the shower will propagate further before the individual particles fall below the critical energy. For vertically incident γ -rays in the STACEE energy regime (~ 100 GeV) shower maximum occurs approximately 10 km above sea level. After shower maximum the number of particles in the cascade will decrease rapidly as they range out and stop. The STACEE detector is located 1.7 km above sea level. At that altitude, as Figure 3.2 shows, very few particles from a 100 GeV γ -ray EAS will reach the ground.

A γ -ray EAS is relatively compact and uniform because all the interactions are electromagnetic processes. Shower-to-shower fluctuations are relatively small for γ -ray EAS, at least compared to cosmic-ray EAS. The purely electromagnetic nature of γ -ray EAS




Figure 3.2: Longitudinal development of number of particles in EAS as a function of depth into atmosphere. On the lower axis depth is shown in terms of radiation length (r.l.); on the upper axis depth is shown in meters. The conversion from r.l. to meters assumes a standard atmospheric density profile and a vertically incident γ -ray. From [107].

also makes modelling them easier.

3.1.1 Cherenkov Light Production in Extensive Air Showers

The key aspect of an EAS, at least from the point of view of this work, is the fact that the electrons and positrons are highly relativistic and are moving faster than the speed of light in the local medium. When a charged particle moves faster than the speed of light in a medium it will emit *Cherenkov radiation*. This effect occurs when the electromagnetic fields that the charged particle induces add coherently. The fields only add coherently when the charged particle is travelling faster than the local phase velocity of light. This process is demonstrated graphically in Figure 3.3. The Cherenkov effect can be thought of as the electromagnetic analogue of a sonic boom [64].

The opening angle, θ_c , of the Cherenkov radiation is related to the index of refraction




Figure 3.3: Production of Cherenkov radiation. The circles denote consecutive electromagnetic wavefronts caused by the charged particle; the wavefronts travel outwards at the local phase velocity of light, c/n . The particle on the left is travelling at less than c/n and is therefore incapable of producing Cherenkov radiation. The particle on the right is travelling at greater than c/n ; the electromagnetic waves can therefore add coherently and Cherenkov radiation can be produced. The Cherenkov angle, θ_c , is the angle between the direction of the Cherenkov radiation and the direction of the charged particle. In air θ_c is 1.3° (at STP); in this diagram θ_c has been exaggerated for clarity. From [111].

of the medium and is given by

$$\cos \theta_c = \frac{1}{\beta n}, \quad (3.1)$$

where $\beta (= v/c)$ is the fraction of the speed of light at which the particle is travelling and n is the index of refraction. Clearly Cherenkov light is not produced if $\beta < 1/n$. For a highly relativistic particle β will be approximately unity. The index of refraction in air (at STP) is 1.00028; the opening angle is therefore $\sim 1.3^\circ$ at ground-level. The index of refraction and hence the Cherenkov opening angle decreases with increasing altitude, because of the lower atmosphere density. At 10 km, near shower maximum, the opening angle is about 0.7° [111]. The Cherenkov photons in an EAS are predominantly produced at ultraviolet (UV) and blue wavelengths.

There are a number of important points about the *light pool* of Cherenkov photons that arrives at the ground. The first point is that the Cherenkov photons form an extended image on the sky. That is to say, if one were looking up at an EAS and were sensitive enough to see it, one would see a flash of Cherenkov light coming from an extended area of the sky near the source direction; the flash has an angular size of $\sim 0.5^\circ$. The




Figure 3.4: Cherenkov light yield as a function of energy at an altitude of ~ 3 km. The different curves show the results for γ -rays and several types of cosmic-rays. The density is calculated using the photons landing within 125 m from the shower axis. Only photons with wavelengths between 300 and 550 nm and landing within 10 ns of the peak arrival time are included. From [107].

extended nature of the Cherenkov flash is the result of seeing photons from throughout the developing EAS. This fact will be important, since in order to collect as many Cherenkov photons as possible, the STACEE detector will have to be sensitive to an extended area of the sky.¹

The second point is that the Cherenkov light pool is calorimetric; that is to say, observations of the Cherenkov light allow us to reconstruct the energy of the incident γ -ray. The lateral density of the Cherenkov light generated by an EAS shower is shown in Figure 3.4. Careful examination of this figure shows that there is an approximately linear scaling between the lateral density of the Cherenkov light pool and the incident photon energy, at least for γ -rays. This is an important point to which we shall return when discussing energy reconstruction. The figure also shows that Cherenkov radiation is a relatively inef-

¹Coincidentally, 0.5° is also the approximate angular size of the sun and the moon. We can therefore use both these sources for optical calibration of the STACEE detector.




Figure 3.5: Lateral timing profile for γ -ray and cosmic-ray induced EAS. The arrival times of the shower wavefront at different locations on the ground are shown for simulated γ -ray (upper panels) and protons (lower panels). Horizontal axes are positions on the ground (in meters), while the vertical axis is arrival time (in nanoseconds). From [111].

ficient means of energy loss for the particles in an EAS. For instance, a 100 GeV shower will only produce $\sim 10^6$ Cherenkov photons. Since these are blue/UV photons (~ 5 eV), they only represent approximately 0.005% of the total energy of the shower. Nevertheless, despite receiving only a small part of the total EAS energy, it is the Cherenkov light pool that allows STACEE to detect γ -rays.

The final point relates to the shape of the Cherenkov light pool at the ground. Given the relativistic beaming and lateral spread caused by the Cherenkov opening angle, the Cherenkov wavefront arrives at the ground looking like a pancake; that is to say the lateral spread of the wavefront is much larger than the height of the wavefront. The lateral spread of the Cherenkov wavefront is circular, with a radius of 120-150 m. The height of the wavefront is about 1-2 m, which means that the duration of the shower is only ~ 5 ns. The lateral spread of the Cherenkov light pool is of great importance. It means that it

is possible to detect an EAS with a detector placed a considerable distance away from where the γ -ray would have landed had there been no atmosphere to interact with. The result, as we shall see, is that the effective area of the STACEE detector is much larger than the physical area of the detector. Without this increase in effective area, γ -ray fluxes at the ground would be too small to detect and VHE γ -ray astronomy would be much more difficult.

In fact, ‘pancake’ is only a first-order description of the shape of the Cherenkov wavefront. The wavefront does in fact have some important curvature. For a 100 GeV γ -ray, the majority of the Cherenkov photons are emitted near shower maximum. The Cherenkov wavefront at the ground is therefore approximately spherical, with the center of the sphere occurring at shower maximum. For a higher energy γ -ray (1 TeV) the shower will penetrate farther into the atmosphere. Indeed, since the electrons are travelling faster than the local speed of light, Cherenkov photons from the more penetrating particles will arrive first at the ground. It is therefore no longer a good assumption that all the photons originate from shower maximum. The Cherenkov wavefront from a higher energy γ -ray is better described as a conical shape [124]. Both a spherical and conical shape result in the edges of the Cherenkov wavefront arriving up to 10ns later than the center; this can be seen in Figure 3.5. These delays are sufficiently large that we need to account for them at the hardware level. As we shall see in Section 4.5.2, we calculate our additional hardware delays based on the assumption of a spherical Cherenkov wavefront.

3.2 Cosmic-rays: the Background

It may seem surprising that cosmic-rays should be a serious background to the study of γ -rays. Cosmic-rays are certainly not, for instance, a major concern for optical astronomy. However, at the high energies in question, cosmic-rays undergo many of the same type of particle physics interactions as γ -rays. In particular, cosmic-rays entering the atmosphere also create extensive air showers, which in turn produce flashes of Cherenkov light. Though they differ in some respects, a cosmic-ray Cherenkov flash is sufficiently similar to a γ -ray Cherenkov flash that the STACEE experiment can trigger on both. Unfortunately, the flux of cosmic-rays is much higher than the flux of γ -rays. For instance, in a circular bin of radius 0.5° around the Crab Nebula the integral flux of cosmic-rays above 1 TeV is 400 times higher than the integral flux of γ -rays [107]. It is therefore of prime importance that the ground-based γ -ray detectors be built with a view to background suppression, both at a hardware and software level. The following description will therefore emphasize the differences between γ -ray and cosmic-ray Cherenkov flashes.

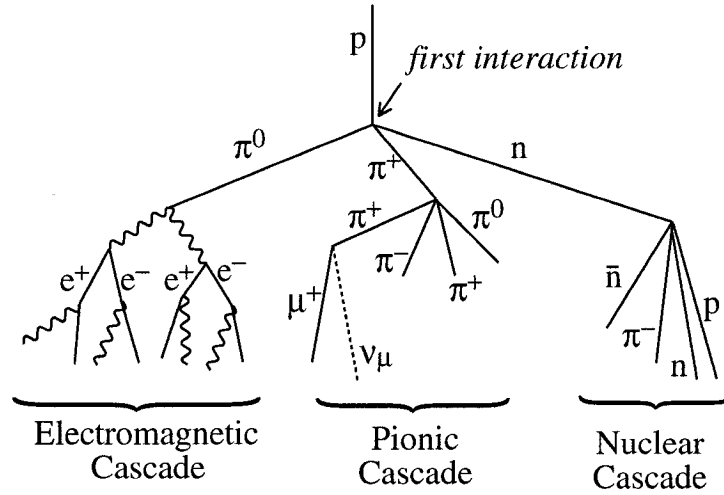


Figure 3.6: Cartoon showing the different components that make up a proton induced EAS. The initial interaction can produce neutral pions, charged pions or nucleons. A neutral pion will decay to two photons and generate an electromagnetic cascade. The charged pions can produce either more charged pions or neutral pions. The nucleon can, like the initial nucleon, produce both pions and nucleons. From [111].

The composition of a cosmic-ray EAS is significantly more complicated than that of a γ -ray, as is shown in Figure 3.6. The initial interaction of a cosmic-ray can produce neutral pions, charged pions or nucleons. The neutral pions will quickly decay to a pair of photons, which then initiates an electromagnetic cascade (and a Cherenkov flash) like a γ -ray EAS. The charged pions can produce more charged pions, neutral pions or muons. The nucleons will interact with atmospheric nuclei and start another chain, just like the initial cosmic-ray [111]. All the highly relativistic charged particles will produce Cherenkov radiation, as in the γ -ray induced EAS.

There are several important consequences of having a large fraction of hadronic particles in a cosmic-ray EAS. The first is that the Cherenkov light pool from a cosmic-ray EAS is much less homogeneous than that of a γ -ray EAS. This is because the nuclear collisions involved in cosmic-ray showers can produce fragments with large transverse momenta. These fragments can initiate sub-showers at large lateral distances from the initial shower axis. The consequence is that the Cherenkov light pool from a cosmic-ray is much more ‘lumpy’, both in its spatial and temporal distribution. This is well demonstrated in Figure 3.5, which shows the simulated arrival times of the Cherenkov wavefront for both γ -ray and cosmic-ray induced EAS. Whereas the γ -ray wavefronts appear very uniform, one can clearly see numerous sub-showers in the cosmic-ray wavefronts. Spherical or conical is a reasonably good description of the shape of the γ -ray Cherenkov

wavefronts, whereas the same is not true for the cosmic-rays. In addition to being more lumpy, a cosmic-ray EAS will also have a greater lateral spread.

The second consequence of the hadronic component is that the Cherenkov photon density for a given energy of cosmic-ray is lower, on average, than for a γ -ray of the same energy. This is shown in Figure 3.4, where one can see that the Cherenkov photon density at the ground from a 100 GeV γ -ray is equivalent to that from a 400 GeV proton. There are several reasons for this. First, the components of a cosmic-ray EAS are on average heavier than those in a γ -ray EAS. There are consequently fewer charged particles in cosmic-ray EAS and hence less Cherenkov photons. Second, as noted above, a cosmic-ray EAS will deposit Cherenkov photons over a wider area, thereby lowering the average lateral photon density [124, 111].

Cosmic-rays comprise nuclei from hydrogen up to iron, as well as energetic electrons. However, the dominant species of cosmic-rays for STACEE are the hydrogen nuclei (protons) and helium nuclei; no other nuclei contribute significantly to our cosmic-ray trigger rate.

Chapter 4

STACEE detector

The previous chapter described the processes that occur when a VHE γ -ray enters the atmosphere. In particular, we have described the flash of UV/blue Cherenkov light that an EAS produces. In this chapter we shall introduce the STACEE experiment which was designed to detect that Cherenkov flash. We shall start with a description of the Solar Farm concept and why it helps us close the ‘high-energy gap’. We shall then describe in detail the construction of the STACEE detector, explaining both the optical and electronics systems. We shall also describe the various steps required to calibrate the STACEE detector.

4.1 The Solar Farm Concept

As noted in Chapter 1, there exists a gap in our coverage of the γ -ray spectrum, between the high-energy limit of satellite detectors and the low-energy limit of ground-based detectors. The STACEE experiment was designed to help close that gap. We must therefore start by explaining how to build a ground-based γ -ray detector with a lower energy threshold.

The energy threshold, E_{th} , of a ground-based Cherenkov detector¹ is set by the presence of Night-Sky Background (NSB) photons. NSB photons are the result of both ambient starlight and anthropogenic light pollution. Blue NSB photons will interact with our detector in the same manner as blue Cherenkov photons. Indeed, even at the darkest observing site there are vastly more NSB photons than Cherenkov photons hitting the detector each second. Cherenkov γ -ray astronomy is only possible because the Cherenkov

¹We shall, in later chapters, define more precisely the term ‘energy threshold’ as it is used in the context of the STACEE experiment. For the present it suffices to consider the ‘energy threshold’ intuitively, ie as the minimum energy necessary to trigger the experiment.

flash is so short. During the ~ 5 ns of a Cherenkov flash the flux of Cherenkov photons can exceed the flux of NSB photons. It is therefore possible to build a detector that triggers on the Cherenkov flash. Nevertheless, at a certain hardware threshold our detector would start triggering on NSB photons alone; the energy threshold of the STACEE detector is set by the requirement that this *not* occur.

This energy threshold depends on many different elements of the detector. It can be shown that the energy threshold of a Cherenkov γ -ray detector scales *approximately* as

$$E_{th} \propto \sqrt{\frac{\Phi \Omega t}{\eta A}} \quad (4.1)$$

where Φ is the flux of NSB photons, Ω is the solid angle viewed by the detector, t is the detector trigger window, η is the detector collection efficiency and A is the mirror area [148]. The crucial inverse square-root dependence on area can be explained as follows. Imagine doubling the mirror area of the detector; there would, on average, be twice as many Cherenkov photons and twice as many NSB photons. But since the NSB photons are a Poisson background, the signal-to-noise improves by a factor of $\sqrt{2}$. Under the assumption that the signal strength varies linearly with the incident γ -ray energy, the energy threshold has therefore been lowered by $\sqrt{2}$.

The majority of the parameters in Equation 4.1 are technically difficult to improve or are constrained by the physics of EAS and the observing site. The mirror area, however, can be increased, thereby pushing down the energy threshold of your detector. In particular, a large collection area can be achieved by using pre-existing solar-power arrays. Using solar-power arrays as the primary mirrors is a cost-efficient scheme, since the primary mirrors, normally one of the most expensive elements of a γ -ray experiment, are already built. We refer to the scheme of using existing solar-power arrays as the primary mirror of a Cherenkov γ -ray detector as the Solar Farm concept.

The Solar Farm concept was first proposed in the early 1980s; at that time, however, the concept was deemed impractical [36]. The concept was revived and improved in the 1990s. In particular the introduction of a secondary optical element solved many of the problems with the original concept [109, 141]. There have been several collaborations formed to exploit the expected benefits of the technique. In addition to the STACEE collaboration, the CELESTE collaboration set up a similar detector at Thémis, in France [113]. Other, later groups included the GRAAL collaboration near Almeira, Spain [9] and the CACTUS collaboration near Barstow, California [95]. At the present time, only the STACEE and CACTUS groups are still taking data.




Figure 4.1: Photograph of the National Solar Thermal Test Facility in Albuquerque, New Mexico. The image shows the array of heliostats that is used, during the daytime, to focus light from the sun onto the central receiver tower (on the left). Courtesy of the NSTTF.

4.2 The STACEE Detector

The STACEE detector was built using facilities at the National Solar Thermal Test Facility (NSTTF) in Albuquerque, New Mexico (34.96° N, 105.51° W, 1705 m above sea level). The NSTTF was established in the 1970s, with the goal of developing the potential of solar energy. The most important part of the NSTTF is the array of 220 large steerable mirrors (referred to as *heliostats*). The heliostats can be used to focus light from the sun onto the tower, allowing for a maximum irradiance of 260 W/cm^2 . The high solar irradiance is used for both solar power generation and materials testing. A picture of the NSTTF heliostats and tower is shown in Figure 4.1.

The STACEE collaboration uses the NSTTF heliostats as our primary mirror. This has proven to be a very satisfactory relationship, since STACEE only uses the heliostats during night time and therefore has a minimal impact on the normal operations of the NSTTF. In addition, since the facility was already in existence, the STACEE collaboration has managed to construct a γ -ray detector at a fraction of the cost that would have been incurred building it from scratch. The current value of the NSTTF is estimated at \$120 million, far in excess of the monetary value of the STACEE project.²

A diagram of the STACEE detector is shown in Figure 4.2; it demonstrates the principle of using the NSTTF heliostats to detect γ -ray induced Cherenkov flashes. The heliostats are used to concentrate the Cherenkov light from an EAS onto a set of secondary

²For details about the NSTTF, see http://www.sandia.gov/Renewable_Energy/solarthermal/nsttf.html




Figure 4.2: Diagram showing the STACEE concept. The heliostats track the position of sources on the sky. The cherenkov photons from EAS are reflected onto secondary mirrors in the tower. The secondary mirrors focus the Cherenkov photons onto a bank of PMTs. From [124].

mirrors on the tower. STACEE uses a total of 64 heliostats to concentrate light onto five secondary mirrors. The secondary mirrors then further concentrate the Cherenkov light onto cameras of PhotoMultiplier Tubes (PMTs). The PMTs are arranged so that each one sees the light from only a single heliostat. The PMTs convert the Cherenkov photons into electrical signals. These signals are then processed and used to determine whether the experiment triggers and thus whether the data from that Cherenkov flash is recorded.

As Figure 4.2 shows, the STACEE experiment is relatively simple in concept. Nevertheless, the short timescale and rarity of Cherenkov flashes³ means care must be taken with each element of the STACEE detector. A great deal of work has gone into the design of these elements. In what follows we shall explain in detail both the optical and electronics components of the STACEE experiment. The description will emphasize how the physics of EAS led to various design choices, as well as covering the extensive calibration that is necessary to ensure the optimal operation of the detector.

4.3 STACEE Optical System

As already noted, the principle advantage of the Solar Farm concept is the increased mirror area compared to the first generation of IACTs. With 64 heliostats, the total primary mirror area of the STACEE detector is $\sim 2400 \text{ m}^2$. This is a factor of 30 larger than the

³The typical STACEE trigger rate is approximately 5 Hz.




Figure 4.3: Schematic of a STACEE heliostat. A heliostat is composed of 25 individual facets with a total surface area of 37 m^2 . The heliostats can be independently moved in the azimuth and elevation directions.

original Whipple 10 m telescope, which results in a significantly lower energy threshold.⁴ However, this low threshold is only realized if the majority of the Cherenkov photons that hit our heliostats are collected onto our PMTs. The goal of the STACEE optical system is to make this collection as efficient as possible.

4.3.1 Heliostats

Each NSTTF heliostat is composed of 25 individual facets, mounted on an azimuth/elevation drive. A diagram of a single STACEE heliostat is shown in Figure 4.3. The total mirror area of each heliostat is 37 m^2 . The individual heliostat facets are aligned so as to direct light onto the tower; in addition, each facet is made slightly concave by applying tension on the back of the mirror structure. The result is that when we reflect the light from the

⁴Though not, unfortunately, the factor of 5.5 lower that one might expect from Equation 4.1. See Section 6.3 for details about the STACEE energy threshold.

sun onto the tower, the resulting image is ~ 2 m across.⁵

The heliostats are moved in azimuth and elevation using motors with 13-bit encoders. During normal operation, the heliostats are set to track a point on the celestial sphere, continually updating their pointing in order to ensure that photons from that direction are always reflected onto the appropriate secondary. The 13-bit encoder corresponds to a pointing precision of 0.044° . Given the 0.5° size of the Cherenkov flash, 0.044° pointing precision is sufficient. The accuracy of heliostat pointing is a more complicated problem and will be described later.

The heliostat facets are back-aluminized mirrors. The transmission through a layer of glass means that the heliostats absorb all the UV photons below a wavelength of 320 nm. This is not ideal for STACEE, but it does make the NSTTF heliostats more durable.

4.3.2 Secondary mirrors

The secondary mirrors serve two goals. First, they further focus the light from the heliostats, making the ~ 2 m image at the secondary mirror into a ~ 15 cm image in the camera plane. A 15 cm image is a much more manageable size for the PMT camera. Second, the secondary mirrors allow the light from many different heliostats to be separated in the camera plane and mapped onto individual heliostats. So while the images from each heliostat overlap at the secondary mirror, they are separated in the camera plane. This separation is of crucial importance. There are large differences in travel time for photons arriving from different heliostats; these time-of-flight differences are much larger than the 5 ns length of the Cherenkov flash. In order to maximize the benefits of the Solar Farm concept, we need to correct for these time-of-flight differences before applying a trigger condition. The mapping of the light from each heliostat onto an individual PMT allows this correction to be performed by our electronics, as we shall explain in more detail later. As noted, the use of secondary mirrors is one of the crucial improvements of STACEE over early Solar Farm experiments [36].

There are five secondary mirrors arranged on the central tower as shown in Figure 4.4. There are three mirrors at the 160' level which reflect the light from 48 heliostats. Each mirror serves a camera of 16 PMTs, which we refer to as the East, North and West cameras. There are two more mirrors at the 120' level which reflect the light from the last 16 heliostats. Each of these mirrors serves a camera of 8 PMTs, which we refer to as the South-East and South-West cameras. The mirrors at the 160' level are each composed

⁵The sun and the moon have an angular size on the sky that is comparable to that of the Cherenkov flash from an EAS ($\sim 0.5^\circ$). It is therefore appropriate to use the sun and moon to calibrate various aspects of the STACEE optical system.




Figure 4.4: Layout of the central tower, showing the positions of the STACEE secondary mirrors on the 160' and 120' levels, as well as the data-acquisition (DAQ) control room. 160' and 120' refer to the height above the ground of those levels. From [124].

of seven hexagonal mirror elements, with a total width of 1.9 m and focal length of 2.0 m. The hexagonal mirror elements are suspended on a spider frame. An example of a 160' secondary mirror (and the corresponding PMT camera) is shown in Figure 4.5. The mirrors at the 120' level are single surfaces with a width of 1.1 m and a focal length of 1.1 m. All the secondary mirrors are spherical. The secondary mirrors are front-aluminized and are therefore reflective to UV photons. The reflectivity of the secondary mirrors is




Figure 4.5: Image of a STACEE secondary mirror and PMT camera. The secondary mirror is in the upper-right and the 16 PMT camera is in the lower-left. A STACEE collaborator is shown for scale.

periodically tested and is found to be stable at around 90%.

4.3.3 PMT Cameras

The individual PMTs are organized into cameras using metal support structures. As noted, the 160' level cameras each contain 16 PMTs, whereas the 120' level cameras each contain 8 PMTs. In principle, the cleanest images would be obtained if the secondary optics and cameras were exactly on-axis. Unfortunately, in that case most of the light from a heliostat would be occulted by the camera itself. We therefore situate the camera such that it makes an angle of 12.8° with respect to the focal axis of the secondary mirror. The cameras are, however, still situated such that the front of the PMT assemblies are a *focal length* away from the center of the secondary mirror.

4.3.4 PMT Can Assembly

The final element in the optical system is the PMT can assembly. The PMT can assembly serves two purposes: it provides protection for the PMT inside and serves to hold together the PMT and an optical concentrator known as a DTIRC (Dielectric Total Internal Reflection Concentrator) [105]. Optical coupling between the DTIRC and the PMT is provided




Figure 4.6: The standard STACEE PMT can assembly. Above are the individual parts; below is the assembled can. The labelled parts are: (a) retaining ring, (b) DTIRC, (c) can, (d) RTV silicone disk, (e) PMT, (f) compliance pad, (g) threaded can base, (h) long Allen screw.

by a thin RTV silicone disk. A schematic of a PMT can assembly is shown in 4.6.

The DTIRCs are the last part of the optical path. The DTIRCs are made of acrylic with an index of refraction of 1.49. The left side of Figure 4.7 is a picture of a DTIRC, showing its unique shape. As the picture shows, the DTIRCs have a spherical front face, conical sides and a flat back face. This special design serves two purposes. First, the DTIRCs serve to further concentrate the incoming Cherenkov photons from the 15 cm, camera-plane image down to the PMT aperture of 5 cm. Second, the DTIRCs define a fixed Field of View (FOV) for the PMT. This dual purpose is achieved by an appropriate choice of materials and concentrator shape, as is shown in the right side of Figure 4.7. This series of plots shows the results of a ray-tracing simulation for a DTIRC with a 24° FOV. Any photons incident on the front face of the DTIRC from angles less than 24° will be directed onto the back face by total internal reflection. Conversely, any photons with angles greater than 24° will be transmitted outside the walls of the DTIRC and consequently suppressed.

The ability to specify the FOV of a PMT is an important consideration for STACEE. Without the DTIRCs, the size of the FOV on the sky that each PMT views is defined by the geometry of the heliostats and secondaries. The FOV of each heliostat is approximately given by the angle subtended by the secondary mirror when viewed from that heliostat. For the heliostats farthest from the tower the corresponding PMT sees a FOV of the sky that is 0.5° across. However, this value increases for heliostats that are closer to the tower, so that for the closest heliostats the PMTs would see a FOV that is 0.9° across.




Figure 4.7: The DTIRC optical concentrator. Left: Photo of a typical DTIRC. Right: ray-tracing simulation of a DTIRC with a 24° field of view, with the incoming light at various angles of incidence. For angles less than 24° the light is internally reflected to the bottom aperture, whereas for angles greater than 24° the light is lost out of the sides. This design was produced for STACEE using a specialized ray-tracing program by Roland Winston and Joe O’Gallagher at the University of Chicago. From [111].

It is undesirable that different heliostat/PMTs would see different sized patches of the sky, since they would then have different sensitivities to the photons from Cherenkov showers, as well as to NSB photons. This problem is solved by the use of DTIRC’s with different fields of view. STACEE uses DTIRC’s with three different cutoff angles: 19° , 24° and 28° . By installing a 19° DTIRC on a PMT that is mapped to a nearby heliostat, we effectively decrease the size of the secondary that that PMT views; by the argument above, this decreases the patch of the sky from which the PMT receives photons. Use of these differently sized DTIRC’s allows us to ensure that there are no gross differences in efficiency between PMT’s mapped to nearby or faraway heliostats [44, 124].

The interaction of light with each of the elements of the optical system involves wavelength dependent transmission or reflection efficiencies. These efficiencies are shown in Figure 4.8; the plot also includes the effect of the PMT quantum efficiency, which is strongly wavelength dependent. The plot shows that STACEE is principally sensitive to blue light, with very little sensitivity to UV light; as noted in Section 4.3.1, the UV photons are lost because of the back-aluminized heliostat mirrors. It should be emphasized




Figure 4.8: Optical efficiency as a function of wavelength. It includes the effects of heliostat and secondary reflectivity, DTIRC transmissivity and quantum efficiency of the PMTs. The STACEE overall throughput curve does not include the effects of light that was lost because of imperfect optics. From [53].

that Figure 4.8 only includes losses due to wavelength dependent effects; it does not deal with losses due to imperfect optics, such as light scattered out of the optical pathway.

4.4 Optical Calibration

The preceding section gave a detailed description of the various components of the STACEE optical system. It should be clear now how each component is *meant* to behave; the challenge is to ensure that that is how they *actually* behave. Two elements are of particular concern: the heliostat pointing and the secondary mirror alignment. The following describes the tests that are done to ensure that these elements are behaving properly.

4.4.1 Heliostat Pointing

As noted, the 0.044° pointing precision of the heliostat should be sufficient for our purposes. An independent system monitors whether the heliostats are actually pointing where they are supposed to. We shall explain later how this information is used to remove peri-




Figure 4.9: Currents for a single channel over the course of a set of drift scans on Iota Cancri. The currents are overlaid on a sky chart of the area covered by the drift scan. The coloured rectangles indicate the current for the PMT when the heliostat was aimed at a given RA and declination; red indicates high currents, blue indicates low currents. The coloured ellipses are a fit to the excess current caused by Iota Cancri. The close alignment between the centroid of the fit and the location of Iota Cancri (blue central dot) indicate that the heliostat for this particular channel is properly aligned. From [27].

ods where the heliostat is not pointing correctly. This ensures that we have confidence in the *precision* of our heliostat pointing. A more difficult question concerns the *accuracy* of our heliostat pointing.

Our principal test of the heliostat pointing accuracy is known as *drift scans*. A drift scan consists of setting the heliostats at a point 1° in right ascension ahead of a bright star and freezing them at that position. The bright star will slowly drift through the field of view of the heliostats, creating a spike in the measured PMT currents. If the spike for a given channel occurs at exactly four minutes after the heliostats are frozen, then we know that the heliostat is pointing properly in the right ascension direction.⁶ If not, then we can adjust the heliostat's pointing. Similar drift scans are then done at offsets from the star of $\pm 0.25^\circ$ and $\pm 0.50^\circ$ in declination. These scans test the pointing accuracy in declination. Figure 4.9 is a graphical representation of the results of a drift scan for a well-aligned heliostat.

⁶An observer watching a fixed azimuth and elevation will see points on the celestial sphere that are separated by 1° in right ascension pass in a four minute period.




Figure 4.10: CCD pictures of the image formed by the moon. The image is projected on a board in the camera plane. The cross-hairs mark the center of the DTIRC and the circle marks the aperture of the DTIRC. The image on the left is from a heliostat near the tower and the image on the right is from a heliostat further away from the tower. Notice that the image of the nearby heliostat is relatively clean; the individual heliostat facets are clearly visible. The image from the faraway heliostat shows longer coma tails, the result of increased off-axis optical aberration. From [124].

4.4.2 Secondary Mirror Alignment

The spider frames holding the 160' secondary mirrors do not have perfect mechanical stability and are therefore susceptible to small changes in pointing. We test for these changes by performing *cameraspots*. Cameraspots consist of setting a given heliostat to track the full moon and examining the resulting image in the camera plane. Figure 4.10 is an example of a CCD picture of the image that the moon forms on a piece of cardboard in the camera plane. The circle in the image marks the extent of the DTIRC/PMT assembly. These images can be analyzed for all the PMTs in a camera; we check that the maximum amount of light is being delivered to each PMT. In the case of suboptimal light collection either the pointing of the secondary mirror or the position of individual PMTs can be adjusted. These cameraspots have shown that the efficiency of this part of the optical system is normally within 5% of optimal. The 120' level secondary mirrors have greater mechanical stability and therefore do not require frequent alignment testing.

Figure 4.10 also shows that the images suffer aberrations from the off-axis arrangement of the secondary/camera assemblage. In particular, this causes the long coma tails that can be seen in the images; this effect becomes more noticeable for more distant heliostats (which are more off-axis). Considerable light is lost because of the imperfection of the images.

4.5 STACEE Electronics System

Though the blue photons that STACEE detects are in the wavelength band of traditional optical astronomy, the electronics used to detect the Cherenkov photons are very different. This is principally because of the short duration and relative faintness of the Cherenkov flash. In order to detect these flashes STACEE uses a specialized set of electronics that owe more to the field of particle physics than to optical astronomy.

In particular STACEE uses PMTs to convert the blue Cherenkov photons into electrical signals. The signals from the PMTs are then used for the following purposes:

1. The signals are discriminated and the time-of-flight differences accounted for. These processed digital signals are used to determine whether or not the experiment triggers. This procedure is handled by a custom STACEE electronics system called MADDOG.
2. The raw analog PMT signals are digitized by our Flash Analog to Digital Converters (FADCs). The resulting digitized signals allow for more sophisticated analysis of the Cherenkov flash during later analysis.
3. The PMT currents are monitored to ensure that NSB light levels are not too high.
4. Various other quantities, such as the PMT discriminator rates and L1 trigger rates (to be explained later), are recorded for offline analysis.

The different elements of the electronics system are controlled by our Data Acquisition System (DAQ). The DAQ handles communication between the different modules, checks if the experiment has triggered and writes out the *event* if it has.⁷ Individual STACEE runs are started by entering commands on the DAQ. The DAQ was written by STACEE collaborators and runs on a single board computer running LynxOS, embedded in a VME crate.

Figure 4.11 shows a simplified schematic of the electronics system for a single *channel*. We shall use the term *channel* to denote the data coming from a particular PMT; so, for instance, the “channel 12 current” is the current measurement for PMT #12. In the following sections we shall describe in detail each different element, as well as describing the various electronics calibrations.

⁷STACEE is similar to particle physics experiments in that data is divided into discrete events, rather than being taken continuously. This contrasts with, for instance, the continuous integration of a CCD camera.




Figure 4.11: Simplified schematic of STACEE electronics for a single channel. The cherenkov photons induce electrical signals in the PMT. The PMT signal is sent into a filter box. The low-frequency component (the DC current) is sent to an ADC for measurement. The high-frequency signal is fanned out to the FADCs and the discriminators. The FADCs digitize the signal. The discriminated pulses from all 64 channels are sent to the MADDOG trigger module, which accounts for expected delays and decides whether to trigger the experiment (L2 triggers). MADDOG also produces local (L1) triggers and vernier hits.

4.5.1 Photomultiplier Tubes and Signal Filtering/Amplification

There are three principal considerations that drove the choice of PMT for STACEE. First, the PMT must be sensitive in the blue/UV wavelength of interest to STACEE. Second, the large FOV of the heliostats means that each PMT receives a very large flux of Night-Sky Background (NSB) photons. The PMT therefore must be capable of handling consistently high anode currents. Finally, the short duration of the Cherenkov flash ($\sim 5\text{ns}$) suggests that we need a PMT with a fast rise time, in order to be able to reasonably study the Cherenkov wavefront.

To satisfy these requirements, STACEE uses Photonis PMTs (model XP2282B). As shown in Figure 4.8, this PMT has a peak quantum efficiency of $\sim 27\%$ near 400 nm, as required. The PMT risetime is 1.5 ns and the transit time spread is 0.5 ns, thereby allowing STACEE to characterize the short Cherenkov pulses amidst the much larger background of NSB light. STACEE typically operates the PMTs at gains of approximately 1.1×10^5 , which is at the low end of the manufacturer's specification. A more detailed discussion of how STACEE calibrates the PMT gains will be presented in Section 4.6.2.

The PMT signals are then run through a low frequency/high frequency filter box. The low-frequency, DC-component of the signal is measured by a set of Joerger scanning ADC modules (Joerger ADC-32, ADC-32A); this is a measure of the DC current being put out by the PMTs. This is important, as the PMTs can be seriously damaged if they are caused to produce too much integrated charge. The currents are continually monitored by the DAQ; if the currents get too high, then the PMT high voltage is automatically shut off. The currents are also stored in the data stream, for use in later analysis.

Cherenkov pulses (because of their short duration) are encoded in the high-frequency, *signal* output of the filter boxes. The signal output is sent to a set of two $\times 10$ amplifiers (Phillips 776) since the pulse amplitudes are still small. The amplified signals are then passed into a linear fanout module (Phillips 748). One output goes to the discriminators and trigger; the other output goes to the FADCs.

4.5.2 MADDOG: The Trigger System

One of the principal challenges of the Solar Farm concept is accounting for the time-of-flight (TOF) differences introduced by the heliostat positions and shower geometry. For instance, consider a vertically-incident Cherenkov wavefront that arrives at the near and far heliostats simultaneously. Since light travels 0.3 m/ns, the photons from the nearby heliostats will arrive at the tower hundreds of nanoseconds before the photons from the faraway heliostats. Making the trigger gate long enough to account for the TOF difference would significantly raise the energy threshold of the experiment. It is therefore important to account for the TOF differences with additional *delays* before applying the trigger condition. Accounting for the TOF differences resulting from the heliostat positions is easy, since the differences are fixed. But accounting for TOF differences due to the expected shower wavefront direction is more difficult, since the wavefront direction (and hence TOF differences) changes as the celestial source moves across the sky. In order to correct for the TOF differences to a 1 ns precision, we need to change our hardware delays every 10 seconds. This is a challenging requirement, but it has been accomplished using our custom-built MADDOG delay system.

The first step is to discriminate the analog signal coming from the fanouts. We use LeCroy discriminators (model 4413), which produce an ECL logic signal if the PMT signal exceeds a certain voltage threshold. When a discriminator produces an ECL signal we call it a *hit*. Converting the signal from analog to digital simplifies the process of accounting for the delays. The discriminator threshold, which is typically on the order of -130 mV, is set by the DAQ. All discriminator channels are nominally set to the same threshold, though, as we shall see, the actual thresholds differ somewhat (see Section




Figure 4.12: Bird's-eye view of heliostat field. The larger squares denote the 64 STACEE heliostats. The numbers inside these squares indicate which L1 cluster each heliostat/channel belongs to; clusters are numbered 0 through 7. This diagram also shows the coordinate system used by STACEE. North is the +y direction and east is the +x direction. For this plot the tower is taken to be at [0,0], though we shall sometimes use the center of the heliostat field as the origin.

4.6.3). We shall also describe later the process of setting the discriminator thresholds during data-taking in Section 5.1.1.

The discriminator outputs are fed into the **MADDOG** system. MADDOG stands for McGill Asynchronous Digital Delays for Observation of Gammas. MADDOG is based on field-programmable gate array (**FPGA**) technology implemented on a series of VME boards. MADDOG is STACEE's solution for the problem of accounting for TOF delays in a fast and adaptable manner. The delays are calculated based on the assumption that the Cherenkov wavefront is spherical, with the center of the sphere located at *shower maximum*. We also assume that the EAS *core position* is at the center of the field (the core




Figure 4.13: Diagram of the discrimination, delay and trigger chain for a cluster of channels. Note the timebase change before and after the ‘Signal Delays’ module. Note also that the widths of the various pulses are not to scale.

position is the point on the ground where the incident particle would have landed had the atmosphere not intervened). MADDOG takes the hits from the discriminators, together with the currently observed azimuth and elevation, and delays the signals appropriately.

MADDOG then uses the delayed hits to check if the trigger condition has been met. STACEE uses a two-level trigger condition. Level 1 (L1) triggers are defined by clusters of STACEE channels. There are 8 clusters, each made up of 8 channels corresponding to heliostats close to one another. Figure 4.12 shows which heliostats were chosen to form each of the clusters. MADDOG produces an L1 trigger if there is a sufficient number of hits within a trigger window. The number of required hits is known as the Local Trigger Condition (LTC). The LTC is normally set to 5 for all clusters; ie, there must be at least 5 hits within a cluster in order to generate an L1 trigger. The window width for the L1 trigger is either 12 or 24 ns, depending on the exact trigger mode used. The narrowness of this trigger window compared to the raw TOF differences of hundreds of nanoseconds makes it clear why we require a system for accurately and actively correcting for delays.

MADDOG then uses the L1 triggers to determine whether the experiment as a whole triggers. When the experiment triggers it is known as a Level 2 (L2) trigger. Like the L1 level, L2 triggers are defined by a simple multiplicity requirement. If a sufficient number of clusters fired within a certain time, then an L2 trigger is generated. The number of required L1 triggers is known as the Global Trigger Condition (GTC). Like the LTC, the GTC is normally set to 5. A diagram of the the complete discriminating/delaying/triggering process is shown in Figure 4.13; the diagram shows the process for only a single cluster. When MADDOG generates an L2 trigger, it stops checking for

more triggers and informs the DAQ. The DAQ then handles reading out all the relevant data-taking modules and writing the information to disk. Further details on the MADDOG system are available elsewhere [96, 120].

It might be asked what the benefit is of having a two-level trigger. One could imagine, for instance, a one level trigger that simply required that at least 25 channels had a hit within 12 ns. It turns out that the two-level trigger has additional benefits in terms of cosmic-ray rejection. The two-level trigger requires that most of the clusters trigger. Consequently, the Cherenkov wavefront must be somewhat uniform, in order to trigger clusters across the whole heliostat field. Since γ -ray induced Cherenkov wavefronts are more uniform than cosmic-ray induced Cherenkov wavefronts, the two-level trigger is more favorable to γ -ray events. The two-level trigger therefore helps to suppress cosmic-rays at a hardware level.

In addition to the global triggers, the discriminator/MADDOG chain also produces several other pieces of data that are of interest to us. First, MADDOG records the hits for each channel that occurred during the trigger window; these are referred to as *vernier hits*. Second, the digital signals from the discriminators are also sent to a set of scaler modules, which count the rate of individual discriminator hits for each channel. We refer to these measurements as *PMT rates*; they are typically in the MHz range. Finally, each L1 trigger is also sent to a scaler module, which counts the rate at which the different clusters trigger. We refer to these measurements as *L1 rates*. The L1 rates, as we shall see, vary over many orders of magnitude. All of these measurements will prove to be useful later for calibration and data quality monitoring purposes.

4.5.3 Flash Analog-to-Digital Converters

In addition to being fanned out to the trigger system, the AC-coupled PMT signals are also sent to our Flash Analog-to-Digital Converters or **FADCs**. These fast waveform digitizers allow us to save an accurate representation of the signal shapes on a nanosecond timescale. The digitized signal waveforms are essential to the offline analysis; most of the sophisticated analysis techniques that will be used in this work would not be possible without the FADCs. The complete installation of an FADC for each STACEE channel in the summer of 2002 represented a great improvement over earlier versions of the STACEE experiment.

STACEE uses a set of four commercial FADC units built by Acqiris;⁸ the model number is DC270. The FADC units each have their own CPU and run a custom-written Linux-

⁸For more information see <http://www.acqiris.com/>




Figure 4.14: Sample FADC trace. This is a real Cherenkov pulse as recorded by our FADCs. The dashed line shows the zero mV baseline. The discriminator threshold is typically around -130mV.

based data acquisition program. Each unit has 16 waveform digitizers. The FADCs have a maximum sample rate of 1 Gigasample/second, which is what we use. The FADCs have an 8-bit precision and an adjustable full scale range. We typically use a 1.0 Volt full scale range, which means we have ~ 4 mV digitizing precision. The FADC range is normally set to be $[-890\text{mV}, 110\text{mV}]$. The FADCs put the data in a 2048 ns circular buffer, which is long enough to account for the same TOF delays that MADDOG must deal with. When the DAQ is triggered, it sends a signal to the FADCs, which then writes the appropriate section of the circular buffer to disk.

We normally save 192 ns worth of FADC data for each channel for each event. We shall refer to each digitized waveform as a *trace*. The 192 ns are chosen such that the Cherenkov pulse occurs near the center of trace. A typical FADC trace is shown in Figure 4.14. This trace shows the pulse from a real Cherenkov event. There are a number of things to note about this sample trace. First, the FWHM of the pulse is ~ 5 ns, an illustration of the short duration of the Cherenkov flash and the speed of the electronics. Next, all the fluctuations other than the central pulse are the results of random NSB photons. Given that our discriminator threshold is around 130 mV, it is clear that the PMT rates, which are entirely driven by NSB hits, will be quite high; the PMT rates are typically

in the MHz range. Finally, note that the principal peak is negative. Despite the clearly negative nature of the peak, this work will describe large pulses as being *above* threshold. That is to say, we shall ignore the negative nature of the pulses. So we would describe the central peak in Figure 4.14 as having a pulse height of 650 mV, instead of -650mV. Hopefully this will cause a minimum of confusion.

The FADCs provide us with other information in addition to the digitized waveform. As noted above, the FADC circular buffer contains far more samples than we actually write to disk. In principle we could write the entire circular buffer to disk. However, this would result in an unacceptable increase in readout time and hence to the deadtime of the system. There is some useful information, however, that can be extracted from the rest of the buffer. In particular, it proves useful to extract the mean and variance of the waveform for the 400 bins before the 192 ns that are written to disk. Each FADC trace therefore has a mean and a variance associated with it. For instance, the FADC sample in Figure 4.14 has a variance of 1700 (mV)^2 (or, equivalently, an RMS of 41 mV). This *channel variance* characterizes the rate of NSB photons; in fact, the channel variance is linearly related to the PMT currents. This information will be very useful later when we come to discuss padding techniques.

4.6 Electronics Calibration

The preceding section gave a detailed description of the various components of the STACEE electronics system. It should be clear now how each component is *meant* to behave; the challenge is to ensure that that is how they *actually* behave. The STACEE electronics therefore require extensive calibration. This section will cover some of the more important elements of electronics calibration. In particular, we shall discuss the PMT gains and the discriminator thresholds. Both these quantities are important to a proper understanding of our data, as well as for defining our energy threshold.

4.6.1 Laser System

We start by introducing the STACEE laser system, which is important for our electronics calibration. The STACEE experiment is typically interested in events where small numbers of photons (~ 100) hit our PMTs in a short time (~ 5 ns). To properly calibrate the PMTs, we require similar bursts of light. This requirement has been satisfied by the installation of a custom laser system. The system allows for an appropriate amount of light to be delivered to the PMTs.




Figure 4.15: Setup of the STACEE laser system. The system allows for a variable amount of light to be delivered to the PMTs. From [54].

A diagram of the STACEE laser setup is shown in Figure 4.15. The backbone of the system is a nitrogen laser that generates UV light pulses. The light pulses are approximately 4 ns long. The laser pulses are passed through a stilbene dye module, which re-emits the light at a more useful blue wavelength ($\lambda = 420$ nm). Part of the beam is then split off to a photodiode; the photodiode signal is used to trigger the experiment in certain calibration modes.

The majority of the laser beam then passes through a set of adjustable filters. The adjustable filters are important, since it is difficult to build a laser/fibre optics system that delivers a precise amount of photons to the PMTs. We therefore start with a strong laser pulse and attenuate it until the light levels at the PMTs are appropriate. The adjustable filters also give the laser system a large dynamic intensity range. The filters can be adjusted by the DAQ, allowing a range of light levels to be sampled during the course of a given calibration run.

After the filters, the beam is then split again, with a portion going to a set of monitor photodiodes. These photodiodes measure the relative pulse-to-pulse intensity. The rest of the beam is then sent through a set of fibre optics to the PMT cameras. There are five fibre optic cables, one for each camera. Each fibre optic cable terminates at a light diffuser, situated near the camera's focal point. The diffused laser light is bounced off the mirror, thereby illuminating each PMT in the camera.

Since it acts as a standard, tunable input to the PMTs, the laser system is a valuable tool for understanding the electronics system. It is used for a variety of purposes beyond the ones that will now be mentioned.

4.6.2 Gain Calibration

Calibrating our PMT gains is an important task, but remains a somewhat unsettled matter despite years of study. The energy threshold that we shall quote for the STACEE

experiment depends strongly on our average PMT gains; it is therefore crucial that we understand them as well as possible. This task is made more difficult by the high-current environment in which we operate our PMTs. The result of this environment is that the PMTs ‘wear out’; that is to say we see a noticeable decrease in gain with time. This decrease in gain may be as large as 1% per night [51], though it is probably not that severe. The decrease means that the gains must be continuously monitored; periodically the PMT high voltage is increased in order to keep the gains approximately constant.

We have a number of different schemes for monitoring PMT gains, of which we shall mention two. The first method is based on calculating the measured charge induced by the laser [146]. The charge is measured using the FADC traces. For instance, in the case of the peak in Figure 4.14 we would integrate over the range [85ns, 100ns] to calculate the charge. The peak in that figure was caused by Cherenkov photons, but a laser peak would look almost identical. The resulting charge is expressed in picoCoulombs (pC). This charge, C , is related to the original number of photoelectrons (pe) in the PMT, N , by

$$C = G_{PMT} G_{amp} N 1.6 \times 10^{-7} \text{ pc pe}^{-1} \quad (4.2)$$

where G_{PMT} is the PMT gain, G_{amp} is the amplifier gain and the factor of 1.6×10^{-7} is the number of pC for a single photoelectron. If we take a series of laser shots at the same intensity, we shall find that there is a spread in the measured charge; this spread is the result of the counting statistics of N , assumed to be Gaussian. The RMS spread in measured charge, σ_C , is therefore

$$\sigma_C = G_{PMT} G_{amp} \sqrt{N} 1.6 \times 10^{-7} \text{ pc}^{1/2} \text{ pe}^{-1/2}. \quad (4.3)$$

Squaring the second equation, dividing by the first and rearranging, we find that the PMT gain is

$$G_{PMT} = \frac{\sigma_C^2}{C G_{amp} 1.6 \times 10^{-7}}, \quad (4.4)$$

thereby giving us a method of calculating the absolute PMT gain.

In fact, this is only true if the fluctuations from the NSB are negligible. In normal running conditions that is not the case and there is therefore an additional term in the above formula for σ_C . However, the additional term is constant for different laser pulse heights. We can therefore extract the necessary information if we vary the laser intensity. The laser intensity can be measured using the monitor photodiode pulse height. Figure 4.16 shows examples of C and σ_C^2 versus photodiode pulse height for a single channel. As noted, there is a constant offset in the σ_C^2 versus photodiode pulse height plot, resulting




Figure 4.16: Absolute gain measurement technique. The left side shows measured charge as a function of monitor photodiode pulse height. The right side shows the variance of measured charge as a function of monitor photodiode pulse height. The slopes of the fitted lines for the two plots can be used to calculate the absolute PMT gain.

from the NSB contribution. It is, however, only the slopes that are of interest to us. Using the slopes of these two curves we can calculate the PMT gain for this channel, which is $(1.13 \pm 0.06) \times 10^5$.

While the preceding method is a relatively clean way of determining the absolute PMT gain, it does have drawbacks. Principally, it requires long special runs that waste time that could be otherwise be used for source observation. We therefore have a second method of calculating the PMT gains. This method is similarly based on the Gaussian photoelectron counting statistics; in this case, however, we are looking at photoelectrons caused by NSB light, as opposed to laser light [52].

We have two quantities related to the NSB flux in our data stream. The first is the current, I , which is related to the rate of NSB photoelectrons, R , by

$$I = G_{PMT} R 1.6 \times 10^{-7} \text{ pc pe}^{-1}. \quad (4.5)$$

The second is the channel variance, ie the variance of the NSB fluctuations as measured by our FADCs. The channel variance, σ^2 , is related to R as

$$\sigma^2 \propto G_{PMT}^2 G_{amp}^2 R. \quad (4.6)$$

Clearly, therefore,

$$\sigma^2 / I \propto G_{PMT}. \quad (4.7)$$




Figure 4.17: Figure shows a profile of the probability of a discriminator hit as a function of the measured FADC pulse amplitude. The fitted line is a modified Error Function. The 50% point defines the measured effective threshold for this channel; in this case 126 mV.

This is excellent, since both the current and the channel variance are available in a standard data run. The value of σ^2/I should be the same for all channels and changes in σ^2/I will tell us about changes in the PMT gain. The downside is that σ^2/I cannot be easily converted into an absolute measure of gain. It is, rather, a relative measure.

These are therefore the two principal methods for determining the gains of the STACEE PMTs. One method is absolute, but requires time-consuming special runs. The other is relative, but can be calculated using regular data. We used a combination of both methods to calculate the PMT gains during the two seasons of data in this thesis. The absolute gain measurements are used to set the scale of the PMT gains, whereas the σ^2/I measurements account for the slow night-to-night gain decrease. Using these methods, we found that the average PMT gains varied between 1.1×10^5 and 1.2×10^5 for the 2002-2004 seasons.

It should be noted that there is still some concern about systematic biases in these gain measurement techniques. These concerns will be partly addressed when we start discussing comparisons between real data and simulated data. The good agreement between real and simulated data will show that our understanding of our PMT gains is solid.

4.6.3 Nominal and Effective Discriminator Thresholds

Setting the discriminator thresholds is a crucial part of our standard operating procedure, since the discriminator thresholds strongly influence the energy threshold of the experiment. We shall discuss in more detail in Section 5.1.1 how we set the discriminator thresholds in order to avoid triggering on Night Sky Background. For the moment, however, we must explain a distinction between nominal and effective thresholds. The *nominal* threshold is the threshold that the DAQ instructs the LeCroy discriminator modules to use (typically around 130 mV). Careful study of the FADC data, however, has shown that the actual or *effective* thresholds are somewhat different from the nominal thresholds. This phenomenon became clear when we noticed that there were often events where MAD-DOG would record a hit but the FADC trace clearly indicated that the Cherenkov pulse did not cross the nominal threshold.

Since most STACEE events only barely trigger the experiment, it became important to properly understand our effective discriminator thresholds. Our standard method is to construct profiles of the probability of finding a discriminator hit for pulses of a given FADC amplitude. An example of this type of distribution is shown in Figure 4.17. Above a certain amplitude we almost always see a discriminator hit.⁹ The fitted form is a modified Error Function. We define the measured effective threshold for the channel to be the point at which the Error Function returns a 50% probability. In the case of Figure 4.17 the effective discriminator threshold was found to be 126 mV, whereas the nominal threshold was 140 mV. We perform this procedure separately for each channel.

We have found that the effective thresholds are typically lower than the nominal thresholds and that they vary from channel to channel. It was, however, also found that the effective thresholds are stable from night to night, as long as the same nominal threshold is used. The effective thresholds only changed when, for instance, we changed the discriminator module [133].

The measurements of our effective thresholds are made every night using Cherenkov data. The resulting information is stored in our database and used in the offline data analysis. Using the observed effective threshold allows for a significantly better understanding of our detector and our data.

⁹Why *almost* always? Because the discriminators have a finite dead-time (~ 7 ns), meaning that they are occasionally dead for the large pulses.




Figure 4.18: Example showing RF noise. This is the FADC data for a channel showing considerable RF noise. This data is taken using fake triggers, with the PMT high voltage turned off. The data shows a clear 10 ns periodicity, corresponding to 100 MHz noise.

4.6.4 Fake Triggers and Radio-Frequency Interference

We have already noted two different modes in which the STACEE detector can be triggered. The first trigger mode is the MADDOG trigger. This is our science trigger, where we look for a given multiplicity of PMT hits and cluster triggers. The second is the laser photodiode trigger. In this mode we trigger on the laser photodiode, for the purpose of electronics calibration.

There is a third trigger, also used for calibration purposes. This is the *fake trigger*. Fake triggers are taken at a 0.5 Hz rate and are used to study the background environment. For instance, the FADC data for fake triggers provides a perfect environment for studying the NSB background, without any worry about ‘contamination’ from Cherenkov pulses.¹⁰ Needless to say, the fake triggers are appropriately tagged and are never used in any real analysis of STACEE data.

One important source of noise that is measured using fake triggers is *Radio-Frequency* (RF) interference. Under certain conditions, parts of the electronics chain act as radio antennae and pick up 100 MHz radio signals. An example of a channel with particularly

¹⁰Cherenkov events are sufficiently rare that you would essentially never see one if you randomly triggered the experiment.

bad RF noise is shown in Figure 4.18. This noise is present even when the PMTs are powered down. We therefore monitor RF noise by taking special runs with the PMT high voltage turned off. We then examine the fake triggers; channels with particularly bad RF noise are flagged. We attempt to reduce the RF noise on these channels by a judicious re-seating of electrical connections.

Chapter 5

STACEE Observations and Basic Data Analysis

The preceding chapter has described the hardware used by the STACEE experiment. It should now be clear how STACEE triggers on and records each EAS-Cherenkov event. However, this is only the start of the process. The typical rate of γ -ray events is dwarfed by the rate of cosmic-ray events; STACEE therefore uses an ON-OFF observing technique to detect γ -ray sources with confidence. This chapter will start with a description of this standard STACEE observing technique.

This chapter will also describe the basics of STACEE data analysis. The simple hardware trigger rates are not normally sufficient for detection of a γ -ray source; the *offline* data analysis is therefore a crucial part of generating credible results with STACEE. The offline data analysis is used to remove periods of bad detector performance, eliminate biases in our raw trigger rates and increase the sensitivity of the instrument. This chapter will focus on the steps required to calculate the raw γ -ray rate. Descriptions of the more advanced analysis techniques will be the subject of later chapters.

These two elements, the observing technique and the data analysis framework, are the basis for all the remaining analysis and results presented in this work. It is therefore sensible that we present an example illustrating these two elements. To that end, this chapter also includes a description of the simplest possible STACEE result; a measurement of the raw Crab γ -ray rate. The example is also important because it will introduce the 2002-2004 Crab data set, which we shall be using throughout this work to calibrate our analysis techniques.

We shall conclude the chapter by introducing our four AGN data sets. These comprise a data set on Mrk 421 for each of the 2002-2003 and 2003-2004 seasons, as well data sets for both 3C 66A and OJ 287 taken during the 2003-2004 season.




Figure 5.1: Rate vs Threshold Curve. The Normal Mode curve is taken with the proper set of MADDOG delays needed for real observations. The Scrambled Mode curve is taken with a random set of MADDOG delays.

5.1 STACEE Observations

There are two primary steps that are taken in a standard night of data-taking. The first step is to choose the nominal discriminator threshold at which to run the experiment. This choice is made by taking rate vs threshold curves. The second step is to take pairs of data runs on a putative γ -ray source; taking pairs will allow us to remove the cosmic-ray background. Together these steps constitute the basic STACEE observing procedure. The following sections describe both steps in detail.

5.1.1 Rate vs Threshold Curves

The first task each night is to determine the discriminator threshold at which to run the experiment. As noted previously, the discriminator thresholds must be set such that the experiment will not trigger on NSB background alone. To ensure that this is the case, we take what are referred to as *rate vs threshold* curves. These curves are exactly as advertised; we measure the L2 MADDOG trigger rate as a function of the nominal discriminator threshold. An example of a rate vs threshold curve is shown in Figure 5.1.

The points in Figure 5.1 marked as ‘Normal Mode’ correspond to running MADDOG with the standard delays. That is to say that the discriminated hits are delayed by the

appropriate amount of time required to account for the expected TOF differences of a Cherenkov wavefront. The normal mode curve shows a *break* at approximately 110 mV. To the left of this *breakpoint* the L2 rate is dominated by triggers caused by *random* coincidences of NSB hits. The NSB events have a Gaussian distribution of pulse-height amplitudes, so the individual PMT rates increase quickly as the threshold is lowered. That is why the Normal Mode curve rises so rapidly below the breakpoint. Above the breakpoint the L2 triggers are the result of *coherent* coincidences of Cherenkov hits. The Cherenkov events have a power law distribution of pulse-height amplitudes, which explains why the Normal Mode curve falls off more slowly above the breakpoint.

The fact that the triggers above the breakpoint are the result of Cherenkov events is confirmed by the ‘Scrambled Mode’ curve in Figure 5.1. In Scrambled Mode the MADDOG delays are set to random values; ie MADD OG is set so that it does *not* correctly account for TOF differences for the Cherenkov wavefront. The result is that there is no breakpoint for the Scrambled Mode curve. EAS-induced Cherenkov events never trigger the experiment in Scrambled Mode and the Scrambled Mode rates quickly go to zero above the breakpoint.

The rate vs threshold curves demonstrate two important points:

- The STACEE experiment does indeed trigger on real EAS-Cherenkov events (or at least some sort of very short light burst that looks exactly like we would expect a Cherenkov flash to look).
- Correcting for the TOF delays is essential for the detection of these Cherenkov events.

As noted in the preceding chapter, setting the operating threshold sets the energy threshold of the STACEE experiment. The goal is to set the operating threshold as close to the breakpoint as possible, in order to minimize our energy threshold; at the same time, we must also ensure that there is a very small chance that the experiment will trigger on NSB alone. We typically choose a nominal threshold approximately 10 - 15 mV above the breakpoint, as a good compromise between these two competing requirements. The operating threshold is normally similar night to night, though it does change somewhat if the amount of NSB changes significantly, for instance if a bright patch of the sky is being observed or if atmospheric conditions change. In addition, slow changes in the average PMT gain over the course of months can also lead to changes in the operating threshold.

5.1.2 ON-OFF Observation Technique

Figure 5.1 also shows that the STACEE L2 trigger rate is approximately 5-10 Hz. The vast majority of these events are caused by cosmic-rays. As we shall see, the detected rate of γ -rays from even the strongest γ -ray source is only about 0.05 - 0.1 Hz. In order to be able to extract convincing detections of γ -ray sources in this high background environment, STACEE uses an ON-OFF observing technique. The ON-OFF observing technique is conceptually simple. We start by setting the heliostats to track the source that we wish to observe. We then start the experiment and record a 28 minute *ON-source* run. When the ON-source run is finished, we take two minutes to reorient the heliostats at a point on the sky that has no likely γ -ray sources. We then restart the experiment and record a 28 minute *OFF-source* run.¹ The OFF-source position is chosen to be exactly 30 minutes greater in RA than our original ON-source position; this ensures that the OFF-source run covers the same range of azimuth and elevation.

A *pair* of ON-OFF runs is the basic unit of STACEE analysis. The difference between the measured ON and OFF L2 trigger rates constitutes our fundamental measurement of whether a given source is an emitter of γ -rays. This rate difference is rarely statistically significant for a single pair; a typical STACEE data set therefore consists of many pairs representing dozens of hours of observations.

There are a number of assumptions that underline the ON-OFF technique. They are:

- That the angular size of the γ -ray emitting region is smaller than our FOV ($\sim 0.5^\circ$).
- That there is no significant source of γ -rays in the OFF field.
- That there are no systematic differences between observations of the ON and OFF fields.

The first two points do not pose any serious concern for the observations described in this work. Ensuring that there are no systematic differences between the ON and OFF observations is more challenging. There are several basic concerns that need to be accounted for. First, the trigger conditions must naturally be the same during the ON and OFF runs; ie the LTC, GTC and nominal discriminator thresholds must be the same. Second, we must ensure that the efficiency of the STACEE detector is the same in both halves of the pair. This means removing periods of time where the detector is malfunctioning in any way. Also, the STACEE detector efficiency depends on the portion of the sky being observed; this is why we take care to ensure that the OFF run covers the exact same range

¹The ON-OFF terminology is very common throughout this thesis. We shall often use a shorthand, describing, for instance, the “ON currents” as opposed to the “currents during the ON-source run”.

of azimuth and elevation as the ON run. Finally, concerns arise because of differences in the brightness of ON and OFF fields; this problem will be covered in more detail in Chapter 7. Once this fairly lengthy list of concerns is accounted for, the ON-OFF observing technique allows us to make a credible statement as to whether or not a given source is an emitter of γ -rays, within the detection sensitivity of STACEE.

5.2 STACEE Software Analysis Framework

As noted, the simplest form of the STACEE result is the difference between the measured ON and OFF rates. If this rate difference were scientifically interesting then we would not need any offline software analysis framework. This is not, however, the case. Firstly, the rate measurements might be compromised by periods when the detector is malfunctioning; these periods need to be removed. Secondly the rate measurements may contain systematic biases. These biases need to be accounted for in the offline analysis. Finally, the raw sensitivity of the STACEE experiment is not optimal; we have therefore developed various analysis techniques for suppressing our cosmic-ray background and improving the final sensitivity. The STACEE analysis framework addresses all these issues.

The standard STACEE analysis framework is the result of work by many different STACEE collaborators over the period of several years. Development of a common analysis framework has allowed for significant efficiencies in terms of reusable software. In addition, a common analysis framework means that the various software elements have undergone extensive double-checking and are therefore less likely to contain errors. Our analysis programs are written in C^{++} and make extensive use of the ROOT analysis tools and the MySQL database package.²

There is a series of distinct steps in a normal analysis of STACEE data. These steps are summarized in Table 5.1. The first step is the `Pass0` program which takes the various hardware dependant quantities in the original data stream and converts them to more useful values, such as mV and MHz. `Pass0` also does some basic analysis of the FADC data. The second step is the data quality step. This is where we cut out periods of data-taking that were taken under conditions that were not optimal. The final step is the `Pass1` program, which is where more sophisticated analysis techniques are applied to the data produced by `Pass0`. These techniques are concerned with correcting systematic biases, suppressing the cosmic-ray background and reconstructing the incident photon properties. It is only after the data have been run through these various steps that scientifically

²For details on ROOT and MySQL see <http://root.cern.ch/> and <http://www.mysql.com/>.

Analysis Step	Purpose
Pass0	Initial calibration of STACEE data. Conversion of hardware units (ie ADC counts) into more sensible units (ie mV). Basic FADC analysis: pulse fitting and charge calculation.
Data Quality Cuts	Removal of periods where the detector was malfunctioning or the weather was poor.
Pass1	Application of event reconstruction techniques, background suppression algorithms, padding.
Final Results	The combination of data from pass1 and data quality programs to generate final science results.

Table 5.1: Sequence of steps for standard STACEE analysis.

interesting information can be extracted.

In the following sections we shall describe the important features of the `Pass0` and `Data Quality` programs. The techniques in `Pass1` are more complex and are the subject of Chapters 7, 8 and 9.

5.2.1 Pass0 Program

The `Pass0` program runs on the raw data file for each STACEE run. In addition to doing basic conversion of hardware quantities, `Pass0` also calculates several quantities that are of relevance to the `Pass1` analysis. In particular, for every event `Pass0` performs the following operations:

1. Calculation of the dead-reckoned Cherenkov wavefront arrival time. The arrival time, t_0 , is essentially the average of the vernier hit times for all channels. We call it dead-reckoned since the channel-to-channel delays still assume that the Cherenkov wavefront was spherical and that the EAS core position was at the center of the field.
2. Calculation of the dead-reckoned charge for each channel. The dead-reckoned charge, q , is calculated using a fixed time window around t_0 . In particular the charge estimate is given by integrating the 16 FADC bins in the range $[t_0-6\text{ns}, t_0+9\text{ns}]$. The calculated charge is in picoCoulombs (pC). The advantage of using a fixed window is that we can calculate a charge for all channels, including those that do not have a significant Cherenkov pulse. This, as we shall see, is important for proper energy reconstruction.
3. Correction of saturated FADC pulses. As noted earlier, the low voltage end of the FADC range is -890 mV. When pulses fall below that level, the FADC trace satu-




Figure 5.2: An example of the important elements of `Pass0` FADC analysis. The black curve is the original FADC trace; notice the saturation that occurs 90 ns into the trace. The dashed red curve is the fitted functional form. The dotted blue line marks the dead-reckoned arrival time, t_0 . The shaded red region marks the area used for the calculation of the dead-reckoned charge, q . This FADC trace is from real data.

rates. This is a particularly serious problem for the calculation of charge; saturation of the FADC trace leads to an underestimate of the measured charge. We therefore perform a simple correction of the saturated part of the FADC trace. A Gaussian function is fit to the saturated portion of the trace, using the unsaturated bins before and after to constrain the fit. It has been found that such a fit does an adequate job of extrapolating the saturated portion of the pulse. The charge calculation uses the ‘corrected’ FADC trace. In addition, `Pass0` saves the corrected version of the FADC trace for use in `Pass1` analysis. The correction of the saturated pulses is important; approximately 50% of our events have at least one channel with a saturated FADC pulse.

These basic elements of FADC analysis are demonstrated in Figure 5.2. The calculated quantities and the corrected FADC trace are stored in a `Pass0` output file. It is this file that is used in the subsequent `Pass1` analysis.

5.2.2 Data Quality Programs

The ‘Data Quality’ step is not a single program, but rather a series of programs designed to identify the different ways in which STACEE data can be compromised. These programs do not work on an event-by-event basis; rather they identify periods of time where the data is compromised.

The data quality programs must treat the data in a pair-wise fashion. That is to say, if a section of an ON run is removed because of a data quality issue, then we need to ensure that the same section of the OFF run is also removed. As noted, the STACEE expected γ -ray rates are less than 1% of raw ON and OFF rates; it is therefore very important to ensure that the process of cutting out periods of data quality problems treat both halves of a pair in a completely equivalent manner. A C++ software framework (called `stoffs_pair`) has been written to ensure that this task can be done in a proper and consistent manner. The following sections will describe the different conditions which can compromise our data; in each case a `stoffs_pair` program has been written to remove the problematic periods. More information on the extensive work done on this topic is available elsewhere [26,27].

High Voltage Shutdowns

As noted in Section 4.5, our PMTs can be damaged if their currents get too high. To protect against this, the DAQ automatically shuts off the PMT high voltage if the currents rise above 100 μ A. This can occasionally happen during the middle of a data run if, for instance, a plane flies through the heliostat’s FOV. Normally the remainder of the run is still useable; we therefore just remove the section of the pair where the HV was turned off.

FADC Driver Failures

As noted in Section 4.5.3, the FADCs are run by data acquisition programs that operate on separate computers. Occasionally the FADC acquisition programs fail. The main DAQ program will, however, continue to run even when an FADC unit has failed. The resulting data is useless since all of our `Pass1` analysis schemes require FADC data. We therefore remove periods where the data for any FADC unit is unavailable.

Frosty Heliostats

Though in a relatively dry climate, the NSTTF does occasionally suffer from frost. Frost on the heliostats is a serious problem, since the heliostat reflectivity will rapidly decrease if frost accumulates. The only 100% certain way of checking for frost is to walk out onto

the NSTTF field and visually inspect the heliostats. This is not feasible during normal data-taking, so we have developed a reasonable alternative. A small weather station at the NSTTF measures the ambient air temperature and the dew point.³ We have empirically found that frost is likely to form when the difference between the air temperature and the dew point is less than ten degrees Fahrenheit. These two quantities are therefore continuously monitored and recorded in our data stream. We remove periods where there is potential for frost on the heliostats.

Low Elevation Data

The response of the STACEE detector changes significantly as a source tracks across the sky. The STACEE energy threshold for a given source is lowest when the elevation is highest. For this reason, and for others that we shall mention when we come to discuss background suppression, we remove any periods of data taken when the source is more than 2.5 hours from transit; that is to say, any data for which the Hour Angle (**HA**) has a value of $|\text{HA}| > 2.5 \text{ hr}$.⁴ It should be noted that low elevation data is not strictly ‘bad’; the γ -ray rate measured using low elevation data is not wrong and could be used for certain types of γ -ray studies. But with STACEE’s low energy goal there is no benefit to using the low elevation data in the present work.

Bad Heliostat Tracking

The NSTTF heliostats were built in the 1970s and have been exposed to the elements since then. Consequently, the heliostats do not always perform as they are meant to. This is clearly a problem; if a heliostat stops tracking halfway through a pair then the decrease in detector efficiency can lead to a spurious rate excess or deficit. We therefore check the exact pointing of each heliostat throughout the night. The actual pointing directions are logged by an automatic computer program every twenty seconds. If a heliostat is found to be pointing more than two bits away (0.088°) from where it should be pointing, then the data is assumed to be corrupted. Periods where any heliostat is malfunctioning are removed.

An exception is made if a heliostat is stowed for both halves of a pair. In that case the malfunctioning heliostat should not contribute to any spurious signals and the data is still useable.

³The dew point is the temperature at which condensation starts to form.

⁴Hour Angle is a measure of how close a source is to its transit point. The transit point is the location in the sky where the elevation of the source is a maximum. $\text{HA} = 0 \text{ hr}$ means that a source is at the transit point; $\text{HA} = +1 \text{ hr}$ means that a source is an hour past its transit point. A source at $\text{HA} = 0 \text{ hr}$ is therefore at a higher elevation than a source at $\text{HA} = +1 \text{ hr}$.




Figure 5.3: L1 rates taken under good and bad weather conditions. The figures show L1 rates for all eight clusters for both halves of an ON/OFF pair. The black curves are for the ON half of the run; the red curves are for the OFF half of the run. The left-hand plot shows L1 rates taken under good weather conditions; the right-hand plot shows L1 rates taken under poor weather conditions. The shaded portions of the right-hand plot show periods which have been removed by the automated `l1_scatter` program.

Unstable Weather Conditions

By far the most challenging data quality problem is identifying data taken under poor weather conditions, which normally means clouds. The other data quality issues have reasonably precise symptoms and are therefore easy to remove. In contrast, identifying data taken under poor weather conditions has proven to be more difficult. It is, however, just as important to account for. Clouds increase reflection of the lights of Albuquerque which leads to increases in the PMT currents. This in turn leads to an increase in *promotion*. Promotion is an increase in the measured L2 rate that results from NSB fluctuations; it will be discussed more completely in Chapter 7. For the moment, it will suffice to understand that large increases in PMT currents are a problem and that periods with bad weather conditions must be removed.

While it is possible to observe current increases that result from worsening weather, it turns out that the L1 cluster rates are the most sensitive indicators of clouds. The L1 cluster rates are driven entirely by and strongly dependant on NSB hits. Even a small increase in the NSB flux can lead to orders of magnitude change in the L1 rates. The left-side of Figure 5.3 shows the L1 rates for a pair where the atmospheric conditions were stable and clear. The L1 rates for all eight clusters for both halves of the pair are stable. The right-side of Figure 5.3, by contrast, shows a pair taken under poor weather conditions; in particular, a cloud caused a noticeable increase in L1 rates during the first part of the ON run.

While it is clear by eye that the L1 rates are less stable in the second plot, the challenge




Figure 5.4: `L1_scatter` Plots. In the left-side plot each point shows the average of the ON and OFF L1 rates for a 20 second period during a pair. The data is for a single cluster and is taken from a full season of Crab data. The red points represent bad weather conditions. The right-side plot shows the distribution of θ_{L1} for the same set of data. The dashed vertical lines represent 3σ deviations from the main distribution. The data outside the 3σ level represent bad weather conditions.

is to devise a program that can automatically characterize the pairs with abnormal L1 rates. This problem is made more challenging because the L1 rates can vary significantly during the night, even when the weather is good. These benign changes are the result of the sources moving across the sky, as well as brightness differences in different fields. The benign changes in L1 rates can be over orders of magnitude and make it impossible to specify a ‘good’ range of L1 rates.

It turns out that the best quantity to remove bad weather is related to the ratio of the ON and OFF L1 rates for a cluster [27]. Though the L1 rates will vary for a given cluster as a source tracks across the sky, the ratio of the ON to OFF L1 rates is more stable. This becomes clear when we look at Figure 5.4; the left-side of the figure shows a scatter plot of the ON and OFF L1 rates for a single cluster. The L1 rates are averaged over twenty second intervals; the plots show the results for an entire season of Crab data. The plot makes it clear that there is a strong correlation between the ON and OFF L1 rates for most of the data; the well correlated data represents periods of good atmospheric conditions. Note that the L1 rates can vary over more than an order of magnitude during good weather conditions. The long ‘tendrils’ extending away from the main distribution correspond to data taken when clouds were present. It is this data that we wish to remove.

In order to do so, we define a quantity related to the ratio of the ON and OFF L1 rates.

The quantity, θ_{L1} , is defined as

$$\theta_{L1} = \text{Log} \left(\frac{L1_{ON}}{L1_{OFF}} \right). \quad (5.1)$$

The distribution of θ_{L1} for a single cluster is shown in the right-side of Figure 5.4. This distribution shows a clear central peak, which corresponds to the periods of good data. All the points outside the peak correspond to periods of unstable weather. We therefore fit a Gaussian to the data and classify all points lying more than 3σ from the central maximum to be ‘bad’. The effectiveness of this algorithm can be seen in the left-side plot of Figure 5.4, where the points in red are what we classified as bad based on the θ_{L1} distribution. The cut on θ_{L1} clearly eliminates all the outlying points. Note that we must generate these distributions separately for each cluster. The average ratio of the ON and OFF L1 rates can be different since each cluster sees a slightly different portion of the sky with different star fields and therefore different NSB levels. If any of the eight clusters fail the 3σ condition then the time interval is removed.

These calculations are automatically performed in a program called `l1_scatter` which takes care of removing sections that fail the 3σ condition defined above. Further proof of the effectiveness of this program can be seen by looking back at the right-side plot of Figure 5.3. The shaded sections are portions of the run which have been automatically cut by the `l1_scatter` program. The program has removed the worst of the data during this pair, data that was ruined by clouds or some other atmospheric disturbance.

5.3 Example Analysis: Crab 2002-2004 Raw γ -ray Rate

In this section, we shall apply the programs described in the preceding sections to the a data set taken on the Crab Nebula during 2002-2004. The Crab is the ‘standard candle’ for γ -ray astrophysics. There are two aspects of the VHE γ -ray emission from the Crab that make it useful for calibration:

1. The Crab is the brightest source (on average) in the VHE γ -ray sky. This brightness explains why the Crab was the first confirmed detection by VHE γ -ray detectors.
2. The Crab VHE γ -ray emission is very steady, with no observable variation in its flux.

Data removed because of...	ON Livetime Remaining (ks)	Fraction of Original Livetime Remaining
... No Data Removed.	163	100%
... High Voltage Shutdowns.	162	99%
... FADC Failures.	154	94%
... Frosty Heliostats.	149	91%
... Low Elevation Data.	149	91%
... Bad Heliostat Tracking.	141	86%
... Unstable Weather Conditions.	131	81%
... All of the Above.	88	54%

Table 5.2: Effect of various data quality cuts on the Crab 2002-2004 data sets. The table shows the ON-source livetime remaining after removing different problems and also shows the fraction of the original livetime remaining after cuts. The final line shows the cumulative effect of all the data quality cuts.

This Crab data set will be used extensively throughout this thesis for testing purposes.⁵ Our goal in this section will be to calculate the *raw γ -ray rate* from the Crab. The raw γ -ray rate is the difference between the ON and OFF L2 trigger rates after all the periods of bad quality data have been removed. All else being equal, a statistically significant excess in the raw γ -ray rate would indicate that a source is an emitter of VHE γ -rays. This simple analysis will therefore be a good example of how the standard data quality programs are applied and what the end result is.

The Crab 2002-2004 data set was composed of 119 ON-OFF pairs, comprising an ON-source livetime of 163 ks. There were some small instrumental changes between the 2002-2003 and 2003-2004 seasons, but nothing that is of importance for the raw γ -ray rate measurements. Table 5.2 summarizes the ON livetime that remains after periods with data quality problems were removed. After all the data quality cuts the final livetime is 88 ks; this is only 54% of the original livetime. As will be seen later, this fraction is typical of the amount of data lost because of data quality cuts. It is clearly unfortunate that we discard such a large fraction of our data; nevertheless, using data without these data quality cuts would seriously compromise our results.

Once the data quality cuts have been applied, we simply have to count up the total ON and OFF livetime and the numbers of ON and OFF events. It should be noted that the total livetime can be systematically different for ON and OFF data.⁶ One consequence of

⁵The Crab is, of course, also of interest in its own right. Details on the scientific results from STACEE observations of the Crab are available elsewhere [44, 110, 111].

⁶This is because the event deadtime can be systematically different for ON and OFF runs. The event deadtime is the time that it takes the DAQ to write a triggered event to disk. Higher NSB rates lead to more hits and hence more deadtime. Hence a systematic difference in NSB levels leads to a systematic difference in livetime.

$t_{ON}(\text{ks})$	$t_{OFF}(\text{ks})$	N_{ON}	N_{OFF}	Raw γ -ray Rate	Significance
88.271	88.245	530709	525932	$3.2 \pm 0.7 \text{ cts min}^{-1}$	4.5σ

Table 5.3: Raw γ -ray rate for the 2002-2004 Crab data set. The table shows the ON and OFF livetime (t_{ON}/t_{OFF}) and ON and OFF number of events (N_{ON}/N_{OFF}) after all data quality cuts have been applied.

Data Set	Time Period	Number of Pairs	Raw ON livetime (ks)	Final ON livetime (ks)
Mrk 421 (2002-2003)	Dec '02 - May '03	70	96	58
Mrk 421 (2003-2004)	Jan '04 - Apr '04	45	62	42
3C 66A (2003-2004)	Sept '03 - Dec '03	87	121	83
OJ 287 (2003-2004)	Dec '03 - Feb '04	28	40	21

Table 5.4: Summary of AGN Data Sets. The table lists the four AGN data sets that are the focus of this work.

having different ON and OFF livetimes is that we need to take more care in calculating the error on the measured rate. Li and Ma (1983) developed a method for calculating the statistical significance of data sets where the ON and OFF livetimes are different [82]; we used their methodology for calculating significances in this work. The difference between the ON and OFF livetimes is normally smaller than 2%, so the correction of our results by using the full Li and Ma formulation is small.

The final results for the raw Crab γ -ray rate are shown in Table 5.3. Our raw Crab γ -ray rate is $3.2 \pm 0.7 \text{ min}^{-1}$, which would suggest that we have detected the Crab at a 4.5σ level during 2002-2004. This is not, however, the final story. As our standard VHE candle, we shall use this Crab data set for the calibration of various analysis procedures. We shall be returning to our measured Crab rate when discussing both library padding and background suppression techniques.

5.4 AGN Data Sets

We shall end this chapter by introducing our AGN data sets. Table 5.4 lists the four AGN data sets that will be analyzed in this work. They are: two seasons of data on Markarian 421, one season of data on 3C 66A and one season of data on OJ 287. The table lists the periods over which the data was taken and the number of pairs taken. The table also lists the raw ON livetime and the final ON livetime after all the data quality cuts have been applied. The same data quality cuts have been applied to all the AGN data sets; the procedure is identical to that used for the Crab data set. We do not, for the moment, present the measured raw γ -ray rates for each source. We will leave that to Chapter 10,




Figure 5.5: Hour Angle Distributions. The left-side plot shows the distribution of livetime spent at different hour angles for Crab 2002-2004 data set. Right-side plot shows same distribution for Mrk421 2003-2004 data set.

after we have introduced and tested additional analysis steps. This will ensure that we are not prejudicing our final results in any way.

We introduce the AGN data sets at this point for two purposes. First, Table 5.4 emphasizes the fact that a large amount of livetime is lost due to the standard data quality cuts. This is regrettable, but, as noted, necessary. Despite this, there remains a reasonably large amount of data in each data set. The two years of Markarian 421 data comprise 100 ks of livetime, which is comparable to the amount of livetime in the Crab data set. The same is true of the 83 ks of livetime in the 3C 66A data set. The OJ 287 data set is rather smaller, with a final livetime of only 21 ks.

Second, while we shall not yet discuss AGN γ -ray rates, there are some aspects of these data sets that are immediately important. In particular, it is useful to present the distribution of Hour Angle (HA) at which the different AGN data sets were taken. As noted earlier, the overall distribution of HA for a data set is important because the response of the STACEE detector changes significantly with elevation. To properly characterize our detector we must therefore know our HA distribution during each data set. We shall often refer to these distributions in later chapters when calculating our average simulated response using simulations done at specific pointings. Figure 5.5 shows the distributions of HA for the Crab (2002-2004) data set and the Mrk 421 (2003-2004) data set. Note that in order to minimize the overall energy threshold of our detector we try to take data as close to transit as possible. For instance, in the Mrk421 (2003-2004) data set a majority of the data was taken within one hour of transit. Note also that there is no data with $|\text{HA}|$

> 2.5 hr since we removed it during our data quality cuts.

Chapter 6

Simulations of the STACEE Detector

With the preceding chapter's description of the simplest STACEE result, our measurement of the raw γ -ray trigger rate, it is now time to introduce the STACEE simulations. Simulations are a crucial aspect of any analysis of STACEE data. Without simulations we cannot place STACEE measurements in a proper scientific context. The different elements that allow STACEE to detect a γ -ray Cherenkov flash are generally too complex to be described in a simple analytical form. As an example we might ask what our Crab raw γ -ray rate tells us about the Crab Spectral Energy Distribution (SED). Answering this question is not simple. How many Cherenkov photons will a 100 GeV γ -ray shower produce? What will be the spatial and temporal distribution of those photons? How many of the Cherenkov photons will be collected onto our PMTs? Will there be sufficient Cherenkov photons to trigger the experiment? To answer these questions we must use simulations. A complete simulation of the EAS/Cherenkov Flash/STACEE Detector chain must be done before our observations can be put in a scientific context that is useful to other astrophysicists.¹

In addition to the larger goal of producing science results, the simulations are valuable for understanding details about our detector. For instance, simulations allow us to understand the consequences of detector malfunctions. We might wonder what effect a malfunctioning heliostat would have on our detector. It is clear that it would degrade the performance of our detector, but it is not clear how seriously. Simulations allow us to answer such questions with precision. Simulations also allow us to create pure samples of γ -rays and cosmic-ray events, thereby providing a test bed for developing background suppression techniques.

For all these reasons, simulations are crucial for STACEE. In the following sections we

¹Indeed, we have already seen simulated results. All the plots used in Chapter 3 to describe extensive air showers were produced using simulations.

shall explain in detail the three different elements of the STACEE simulation chain. We shall also describe some of the particulars of the simulations that were done for this work and will present important simulated results, such as the determination of the STACEE energy threshold. Finally we shall provide a couple of basic comparisons between simulated and real quantities. Since simulations are only useful if they can be trusted to faithfully reflect reality, successful comparisons give us increased confidence in our understanding of our detector.

6.1 Simulation Overview

There is a great deal of variability inherent in the type of event that STACEE observes. This is particularly true of the processes involved in the extensive air showers. Even if we specify that we are only interested in the Cherenkov light pool of a 100 GeV γ -ray, there will be event-to-event variability as a result of the fluctuations involved in every step of the electromagnetic cascade. When the particle energies, type, directions and core positions are also allowed to vary, the variety of possible events becomes very large. It is fair to say that no two STACEE events are exactly the same. Given this randomness, STACEE uses a *Monte Carlo* type of simulation. In a Monte Carlo simulation we specify the physics that particles are required to follow and the environment that these particles will interact with. The programs then simulate a single incident particle, using random number generators to specify its energy, direction and core position. The particle is propagated through the atmosphere and the EAS/Cherenkov Flash is simulated; the interaction of the Cherenkov Flash with our detector is then modelled. This process is repeated until a sufficiently large number of Monte Carlo events have been generated. The goal of Monte Carlo simulations is not to simulate exactly every possible type of event, but rather to create a set of events that adequately spans the expected range of event properties.

Simulations of the STACEE experiment are divided into three sequential programs. They are:

1. `CORSIKA`: Simulation of the physics of extensive air showers and their production of Cherenkov photons. The output is a list of Cherenkov photons on the ground.
2. `Sandfield`: Simulation of the STACEE optical system. It takes the list of Cherenkov photons and traces their paths through the optical system until they reach the PMTs or are lost. Output is the number of photoelectrons detected in each PMT.
3. `Elec`: Simulation of the STACEE electronics. It takes the list of photoelectrons,

runs them through an electronics simulation and decides whether the event should trigger. Output is a data file that is almost identical to real STACEE data files.

Detailed descriptions of each of these programs are given in the following sections.

6.1.1 CORSIKA

The first step of the simulation chain is CORSIKA. Understanding extensive air showers is a necessary task for all ground-based γ -ray and cosmic-ray detectors. The package STACEE uses is called CORSIKA; it was originally developed for the KASCADE air-shower experiment [57]. CORSIKA is useful because it gives the option to use different treatments of the electromagnetic and hadronic interactions that occur in EAS. This is particularly important for very high energy (>100 GeV) hadronic interactions, where there is limited experimental data to constrain the interaction models. In the case of STACEE, we use the QGSJET model for high energy hadronic interactions [66] and the FLUKA model for low energy hadronic interactions [41,42]. For electromagnetic interactions, we use the well-tested EGS4 model [104]. For this thesis we use CORSIKA version 6.200.

To start a simulation, we specify the primary particle, its energy and its incident direction. We must also specify the altitude of our detector and the direction and strength of the local magnetic field. CORSIKA propagates the incident particle through the atmosphere, following it until it initiates an extensive air shower. CORSIKA then tracks the secondary particles, creating and destroying them as required by the appropriate physics model; the Cherenkov photon production by charged particles is also simulated. The output from CORSIKA is a file listing all the Cherenkov photons that have reached the altitude of STACEE.

Using CORSIKA also requires that we choose an atmospheric profile. We use the ‘US standard atmosphere’. Using this single standard atmosphere does not account for the fact that the atmosphere changes over the course of the year. Experience with STACEE has shown that our average Cherenkov trigger rate is reduced during summer time, a fact that we attribute to changes in the atmospheric profile. Work by members of other collaborations confirms this. Bernlohr (2000) found that there was a 15-20% reduction in the Cherenkov yield between winter and summer at mid-latitudes for simulated 100 GeV γ -rays. Their results are plotted in Figure 6.1. The fact that we use a single atmospheric profile for all our simulations will therefore add to our overall systematic error.

Further details about STACEE’s implementation of CORSIKA are available elsewhere [102].




Figure 6.1: Average lateral profile of Cherenkov photon density for different atmospheric profiles. The plots were derived from 100 GeV γ -rays generated with Corsika 5.71, at an altitude of 2200 m. The right-side plot shows the 15% difference in Cherenkov yield between summer and winter profiles at mid-latitude. From [15].

6.1.2 Sandfield

The next step of the simulation chain is *sandfield*. *sandfield* is a ray-tracing program designed to account for the optical component of the STACEE experiment. *sandfield* takes the list of Cherenkov photons on the ground generated by CORSIKA and propagates them through the heliostats, secondary mirrors and DTIRCs. The program also accounts for all the wavelength dependent absorption and transmission coefficients (and PMT quantum efficiency) that are shown in Figure 4.8.² The output of *sandfield* is a file listing the number of photoelectrons incident on each PMT for each event. Unlike CORSIKA, *sandfield* was written and is used solely by STACEE collaborators.

One of the goals of simulating STACEE is to determine how much difference there is in a given event depending on the *core position* of the EAS. The core position is the point on the ground where the incident particle would have landed had there been no atmosphere. An EAS might trigger the experiment if its core position is at the center of

²Actually, for reasons of computational efficiency, some of the wavelength dependant factors are applied in CORSIKA; there is no benefit in ray-tracing photons whose wavelength is outside the reflectivity curve.

the heliostat field or if the core position is 100 m away from the center of the field. But the distribution of charge across the different channels will be quite different; for instance, the total number of Cherenkov photoelectrons measured for an event will be lower if the core position is far from the center of the field. This dependence of the measured event properties on the core position will recur throughout this work, particularly when we come to discuss energy reconstruction in Chapter 9.

In order to quantify this effect, we need to simulate EAS with a variety of different core positions. We therefore scatter our core positions uniformly within a disk of radius 250 m for γ -ray EAS; we refer to this value as the scattering radius, R_{scat} . The Cherenkov light pool for cosmic-ray events is more spread out, so they can trigger the experiment with core positions at larger distances; we therefore use R_{scat} of 400 m for simulated cosmic-rays. Scattering the EAS core positions over such a wide area is somewhat inefficient since while some EAS might trigger the experiment with a core position at 200 m, most will not. To alleviate this problem we reuse each CORSIKA shower at several different core positions. CORSIKA is significantly more time-consuming than *sandfield* and we therefore save a great deal of computing power by reusing each CORSIKA-simulated EAS. However, we should not overdo this. If we reuse a simulated EAS too many times, then we will introduce too many correlated events into our data. Through a variety of tests, we have arrived at a good compromise of 5-10 re-uses of each CORSIKA EAS. In this work, each EAS is used (or *thrown*) at 5 different core positions.

The *sandfield* program has been extensively tested and calibrated by various STACEE collaborators. For instance, *sandfield* has been used to simulate *sunspots*, where we use individual heliostats to reflect the sun onto the solar tower; the simulations correctly predict the size and shape of the sunspots [60, 124]. We therefore have some confidence in our simulation of the STACEE optical system.

6.1.3 Elec

The final step of the simulation chain is the *elec* program. *elec* simulates the STACEE electronics system. Like *sandfield*, the *elec* program has been developed by STACEE collaborators. The *elec* program aims to reproduce as faithfully as possible all the important elements of the electronics chain, such as the PMTs, the discriminators, MADDOG and the FADCs. There are a number of important input parameters that must be specified before running *elec*. These are:

- The average PMT gain for each channel.
- The effective discriminator threshold for each channel.

- The average current for each channel.

As noted in Section 4.6.2 and 4.6.3, we have put a fair amount of work into understanding the PMT gains and effective discriminator thresholds. Before starting `elec` we estimate average values of all the input parameters using the real data set that we wish to simulate.

Using the input parameters, `elec` then creates the simulated electrical signals for each channel. The first step of the process is to add the Cherenkov photoelectrons. As noted the `sandfield` output is a list of photoelectrons and their arrival times at each PMT. For each photoelectron, `elec` creates a simulated electrical pulse with the following functional form:

$$F(t) = Ct^{2.37}e^{-1.32t}. \quad (6.1)$$

This functional form is inspired by measurements of real single-photoelectrons [127]. The height of each individual pulse is determined by the specified PMT gains.

Reading the list from the `sandfield` output only allows us to build up the electrical signal created by the Cherenkov photons. In order to add NSB background to this signal we use the input currents and PMT gains; with these we can calculate the average rate of NSB photoelectrons for each channel. `elec` then adds to the simulated signal enough single photoelectron pulses to match this flux of NSB photoelectrons. The NSB pulses have random hit times and use the same single photoelectron pulse shape as the Cherenkov pulses. Once the simulated electrical signal is completed (ie Cherenkov and NSB photoelectrons have been added) we simulate our discriminator modules using the specified effective discriminator thresholds. A list of discriminator hits and hit times is thereby generated. If the list of hits satisfies the requirements of the simulated MADDOG system, then the event is said to have triggered and is saved to the `elec` output file.

The output of `elec` is a data file in the same format as a standard file generated by the STACEE experiment. A significant amount of work has been done to ensure that the simulated data product is as similar as possible to real data product. The net result is that a simulated `elec` data file can be run through the `Pass0` and `Pass1` programs in a way that is almost the same as for a real data file. This greatly simplifies comparisons between real and simulated data. Figure 6.2 shows examples of real and simulated FADC traces. The figure shows that the simulated data is capable of reproducing at least the gross features of the FADC traces.

The `elec` output also contains additional ‘Monte Carlo Truth’ information that does not exist in real data. The most important of these MC-Truth quantities are:




Figure 6.2: Real and Simulated FADC traces. The trace on left is from real data; the trace on right is from simulated data.

- Type of incident particle (ie γ -ray, proton, helium nuclei).
- Energy of incident particle.
- EAS core position.
- Number of photoelectrons for each channel.

The knowledge of the actual energy of a given event is very important when, for instance, we come to developing background suppression and energy reconstruction schemes.

6.2 Simulation Details

The preceding sections described the basics of the STACEE simulation packages. In the following sections we touch upon a couple of details about the simulations that were performed for this particular work.

6.2.1 Type and Energy Range of Simulations

The choice of the type and energy for our simulations is determined by what we want to do with them.

Our simulations of γ -rays have a number of important purposes. The first is to describe how the likelihood of a γ -ray triggering the experiment varies with energy. This leads to an understanding of our effective area curves, which we shall describe in Section 6.3. The second purpose of γ -ray simulations is to allow the development of background suppression techniques. By comparing γ -ray and cosmic-ray simulations, we can determine quantities that allow the suppression of the latter. Finally, γ -ray simulations will

allow us to do more sophisticated spectral analyses for sources that we detect. Rather than simply calculating the integral flux of a source, we shall be able to say something about its spectral energy distribution.

There are therefore a number of different purposes for our γ -ray simulations. In order to have optimal results for the different purposes, we use two different schemes for choosing the energies at which the simulations are performed:

- Discrete energies: we perform a series of simulations at discrete energies. For instance, we perform separate simulations at 100 GeV, 110 GeV, 120 GeV, 130 GeV, 140 GeV, 150 GeV, 170 GeV and 200 GeV for the range 100 - 200 GeV, where we are most interested in our detector response. The discrete energy simulations are used for generating our effective areas, energy thresholds and integral flux measurements.
- Continuous energies: we perform a single simulation with a continuous distribution of energies. Specifically, the distribution of energies are defined by a power law,

$$dN/dE \propto E^{-\alpha}, \quad (6.2)$$

where α is the spectral index. This distribution of energies is close to that of many real VHE γ -ray sources, which means that a simulation with a continuous, power-law distribution of energies will be a good match to real data.³ We therefore use continuous energy simulations for both our background suppression studies and our energy reconstruction work. Unless otherwise noted, we use a spectral index of 2.4 for continuous energy γ -ray simulations.

Simulations of cosmic-rays are somewhat less important. As noted, they will prove useful for calibrating our background suppression techniques. Cosmic-ray simulations will also be compared to distributions of real cosmic-rays. Since our OFF-source data is a pure sample of cosmic-rays, we have a set of precise cosmic-ray measurements against which to compare our cosmic-ray simulations. We shall often do these comparisons between real and simulated cosmic-rays as a test to give us additional confidence in our simulations. There is, however, no benefit in generating effective area curves for cosmic-rays; it therefore suffices to only simulate cosmic-rays with a continuous power-law distribution of energies. We use spectral indices of 2.71 and 2.65 for our proton and helium nuclei simulations, respectively.

³To be clear: the power law describes the incident spectral energy distribution of the source, not the distribution of the energies of γ -rays that trigger the experiment.

The other element that we must choose is the range of energies that our simulations will cover. The low end of our energy range is chosen to be an energy at which the chance of triggering the experiment is negligible. For γ -rays the low end of our simulated energy range is typically 20 GeV. For cosmic-rays the Cherenkov light pool density is lower, so the low end of the energy range is higher; for protons we start simulations at 80 GeV and for helium nuclei we start simulations at 150 GeV. Choosing the high energy end of the simulations is less crucial since there will always be very few events at those energies. The high end of the energy range is 10,000 GeV for the discrete γ -ray simulations and 30,000 GeV for the continuous γ -ray simulations; it is 1,000,000 GeV for proton simulations and 3,000,000 GeV for helium nuclei simulations.

6.2.2 Pointing of Simulations

The performance of the STACEE experiment depends strongly on the elevation and azimuth of the source we are observing. This is partly because the Cherenkov light from an EAS that arrives from a lower elevation must pass through a larger length of the atmosphere. The Cherenkov photon density at the ground is therefore reduced, both by the spreading of the light pool and by increased atmospheric attenuation. A 100 GeV γ -ray that would have triggered STACEE if it came from zenith would probably not trigger if it came from a lower elevation. In addition, the response of STACEE also changes due to our optics. STACEE has a mirror area of 2400 m², but only if we are observing a source that is right behind the tower. If we observe sources anywhere else, our primary mirror area is effectively reduced by the off-axis alignment of the heliostats.

In order to account for these effects, it is necessary to do simulations at several points on the sky. For this work, we have chosen to do simulations at the transit point for each source and points +1 and +2 hours after the transit point. The transit point is where the elevation of a given source is maximized and consequently where the energy threshold for that source is lowest. The overall detector response is equal on either side of the meridian, so that data taken an hour before transit can be understood by looking at simulations done an hour after transit. Table 6.1 summarizes the azimuth and elevation of the points chosen for the Crab, Mrk 421, OJ 287 and 3C 66A simulations. Note that we shall normally refer to the different pointings by their hour angle, **HA**, location. So we shall refer to the ‘Mrk 421 +1HR’ simulation, instead of saying ‘the simulation with an azimuth of 289.5° and an elevation of 77.5°’.

In addition to the large scale pointing of the simulations, CORSIKA also allows us a smaller scale scattering around the nominal direction. This is important for simulating cosmic-rays, whose arrival directions are isotropic. For cosmic-rays we scatter the inci-

HA location	0HR (Transit)		+1HR		+2HR	
	Az	El	Az	El	Az	El
Crab	180.0	77.1	229.4	71.6	251.7	60.8
Mrk 421	360.0	86.8	289.5	77.5	286.6	65.8
OJ 287	180.0	75.1	225.6	70.1	248.6	59.7
3C 66A	360.0	81.9	309.3	75.9	298.3	65.5

Table 6.1: Simulation Pointings. The table lists the azimuth and elevation of the different points on the sky that were simulated for the Crab, Mrk 421, OJ 287 and 3C 66A. The Hour Angle locations are listed in hours, the azimuth and elevation are listed in degrees.

dent particles in solid angle out to a maximum angle of 3.0° from the nominal direction. We have found that it is very unlikely that a cosmic-ray will trigger the experiment if its arrival direction is more than 2.0° from the nominal direction of the source, so our solid-angle scattering adequately covers the possible range of incident directions. The γ -ray sources studied in this work are all point-sources, so the γ -ray simulations are not scattered in solid-angle.

6.2.3 Computational Requirements of Simulations

The simulations that are required to properly characterize STACEE are time-consuming, particularly for simulations of cosmic-ray EAS. This is partly because the hadronic processes involved in cosmic-ray EAS are more complicated and hence more time-consuming to simulate. It is also because the Cherenkov yield from hadronic showers is lower and we must therefore simulate more showers in order to get an equal number of triggered events.

For example, the Markarian 421 cosmic-ray simulations required almost a year and a half to run. The simulations were continuously running on an average of five computers with 1 GHz CPUs. Given this large investment of CPU time, it makes sense to reuse previous sets of simulations. For my simulations of the Crab, I used the set of `sandfield` simulation files done by Pascal Fortin for his PhD thesis [44]. It should be noted these simulations do not use the exact same versions of the simulation programs, nor do they necessarily use the same `CORSIKA/sandfield` input parameters. This is because they were performed some time ago and the STACEE simulation packages have undergone small upgrades since then. The modifications are, however, very small and not expected to change the final results in any significant ways. I ran my own version of `elec` on Pascal's `sandfield` files, so everything I describe about the electronics simulations apply as exactly to the Crab data as much as to the other data sets.




Figure 6.3: Effective Area for Mrk 421 (2003-2004). The plot shows the different effective areas for the transit, +1HR and +2HR simulations.

6.3 Effective Areas and Energy Threshold

Having described the parameters of our simulations, we can now introduce some of their most important results. In particular, we shall now describe the STACEE *effective area curves*. Our effective area curves allow us to define the energy threshold of the experiment, as well as to transform a measured trigger rate into a source flux measurement or source flux upper limit. The effective area $A_{eff}(E, HA)$ is calculated as

$$A_{eff}(E, HA) = \pi R_{scat}^2 P(E, HA), \quad (6.3)$$

where R_{scat} is the scattering radius and $P(E, HA)$ is the ratio of the number of γ -rays that triggered the experiment to the number that were simulated (ie. $P(E, HA)$ is the probability that a shower triggers the experiment). The effective area is a useful way to describe the ‘efficiency’ of STACEE for detecting γ -rays as a function of energy; a large effective area means a high ‘efficiency’.⁴ Note that the effective area is dependant on both the incident γ -ray energy, E , and the source hour angle, HA .

In order to calculate the γ -ray effective area, we ran simulations at a series of discrete energies. The results are shown in Figure 6.3 for a set of Mrk 421 simulations. The figure

⁴Effective area isn’t exactly like efficiency, since there is no γ -ray energy at which the STACEE experiment can be said to be 100% efficient. This, in turn, is because of there is no fixed limit to the core position of the γ -rays that STACEE detects.

shows the effective area curves for γ -rays coming from the Mrk 421 transit point and points +1 and +2 hours past transit. For each pointing, the simulations used a slightly different version of the `elec` input parameters, such as PMT gain, current and effective discriminator threshold. The input parameters were average values calculated from real Mrk 421 data taken near the HA of each pointing during the 2003-2004 season. There are a number of features of the effective area curves that are worth emphasizing:

First, the effective area increases with γ -ray energy. This is intuitive; higher energy EASs produce more Cherenkov photons and hence are more likely to trigger the experiment. Near 100 GeV the increase is dramatic, with the effective area increasing by many orders of magnitude; above that the increase is slower.

Second, there is a noticeable difference in the effective area curves for different source positions. For the +2HR pointing the effective area has decreased significantly at energies below 200 GeV. This is because Cherenkov photons from the γ -ray induced EAS must pass through a greater length of atmosphere to reach the detector. The Cherenkov light pool density at the detector is therefore reduced, both by the spreading of the light pool and by increased atmospheric attenuation. Lower energy γ -rays from the +2HR pointing are therefore unable to trigger the experiment. At higher energies an opposite effect is seen. Again, the greater path length spreads out the light pool. But since there are more than enough Cherenkov photons to trigger the experiment, the only consequence of a spread out light pool is that γ -rays whose core position is farther from the center of the field can trigger the experiment. Consequently, the effective area is even greater at higher energies for γ -rays from the +2HR pointing.

Figure 6.3 therefore demonstrates why we must do simulations at different pointings. In the end, however, we wish to extract information from our data set as a whole. It is therefore useful to start creating an average simulated curve, whereby the simulations at different pointings are weighted by the amount of time spent near each pointing in the real data set. For instance, during the 2003-2004 Mrk 421 data set 39% of the data was taken near $|HA| = 0HR$, 44% of the data was taken near $|HA| = 1HR$ and 17% of the data was taken near $|HA| = 2HR$.⁵ These percentages are then used to weight the three curves in Figure 6.3. The result is shown in the left-side plot of Figure 6.4, which is the HA-weighted raw effective area curve for the 2003-2004 Mrk 421 data set.

We can now use the HA-weighted effective area curve to properly define our energy threshold. It should be clear from the effective area curves that the STACEE detector does not have a sharp low energy cutoff. We must therefore use a more complicated definition

⁵Remember, we assume that the response of detector is symmetric about the local meridian, so that data taken at $HA = -1HR$ can be analyzed using the simulation at $HA = +1HR$.




Figure 6.4: HA-Weighted Simulation Results for Mrk 421 during 2003-2004 season. The left-side plot shows the HA-weighted raw effective area; the right-side plot shows the HA-weighted average response curve. The response curve is a convolution of the raw effective area curve with a typical γ -ray spectral energy distribution with a spectral index of 2.4.

of energy threshold. To do so, we must calculate the *response curve* for our detector. The response curve is the convolution of our effective area curve with an assumed spectral energy distribution of our γ -ray source. The response curve is therefore a plot of the differential trigger rate (counts $\text{GeV}^{-1} \text{s}^{-1}$) as a function of energy. An example is shown in the right-side plot of Figure 6.4 where the Mrk 421 effective area has been convolved with an assumed source spectrum with a spectral index of 2.4.

The convention in VHE astrophysics is that the energy threshold of a detector is the energy at which the differential trigger rate is maximized. From the plot in Figure 6.4 we can therefore say that the raw energy threshold for the Mrk 421 (2003-2004) data set was 150 GeV. This definition of energy threshold is clearly flawed in some ways. First, $\sim 15\%$ of the γ -rays that trigger STACEE are below the energy threshold. Second, our energy threshold depends on the form of the underlying γ -ray spectrum that we assume; if we had assumed a form with a spectral index of 1.4 instead of 2.4 then the energy threshold would have been higher. This entanglement of energy threshold with the assumed form of the γ -ray spectrum is unfortunate. Nevertheless, in the absence of any better definition of energy threshold, we must continue to use the conventional one. We shall be returning to the effective areas and response curves in later chapters, since the various event cuts that we apply are energy dependant. We shall therefore see that our final effective area curve is different from the raw effective area curve shown in Figure 6.4; in particular, our final

energy threshold will be higher than 150 GeV.

Finally, it should be noted that the integral of the response curve gives the predicted STACEE raw γ -ray trigger rate for a source with this particular energy spectrum. The response curve therefore allows us to convert between a measured γ -ray trigger rate and an integral source flux or flux upper limit, which we shall be doing in Chapter 10.

6.4 Tests of Simulations

As noted, our simulations provide an excellent tool for understanding the STACEE experiment. As we shall see, simulations are extremely useful when it comes to investigating background suppression techniques and energy reconstruction. However, this is only true if our simulations are a good match to our real data. If this is not the case then nothing we learn from the simulations is of value. For instance, it should be clear from the description of `elec` that the simulated results depend strongly on a good understanding of the PMT gains and effective discriminator thresholds. If our estimates of the real PMT gains are wrong, then the simulated data will not match the real data. It therefore behooves us to perform some tests in order to ensure that the simulations are a good match to our real data. It should be noted that both of the following tests are based on earlier work by other STACEE collaborators [111, 124].

6.4.1 PMT Rate Comparison

As noted earlier, the *PMT rate* is the rate at which the discriminator for each channel produces hits. The PMT rate is typically in the MHz range and is driven entirely by hits from NSB photons. Simulating the PMT rates depends only on elements of `elec`. The PMT rates are strongly dependant on the PMT gains, the currents and the effective discriminator threshold; the test is very sensitive to changes in each of these quantities. A successful comparison will therefore give us increased confidence in our electronics simulation. For this comparison, we run `elec` in a special mode, where we produce simulated electrical signals using only NSB photons. There is no need to use a `sandfield` input file for this mode, since the Cherenkov photons have a negligible impact on the PMT rates.

The comparison is done as follows: we calculate the average real and simulated PMT rates for a real individual data run. To get the predicted PMT rate `elec` uses the average currents, PMT gains and effective discriminator thresholds for that particular run. Once the simulations are completed, we calculate the ratio of the real to simulated PMT rates for each channel. We then repeat the procedure for every run in the 2002-2004 Crab and




Figure 6.5: PMT rate comparison - ratio distribution. The plot shows a histogram of the ratio of real-to-simulated PMT rates. Each entry is for a single channel and a single run. The plot shows the results for Crab and Mrk 421 data for 2002-2004 seasons.

Mrk 421 data sets. A histogram of the ratios of real to simulated PMT rates is shown in Figure 6.5. The plot shows that the peak of the distribution of ratios is centered near 1; this indicates that our electronic simulation is fairly good.

There is, however, a noticeable tail to the distribution where the real PMT rate is more than twice the simulated PMT rate. The explanation for this tail becomes more clear when we present our results in a different fashion. Figure 6.6 shows a profile plot of the real and simulated PMT rates as a function of current. The first point to note is that, as expected, the PMT rates increase with increasing current. The second point is that there is generally excellent agreement between the real and simulated curves. The only place where the agreement breaks down is for currents that are smaller than $30 \mu\text{A}$; in the low current/low PMT rate regime, the real PMT rates are noticeably larger than the simulated PMT rates. The most logical explanation for this discrepancy is the presence of Radio-Frequency (RF) noise in the real data. As described in Section 4.6.4, RF noise causes additional background fluctuations in the electrical signals that are unrelated to the PMT currents. Large amounts of RF noise can lead to increases in our real PMT rates; the relative increase is most significant when the currents are small. The `elec` simulation, by contrast, does not contain significant RF noise. It is therefore reasonable to claim that RF noise is the cause of the discrepancy between real and simulated PMT rates in the low




Figure 6.6: PMT rates as a function of current. The plot shows profiles of PMT rate vs current for both real and simulated data. As before, the plot shows the results for Crab and Mrk 421 data for 2002-2004 seasons.

current regime.

Even excluding the large-ratio tail, one might object that there is a considerable spread of the ratio around 1.0. This is the result of how strongly the PMT rate depends on the various electronics parameters. Even small errors in our estimate of our currents or effective thresholds can lead to a noticeable difference between the real and simulated PMT rates. What is most important, however, is that *on average* our description of the electronics is good enough that the ratio distribution is peaked at 1.0. This comparison therefore shows that our simulation of our electronics chain is working properly. The runs used for this comparison span the 2002-2004 seasons, which indicates that we understand our PMT gains and effective discriminator thresholds well over this entire period. It has been found that changing the PMT gains used in the simulation by as little as 10% results in a noticeably worse match between real and simulated PMT rates.

6.4.2 Mrk 421 Cosmic-ray Rate Comparison

The second test of our simulation chain is a comparison of our real and simulated OFF-source cosmic-ray rate. Our OFF-source trigger rate is entirely due to cosmic-ray induced Cherenkov events. Simulating the OFF-source cosmic-ray rate therefore tests our entire




Figure 6.7: Measured spectral energy distributions of proton and helium nuclei cosmic-rays from several different experiments. The dashed red line is the fit that will be used in the cosmic-ray simulations in this work. The plot shows that there is a substantial amount of disagreement between the results of different groups. From [124].

simulation package: CORSIKA, sandfield and elec. A successful comparison will provide increased confidence in all elements of the simulation chain.

The first thing we require for this test is an estimate of the spectral energy distribution of cosmic-rays in our energy regime. The cosmic-ray spectrum has been measured by several other experiments; Figure 6.7 summarizes these measurements. The dashed lines show the best fit for the results from several different collaborations (it is clear that the agreement between the different experiments is not perfect). The distribution of particle energies is found to be adequately described by a power-law distribution. The measured spectral indices are 2.71 and 2.65 for proton and helium nuclei, respectively, which was

Position	Proton Rate (Hz)	Helium Rate (Hz)	Total Predicted Rate (Hz)	Total Real Rate (Hz)
Transit	3.96 ± 0.06	1.80 ± 0.05	5.76 ± 0.08	5.29 ± 0.02
+1 HR	4.18 ± 0.09	1.80 ± 0.05	5.98 ± 0.10	5.50 ± 0.02
+2 HR	3.78 ± 0.09	1.70 ± 0.07	5.48 ± 0.11	4.80 ± 0.03

Table 6.2: Cosmic-rate rate comparison. The table shows the predictions for the proton, helium and total trigger rate for three different pointings of Mrk 421. The table also shows the real measured cosmic-ray trigger rate for comparison.

used in our cosmic-ray simulations.

The simulations are performed at specific points along the track that our source takes across the sky. The input parameters for the `elec` (PMT gains, currents, effective discriminator thresholds) are estimated by averaging the values for real data taken near each pointing. Using these parameters, full simulations are performed for three different source points: transit, +1HR and +2HR. The predicted cosmic-ray trigger rate for each individual species (protons and helium nuclei) is calculated as

$$\text{Trigger Rate} = \pi R_{\text{scat}}^2 \Omega \Phi P, \quad (6.4)$$

where R_{scat} is the radius over which the core positions were scattered, Ω is the solid angle over which the incident directions were scattered and Φ is the integral flux of the cosmic-ray spectrum above a low energy cutoff, E_0 , at which the probability of triggering the experiment is negligible. The final quantity, P , is calculated using the cosmic-ray simulations. P is the probability that one of our simulated cosmic-ray events will trigger the experiment; that is to say, P is the ratio of the number of simulated events that trigger the experiment over the total number of simulated EAS.

Let us give a specific example. We simulated a total of 20,000,000 proton events coming from around the Mrk 421 transit point; of these, only 5001 events actually triggered the experiment. The energies of the simulated protons were chosen based on a power law spectrum, with a spectral index of 2.71 and a low energy cutoff, E_0 , of 80 GeV. The integral proton flux above this cutoff is believed to be $3.66 \text{ cts m}^{-2} \text{ sr}^{-1} \text{ s}^{-1}$ (from Figure 6.7). The simulations were done by scattering the core positions over a circle of radius 400 m and by scattering the incident direction over a solid angle of $8.6 \times 10^{-3} \text{ sr}$. The estimated rate of triggers from cosmic-ray protons is therefore $3.96 \pm 0.06 \text{ Hz}$.

Table 6.2 summarizes the results for both species for all three pointings. The table also lists the total predicted cosmic-ray rates and the real measured cosmic-ray rates from the Mrk 421 OFF-source. The simulated cosmic-ray rates are very close to the measured rates. The agreement is not perfect; it is apparent that the simulations over-

predict the cosmic-ray rate by approximately 10%. Nevertheless, given the uncertainties in the underlying cosmic-ray spectral energy distributions, our lack of understanding of the atmospheric conditions and the general complexity of the STACEE experiment, the comparison should be considered successful.

These two tests provide convincing evidence that our simulations are a good match to our real data. This should provide confidence when we start to use the simulations to develop new analysis techniques and extract meaningful scientific results from our STACEE measurements.

Chapter 7

Library Padding: Accounting for the Promotion Effect

We described in Chapter 5 how the raw γ -ray rate, the difference between the ON and OFF trigger rates, is calculated. If the ON and OFF observations were exactly equivalent then the raw γ -ray rate would be an accurate measure of whether a source was an emitter of VHE γ -rays. We have already explained most of the steps that we take to make the ON and OFF observations equivalent. These include taking care that ON and OFF runs cover exactly the same part of the sky and ensuring that the detector is not malfunctioning in any way for half of a pair. There is, however, one additional bias that we must correct for during the off-line data analysis; this is the effect of promotion.

Promotion couples our cosmic-ray trigger rate to the amount of night sky background (NSB) photons. The coupling occurs as follows. NSB photons cause additional fluctuations of the electrical signal at our discriminators. The fluctuations can push above threshold (promote) an event that would not have triggered by itself; they can also push below threshold (demote) an event that would have triggered by itself. However, there are more low energy events than high energy events; there will therefore always be more events to promote than demote. Consequently the presence of NSB photons results in a net excess of promoted events and hence an increase in our triggered cosmic-ray rate.

The fact that promotion couples the cosmic-ray trigger rate and the NSB is not *per se* a problem. The problem occurs when there are systematically different amounts of NSB light for the ON and OFF runs. For instance, if there is a bright star in the ON-source field-of-view, then there will be systematically more ON-source promoted events. The result will be a systematic bias of our raw γ -ray rate. The goal of this chapter is to eliminate this spurious excess and ensure that any remaining signal is the result of γ -rays.

Promotion is a generic feature of all γ -ray Cherenkov telescopes, since any PMT-

based detector will be sensitive to the changes in NSB levels. This was noted during the early work on detecting Cherenkov light from cosmic-rays in Jelley and Galbraith (1958) and Jelley and Porter (1963) [64, 65]. These early experiments accounted for promotion at the hardware level, by attaching a lamp near the front of the PMT. The lamp was connected to a feedback circuit that ensured that the PMT always saw a constant light level, no matter which stars were passing through the field of view. A similar hardware based technique was also used during the first successful Whipple detection of γ -rays from the Crab Nebula [150]. For the Whipple work this typically meant adding 30% more current than the dark night sky caused by itself. Later Whipple results corrected for promotion in off-line software analysis [30].

There have been a number of different STACEE attempts to account for promotion. Boone (2002) developed a method of estimating the excess hadron rate from promotion as a function of the ON/OFF current difference [22]. With the introduction of the FADCs, STACEE has been able to develop more sophisticated techniques to account for promotion. Scalzo (2003) developed a ‘Dynamic Thresholds’ technique of incrementally raising the analysis thresholds in order to account for promotion [124].

Library Padding is a third scheme to account for promotion. Library Padding involves adding extra noise to the FADC traces and re-imposing a trigger condition using a higher threshold. Library Padding is the procedure that comes closest to producing the data set that we would have acquired if we never saw any field brightness asymmetries. It accomplishes this at the cost of having a higher energy threshold than the other two techniques. Library Padding is a particular implementation of a padding-type scheme, that is to say a scheme where extra ‘noise’ is added to the data to account for promotion. The scheme is therefore descended from earlier Jelley/Whipple efforts. It is *Library* Padding because we use a library of sample traces derived from real data, as opposed to just adding random Gaussian noise or Monte Carlo generated photoelectrons.

It is the Library Padding scheme that will be used in this work to account for the promotion effect. In this chapter we will explain and verify the scheme. We start by introducing our star data sets in order to quantify the scale of the promotion problem. We then describe the Library Padding procedure in detail; this will include several tests of the variance equalization method. Finally we will prove that Library Padding works by showing that it successfully removes the promotion excess for the star data sets. With the application of Library Padding we claim that all the systematic biases of our γ -ray rate have been accounted for.

Table 7.1: Summary of Star Data Sets. The table shows the star name, B Band magnitude, as well as the livetime, average current difference and raw excess rate of the STACEE data sets.

Star name	B Band Magnitude	Livetime (ks)	$\langle \Delta I \rangle$ (μA)	Raw Excess Rate (counts min ⁻¹)
Hip 89279	5.98	8.1	1.80 ± 0.04	8 ± 2
Her 59	5.29	5.0	3.7 ± 0.2	22 ± 3
Iota Corona Borealis	4.93	10.5	4.9 ± 0.1	20 ± 2

7.1 Promotion and Star Data

The introduction gave a purely qualitative description of promotion. It is easy to understand that NSB light would occasionally influence the triggering of a Cherenkov event; but it is not clear how often this would occur. In order to better quantify promotion STACEE has undertaken a series of *star observations*. During star observations we use our regular ON/OFF technique; but instead of being a potential γ -ray emitter, the ON source is just a bright star. The bright star results in more promotion. The resulting difference between the ON and OFF rates is therefore entirely due to the effect of promotion. This rate difference is what we defined as the raw γ -ray rate; but since it is clearly not caused by γ -rays we shall, to avoid confusion, call it the ‘raw excess rate’ for the star data sets.

We took a total of three data sets on bright stars during the 2002-2004 seasons. We ran these data sets through the standard data quality programs in order to remove bad data and then computed the raw excess rate. A summary of the livetimes and raw rates for the three data sets is shown in Table 7.1. The table also lists the average current difference for each data set. The average current difference, $\langle \Delta I \rangle$, characterizes how much brighter the ON source field is than the OFF source field. The average current difference is defined as

$$\langle \Delta I \rangle = \frac{\sum_{i=0}^{64} (I_i^{ON} - I_i^{OFF})}{64}, \quad (7.1)$$

where I_i^{ON} and I_i^{OFF} are the PMT currents for channel i for the ON and OFF fields.

A larger current difference means a larger difference in the level of NSB fluctuations. We would expect that this would lead to a larger rate of promoted events and this is indeed what we find. This is clearly seen in Figure 7.1, which is a plot of the raw excess rate versus the average current difference. There appears to be an approximately linear relationship between the promotion rate and the average current difference. In order to give a sense of the scale of the promotion effect, Figure 7.1 also shows the raw excess rate from the Crab Nebula that we calculated in Chapter 5. The raw Crab excess rate




Figure 7.1: Raw excess rate as a function of the average current difference. The plot shows excess rates for the three star data sets, as well as for the 2002-2004 Crab data set. The line is constrained to go through the origin; only the star data sets are used in the fit. The dashed lines indicate the 1σ errors on the slope measurement.

is completely dwarfed by the excess rate from the star data sets. The average current difference for the Crab happens to be small, so that the raw Crab excess rate is mostly the result of γ -rays. But for other sources the average current difference can be much larger. In particular, the average current difference for Mrk 421 is approximately $3 \mu\text{A}$, indicating that the Mrk 421 raw γ -ray rate will be massively distorted by promotion.

Figure 7.1 makes it clear therefore why it is so important that the promotion effect be accounted for. The figure also suggests one simple technique for doing so. As noted, the relationship between the promotion excess and average current difference appears to be linear. The average current difference for a data set is easy to measure, so we could use the empirical fit shown in Figure 7.1 to estimate how much of excess rate is due to promotion. For instance, the average current difference for the Crab is $\approx -0.3 \mu\text{A}$, which means that the excess from promotion is $\approx -1 \text{ count min}^{-1}$. We could therefore estimate that the ‘true’ Crab rate is $4 \text{ counts min}^{-1}$. The earliest STACEE attempts at accounting for promotion were based on just such an empirical correction [22, 23].

The problem with this method is that our final measurement is then only as good as our measurement of the slope shown in Figure 7.1. If we simply subtract off a correction based on that slope, then our final result picks up a systematic error that is proportional to the error on the slope. This might be a relatively small effect for strong γ -ray sources or sources with very small average current differences. But if we wish to observe a relatively

weak γ -ray source with a large current difference, the systematic error from the promotion correction will quickly dominate over the statistical error on the raw rate measurement. In addition, the linear relationship shown in Figure 7.1 is probably only true for data taken under exactly the same hardware conditions; for instance, it might not hold if our nominal discriminator threshold changes from night to night, which it often does. For this reason the STACEE collaboration set about devising a more sophisticated correction for promotion, where the method does not (or at least need not) add a systematic error to our rate measurement. The Library Padding scheme is my particular solution to this problem. It is worth noting that Library Padding is only possible because the FADCs provide us with the digitized PMT signals.

A comment on terminology: for simplicity's sake, we assume throughout this chapter that the ON source field is brighter than the OFF source field. This makes descriptions simpler, since we can then refer to the 'ON-source variance' rather than having to say the 'variance of the run with a brighter star field' or the 'variance of the run with higher currents'. In reality it is of course possible that the OFF source be brighter than the ON source, as it indeed is for the Crab.

7.2 Description of the Library Padding Scheme

Applying the Library Padding scheme involves two distinct steps:

1. Equalizing the NSB fluctuations in the ON and OFF FADC traces; to do this we *pad* the OFF FADC traces for each channel with a set of sample NSB traces.
2. Applying a *padding cut*. This cut is similar to the online trigger but is applied at a higher threshold.

In the following sections we shall explain each of the two steps in more detail.

7.2.1 Background Equalization

The goal of this step is to equalize the NSB fluctuations seen in the ON and OFF FADC traces. Remember we are assuming that the ON source is brighter and has larger NSB fluctuations. The NSB photons are a random background and there is therefore no way to *decrease* the ON-source NSB fluctuations. We must therefore *increase* the OFF-source NSB fluctuations by adding additional noise to the OFF FADC traces. Specifically, we shall add a sample FADC trace which contains only NSB noise. The extra noise is added

early in our analysis chain, during the `Pass0` program, so that the fluctuations are equalized in all the subsequent analysis.

We must start by properly quantifying the background fluctuations. The most appropriate method of characterizing the scale of NSB fluctuations is by computing the *variance*, σ^2 . As noted in Section 4.5.3, the FADC acquisition program calculates a variance for each channel for each event using the 400 samples before the main trace. Specifically, it calculates

$$\sigma^2 = \frac{\sum_{i=1}^{400} (V_i - \langle V \rangle)^2}{N - 1} \quad (7.2)$$

where V_i is the voltage for bin i before the main FADC trace and $\langle V \rangle$ is the mean voltage for the 400 ns. When we mention the variance of a given channel, we will almost always be referring to this 400 sample measurement.¹

To start with, therefore, the average variance of the ON traces will be larger than the OFF traces. We must add a sample trace to the OFF-source FADC data such that the ON and OFF variances will be equalized. A cartoon of this procedure is shown in Figure 7.2. The cartoon shows how we start with an OFF trace, choose an appropriate sample trace that contains only NSB noise and add the sample trace to the original trace. Choosing a sample trace is made simpler by remembering a statistics fact: if we add two Gaussian distributions with variances σ_a^2 and σ_b^2 then the resulting distribution will have variance $\sigma_a^2 + \sigma_b^2$. So, if we assume that our background fluctuations are Gaussian then we can simply pick a sample trace by requiring that the variance of the sample trace be

$$\sigma_{sample}^2 = \Delta\sigma^2 = \sigma_{ON}^2 - \sigma_{OFF}^2 \quad (7.3)$$

Note that the ON-OFF variance difference, $\Delta\sigma^2$, is normally small compared to the original variance level. The typical ON or OFF NSB variance is approximately 2000 (mV)², whereas the NSB variance difference is approximately 100 (mV)². The cartoon in Figure 7.2, which shows that the sample trace has approximately the same NSB variance as the original trace, is therefore an exaggeration of what occurs in reality.

The following sections describe some of the details about the creation and use of sample libraries and the calculation of variance differences. It also includes two tests to show that the padding works correctly.

¹As noted in Section 4.6.2 the variance and current are linearly related. So we could also quantify differences in NSB fluctuations in terms of average current differences, as we did in Section 7.1. But the average variance difference is more directly relevant to the process of padding.




Figure 7.2: Cartoon of procedure of padding real FADC data with sample NSB traces. Only the peak at 18 ns is the result of Cherenkov photons; all other fluctuations are the result of NSB.

Library of sample traces

The following describes the process used to create the library of sample traces. We shall start with a series of useful definitions:

- *sample trace*: a single FADC trace which is to be used to pad our FADC data.
- *sample set*: a set of 1000 sample traces, all having (approximately) the same average variance. Each set is characterized by a certain *variance level* measured in $(\text{mV})^2$.
- *sample library*: the combination of many sample sets at a series of different variance levels. Padding requires a complete sample library.

The sample traces are generated from real FADC data, rather than from simulations. A simulated NSB trace can be relatively easily created by adding a large number of overlapping single photoelectrons. However the final variance will be very sensitive to the assumed form of the single photoelectron pulse shape. Since we do not understand our

single photoelectron pulse shape well, using simulated NSB traces might introduce subtle systematic errors into the padding process. Using sample traces derived from real data is the preferable solution.

In order to create the sample library we took a series of special runs. For these runs we used the South PMT cameras at the 120' level of the tower. The 120' level has a garage door that can be closed; when closed, no NSB photons hit the PMTs. This provides a low variance baseline ($\sigma^2 < 200 \text{ (mV)}^2$) for making the sample traces. This was important since the typical ON-OFF variance difference was approximately 100 (mV)^2 ; we needed sample traces at these small variance levels in order to properly pad the OFF traces. We mounted an LED at the focal point of the secondary mirrors. By varying the LED brightness we could create sample traces at a variety of different variance levels. Since we are only interested in creating a random background of photoelectron fluctuations, it does not matter whether the photoelectrons are caused by an LED or by starlight.

At each different LED light level we would take a run composed of 1000 fake events. The 1000 fake FADC traces for a single channel would constitute a sample set. We characterize each set by the average variance. Since each channel would see slightly different amounts of light we could build up a library of sample sets with different characteristic variances. In principle we would like to have a sample set at each, say, 1 (mV)^2 , ie 1, 2, 3, 4... 1499, 1500 (mV)^2 . In reality, it would require too much effort to produce such a library. We found that it was sufficient to have a sample library with variance level steps of 20 (mV)^2 below 200 (mV)^2 and steps of 50 (mV)^2 above.

Note that our procedure for generating the sample library makes the assumption that a sample FADC trace generated for channel 48 can be used to pad FADC data for channel 49. This should be a reasonable assumption. Problems might occur only if the average PMT pulse shapes are significantly different channel to channel; we have checked that this is not the case.

Calculation of variance differences

When characterizing the variance of a trace, we use the 400 sample measurement of the variance. 400 samples is more samples than we have available if we only use the portion of the FADC trace that is actually written to disk. This statistical power is important because each variance measurement is quite imprecise. This is shown in Figure 7.3, a plot of the calculated variance for each event during a run. There is a 15% spread in the calculated variances for this channel, despite the fact that the current during this run never changed. We may therefore conclude that the error on a single measurement of the variance is approximately 15%.




Figure 7.3: Left: series of variance measurements for a single STACEE channel for a run. Right: current measurements for same channel during same run; indicates that current (and hence the NSB rate) was stable for this run. The spread in the measured variance distributions is therefore the result of measurement error.

The fact that the variance has a relatively wide distribution is not unexpected. The problem is that the error on a single measurement of the variance is large compared to the precision we desire for the ON-OFF variance differences, $\Delta\sigma^2$. In the previous section, we noted that our variance samples were quantized at the level of 20 (mV)^2 ; we would therefore like to know $\Delta\sigma^2$ to approximately that level of precision. On the other hand, Figure 7.3 shows that the error on a single variance measurement can be of the order of 300 (mV)^2 . It is clear that to get a precise measurement of $\Delta\sigma^2$, we must average the variances for many events. It has been found that the required precision is achieved if the ON and OFF variances are averaged over two minute intervals. We therefore use the following formula to choose the variance level for the sample traces

$$\sigma_{\text{sample } ij}^2 = \langle \Delta\sigma_{ij}^2 \rangle = \langle \sigma_{ON_{ij}}^2 \rangle - \langle \sigma_{OFF_{ij}}^2 \rangle \quad (7.4)$$

where $\langle \sigma_{ON_i}^2 \rangle$ is the average ON variance for channel i for the two minute interval j . This measure of the average variance difference for each run is stored in a database at UCLA. During the `Pass0` program the average variance differences are downloaded and used to choose the sample set with which to pad each channel. The result is that the FADC traces that are written to disk by `Pass0` have equal levels of NSB fluctuations in the




Figure 7.4: Histograms of FADC bin values using fake events. The channel variances are the square of the RMS of these distributions. Left: A channel with a relatively low variance. Right: A channel with a relatively high variance. Note the significant amount of clipping near the 100 mV bin.

ON and OFF data. It should be emphasized that this padding is done separately for each channel. Some channels might have a large variance difference and need to be padded with 500 (mV)^2 traces; other channels might have minimal variance difference and need no padding.

Two tests of variance equalization

Equation 7.4 defines how we choose the sample traces with which to pad our data. But it is reasonable to ask whether the equation is an accurate description of how the variances of real FADC traces add. To understand why this might be a concern, we need to take a step backwards and look at the distributions of background fluctuations from which we derived our variance. Figure 7.4 shows histograms of every FADC bin values for fake events for a typical STACEE run; the results are shown for two different channels.

Using fake events ensures that we are not ‘contaminated’ by Cherenkov pulses; the histograms are therefore a clean view of our NSB fluctuations. The square of the RMS of these distributions is equal to the average variance we calculate using the 400 ns measurements of variance. Throughout this section we shall be manually calculating the variance levels, rather than using the 400 ns measurements. This is because the 400 ns measurements do not change when the FADC trace is subsequently modified. Since we wish to




Figure 7.5: Padding linearity test. The plots show the final variance of our FADC data vs the variance of the sample set with which it was padded; the results are shown for five different channels. The straight lines are *not* fits: their slopes are fixed to 1.0 and their y-intercepts are set to the variances of the original unpadded FADC data.

explicitly check the consequences of padding, we cannot use the 400 ns measurements.

There are two features of the distributions in Figure 7.4 that are cause for concern:

1. The distributions of NSB fluctuations are not quite Gaussian. Equation 7.4 is only strictly true for a Gaussian distribution.
2. The right-hand distribution shows that *clipping* can be a significant effect for channels with large variances. Clipping occurs when the FADC trace gets saturated at the high end of the scale at 100 mV. The result of clipping is a truncated distribution and hence an underestimate of the true variance.

Both these effects clearly occur; but how significant are they? Is Equation 7.4 truly an accurate description of how FADC variances add? To decide, we do two tests of our padding procedure.

For the first test we start with a typical STACEE run. We process this run through `Pass0` several times, each time padding the FADC data for all channels with sample sets of progressively larger variance; we pad using the 100, 200, 400, 600, 800, 1000, 1250 and 1500 (mV)² sample sets. We then calculate the final variance for several channels




Figure 7.6: ON-OFF variance differences for iota Corona Borealis data set. The variance differences are shown for each channel, for both padded and unpadded FADC traces. Note that these points do have error bars; they are just too small to see.

for each sample set. We use the padded FADC traces that `Pass0` writes out to manually calculate the final variances (as described earlier).

The results are shown in Figure 7.5. The plot shows the calculated variances for several different channels versus the variance level of the sample set used to pad that run. The relationship between the two is expected to be linear with a slope of 1.0. In general this is indeed what we find. This is particularly true for channels that have a lower overall variance such as channel 34; the curve for this channel is confirmation that equation 7.4 is generally valid.

However, there is a breakdown of linearity at larger variance levels. For instance the final variances for channel 30 diverge significantly for large variance sample sets. This is the result of clipping. The more noise that is added to the FADC trace the more clipping there is; consequently the more the variance is underestimated.² In general the effect of clipping is fairly small. The average variance difference for real data is less than 600 (mV)^2 , even for bright sources. We make the assumption that this effect is negligible and do not try to correct our measurements of the variance.

A more important test of variance equalization is shown in Figure 7.6. This plot shows the difference between the ON and OFF variances for each channel for the iota Corona

²It should be noted that the process of padding *respects the limits of the FADC*. So if a sample trace is added and the resulting trace is outside the normal range of the FADCs then these values are *clipped* to the maximum/minimum FADC values.

Borealis data set. We use the padded FADC traces that `Pass0` writes out to manually calculate the ON and OFF variances (as described earlier). Figure 7.6 shows that without padding there is a large difference between the ON and OFF variances. After applying the Library Padding procedure the variance difference has been eliminated, as expected.

Together these two tests prove that Equation 7.4 is a good description of the variances of our NSB fluctuations. They also show that the procedure of adding sample traces does eliminate any difference between the ON and OFF NSB fluctuations as seen in our FADC traces. This padding procedure is therefore applied to *all* our real data sets.

7.2.2 Padding Cut

The second step of Library Padding is applying a *padding cut* using our padded (ie variance equalized) FADC data. We choose to apply a cut that closely resembles the online trigger. We analyze each FADC trace and look for ‘hits’ where the pulse crosses the analysis threshold. These hits are used to check for the trigger condition, which, like the hardware trigger, is a two-level trigger condition requiring 5/8 hits in a cluster and 5/8 clusters in the experiment in order to satisfy the cut. The crucial difference is that the analysis threshold for the padding cut is at a higher value than the hardware threshold.³

We apply a cut at a higher threshold because we cannot get back OFF-source events that would have triggered the experiment had the online background fluctuations been larger; these events are not part of our data set. Suppose we applied our padding cut at the same threshold as used by the hardware. In that case essentially all the events would pass the cut, in spite of all the extra NSB noise that we added. Consequently we would still have a significant excess. It is clear, therefore, that we must raise the threshold when applying our padding cut. The relevant quantity is the *threshold raise*, X ; the FADC data for each channel is analyzed at a threshold of

$$\text{Analysis Threshold} = \text{Effective Threshold} + X \quad (7.5)$$

One problem with the Library Padding scheme is that it is not *a priori* obvious how much it is necessary to raise the threshold. We can, however, empirically estimate the value of X using our real star data sets. This procedure is described in the following section.

³The effective discriminator thresholds.




Figure 7.7: Final excess rate as a function of X for iota Corona Borealis. The curves are for both padded and unpadded data. The padded curve shows that the excess is successfully eliminated for $X > 24$ mV.

7.3 Real Tests of the Library Padding Scheme

This section describes the application of Library Padding to the star data sets that were introduced earlier. For this analysis we shall introduce a new measurement, the STACEE *padded γ -ray rate*. The padded γ -ray rate is the difference between the ON and OFF rates for events that satisfy the padding cut. *The goal of Library Padding is therefore that the padded γ -ray rate be consistent with zero for the star data sets.*

The first star data set that we will analyze with Library Padding is iota Corona Borealis. We shall use this data set to determine the value X at which the promotion excess is eliminated. We wish to determine X for iota Corona Borealis because this is the star data set for which there is the largest average variance difference. It is reasonable to assume that we must raise the analysis thresholds more for a data set with a larger average variance difference. By defining X using the data set with the largest current difference we can have confidence that Library Padding would also work for any other data set.

In order to determine how much to raise the analysis threshold we recalculate the ON and OFF rates after a padding cut has been applied at a variety of different values of X . The results for iota Corona Borealis are shown in Figure 7.7. The plot shows the difference between the ON and OFF rates as a function of X . For unpadded data there is a significant excess for all values of X . This indicates that you cannot solve the problem of promotion simply by raising the analysis threshold; it is essential that the NSB




Figure 7.8: Padded γ -ray rates from several star sources. The dashed line denotes the zero rate level (it is not a fit).

fluctuations be equalized before a padding cut be applied. For padded data the excess goes away by $X = 24$ mV. We shall therefore set $X = 24$ mV for our proper padding cut.⁴ The padded γ -ray rate for iota Corona Borealis is therefore -0.9 ± 1.7 counts min^{-1} , which is consistent with zero.

As additional tests, we then apply the same Library Padding procedure to the Her 59 and Hip 89279 star data sets. We apply the padding cut (with $X = 24$ mV) and compute the padded γ -ray rate. For Her 59 the rate is 0.7 ± 2.3 counts min^{-1} and for Hip 89279 the rate is -0.9 ± 1.7 counts min^{-1} . Library Padding successfully eliminates the promotion excess from both these star data sets.

These tests on the star data sets clearly demonstrate that the Library Padding procedure is working properly. These results are shown graphically in Figure 7.8, which shows the padded γ -ray rates for our three star data sets as a function of the current difference; the promotion excess has been eliminated. The figure also shows that Library Padding does *not* remove the excess for the Crab. This is good and expected; the excess from the Crab is the result of Crab γ -rays. After Library Padding, the padded Crab γ -ray rate is 2.78 ± 0.58 counts min^{-1} . *Since Library Padding has removed the last source of systematic error, we can therefore meaningfully say that we have detected the Crab at the 5σ level.*

⁴Remember, our effective thresholds are typically about 120 mV. So we are raising the thresholds by an additional 20%.




Figure 7.9: Effective areas for 2003-2004 Mrk421 data set. The plot shows the raw effective area and the effective area after the padding cut has been applied.

7.3.1 Library Padding and the Energy Threshold

We have shown that the Library Padding satisfies its stated goal: the removal of the spurious excess due to promotion. It must be added, however, that there is a serious drawback to the Library Padding procedure; it increases the STACEE energy threshold. This is unsurprising. It was noted before that we run the discriminator thresholds as low as possible in order to reduce our energy threshold. Clearly, therefore, increasing the discriminator thresholds in the offline analysis will increase our energy threshold. This is shown in Figure 7.9, which shows the STACEE γ -ray effective area curve before and after the padding cut has been applied. The raw STACEE effective area curve is what we originally calculated in Section 6.3. This particular effective area curve is for the 2003-2004 Mrk 421 data, but our conclusion is true for any source: the padding cut preferentially removes low energy γ -rays.

This is clearly unfortunate and is contrary to the stated low energy goal of STACEE. There was therefore some resistance to using the Library Padding scheme within the STACEE collaboration. Several different alternative schemes were proposed for accounting for promotion. In particular, there was the Dynamic Thresholds scheme which involved using different thresholds for the ON and OFF data [124], which resulted in a lower final energy threshold. Nevertheless, the Library Padding scheme has become the standard STACEE method for eliminating promotion. I feel that this is principally be-

cause the scheme most closely recreates the effects that actually happen at the hardware levels; hence it is the most intellectually transparent solution. Since the promotion effect is so large compared to our typical γ -ray rates, it is essential that the scheme that corrects for promotion be as trustworthy as possible. Library Padding provides the results in which we can have the most confidence.

Chapter 8

Cosmic-Ray Background Suppression

The previous chapter has described how we correct our measured γ -ray rate to account for the effect of promotion. After having applied the padding cut we can say with confidence that our measurement of the Crab *padded γ -ray rate* of 2.78 ± 0.58 counts min^{-1} means that we have detected this source at an approximately 5σ level. Despite this clear detection, the raw sensitivity of the STACEE experiment is not great. A 5σ detection for a full two years of observation of the Crab is not competitive with the results of most other VHE γ -ray experiments.

It is therefore desirable that we improve the sensitivity of STACEE with the use of offline analysis techniques to suppress our cosmic-ray background. It was the development of powerful background suppression techniques that led to the first credible successes of ground-based γ -ray astronomy. The background suppression techniques developed for STACEE are different from those that made the imaging Cherenkov telescopes successful. Nevertheless, they are inevitably based on the same inherent differences in the Cherenkov light pools of γ -ray and cosmic-ray induced EAS. In Chapter 3 we noted several differences between the γ -ray and cosmic-ray induced events; the most important of these differences are:

- The arrival direction of cosmic-rays is isotropic, whereas the γ -rays come from point sources (at least for the sources of interest to this work).
- The lateral and temporal distribution of the Cherenkov light pool from a γ -ray EAS is, on average, more homogeneous than from a cosmic-ray EAS. Specifically, the Cherenkov wavefront of a low energy γ -ray is approximately spherical.
- The Cherenkov wavefront from a γ -ray EAS is, on average, shorter in duration than from a cosmic-ray EAS.

A great deal of work has been done by STACEE collaborators over the years on developing techniques for suppressing cosmic-rays. The original STACEE detection of the Crab used a background suppression cut that was based on how well a spherical wavefront could be fit to the measured discriminator hit times. This technique resulted in an approximately 25% increase in the sensitivity of the STACEE experiment [110]. Later work concentrated on reconstructing the direction of the incident Cherenkov wavefront. Since the arrival of the cosmic-rays is isotropic, precise direction reconstruction for a point-source of γ -rays would provide for powerful background suppression. Though promising, this work suffered from large differences between real and simulated data [124], which made the technique difficult to use.

Extensive study, much of which is described in this chapter, has shown that a different method has even greater potential for eliminating cosmic-rays. The technique involves realigning our FADC traces and characterizing the sharpness of the resultant peak; the greater smoothness and sphericity of a γ -ray Cherenkov wavefront allows for the calculation of a quantity, ξ , which can be used for background suppression. This technique was originally developed by members of the CELESTE group and is documented in Manseri (2004) [91]. In addition to the overall concept, several smaller investigations in Sections 8.3.2 and 8.3.3 of this chapter were directly or indirectly inspired by Manseri's work. The present study therefore owes a great deal to this CELESTE work, as well as to the original STACEE work on this subject in Kildea (2005), another STACEE collaborator [71, 72].

We shall, for reasons that will soon become clear, refer to this scheme as the Grid Alignment technique. It is this technique that we shall use in this work in order to suppress our cosmic-ray background. As we shall show, the technique works well in the context of STACEE, leading to an overall improvement in sensitivity of approximately 70%. This chapter starts with a detailed description of the Grid Alignment technique and an explanation of why it can be used for background suppression. We then investigate some of the particularities of this technique and discuss ways in which it can be improved. For these sections, we shall use extensively our γ -ray and cosmic-ray simulations of the Crab. Finally, we apply the cut to our 2002-2004 Crab data set, in order to prove that the technique works almost as well with real data as our simulations predict.

8.1 Description of Grid Alignment Technique

The essence of this technique is based on the fact that the Cherenkov wavefront of a γ -ray is approximately spherical, whereas that of a cosmic-ray is not. The goal, therefore, is to define a good way of measuring the 'sphericity' of a given event. Among the




Figure 8.1: FADC traces for a single triggered cluster. The traces demonstrate the challenge of distinguishing between Cherenkov and NSB pulses for events that only barely triggered the experiment.

many challenges to effectively measuring the sphericity, two in particular stand out. The first is the fact that the Cherenkov signal in an individual channel can be negligible or non-existent; the second is that in order to properly define the sphericity of a Cherenkov wavefront we must have an estimate of the core position. In this section we shall describe how the Grid Alignment technique simultaneously accounts for both these problems; in the following sections we shall then show how we can use the resulting information to distinguish between γ -rays and cosmic-rays.

The first challenge is that it is difficult to characterize the Cherenkov wavefront based on the individual FADC traces alone. This is because the Cherenkov pulse for an individual channel is often not significant compared to the underlying NSB fluctuations. This is natural; we wish to run STACEE with the lowest possible energy threshold. This means that while the event as a whole is safely above the level of NSB, the individual FADC traces may not be. This is shown in Figure 8.1, which shows eight real FADC traces of a STACEE cluster. There were discriminator hits on five out of the eight channels in this cluster and the cluster therefore contributed to triggering the experiment. Yet, looking at each trace it is difficult to say whether the individual hits were the result of Cherenkov photons or NSB photons. It is only by considering the event as a whole that we have confidence that it is the result of an EAS. This illustrates the difficulty of characterizing a given EAS event only on the information from individual channels. For a higher energy event the pulse heights would be larger and it is easier to characterize the arrival time and charge of each FADC pulse. But since STACEE is principally interested in low energy events we cannot use techniques that only work for high energy γ -rays.

It is clear therefore, that we are better off treating the FADC information in some sort of overall, aggregate manner. In particular, if we sum together all 64 FADC traces for




Figure 8.2: Sum of all FADC traces. The FADC trace has been inverted, so that the main peak is positive. The trace is for a simulated Crab γ -ray event. The plot also shows the height, H , and width, W , of the summed pulse.

a given event then we shall significantly increase the signal to noise of our Cherenkov pulses. But how should we align the individual FADC traces before we add them together? To start with we simply use the same choice as the hardware MADDOG system made; namely, that the FADC traces should be aligned based on the assumption that the Cherenkov wavefront is a sphere centered at shower maximum, whose core position is at the center of the heliostat field.¹

An example of this type of summed FADC trace is shown in Figure 8.2. This is the summed FADC trace for the same event as shown in Figure 8.1. It is clear that while the Cherenkov pulse was difficult to distinguish in the individual FADC traces, the Cherenkov pulse in the summed trace is clear. With a clear Cherenkov pulse we can now try to characterize the sphericity of a given event. To do so, we measure the *height over width* value, H/W of the summed pulse. The height, H , and width, W , of the summed pulse are shown graphically in Figure 8.2; note that the pulse width is calculated as the Full

¹A clarification is necessary here. The reader will note that shower maximum varies for different γ -rays; a high energy γ -ray is more penetrating and will therefore have a shower max that is deeper into the atmosphere. However, the MADDOG system does not, a priori, know what the energy of a γ -ray will be. MADDOG therefore assumes a fixed value for the distance to shower maximum when calculating the hardware delays. This fixed value is based on the assumed shower max being at a fixed number of radiation lengths into the atmosphere. The height of shower maximum therefore varies depending on the elevation of a source, something that we shall return to later in this chapter, but does not vary for each event. For the rest of this chapter, 'shower maximum' will refer to a point that is a fixed number of radiation lengths into the atmosphere.

Width at Half Maximum (**FWHM**). H/W is a measure of how well aligned our FADC traces are. If the Cherenkov wavefront was not spherical then the summed FADC trace would be shorter and wider and the value of H/W would be smaller since the alignment of the traces would be incorrect.

However, the preceding statement is only true if the EAS core position actually was at the center of the heliostat field, which is generally not the case. As has already been noted, the core positions of triggered γ -rays events can be spread out over a large area; some γ -ray core positions are over 100 m from the center of field. A γ -ray whose core position is in the west of the heliostat field will result in a set of FADC traces where the Cherenkov pulses arrive early for the west camera channels and late for the east camera channels. If we blindly add the FADC traces based on our previous assumption then the summed pulse will be broader and have a lower H/W value than if the same γ -ray had landed at the center of the heliostat field. Accounting for different core positions is therefore the second challenge to characterizing the wavefront sphericity.

In order to correct for this effect we need to redo the alignment of the FADC traces based on different possible core positions. In fact, since we don't know the real core position, we create a whole grid of different possible core positions. We use a 30x30 grid of points, with 15 m spacing, centered on the middle of the heliostat field. For each of these grid points, we recompute the FADC alignment using the point as the assumed core position.² Other than the different core position, everything else about the Cherenkov wavefront is the same; we still assume a spherical Cherenkov wavefront centered at shower maximum. We then sum the FADC traces for each different alignment and calculate the H/W value for that particular point. In principle the alignment of the FADC traces will be best for the assumed core position which is closest to real core position; this is the point that will have the maximum value of H/W .

This method is graphically demonstrated in Figure 8.3. The plot shows the H/W value versus the different assumed core positions. The plot is for a simulated γ -ray event. There is a clear peak in the H/W distribution near to the origin (the center of the heliostat field). The peak is where the alignment of the FADC traces is best; it therefore appears that the core position for this event is near to the center of the heliostat field. It should be emphasized that this procedure is repeated *for each event*. For each event a distribution like that shown in Figure 8.3 is generated; this in turn means that for each event there are 900 different summed traces, based on the different assumed core positions, for which an H/W value must be calculated. The Grid Alignment technique is therefore computation-

²This explains the name of our method. We need a *Grid* of possible core positions upon which to do different *Alignments* of FADC traces.




Figure 8.3: Distributions of H/W as a function of assumed core position. Each point in the distribution results from a different alignment of the FADC traces and hence a different H/W value. In the upper left is the summed FADC trace for an assumed core position near the peak of the distribution; in the upper right is a summed FADC trace for an assumed core position 200 m away from the peak. Note that the pulse in upper-left trace is higher and narrower than the pulse upper-right trace. Plot is for a *single* simulated γ -ray event.

ally intensive, though not prohibitively so.

8.2 Using the Grid Alignment Technique for Background Suppression

The preceding section has introduced the Grid Alignment technique. It is now time to explain how we shall use this information to suppress our cosmic-ray background.

The quantity that we use to suppress cosmic-rays is called ξ . To calculate ξ , we follow the following procedure:

1. Locate the assumed core position that maximizes H/W . We call the position where this occurs \mathbf{r}_{grid} and the value at that point H/W_{max} . The value \mathbf{r}_{grid} is a good estimate of the true core position, which we shall call \mathbf{r}_{true} . It is clear that for the

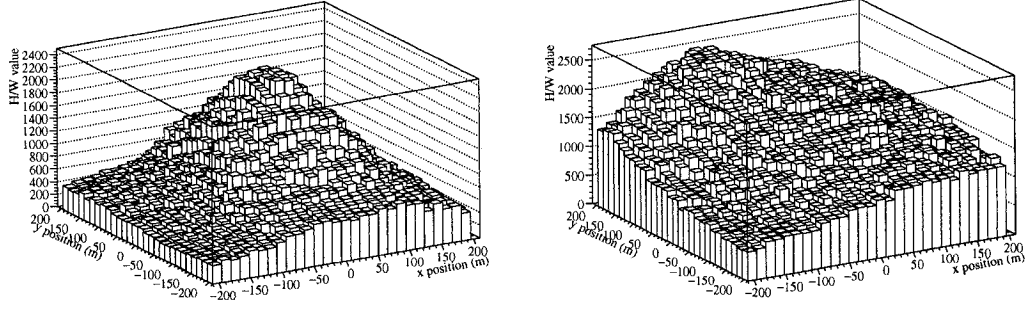


Figure 8.4: Distributions of H/W as a function of assumed core position. Left side: H/W distribution for a simulated γ -ray EAS. Right side: H/W distribution for a simulated proton EAS.

event shown in Figure 8.3 \mathbf{r}_{grid} is very near to the origin.

2. Calculate the value of H/W at four assumed core positions that are 200 m away from \mathbf{r}_{grid} . The four points are referred to as the *off-alignment points*. The average of the H/W value at four off-alignment points is called H/W_{200m} .
3. Calculate the ratio of H/W_{200m} to H/W_{max} , a quantity that we call ξ . Specifically,

$$\xi = \frac{H/W_{200m}}{H/W_{max}}. \quad (8.1)$$

Why is ξ a good discriminator between γ -rays and cosmic-rays? The principal reason is as follows: a γ -ray Cherenkov wavefront is well described as being a sphere centered at shower maximum. Aligning the FADC traces based on this assumption will therefore result in an almost optimal (ie maximized) value of H/W . Since the FADC traces are well aligned at \mathbf{r}_{grid} , moving to a point 200 m away will therefore result in a significant decrease in the calculated H/W and hence a very small value of ξ .

A cosmic-ray Cherenkov wavefront, by contrast, is only poorly described as being spherical. Consequently, even at \mathbf{r}_{grid} the alignment of the FADC traces will not be optimal. There is therefore less decrease in H/W when it is recalculated at a point 200 m away from \mathbf{r}_{grid} . Values of ξ will therefore be, on average, significantly larger for cosmic-rays than for γ -rays. This fact is demonstrated in Figure 8.4, which shows simulated H/W distributions for a γ -ray and a proton. The γ -ray distribution has a clear peak and H/W falls off quickly away from that peak. The proton distribution has only a weak peak, which falls off slowly. In this case, the value of ξ is 0.27 for the γ -ray and 0.61 for the proton.

It should be noted that the discriminating power of ξ is not only based on the greater sphericity of a γ -ray Cherenkov wavefront; it is also based on the fact that a γ -ray wavefront has a shorter duration. It can be shown that the quantity ξ depends on the intrinsic duration of the Cherenkov wavefront [84]. γ -ray wavefronts are, on average, shorter in duration than cosmic-ray wavefronts; this is therefore an additional reason why the values of ξ are lower for γ -rays. Both the differences in sphericity and duration of the Cherenkov wavefront contribute to the discriminating power of ξ , though of the two the dependence on sphericity is probably more important.

That then, is the basic premise of this technique; we shall look for events with small values of ξ in order to preferentially select γ -rays over cosmic-rays. The devil, as always, is in the details. In the following sections we will use simulated data to show the various particularities of ξ , as well as testing how well the technique works for real data. We will start, however, by presenting the simplest predictions that simulations make about using a cut based on ξ . We need to define a figure of merit with which to evaluate the effectiveness of this technique. We shall use the *quality factor*, Q , to define the effectiveness of the ξ cut at suppressing cosmic-rays and retaining γ -rays. The quality factor is defined as

$$Q = \frac{N'_\gamma/N_\gamma}{\sqrt{N'_h/N_h}}, \quad (8.2)$$

where N_γ and N'_γ are the numbers of γ -rays in the data set before and after the cut has been applied and where N_h and N'_h are the numbers of cosmic-rays before and after the cut has been applied. A high Q -factor therefore implies that a cut is effective at removing cosmic-rays and retaining γ -rays; a Q -factor less than 1.0 means the cut makes our sensitivity worse than it was before.³

Throughout this chapter we shall be using simulations of γ -rays and cosmic-rays coming from different points along the path of the Crab through the sky. We shall be using the γ -ray and cosmic-ray simulations with a continuous spectrum of energies, as described in Section 6.2.1. We do this because at the end of the chapter we shall be comparing our simulated predictions about ξ with the results of the real Crab data set; it is desirable that the two match as closely as possible.

We start by presenting the results at the Crab transit point; the distributions of ξ for γ -rays and protons are shown on the left-side of Figure 8.5. The plot shows that there is a clear separation between the γ -rays and protons. The best Q -factor comes from a cut at

³It should be noted that the denominator in Equation 8.2 should really be of the form $\sqrt{(N'_h + N'_\gamma)/(N_h + N_\gamma)}$. But since the number of cosmic-rays always greatly exceeds that of γ -rays the approximation of Equation 8.2 is appropriate.




Figure 8.5: Basic ξ results. The left-side plot shows the values of ξ for both γ -rays and protons, for simulations done at the Crab transit point. The two curves have been normalized to have the same maximum height. The right-side plot shows the Q -factor that results from a cut at a given value of ξ .

$\xi < 0.325$, which gives a Q -factor of 2.6, but only retains 40% of the original γ -rays. A slightly looser cut at $\xi < 0.35$ has a Q -factor of 2.3 but retains 60% of the original γ -rays. We shall be using the value of 0.35 throughout this chapter as our standard cut value.

This simple test demonstrates a key problem with using a cut based on ξ ; any cut with a reasonable Q -factor requires that we cut out a significant fraction of γ -rays. If all the γ -rays that we observe were exactly the same then this would not matter; any cut that improved our signal-to-noise ratio would be desirable, even if the cut only retained 1% of the original γ -rays. However, the γ -rays we observe have a spectrum of different energies and, as we shall see, the quantity ξ has a dependence on the energy of the incident γ -ray. This means that a cut based on ξ will preferentially retain or reject γ -rays based on their energy. This is unfortunate since we do not, a priori, know exactly what the energy spectrum of γ -rays for a given source might be; we do not, therefore, know what Q -factor we can expect. It is for this reason that we use 0.35 as our standard ξ cut value, even though a slightly tighter cut would give an even better Q -factor. For this reason, most of the modifications that we shall be describing in the remainder of this chapter concern efforts to maximize the fraction of γ -rays that are retained, while still maintaining a high Q -factor.

8.3 Biases of the Technique

In a perfect world, the quantity ξ would not depend on any other quantities such as the initial γ -ray energy, core position or direction of incidence. We could then apply a cut on




Figure 8.6: Dependence of ξ on the incident γ -ray energy for Crab transit simulations. The left-side shows results as a profile; that is to say, it shows the average value of ξ as a function of energy. The vertical error bars are errors on the mean. The right-side shows results as a 2-dimensional histogram.

ξ based on the results in the previous section to any data set and expect the same response. Unfortunately this is not the case. As we shall show, the quantity ξ does depend on many of these different quantities. It is important to examine each of these dependencies in some detail. In some cases, the examination suggests ways in which the raw ξ value can be modified in order to minimize a dependence. Even in the case where a correction is not possible, it is still important to understand the dependence. The following sections detail the most important dependences of the quantity ξ .

8.3.1 Dependence of ξ on Energy

Probably the most worrisome aspect about the Grid Alignment technique is that ξ has a noticeable energy dependence. This is clearly seen in Figure 8.6, where we plot ξ vs the actual energy for Crab transit γ -ray simulations. It is clear that ξ increases with increasing energy.

It is not entirely clear why ξ depends on the incident γ -ray energy. The explanation may be in part related to the shape of the Cherenkov wavefront for high energy γ -rays. As noted in Chapter 3, for higher energy γ -rays the Cherenkov wavefront is better described as a cone. Our assumption of a spherical shape to the wavefront would therefore be incorrect and would lead to a non-optimal alignment of the FADC traces for high energy events. This, in turn, would result in a larger value of ξ [84]. The energy dependence may also be related to higher energy γ -rays having a longer duration; as noted before, a longer intrinsic wavefront duration results in a larger value of ξ .

Nor, unfortunately, is there any clear way to correct for this dependence. We shall,




Figure 8.7: Position of the four off-alignment points. The origin of this coordinate system is \mathbf{r}_{grid} , the position at which H/W is maximized. As usual, however, north is in the $+z$ direction and east is in the $+x$ direction. The position of the shower maximum is denoted by the blue sphere; the distance to shower maximum is L . The four off-alignment points are marked by the four smaller red spheres; each is a distance of D from \mathbf{r}_{grid} .

in the following chapter, present a method for reconstructing the incident γ -ray energy. But experience has shown that there is no way of using this information to correct for the energy dependence without seriously degrading the Q -factor of the ξ cut. Our inability to correct for this dependence means that we must always be conscious of the underlying γ -ray energy spectrum. For instance, the γ -ray simulations used in this chapter assume that the spectral index of the Crab is 2.4. The increase in sensitivity from using the ξ cut will therefore be different if we are looking at a source with a spectral index of 3.4.

8.3.2 Dependence of ξ on Source Position

Another important concern is how the quantity ξ depends on the position of the γ -ray source on the sky. There are a number of obvious aspects of the technique that need to be considered, particularly with regards to how we choose the four off-alignment points. Our initial decision that these points are simply 200 m horizontally from \mathbf{r}_{grid} is too simplistic and should be modified.

Figure 8.7 demonstrates how we choose the location of the four off-alignment points.

There are two important elements to note about our choice. Both of these elements are related to an important principle about the off-alignment points; namely that the points must represent *an equivalent change in the relative delays as the source position moves across the sky*.⁴ The calculation of H/W at the off-alignment points is meant to provide a measure of how quickly the alignment of the FADC traces moves away from optimal alignment. Another way of thinking about the off-alignment points is that they provide a normalization of the H/W_{max} value. However, this normalization requires that the points be placed such that the change in the relative delays at the off-alignments point be the same as the source position moves across the sky. This, as we shall explain, is not as simple as requiring that the distance to the off-alignment points be always 200 m from \mathbf{r}_{grid} . There are two elements of Figure 8.7 that show why this is not the case.

The first element to note in Figure 8.7 is that the four off-alignment points are not located on a horizontal plane; instead they are located on a plane that is perpendicular to the source direction. Placing the points on a horizontal plane (as we initially did) is a flawed procedure, since only differences in position that are in a direction that is perpendicular to the EAS axis will contribute to changes in the relative FADC alignment. This problem is solved by simply placing the four off-alignment points in the plane perpendicular to the EAS direction. It should be clear that the four off-alignment points will be located in the horizontal plane only if the source is exactly at the zenith position.

The second element to note in Figure 8.7 is that the distance from the off-alignment points to \mathbf{r}_{grid} , D , must change as the source moves across the sky. The realignment of the FADC traces is based on the assumption of a spherical wavefront centered at shower maximum. At the Crab transit point that means that shower maximum is 11.3 km away and hence the spherical wavefront is assumed to have a radius of 11.3 km. But as the source moves further from the transit point, the distance to shower maximum, L , increases. For γ -rays coming from a point two hours past the Crab transit, L has increased to 13.6 km. This is expected, as γ -rays incident at lower elevations will have a shower maximum that is *higher* in the atmosphere; hence the distance from shower maximum to the heliostat field will be greater. For data taken two hours past transit, we therefore assume that the spherical wavefront has a radius of 13.6 km.

This, however, has an effect on ξ . If the assumed radius is 13.6 km, then the change in the relative delays at a point 200 m from \mathbf{r}_{grid} will be less than if the assumed radius is 11.3 km. Since the change in the relative delays is smaller, ξ will be, on average, larger. To correct for this, we must use a value of D that changes with the source position; in

⁴By ‘relative delays’, we mean the difference between the delays for different channels. So if we changed the delays for all the channels by 1 ns, then there would be no change in relative delays; but if we only change the delays for half the channels by 1 ns then there would be a change in the relative delays.




Figure 8.8: Dependence of ξ on source position after the corrections described in the text. The plot shows distributions of ξ for γ -ray simulations at Crab transit and +1, +2, +3 hours after transit.

particular, we must have

$$D \propto L. \quad (8.3)$$

D will only be equal to 200 m when the source is at the Crab transit point; as the source moves away from that point, D will increase.

These corrections eliminate the majority of the dependence of ξ on elevation. Nevertheless, some dependence still remains. This is shown in Figure 8.8, which shows the ξ distributions for γ -ray simulations at various pointings. The plot shows that there remains some increase in ξ as we get further away from transit. The cause of this residual dependence is not clear. Nevertheless, Figure 8.8 demonstrates an additional reason why, as noted in Section 5.2.2, we do not use data that is more than 2.5 hours past transit. The figure makes it clear that the distribution for the +3HR pointing is shifted significantly to the right of the other distributions. Even after eliminating data taken more than 2.5 hours from transit, it is clear that to properly understand a cut on ξ we must combine simulated data from several different pointings, in much the same manner as we did when calculating the HA-weighted effective area curves.

These HA-weighted ξ distributions are shown in Figure 8.9. This plot shows the same information as Figure 8.5, except that the γ -ray distribution is now a combination of the distributions from several different pointings; the weighting for the different pointings is given by the HA-distributions for the 2002-2004 Crab data set. The same is true for the cosmic-ray distribution, except that the distribution also combines the results for both proton and helium nuclei together. The proton and helium nuclei distributions are normalized

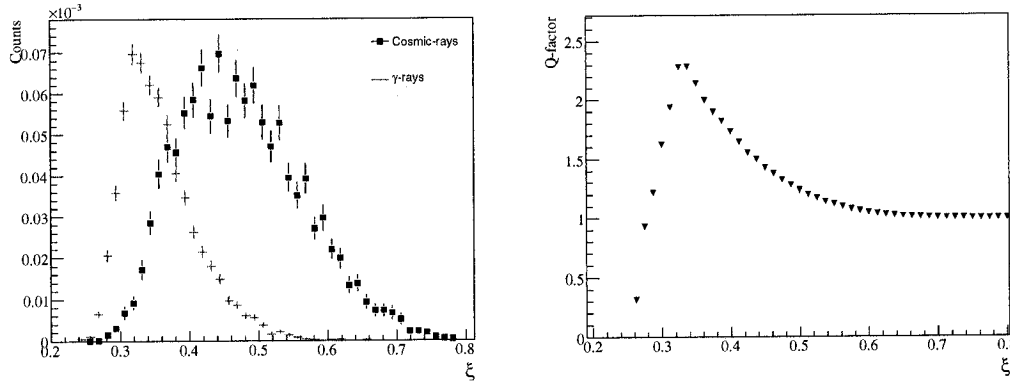


Figure 8.9: HA-weighted ξ results. The left-side plot shows the distributions of ξ for both γ -rays and cosmic-rays. The right-side plot shows the Q -factor that would result from a cut at a given value of ξ .

based on their assumed SEDs.

There are a couple of important differences to note between the transit results in Figure 8.5 and the HA-weighted results in Figure 8.9. First, the Q -factor for a cut of $\xi < 0.35$ has decreased slightly from about 2.6 to 2.2. Second, the value of ξ that results in the maximum Q -factor has shifted to the right slightly. Nevertheless, the two distributions are still quite similar, suggesting that the conclusions we drew from studying transit simulations alone are still useful.

This is also an appropriate place to note an interesting dependence of ξ on whether a γ -ray emitter is a point source. Figure 8.10 shows a plot of the distributions of ξ for two sets of simulations of 200 GeV γ -rays. In the first set of simulations, the γ -rays come exactly from the transit point. That is to say, we assume that the Crab is a point-source emitter of γ -rays; this is our standard procedure for our γ -ray simulations. For the second set of simulations, however, we scatter the γ -rays in solid angle out to a maximum angle of 1.0° from the transit point; the simulation therefore reproduces the expected results from an extended γ -ray source. The distribution of ξ from the extended source is noticeably shifted to the right of the distribution from the point source. It is clear, therefore, that if an otherwise identical γ -ray came from 0.5° away from the nominal source direction then it will have a higher ξ value. This has two important implications. First, it shows that a cut on ξ will also be, to some extent, a cut on the direction of the Cherenkov wavefront; at least part of the reason for the effectiveness of the ξ cut is therefore because the γ -rays come from a point source whereas the cosmic-ray background is isotropic. The second point follows from the first: namely that this cut will be much less effective if it is applied to an extended source of γ -rays. This is not a concern for the present work, where





Figure 8.10: Dependence of ξ on wavefront direction. Both of these curves are for simulations of 200 GeV γ -rays. The first set of simulations assumes the γ -ray emitter is a point source; the second set assumes that the γ -ray emitter is an extended source with angular radius of 1.0° .

all the γ -ray sources are known to be point sources; but it may be relevant for future work.

8.3.3 Dependence of ξ on South Camera Fraction

The final dependence of ξ , the dependence on the *South camera fraction*, is also the most complicated. Understanding this dependence requires us to delve into the particularities of the STACEE detector. In addition, the method of correcting for this dependence is a purely empirical correction, which may seem *ad hoc*. Nevertheless, this dependence is sufficiently important that it is worth the effort to understand and correct.

The dependence is as follows: *the value of ξ is larger for events where the majority of the Cherenkov photons landed on the South camera heliostats.* To understand this dependence, we must understand more about the particular geometry of the South camera heliostats. The South camera heliostats are marked in Figure 4.12 as being the heliostats in clusters 6 and 7. The crucial point is that the heliostats in these two clusters are much more closely spaced than the heliostats of the other six clusters. Why does this matter? It matters because the relative changes in FADC delays between two heliostats depend on how far apart the heliostats are. If the heliostats are close together then changing the assumed core position will result in a smaller change in the relative delays than if the heliostats are far apart. Since the South camera heliostats are so densely spaced this

Figure 8.11: Distribution of S_f for Crab transit γ -ray simulation.

causes a problem. The Grid Alignment technique assumes that the relative change in the delays introduced by moving away from \mathbf{r}_{grid} are the same for all clusters. This is not the case for an event where the majority of the Cherenkov photons are deposited in the South camera.

That is the qualitative explanation of why ξ is larger for events where the majority of the Cherenkov photons ended up in the South camera. It is now necessary to describe this dependence in a more quantitative manner. To characterize the importance of the South camera, we shall use our measurements of the dead-reckoned charge measured using the FADC data, q_i ; the calculation of this quantity was described in Section 5.2.1. Using these measurements, we can calculate the fraction of the total charge for a given event that is seen in the South camera channels. Specifically, we calculate the South camera fraction, S_f , defined as

$$S_f = \frac{\sum^S q_i}{\sum^{all} q_i} \quad (8.4)$$

where the sum in the numerator is over the South camera channels and the sum in the denominator is over all the channels. A large value of S_f indicates an event where the South camera dominates.

Let's start by investigating S_f itself, before trying to understand the dependence of ξ on S_f . Figure 8.11 shows the distribution of S_f for the Crab γ -ray simulation. The mean of the distribution is 0.29, which suggests that the South camera collects, on average, slightly more than a quarter of the Cherenkov photoelectrons for a given event. This is expected, since studies have shown that the South camera tends to slightly dominate the other cameras.




Figure 8.12: Dependence of ξ on S_f for Crab transit γ -ray simulation. The left-side plot shows a profile of ξ dependence on S_f . The line is a fit to the range $S_f > 0.3$. The right-side plot shows same data, presented as 2-dimensional histogram.

The dependence of ξ on S_f is shown in Figure 8.12. It is clearly a somewhat complicated dependence. For the range $0.2 < S_f < 0.3$, where most of the events are, the dependence is flat. There are very few events in the range $S_f < 0.2$, so we shall not attempt to correct them. The important area is the range $S_f > 0.3$ where the dependence is strong. The dependence for $S_f > 0.3$ is consistent with the qualitative description given above; ξ tends to be larger for events where S_f is larger. It is the dependence of ξ on S_f in this range that we wish to empirically correct; the correction will therefore affect about 40% of the γ -rays.

Figure 8.12 clearly demonstrates the dependence of ξ on the South camera fraction. What is less clear is how to correct for this dependence. Our understanding of the problem suggests that the Grid Alignment technique would work better if our heliostats were all equally spaced; this is not, however, a practical suggestion. Nor does there appear to be any way of altering the Grid Alignment method to naturally account for the importance of the South camera. It appears that the only feasible option is to use the data shown in Figure 8.12 to make an empirical correction to our ξ value. To derive the correction, a line has been fit to the portion of the ξ vs S_f profile where $S_f > 0.3$.⁵ The slopes of the line is 0.5. Additional tests have shown that this slope is similar for the γ -ray simulations at other detector pointings. Rewritten as pseudocode, our corrected quantity, called ξ_{cor} , is therefore defined as follows:

```
if ( $S_f > 0.3$ ) then
```

⁵We do not want to correct the section $S_f < 0.3$, where there is no strong dependence on S_f .




Figure 8.13: The upper plot shows the distributions of ξ_{cor} for HA-weighted γ -ray and cosmic-ray simulations. The lower-left and lower-right plots compare the Q -factor and γ -ray retention fraction for different values of the ξ and ξ_{cor} cuts.

$$\xi_{cor} = \xi - 0.5(S_f - 0.3)$$

else

$$\xi_{cor} = \xi$$

How does this correction change our expected results? The upper plot of Figure 8.13 shows the γ -ray and cosmic-ray distributions of ξ_{cor} . These distributions are similar to the distributions in Figure 8.9, except that the distributions are slightly tighter and shifted to the left. The important information, however, is shown in the other plots of Figure 8.13, where we compare the Q -factor and γ -ray retention fraction for the standard and S_f -corrected ξ . As the lower-left plot shows, there is a small improvement in the Q -factor from using the corrected ξ value. The real benefit, however, is seen when comparing the fraction of γ -rays that are retained at different cut values, which is shown on the lower-right plot. As expected, we retain more γ -rays at a given cut value. These plots suggest that we can use our normal cut at $\xi_{cor} < 0.35$ for an improved Q -factor of 2.3, while




Figure 8.14: Final effective area for 2003-2004 Mrk 421 data set. The plot shows the effective area for four different cases. They are the raw effective area (no event cuts); the effective area after the padding cut alone; the effective area after the padding cut and the S_f -corrected ξ cut; and the effective area after the padding cut and the original ξ cut. Both the original and S_f -corrected ξ cuts were made at 0.35.

retaining a greater fraction of the original γ -rays (62%). *Based on these results, we define $\xi_{cor} < 0.35$ as the standard cosmic-ray background suppression cut for the remainder of this work.*

One final aspect of the S_f correction should be mentioned. Not only does the S_f correction increase the fraction of γ -rays that pass our cut; this benefit occurs predominately for the higher energy γ -rays. This means that the S_f correction corrects, at least partially, the energy dependence that was described in Section 8.3.1. This fact is best shown by observing the final effective area after all our event cuts have been made. The final effective area is shown in Figure 8.14; this plot also shows the raw effective area, as well as the effective area after the padding cut alone. While it is clear that both of the cuts based on the Grid Alignment technique decrease significantly the effective area, the S_f -corrected scheme is clearly better, especially at higher energies. Like the earlier effective areas, this figure shows the results for the 2003-2004 Mrk 421 data set.




Figure 8.15: Cosmic-ray S_f distribution. The distributions are for 2002-2004 Crab data set. The simulated distribution is HA-weighted and normalized to have equal area.

8.4 Application of ξ_{cor} to Crab Data

The previous sections have introduced the Grid Alignment technique and the cosmic-ray background suppression quantity ξ_{cor} . Our simulations suggest that the application of this quantity will lead to a significant increase in the sensitivity of the STACEE experiment (though at some cost in terms of a reduced effective area for γ -rays). The crucial test, however, is how well a cut on ξ_{cor} works for real data. As noted, earlier STACEE work on background suppression has often been stymied by large differences between simulations and real data. The differences made interpretation of the results difficult; lack of faith in the background suppression techniques meant they were not applied to real data [124].

We do not, however, wish to test our Grid Alignment technique on our AGN data sets. Any such test would suggest that we were optimizing our cuts and would compromise our AGN measurements. We will therefore, once again, use our Crab 2002-2004 data set as a testbed. As a well understood and constant γ -ray source, the Crab is perfect for testing our background suppression capability. As will be seen, application of the Grid Alignment technique to the real Crab data set shows that the cut does indeed work and behaves, for the most part, as expected.

8.4.1 Crab Cosmic-ray Comparison

Using the Crab data set actually admits two separate tests of the Grid Alignment technique. The first involves comparisons between the simulated and real distributions for cosmic-rays; the second between simulated and real distributions for γ -rays. We shall




Figure 8.16: Cosmic-ray ξ_{cor} distribution. The distributions are for 2002-2004 Crab data set. The simulated distribution is HA-weighted and was normalized to have equal area.

start by observing the distributions for cosmic-rays. It should be emphasized, once again, that our cut on ξ_{cor} is in addition to the padding cut; all the distributions shown in this section are for events that have already passed the padding cut.

Since the ξ_{cor} depends on S_f it is sensible to first test that the distribution of S_f is similar for real and simulated cosmic-rays. Figure 8.15 shows the distributions of S_f for the OFF-source Crab data and for the cosmic-ray simulations. The cosmic-ray simulations are HA-weighted averages of the proton and helium nuclei simulations at different pointings, in the same manner as previously described. The figure shows that there is a fairly good match between the distributions of S_f for real and simulated data. The match is not perfect: the mean of the real data distribution is 0.33, whereas the mean of the simulated distribution is 0.32. This indicates that the South camera is slightly more dominant in real data than it is in simulations. This difference is, however, small and it is fair to say that the match between real and simulated data is good enough that we can proceed to comparing the cosmic-ray distributions of ξ_{cor} .

The comparison is shown in Figure 8.16, which shows the distributions of ξ_{cor} for the real Crab OFF data and the cosmic-ray simulations. As before, the simulations are HA-weighted averages of individual proton and helium nuclei simulations. The comparison shows that we clearly have some problems. The real data distribution has a larger spread than the simulations. The real data has a large tail of events with $\xi_{cor} > 0.7$, which we do not see in simulations. In addition, the real data distribution is shifted to the left of the simulations for $\xi_{cor} < 0.35$. Of these two problems, the disagreement at small values of ξ_{cor} is a greater concern, even though the difference is less obvious. This is because our

expected cut value is at $\xi_{cor} = 0.35$; this comparison therefore shows that this cut value will let pass significantly more cosmic-ray events than the simulations had predicted.

Significant work has been done on understanding these disagreements between the distributions of ξ_{cor} for real and simulated cosmic-ray data. There are two separate aspects of the simulations that are of concern. On one hand, there may be elements of our modelling of the STACEE detector that are causing this discrepancy. For instance, studies have shown that the quantity ξ_{cor} is sensitive to the precise form of the single photoelectron pulse shape used in the simulation [84]. It is possible that small differences between the real and simulated single photoelectron pulse shape may be part of the problem. This type of problem is unfortunate, since it would apply to our γ -ray simulations as much as to our cosmic-ray simulations.

On the other hand, it is possible that the problems arise from the simulations of the cosmic-ray EAS themselves, rather than from problems with the simulation of the STACEE detector. As noted earlier, the simulation of hadronic interactions above 100 GeV is problematic, since the simulations are often poorly constrained by experimental data (see Section 6.1.1 for details). While we believe that our cosmic-ray simulations are generally satisfactory, there may be smaller differences between reality and simulation that contribute to the differences in the distributions of ξ_{cor} . In addition, as Figure 6.7 makes clear, the spectral energy distribution of the cosmic-rays is only approximately known. Mistakes in the assumed form of the cosmic-ray energy spectrum will lead to changes in the distribution of ξ_{cor} , since, as noted, ξ_{cor} depends on the incident particle energy. If the disagreements between the distributions shown in Figure 8.16 result from deficiencies in our simulation of cosmic-ray EAS, then the disagreements are of less concern. This is because the disagreements in these cosmic-ray comparisons would not necessarily imply that there are any problems with our γ -ray simulations. The simulation of γ -ray EAS is significantly simpler⁶ and therefore less likely to be flawed. As long as this is the case, our simulations can still be used to accurately predict the γ -ray effective areas and hence to properly convert our measured γ -ray rates into source flux measurements. The test of whether this is the case will be whether the ξ_{cor} distributions for real and simulated γ -rays, shown in the following section, are a good match.

8.4.2 Crab γ -ray Comparison

We now turn to the most important test of the Grid Alignment technique; its application to Crab γ -rays and its ability to increase the detector sensitivity. The left-side plot of Figure

⁶Since interactions in a γ -ray induced EAS are purely electromagnetic and therefore better understood.




Figure 8.17: Real data γ -ray ξ_{cor} distributions. Distributions are for 2002-2004 Crab data set. The left-side plot shows the distributions of ξ_{cor} for the Crab ON and OFF data; the distributions are corrected for the small live-time differences between the ON and OFF data. The right-side plot shows the difference between the ON and OFF distributions. The distribution is overlaid by HA-weighted Crab γ -rays simulations.

8.17 shows the ON and OFF distributions of ξ_{cor} for the 2002-2004 Crab data set; the distributions have been corrected for the small difference between the ON and OFF live-times. The right-side plot shows the difference between the ON and OFF distributions. The difference distribution is, in principle, the distribution resulting from the Crab γ -rays. This interpretation is confirmed by the fact that the difference distribution is shifted to the left of the original ON and OFF distributions. Indeed, as the right-side plot shows, the difference distribution is well matched by the distribution of ξ_{cor} for the Crab HA-weighted γ -ray simulations. The ξ_{cor} excess in the real data occurs exactly where it should.⁷

These results are encouraging evidence that the Grid Alignment technique is working properly. Remember that the prediction from our simulations was that a cut where we only accept events where $\xi_{cor} < 0.35$ would lead to a significant increase in our sensitivity. It is clear from the two plots of Figure 8.17 that this is indeed the case. We are therefore now in a position to calculate the *final γ -ray rate* for the Crab Nebula. The final γ -ray rate is defined as the difference between the ON and OFF rates after both the padding cut and a cut on ξ_{cor} are applied. Figure 8.18 shows the final γ -ray rate as a function of the exact cut value of ξ_{cor} . *If we retain only events where $\xi_{cor} < 0.35$, we find that our final Crab γ -ray rate is $1.62 \pm 0.20 \text{ min}^{-1}$, which represents a 8.1σ detection.*

⁷There may be some differences between the real and simulated γ -ray distributions, but the large error bars on the real ON-OFF difference distribution makes this impossible to test with any certainty.




Figure 8.18: Final γ -ray rate after ξ_{cor} cut. The plot shows the final γ -ray rate as a function of the cut value of ξ_{cor} . The plot also shows the predicted final γ -ray rate.

There are several things to note about this measurement. First, remember that the padded Crab γ -ray rate is $2.78 \pm 0.58 \text{ min}^{-1}$, a detection of 4.8σ . The measured Q -factor for our cut is therefore 1.7 ± 0.4 . This is less than, but almost consistent with, our predicted Q -factor of 2.3. Also, our final γ -ray rate is $58 \pm 14\%$ of the padded rate, which is consistent with the predicted fraction of 62%. It should also be noted that there is an excellent agreement between the real and simulated final γ -ray rates shown in Figure 8.18. In fairness, the predicted final rate is not a full prediction; the normalization of the predicted curve is simply set so that the two curves agree at $\xi_{cor} = 1.0$. The figure therefore only shows that the shapes of the two curves agree; nevertheless, it is encouraging to note that the agreement is excellent. Notice also that there appears to be better agreement between the real and simulated γ -ray distributions of ξ_{cor} than there was for the cosmic-rays; this suggests that the differences seen in Figure 8.16 were principally the result of deficiencies in the simulation of EAS rather than in the simulation of the STACEE detector.

The final test of the ξ_{cor} cut is shown in Figure 8.19. This is a plot of the final significance of the detection as a function of the amount of live-time accumulated. The curve is well fit by a function with the form $f(t) \propto \sqrt{t}$, which is what we would expect for a constant γ -ray source like the Crab. If our data set was being thrown off by, for instance,




Figure 8.19: 2002-2004 Crab significance vs observation time after ξ_{cor} cut.

a single night of bad data, then the curve would show a noticeable discontinuity. There is no discontinuity, thereby proving that our data set is clean and our analysis techniques are well-behaved. This plot therefore shows that in addition to being *effective*, the ξ_{cor} cut is also *stable*. The plot also shows that after the ξ_{cor} cut our Crab sensitivity is $1.6 \sigma / \sqrt{\text{hour}}$.

The test of ξ_{cor} using our Crab data has been largely successful. After a cut on ξ_{cor} we have measured a final γ -ray rate of $1.62 \pm 0.20 \text{ min}^{-1}$, which represents a significantly stronger detection than the padded rate. The evidence suggests that the cut is operating properly. The ξ_{cor} and padding cuts together compose our standard *event cuts* that will be applied to all data. We shall be applying these cuts to the Crab data before the further energy reconstruction analysis presented in the following chapter. The event cuts will also be applied to the AGN data sets when we come to our final analysis in the subsequent chapter.

Chapter 9

Energy Reconstruction

With the preceding descriptions of the Library Padding and Grid Alignment techniques we are now in a position to make accurate and precise measurements of the rate of γ -rays from a given source. We can determine whether or not a given source is a γ -ray emitter; if it is, we can investigate whether the flux is constant. However, these questions only deal with the integral flux results; that is to say, we don't differentiate between a 200 GeV γ -ray and a 500 GeV γ -ray. A more interesting, though also more challenging, problem is to measure the energy of each γ -ray and to estimate the overall distribution of energies of the γ -ray source. The distribution of γ -ray flux as a function of energies is referred to as the Spectral Energy Distribution (SED). This chapter will explain the process of reconstructing the SED of a γ -ray source using STACEE data.

Chapter 2 explained the reasons why a measurement of the SED of a γ -ray source is of interest. One point should be emphasized: most sources of VHE γ -rays have flux distributions that are well described as being power laws of the form

$$dN/dE \propto E^{-\alpha}, \quad (9.1)$$

where α is known as the spectral index. This, as noted in Section 6.2.1, is the assumed spectral form for our simulations. There are some known deviations from this rule; the SED of Mrk 421, for instance, appears to have an exponential cut-off above ~ 4 TeV [1,77]. However, this exponential cutoff has a minimal effect on the spectrum of Mrk 421 in the energy range that is of interest to STACEE. We will therefore make the reasonable assumption that the SED of the Crab and Mrk 421 are power laws in the STACEE energy range. The problem of spectral reconstruction is therefore reduced to a problem of determining the spectral index. Once we have determined the spectral index it is simple to measure the overall flux normalization using our effective area curves.

It should be noted that this work builds on work by other STACEE collaborators,

particularly the work of Scalzo (2004) [124]. In addition, work on the spectral reconstruction of Mrk 421 using STACEE data has already been presented in Boone (2002) and Carson (2005) [22, 29]. The work in Carson (2005) investigates some of the same Mrk 421 data set which is the subject of this work. However, the methodology presented in Carson's work is different from that of this work; each work represents the development of a parallel technique. Having parallel analysis techniques is worthwhile since energy reconstruction for VHE γ -ray experiments is complicated and prone to systematic errors; the development of an alternative measurement provides confidence in our final results.

We shall start the chapter by explaining the challenges of energy reconstruction for STACEE. In particular, we shall describe the degeneracy between core position and γ -ray energy. The method for breaking the degeneracy required a modification of the heliostat canting scheme. Next, we shall describe the Template Fitting method, which is our algorithm for determining the energy of each γ -ray, as well as the core position. We shall then show how this information can be used to extract the spectral index for a source; the problem of accounting for our detector efficiency is non-trivial. Finally, we shall apply our spectral reconstruction to the Crab Nebula and compare our results with expectations based on results from other experiments.

9.1 Introduction to STACEE Energy Reconstruction

At first sight, the question of reconstructing the energy of an incident γ -ray might seem straight-forward. As shown back in Figure 3.4, the Cherenkov photon density of a γ -ray EAS appears to have a simple linear relation with the incident energy. One might imagine, therefore, a simple energy reconstruction scheme: a) calculate the total measured charge for each event by, for instance, summing up the dead-reckoned charge in each channel, b) use a simple linear relationship to convert the total charge into a γ -ray energy.

Unfortunately, this scheme is flawed. Figure 3.4 obscures an important point, which is that the Cherenkov photon density varies with the core position. The consequence is that a Cherenkov telescope will detect less Cherenkov photons from a 200 GeV γ -ray EAS that lands at a larger distance from the center of the detector. This problem is particularly acute for STACEE, since our restricted field of view makes us more sensitive to the EAS core position.

This phenomenon is demonstrated in Figure 9.1. The figure shows the total number of photoelectrons produced in the STACEE PMTs for a given γ -ray EAS as a function of the γ -ray energy; the different curves show how the relationship changes for different core positions. The plot was generated using γ -ray simulations at the Crab transit point.




Figure 9.1: Number of photoelectrons vs γ -ray energy. The data is presented as a profile and is split into three different sets based on the core position for a particular event. The data is from Crab γ -ray simulations done at transit point. All the information in this plot is derived from Monte Carlo Truth quantities.

This plot shows the same overall trend as Figure 3.4; namely that higher energy γ -rays result in a larger number of photoelectrons. But the figure also shows how strongly the core position distorts this relationship. The distance between the center of the heliostat field and the true core position is denoted as r_{true} . In Figure 9.1 we have subdivided the data set into events where the r_{true} is in the range 0 - 50 m, 50 - 100 m and 100 - 150 m. While there remains a clear dependence of total number of photoelectrons on energy (except, perhaps, for events where $r_{true} > 100$ m), the slope of the relationship is quite different. There is therefore a degeneracy between low energy showers that land near the center of the heliostat field and high energy showers that land near the edge of the heliostat field. Figure 9.1 demonstrates an important feature of energy reconstruction for STACEE: *accurate reconstruction of the γ -ray energy requires a good estimate of the EAS core position.*

There are also several other points to note about this plot. The first is that this plot is generated using our standard γ -ray simulations. Unlike Figure 3.4, the events in Figure 9.1 must actually trigger the experiment. The requirement that an EAS trigger the experiment results in some distortion of linearity of the relationship. This is clear at low energies, where the trigger requirement results in a larger number of photoelectrons than expected.




Figure 9.2: Dead-reckoned charge distributions for cosmic-rays. The solid lines are the real data; red points with error bars are the simulations. The left-side plot is for channel 32; The right-side plot is for channel 58. The plots of q_i are for all events near the Crab transit points.

The second point is that the total number of photoelectrons is a ‘Monte Carlo Truth’ quantity, rather than being a real measured quantity. It is calculated as the sum of the number of Cherenkov photoelectrons actually created on each channel, a quantity stored by the simulations. In real data we must estimate this quantity using the FADC information; but this measurement will be subject to a number of different errors. For instance, it is difficult to accurately measure the number of photoelectrons for very large pulses, where saturation of the FADC traces occurs. The clean relationships in Figure 9.1 therefore represent a best-case scenario, what we might attain in real data if our charge measurement and core determination were perfect.

In the following subsections we shall be examining individually the two elements of energy reconstruction: the charge measurement and the core position determination. Once the individual elements are in place, we will combine them together in our Template Fitting method.

9.1.1 Dead-Reckoned Charge Distributions

The measurement of the charge for each channel is conceptually simpler than the determination of the core position. We have already described in Section 5.2.1 how we calculate the dead-reckoned charge, q_i , for each channel. The dead-reckoned charge was calculated using a fixed 16 ns window around t_0 , our best estimate of the shower arrival time. This

charge calculation includes the pulse fitting scheme that accounts for saturated pulses. It is not absolutely essential that our measured charge distribution accurately match the true charge distribution, though it is desirable that the match be as good as possible. What is important is that the real and simulated measured charge distributions be a good match. To that end, we present a comparison between the dead-reckoned charge for real and simulated cosmic-ray data. Examples of this type of comparison are shown in Figure 9.2; the plots show the distributions of q_i for two channels for data taken at or near transit. A thorough examination of real and simulated q_i distributions shows a range of behavior, of which these two channels demonstrate the extremes. At one end we find channels, such as channel 32, where there is an excellent agreement between the real and simulated charge distributions. At the other end are channels, such as channel 58, where the real data has significantly more charge than the simulations. The comparison for most channels falls somewhere between these extremes.

These results suggest that real data events have slightly more charge, on average, than simulated events. The most probable explanation for this discrepancy is related to the variation in the optical throughput of the STACEE detector. Previous studies using drift scans suggest that there is more channel to channel variation in our optical throughput than our simulation would lead us to expect [83] (see Section 4.4.1 for a description of drift scans). Given the multiplicity requirement of the STACEE detector, a larger variation in the real optical throughput should lead to slightly more charge in the real data than was predicted. This throughput variability would also explain the channel to channel differences seen in Figure 9.2. Though this explanation is reasonable, some discrepancy remains between the real and simulated charge distributions. This is of concern since the spectral reconstruction that we shall present is predicated on there being a good match between real and simulated data.

It should also be noted that we shall sometimes be describing the charge for a given channel in terms of the number of photoelectrons (as in Figure 9.1) and other times in terms of pC (as in Figure 9.2). When we initially measure the charge for a given channel using the FADCs the result is in pC. But the conversion to number of photoelectrons only involves a factor related to the PMT and amplifier gains, which has been extensively calibrated. With the exception of Figure 9.2 we shall be describing charge in terms of number of photoelectrons, since the error on the real charge measurement is based on the number of photoelectrons.




Figure 9.3: Cartoon of the monocanted heliostat pointing scheme, in which all of the heliostats are aimed towards the canting point. The gray ellipses correspond to different possible positions of shower maximum and hence to different possible core positions. The sensitivity of STACEE is maximized for γ -rays with core positions near the center of the heliostat field (marked as x_0). The cartoon is not to scale. From [124].

9.1.2 Core Position Determination and STACEE Canting Schemes

The second element of energy reconstruction is an accurate determination of the core position. It turns out that the problem of core determination is related to our choice about where exactly the STACEE heliostats are pointed. We must therefore step back and describe the STACEE heliostat canting scheme. It was earlier noted that during normal operation the heliostats were set to track the putative γ -ray source. This is not exactly true. The heliostats actually track the position of expected shower maximum for a 100 GeV γ -ray coming from the source position and landing at the center of the heliostat field. The point that the heliostats track is referred to as the *canting point*. For a source that is exactly at Zenith, the canting point is approximately 10 km above the center of the field. The canting scheme where all heliostats are aimed at the canting point is called *monocanting*. Monocanting is demonstrated graphically in Figure 9.3.

The purpose of canting the heliostats is to maximize the amount of Cherenkov photons gathered from a γ -ray landing at the center of the field. The amount of canting required is not small; a heliostat that is 100 m from the center of the field points 0.6° away from the




Figure 9.4: Cartoon of the paracanted heliostat pointing scheme. Only 48 heliostats are aimed at the canting point for γ -rays landing at the center of the field. The other 16 heliostats are aimed straight up in the γ -ray source direction. Paracanting results in a better estimate of the core position. The cartoon is not to scale. From [124].

nominal direction in order to be looking at the canting point. Given that the FOV of our heliostats is restricted to be approximately 0.5° , canting is clearly important in order to collect the most Cherenkov photons coming from shower maximum. By allowing the detector to trigger on low energy γ -rays that land at the center of the field, the monocanting scheme minimizes our overall energy threshold.

It is, however, because of this canting scheme that there is such a strong degeneracy between the γ -ray energy reconstruction and the EAS core position. As Figure 9.3 makes clear, while monocanting maximizes Cherenkov photon collection from γ -rays that land at the center of the field, it also means that our Cherenkov collection efficiency decreases rapidly for showers that land away from the center. An alternative to the monocanting scheme would be to have all the heliostats be *parallel canted*, which means they would not be canted at all. All 64 heliostats would then be pointed exactly in the source direction. In this scheme each heliostat would be most sensitive to a γ -ray EAS that landed exactly on the heliostat. It would be easy in this scheme to determine the core position of a high energy γ -ray, since the heliostats nearest to the core position would have the greatest measured charge. However, STACEE would be unable to detect low energy γ -rays with




Figure 9.5: Choice of 16 parallel canted heliostats. Two heliostats are chosen from each cluster to be parallel canted; these heliostats are circled. The numbers refer to the cluster that each heliostat is in. From [124].

such a canting scheme; such a scheme would therefore defeat the low energy goal of STACEE.

We have developed a compromise between these two alternatives. This modified canting scheme is called *paracanting* and was devised principally by Richard Scalzo [123, 124]. In paracanting, 16 out of the 64 heliostats are aimed directly at the source (they are *parallel canted*); the other 48 heliostats are aimed at the canting point. This scheme is demonstrated graphically in Figure 9.4.

The exact choice of parallel canted heliostats is shown in Figure 9.5. Two parallel canted heliostats were chosen from each cluster. The paracanted heliostats therefore cover the entire heliostat field, which is important for accurate core position determination. The parallel canted heliostats become more sensitive to γ -ray EAS whose core position is *not* at the center of the heliostat field. By observing which parallel canted heliostats have

significant Cherenkov pulses, we can accurately determine the core position for a given event. It should be clear that we can only exploit the additional information of paracanting because the FADCs provide access to the digitized waveforms. Paracanting does raise the energy threshold of the STACEE experiment slightly; but the increased capacity to determine the core position is judged to be an adequate compensation.

However, the advantages and disadvantages of paracanting took some time to understand. The monocanting scheme was therefore used by STACEE for many years. The decision to change to paracanted heliostats only occurred near Christmas of 2003. Our data sets straddle this change. The following data sets were taken with monocanted heliostats: Crab 2002-2003, Mrk 421 2002-2003, 3C 66A 2003-2004. The following data sets were taken with paracanted heliostats: Crab 2003-2004, Mrk 421 2003-2004, OJ 287 2003-2004. Without paracanting our ability to make proper energy estimates is seriously reduced. We shall therefore only be doing spectral analyses for those data sets taken with the paracanted heliostat scheme.

It should be emphasized, however, that the change from monocanted to paracanted heliostats did not affect all elements of STACEE data analysis. The change only resulted in a small increase in the energy threshold of the experiment and hence a small decrease in the observed γ -ray rate excess. The decrease is sufficiently small that it cannot be seen conclusively when comparing the measured final γ -ray rates for the 2002-2003 and 2003-2004 Crab seasons. In addition, it can also be shown that the change from monocanted to paracanted data did not noticeably change the results relating to the Grid Alignment technique. So our conclusions regarding a cut on ξ_{cor} apply equally well to monocanted and paracanted data and our decision to analyze the two data sets together is reasonable. Only matters related to determination of the core position and energy reconstruction were significantly affected by changes in the heliostat pointing scheme.

9.2 Spectral Reconstruction

With an understanding of charge measurements and core position determination, we can now turn to making a measurement of the spectral index of the γ -ray sources that we detect. This process of *spectral reconstruction* actually involves two distinct steps. The first is making an estimate of the energy of each individual event; this step uses a method called *Template Fitting*. The second is using the energy estimate to extract information about the spectral index; this step uses a method called *Forward Folding*.

9.2.1 Core/Energy Template Fitting

Though the preceding sections discussed core position determination as a preliminary step before the energy estimation, the two tasks will actually be performed simultaneously as a part of our *Template Fitting* method. The Template Fitting method was also originally developed by Richard Scalzo [123, 124]. The following section will explain the method in detail.

The method is based on the comparison of the measured dead-reckoned charges for each channel of an event, q_i , to a predicted charge, $q_i^s(r_{true}, E, \theta, \phi)$. The predicted charge for each channel depends on the true core position, r_{true} , the incident γ -ray energy, E , and the source position, (θ, ϕ) . Given the details of complex EAS physics and changing optical throughput with source position, the dependence of the predicted charge on these quantities is quite complicated. It is therefore not possible to define a simple empirical formula for the predicted charge. This problem led to the development of *templates*. A template is a look-up table, defining the expected charge that would be seen for each channel for a γ -ray EAS with a given energy, core position and source direction.

Production of Templates

The procedure for generating templates uses special γ -ray simulations. For a given source position, twenty-five 200 GeV γ -ray EAS are simulated using CORSIKA. A 40x40 grid of points is then defined, with 10 m spacing, centered on the middle of the heliostat field. The first of these points is used as the core position for propagating the twenty-five CORSIKA EAS through *sandfield*; the output of *sandfield* is then used to compute the average number of photoelectrons for each channel. The process is repeated for every other grid point, resulting in charge predictions for every channel and every core position. This information constitutes our template. The templates are stored in a separate file that is used during the Template Fitting process.¹

Figure 9.6 shows examples of the templates for channels 0 and 1 for the Crab transit point. The plots show the number of photoelectrons that are expected as a function of core position for a 200 GeV γ -ray EAS. These templates demonstrate why paracanting allows us to determine the core position. Channel 1 is a regularly canted heliostat, which means that it is aimed at the canting point. As expected the template for channel 1 shows that this heliostat is most sensitive to γ -rays that land near the center of the field. Channel 0, however, is a parallel canted heliostat and hence is aimed directly at the source. The template

¹It should be emphasized that the grid of assumed core positions is different from the grid used in the Grid Alignment method. The similarity of the techniques simply demonstrates the importance of core position determination in event reconstruction.

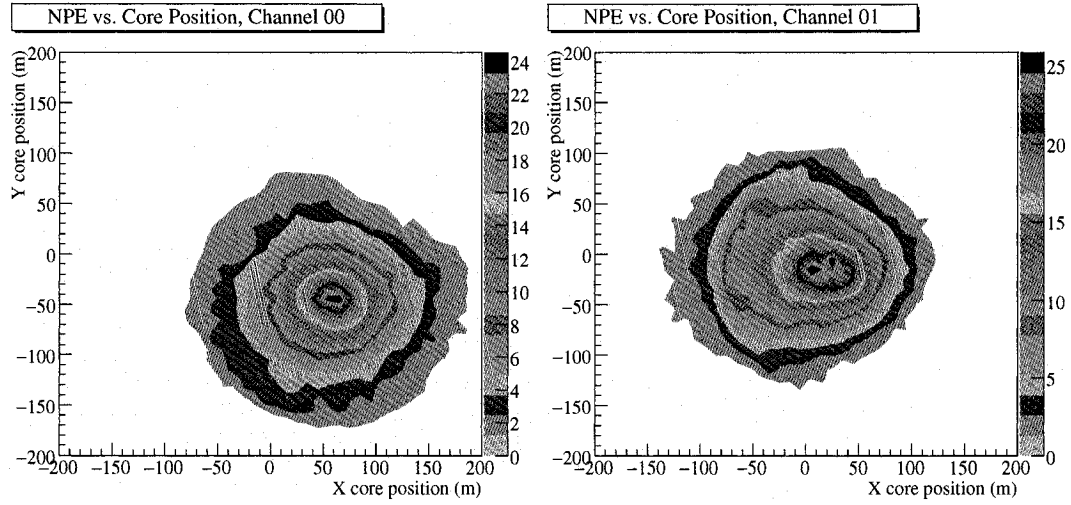


Figure 9.6: Examples of templates for channels 0 and 1 for the Crab transit point. The plots are 2-dimensional histograms. The contours represent the number of photoelectrons that are expected for a given core position. This set of templates was generated for the paracanted heliostat pointing scheme. Channel 0 is parallel canted (pointing at the source) and channel 1 is regularly canted (pointing at the canting point). The origin of the plots corresponds to the center of the heliostat field.

for channel 0 shows that the maximum sensitivity of this heliostat is to γ -ray EAS that land to the south-east of the center of the field. This discrepancy allows for the accurate determination of the core position of a given γ -ray. If the templates had been generated for the monocanted heliostat canting scheme, the template for channel 0 would, like that of channel 1, show the maximum sensitivity near the center of the field.

Several details about template generation are listed below:

- The fact that the templates were generated using mono-energetic γ -rays is simple to account for. We allow a simple linear scaling during the fitting process, such that, for a fixed core position and source direction, the predicted charge varies as

$$q_i^s \propto E. \quad (9.2)$$

The assumption of linear scaling for a fixed core position is not perfect, but it is reasonable given the evidence shown in Figures 3.4 and 9.1. We can therefore use our 200 GeV templates to extrapolate the expected charge for γ -rays of any energy.

- There is no attempt to actually model the STACEE trigger when generating the templates. For instance, the templates would define the number of Cherenkov photoelectrons for a 200 GeV shower landing far from the center of the field, even

though that number of photoelectrons would have no chance of triggering the experiment. The lack of a trigger in the templates will have consequences that will be noted later.

- It is inefficient to generate templates for every point on the sky. The templates are therefore produced at discrete points along the path of a given source. When comparing the templates to real data we generated an average template using linear interpolation between the discrete templates that bracket the real source position. The average template is re-interpolated every time the source moves 0.083° in HA (every 20 seconds of data).

Template Fit and Expected Performance

With the templates in place, performing the Template Fit is relatively simple. Our dead-reckoned charge measurements, q_i , will be compared to the predicted charge, q_i^s , derived from the templates. The comparison will use a χ^2 figure of merit. As always with a χ^2 comparison, care must be taken in defining the error on the charge measurement. The principal source of error is based on the Poisson fluctuations in the total number of photoelectrons produced at each PMT during the dead-reckoned charge window (which is 16 ns long). There are contributions to the total number of photoelectrons, q_i^{tot} , from both Cherenkov and NSB photoelectrons, denoted as q_i^C and q_i^{NSB} respectively, so that

$$q_i^{tot} = q_i^C + q_i^{NSB}. \quad (9.3)$$

Because of the AC-coupling of the PMT signal, our FADC measure of charge is only an estimate of the Cherenkov charge, so that we shall set $q_i^C = q_i$. To get an estimate of q_i^{NSB} we must use our measurement of the PMT current. The PMT current is entirely due to NSB photons, so the current together with our electronics gain allows an estimate of the rate of NSB photoelectrons. This rate is multiplied by 16 ns to get our estimate of q_i^{NSB} .

With these estimates of the number of Cherenkov and NSB photoelectrons, we can write the error on q_i as

$$\sigma_{q_i} = \sqrt{(0.30)^2 q_i^{tot} + q_i^{tot}} = \sqrt{1.09 q_i^{tot}}, \quad (9.4)$$

as long as the charge is measured in photoelectrons. The factor of 0.30 accounts for our single-photoelectron charge resolution being approximately 30%. That means that if the PMT only ever produces one primary photoelectron, there would still be a 30% spread in the final charge. The total number of photoelectrons is always sufficiently large to warrant




Figure 9.7: Reconstructed vs real energies for Template Fitting. The plots are for γ -ray simulations for Crab transit point; all distributions are after event cuts.

the assumption of a Gaussian error distribution. The error on q_i is dominated by the NSB fluctuations for low energy γ -rays and by the Cherenkov fluctuations for high energy γ -rays. Given the relatively large number of γ -ray EAS used to generate the templates, the predicted charge is assumed to have no error.

With the errors in place, the χ^2 figure of merit for this fit can now be defined:

$$\chi_{rec}^2 = \sum_{i=0}^{64} \frac{(q_i - q_i^s(\mathbf{r}_{rec}, E_{rec}))^2}{(\sigma_{q_i})^2}, \quad (9.5)$$

where \mathbf{r}_{rec} and E_{rec} are the reconstructed position and reconstructed energy. Note that NSB contributes solely to the error; the difference in the numerator depends on Cherenkov photoelectrons alone. The only free variables are the core position and the energy; the source direction is known and so is not allowed to vary. During the fitting process the core position and energy are varied until χ_{rec}^2 is minimized. All 64 channels are included in every fit. We do not wish to exclude channels with small or non-existent Cherenkov pulses since they are essential for determining the core position.

Template Fitting has been applied to Crab transit γ -ray simulations to demonstrate the effectiveness of the energy estimate. Figure 9.7 shows a distribution of reconstructed




Figure 9.8: Resolution and bias of reconstructed energy. The left-side plot shows the fractional resolution of the reconstructed energy, as a function of the true energy. The right-side plot shows the fractional bias in the reconstructed energy, as a function of the true energy. The plots are for γ -ray simulations for Crab transit point; all distributions are after event cuts.

versus real energies. Overall, the method works as expected, with a generally good correlation between the real and reconstructed energies. Figure 9.8 quantifies the effectiveness of the energy reconstruction; the figure shows the fractional resolution and fractional bias of the reconstructed energy compared to the true energy. The fractional energy resolution is defined as the RMS of the distribution of $(E_{rec} - E_{true})/E_{true}$ for each energy bin, whereas the fractional energy bias is defined as the mean value of the distribution of $(E_{rec} - E_{true})/E_{true}$ for each energy bin. The first observation is that the energy reconstruction is quite precise. The energy resolution is approximately 20% over the majority of the STACEE energy range. This provides justification for the effort that was put into properly understanding energy reconstruction.

The results are less satisfying when we investigate the distribution of the fractional energy bias. The right-side plot of Figure 9.8 shows a profile of the fractional bias as a function of the true energy. There are two notable features of this curve. First, at low energies (below 150 GeV) there is a tendency to *overestimate* the γ -ray energy. This bias is the result of our assumption that the density of Cherenkov photoelectrons scales linearly with the incident energy. This assumption fails because of the effect of the STACEE trigger. Figure 9.8 only includes events that actually satisfied our trigger condition, which means they must have a certain number of photoelectrons. For instance, a 100 GeV γ -ray that triggers the experiment must produce a larger charge than an average 100 GeV




Figure 9.9: Reconstructed core resolution. The solid line shows the difference between the real core position and our template estimate. The dashed line shows the difference had we assumed that the shower core was at the center of the field. The plots are for γ -ray simulations for Crab transit point; all distributions are after event cuts.

γ -ray, which makes our linear energy scaling of photoelectrons incorrect.

The second feature is that at high energies (above 700 GeV) there is a tendency to *underestimate* the γ -ray energy. This is the result of deficiencies of our pulse fitting routine. A γ -ray above 700 GeV will tend to saturate most of the FADC traces. The pulse fitting routine attempts to reconstruct the saturated pulses. However, the routine is not perfect and the measured charge tends to slightly underestimate the actual charge.

But even if we understand the biases, they can still cause systematic errors when converting our reconstructed energies into spectral measurements. As we shall show in the following section the method of using the reconstructed energies has been carefully chosen to remove the consequence of our biased energy estimate.

In addition to an estimate of the shower energy, the template fitting method also provides an estimate of the core position. We refer to this estimate as \mathbf{r}_{rec} . It should be emphasized that this estimate of the core position is different from the estimate from the Grid Alignment scheme, \mathbf{r}_{grid} , though the two quantities are, of course, measurements of the same thing. The Template Fitting method allows for a much improved ability to determine the EAS core position. This is shown in Figure 9.9, which plots core resolution results for the same Crab transit γ -ray simulations. It is clear that the Template Fitting allows us to find the core with better than 10 m resolution in each of the x and y directions. The plot also shows the core resolution had we simply continued our assumption that the core position was at the center of the heliostat field; the Template Fitting scheme is a vast




Figure 9.10: Crab 2003-2004 cosmic-ray E_{rec} distributions. The black and blue distributions mark the real ON and OFF distributions. Note that the excess of the ON data results from Crab γ -rays. The red triangles mark the distributions for HA-weighted simulated cosmic-rays. All distributions are after event cuts.

improvement.

Crab 2003-2004 Reconstructed Energies

The Template Fitting method can now be applied to our real 2003-2004 Crab data set. Figure 9.10 presents the ON and OFF distributions of reconstructed energy for the 2003-2004 Crab data set. The real ON and OFF distributions in Figure 9.10 have been corrected for the small difference in ON and OFF live-times. There is an excess of ON source events compared to OFF source events; this difference is the distribution of reconstructed energies for Crab γ -rays.

Before studying the distribution of Crab γ -ray energies, we shall perform another consistency check. Figure 9.10 also shows the distribution of E_{rec} for HA-weighted cosmic-ray simulations. As noted, our spectral reconstruction only works if the simulations are a good match to the real data. We do not expect our energy reconstruction to get the right scale when reconstructing cosmic-rays, since the Cherenkov photon density of cosmic-rays at a given energy is much smaller than that of γ -rays and our reconstruction assumes that the particles are γ -rays. We do, however, expect that the real and simulated distributions of E_{rec} would match. The plot shows that the match is decent, though not




Figure 9.11: Crab 2003-2004 γ -ray reconstructed Energy Distribution. The plot shows the difference between the ON and OFF distributions of E_{rec} , which is the result of Crab γ -rays. The ON and OFF distributions were corrected for the small difference in live-times. All the distributions are after event cuts.

perfect. There is some evidence that the simulated data underestimates the cosmic-ray energy. This is not surprising, given that the results in Section 9.1.1 showed that the dead-reckoned charge was slightly larger in real data than in simulations. Given the close connection between total charge and reconstructed energy, the slight mismatch in Figure 9.10 is expected. The caveats about this comparison are the same as noted in Section 8.4.1 when comparing the distributions of real and simulated ξ_{cor} . As with the Grid Alignment method, the most important test will be comparing the real and simulated γ -ray distributions.

Figure 9.11 shows the difference between the ON and OFF distributions in Figure 9.10; this is the distribution of reconstructed energies for Crab γ -rays. From this distribution we will extract a measurement of the Crab spectral index. The bins for these histograms are equally sized in $\log(E)$. The specific ranges for the energy bins are given in Table 9.1.

It is worth noting at this point the differences between the energy reconstruction presented above and the method presented in Carson (2005) [29]. The Carson method uses simulations to define the number of photons expected at each heliostat for a γ -ray shower landing at a given core location. This information, in combination with the core position estimate from the Template Fitting, is used to convert the measured charge into the

Table 9.1: Energy Bins for Spectral Reconstruction.

Bin Number	Bin Range (GeV)
1	63 - 100
2	100 - 158
3	158 - 251
4	251 - 398
5	398 - 631
6	631 - 1000
7	1000 - 1585
8	1585 - 2512

number of Cherenkov photons at each heliostat. An empirical relationship allows for the conversion of the mean number of Cherenkov photons at the heliostats into a measured energy. The Carson energy measurement has a smaller bias in its energy reconstruction, but a worse energy resolution.

9.2.2 Forward Folding

The Template Fitting method allows an estimate of the γ -ray energy on an event by event basis. We have shown the γ -ray distributions of E_{rec} for the Crab 2003-2004 data set in Figure 9.11. It might be thought that with this information the problem of spectral reconstruction is essentially solved. This is, however, not the case. The other half of the process involves accurately quantifying the detector efficiency as a function of energy, in order to convert the energy spectrum of γ -rays *detected by STACEE* into the energy spectrum of γ -rays *incident at the top of the atmosphere*. There are considerable subtleties involved in this process and mistakes can easily lead to systematic errors in the spectral measurements.

These problems are not unique to STACEE. For instance, one of the standard works on Whipple spectral reconstruction is Mohanty et al. (1998) [100]. In that work they describe two different methods of arriving at a spectral measurement of the Crab Nebula. Two methods were required because of the large potential for systematic biases. In that work, the results of both methods agree within errors, demonstrating the robustness of the methods. The spectral reconstruction method used in the present work is inspired by Mohanty Method 2. However, in order to understand the benefits of Method 2, it is worth explaining how Mohanty Method 1 would be applied to STACEE.

A simplified description of Mohanty Method 1 as applied to STACEE is as follows: The first step is simple: make histograms of ON and OFF reconstructed energies. The

difference between the ON and OFF distributions is the distribution from the γ -rays. This distribution is already shown in Figure 9.11. The content of each histogram bin is then divided by the bin width and total live-time; the measurement for each bin is therefore counts $\text{sec}^{-1} \text{ GeV}^{-1}$.

The next step is more troublesome. The energy distribution must be corrected by a factor accounting for the sensitivity of the experiment as a function of energy. The obvious solution would be to divide the reconstructed energy distribution by the final effective area that we presented in Figure 8.14. The units of the resulting distribution would be counts $\text{m}^{-2} \text{ sec}^{-1} \text{ GeV}^{-1}$, which is a standard unit for describing source fluxes in γ -ray astronomy.

Dividing by the final effective area is approximately the correct step; there are, however, complications with this procedure. The first complication is that the effective area curve changes rapidly, especially near the energy threshold. Given the low significance of the STACEE measurements it is necessary for our energy bins to be relatively large in order that the excess in most bins be significant. For instance, the second to lowest energy bin in Figure 9.11 covers the range 100-158 GeV; the effective area changes by more than an order of magnitude over this range. Assigning an average effective area to this bin is a challenge, especially since the average value of the effective area would inevitably be based on some assumption about the form of the underlying spectrum.

The second complication is that the effective area is plotted with respect to the *actual* γ -ray energies, whereas the E_{rec} distribution is a plot of the *reconstructed* γ -ray energies. This is a problem for two reasons. The first occurs if the reconstructed energies are biased with respect to the actual energies. For instance, from the right-side plot of Figure 9.8 it is clear that our reconstruction systematically overestimates the energy of low energy γ -rays and systematically underestimates the energy of high energy γ -rays. The second problem is what is referred to as *migration*. Even if the energy reconstruction has no bias, the finite energy resolution leads to a flattening out of the reconstructed energy distribution. Both of these effects, if uncorrected, will systematically distort the reconstructed energy spectrum. In Mohanty Method 1 this is accounted for by modifying the effective area curve to correct for bias and migration. But this correction is not simple, not least because it depends on the underlying γ -ray spectrum; for instance, migration is a worse problem for a γ -ray source with a steeper spectrum (larger value of α). So the correction to the effective area curve depends on the form of the SED that we are trying to measure. The dependence is weak, which makes the problem tractable using an iterative procedure. Nevertheless, the correction is susceptible to systematic bias.

Once we have a corrected and properly averaged effective area, we can divide the

energy distribution by the effective area. The final result should be a correct measurement of the γ -ray flux as a function of energy. This method is the standard one used by most VHE γ -ray experiments. It is also a variant of this method that has previously been used in STACEE investigations of Mrk 421 in Carson (2005). The Carson method differed from Mohanty Method 1 principally in that the energy estimate itself was corrected for biases and migration.

Yet despite the common use of Mohanty Method 1, there remains the aforementioned concerns about systematic errors. It is therefore useful that STACEE spectral results be confirmed using an alternate method based on Mohanty Method 2; we shall refer to this method as *Forward Folding*. The idea of Forward Folding is simple. We already have a process for calculating the distribution of reconstructed energies for the real γ -rays; in the case of the 2003-2004 Crab data, this distribution is shown in Figure 9.11. The trick is to produce an equivalent distribution for simulated γ -rays. The simulated distribution will depend on the assumed energy distribution of the simulated γ -rays; the method therefore involves varying the spectral index of the γ -ray simulations in order to get the best match with the distribution of E_{rec} for real data. The concept is simple; the details will be shown in the following section.

The principal advantage of Forward Folding is that it naturally accounts for the biases of our original energy estimate, as well as accounting for the effect of migration. There is no need for any complicated averaging of the effective area curves. This is because both the real data and the simulated data are subjected to a completely equivalent treatment; the energy of every real 100 GeV γ -ray might be overestimated, but so will the energy of every simulated 100 GeV γ -ray. The only requirement for getting an accurate spectral reconstruction is that the simulations be a good match to the real data; as long as this requirement is satisfied the Forward Folding technique will provide an accurate measurement of the spectral index. Contrast this with a spectral reconstruction using Method 1, where an accurate measurement requires both that the simulations be a good match to the real data and that the resolution and biases of the energy estimator be properly accounted for. We therefore claim that there is less chance of systematic errors in the method of Forward Folding (or at least that the systematic errors will be different). Techniques similar to Forward Folding has been used successfully by Whipple, as well as by large particle physics experiments, as a means of measuring basic physics quantities.




Figure 9.12: Distributions of γ -ray reconstructed energies for different spectral indices. The different *synthetic* curves were generated from a set of $\alpha = 2.4$ simulations by re-weighting the entries. These histograms are normalized to have equal integrals. The curves are only to guide the eye. All the distributions are after event cuts.

9.3 Crab 2003-2004 Spectral Reconstruction

The best way of explaining the details of the Forward Folding method is by using an example. We shall use the 2003-2004 Crab data set for this test. In addition to a demonstration, the test will also serve to validate the method, since the results can be compared to expectations for the Crab spectral index. It should be emphasized that, unlike the tests of the Grid Alignment method that used both the 2002-2003 and 2003-2004 data sets, this spectral reconstruction test will only use the 2003-2004 data set. This is because only the 2003-2004 data set used the paracanted heliostat scheme and therefore has the potential for accurate core determination.

The first step of the process, the generation of the distribution of E_{rec} for real Crab γ -rays, has already been completed. The second step is to generate the simulated distributions of E_{rec} that we would expect for different spectral indices. We start with our Crab γ -ray simulations; the simulations have a continuous power-law distribution of energies with a spectral index of 2.4. As always, we combine simulations from the transit, +1HR and +2HR pointings in an appropriate manner. After running these simulations through the standard Template Fitting scheme we can construct a distribution of expected




Figure 9.13: Reduced χ^2 for comparisons between real and simulated Crab distributions of E_{rec} . The reduced χ^2 are plotted versus the spectral index of the γ -ray simulations. The solid line is a parabolic curve, fitted to the region near the minimum χ^2 value. The dotted line indicates the χ^2 value at which the error on the estimate of α is calculated.

E_{rec} for Crab γ -rays. However, this histogram of reconstructed energies is dependent on the assumption that $\alpha = 2.4$. In order to calculate the predicted distribution for different spectral indices we reuse the same γ -ray simulations, but re-weight the histogram entries. So, for instance, in order to simulate the E_{rec} distribution for $\alpha = 2.2$, we weight each entry in the histogram with a factor, w , defined as

$$w = \left(\frac{E_{true}}{100 \text{ GeV}} \right)^{0.2}. \quad (9.6)$$

This procedure is more efficient than actually redoing the complete simulation chain for each different spectral index. Note that the re-weighting is done on the basis of individual events, before being binned in histogram. The results of these re-weighted *synthetic* E_{rec} distributions are shown in Figure 9.12. The distributions of reconstructed energy look sensible, with the results for $\alpha = 1.4$ having the highest mean reconstructed energy. The curves are naturally only for events that passed the padding and $\xi_{cor} < 0.35$ cuts.

With a proper set of synthetic γ -ray simulations for different spectral indices, we can now compare to these predictions to the real Crab γ -ray results shown in Figure 9.11. We shall use a χ^2 estimator to determine the goodness-of-fit between the real data and the different synthetic distributions. Specifically, for each different spectral index and simulated




Figure 9.14: Real and simulated Crab 2003-2004 γ -ray reconstructed Energy Distribution. The plot shows the difference between the ON and OFF distributions of E_{rec} , which is the result of Crab γ -rays, as well as the distribution for the $\alpha = 2.2$ simulation.

prediction, we calculate the estimator using the eight bins of the E_{rec} histograms as

$$\chi^2 = \sum_{i=1}^8 \frac{(N_i - N_i^s(\alpha))^2}{\sigma_{N_i}^2}, \quad (9.7)$$

where N_i is the measured number of events in each energy bin for the real distribution and $N_i^s(\alpha)$ is the number of events for each energy bin in the synthetic simulated distributions. Note that the predicted distributions are always normalized to have the same integral as the real data. Note also that the statistical error on the predictions is always much smaller than the error on the real measurements, so the estimator uses only the error on the real measurements.

The results of this comparison are displayed in Figure 9.13, which shows the reduced χ^2 for the comparison between the real data and the different simulated predictions. The curve is minimized near $\alpha = 2.2$. Figure 9.14 shows the real γ -ray distribution of E_{rec} , along with the prediction for a $\alpha = 2.2$ simulation. The figure shows that $\alpha = 2.2$ is indeed a good fit to the real data.

The distribution of χ^2 near a minimum is expected to be approximately parabolic [49]; a second-order polynomial is therefore fit to the data in Figure 9.13. The parabolic fit gives us a refined estimate of the best-fitting spectral index; in this case it allows us to state that

$\alpha_{crab} = 2.29$ for the 2003-2004 Crab data set. The minimum χ^2 is 3.6 for 6 degrees of freedom, indicating a good fit. One of the drawbacks of the Forward Folding technique is that the power law form of the γ -ray spectrum must be assumed a priori, whereas the true spectrum might be more complicated. The only indication of a poor choice of the assumed spectrum would be a larger value of χ^2 . In this case, the reasonable χ^2 value indicates that a power law description of the Crab SED is appropriate, at least at the sensitivity of our measurements.

The error on the fitted parameter is defined by where the $\chi^2 = \chi^2_{min} + 1.0$ [49]; in our case this is where the reduced χ^2 has increased by 0.167. This increase is shown by the dotted line in Figure 9.13. The error on α_{crab} turns out to be ± 0.17 . So our full measurement of the Crab spectral index is $\alpha_{crab} = 2.29 \pm 0.17$. This is the first STACEE measurement of the Crab spectral index.

How does our measurement of α compare with expectations? Whipple observations of the Crab Nebula lead to a measurement of $\alpha = 2.49 \pm 0.06_{stat} \pm 0.04_{sys}$ for γ -rays in the energy range of 500 GeV to 8 TeV [58]. However, the spectrum of the Crab Nebula is expected to flatten near 100 GeV, which suggests that STACEE would see a spectral index that is smaller than 2.49 [44]. Preliminary results from the MAGIC telescope seem to confirm these predictions. MAGIC measured the spectral index in the range 300 - 3000 GeV to be 2.58 ± 0.16 ; the data points they measured below 300 GeV lay below the extrapolation of the power law, as expected [147] Our estimate of $\alpha_{crab} = 2.29 \pm 0.17$ for the Crab Nebula in the range 100 GeV to 1000 GeV therefore seems reasonable.

Our measurement of the spectral index for the Crab Nebula is therefore consistent with the expectations from other groups. Despite some concerns about the distributions of E_{rec} for cosmic-rays, this Crab comparison demonstrates the fundamental soundness of our techniques for energy reconstruction and the extraction of the spectral index. We can therefore have confidence in our results when we apply the same technique to the Mrk 421 data set.

Chapter 10

AGN Observations and Results

The preceding chapters have completed the description of the STACEE analysis techniques. These techniques have been extensively tested using simulations, real Crab data and star observations. In this chapter we can finally apply these techniques to the three AGN data sets that are the subject of this thesis. The AGN data sets underwent exactly the same data quality and event cuts as the Crab data set that was examined in the previous chapters. The same data quality cuts were used to remove periods where the detector was malfunctioning or the weather was bad. The same Library Padding procedure was applied to remove any excess due to promotion and the same Grid Alignment scheme was used with a cut of $\xi_{cor} < 0.35$ to improve the signal-to-noise of the data set. The cut value was fixed using the simulations and was not tuned based on any AGN results.

In this chapter we show that STACEE observations result in a detection of Mrk 421 and non-detections of both 3C 66A and OJ 287. With the detection of Mrk 421 additional analysis can be applied to the data. The discussion of Mrk 421 results is therefore much longer than those of 3C 66A and OJ 287. For the latter two sources we simply compute the STACEE flux upper limits and compare them to results from other experiments and to model predictions.

10.1 Markarian 421 Observations and Results

STACEE observed Markarian 421 during the 2002-2003 and 2003-2004 seasons. A total of 70 pairs of data runs were taken during 2002-2003 and 45 pairs during 2003-2004. We used the standard data quality cuts to remove periods where the detector malfunctioned or the weather was poor. 40% of the data was removed in 2002-2003, resulting in a final ON source livetime of 58 ks. 32% of the data was removed in 2003-2004, resulting in a final ON source livetime of 42 ks.




Figure 10.1: Mrk 421 ξ_{cor} distributions. The upper plots are for 2002-2003; the lower plots are for 2003-2004. The left-side plots show the distributions of ξ_{cor} for the ON and OFF data; the distributions are corrected for the differences between ON and OFF live-times. The right-side plots show the difference between the ON and OFF distributions. The dashed red line shows our cut value.

The use of the Library Padding procedure was essential for this source, because of a bright star¹ in the ON field of view. The raw ON-OFF excess rate for the entire 2002-2004 was 17.6 ± 0.6 counts min^{-1} , a ‘detection’ of 27.5σ significance. The majority of this excess is the result of promotion. Only after applying the padding cut do we have confidence that an excess is the result of Mrk 421 γ -rays.

After the padding cut we apply our background suppression technique. The results of the Grid Alignment technique are shown in Figure 10.1. The left-side plots show the

¹HD 95934, B band magnitude of 6.2.




Figure 10.2: 2002-2004 Mrk 421 significance vs cumulative observing time. The significance is after ξ_{cor} cut. The solid line is a $f(t) \propto \sqrt{t}$ fit to the data. The dashed vertical line shows the boundary between data from 2002-2003 and 2003-2004 seasons.

ON and OFF distributions of ξ_{cor} after the padding cut. The right side plots show the differences between the ON and OFF distributions (ie the γ -ray distributions). The upper plots are results from 2002-2003 data and the lower plots are results from 2003-2004 data.

Both data sets show evidence for an excess of ON events. We apply the Grid Alignment cut at the value defined in Chapter 8, $\xi_{cor} < 0.35$. For 2002-2003 data this results in a final γ -ray rate of 1.01 ± 0.26 counts min^{-1} , with a detection significance of 3.9σ . For 2003-2004 data this results in a final γ -ray rate of 2.88 ± 0.27 counts min^{-1} , with a detection significance of 10.9σ . The detection is much stronger in 2003-2004, despite having less ON-source time; it is clear therefore that Mrk 421 was in a more active state that year.

Figure 10.2 presents a plot of the final significance of the data set versus time for both seasons of Mrk 421 data. Note that, in contrast to the equivalent plot for the Crab data set (Figure 8.19), the significance vs time for Mrk 421 is not well described as $f(t) \propto \sqrt{t}$. This reflects the fact that the Mrk 421 flux is not constant; it is clear that the significance rises faster during 2003-2004, when Mrk 421 was in a higher flux state. Figure 10.2 also shows that there is no single pair that is significantly affecting the final results. This gives us confidence that our data quality programs have successfully removed any obviously corrupted pairs.

With the detection of Mrk 421, there is a great deal more analysis that can be done.

We shall start by comparing our measured rates to those of other experiments and looking for correlations. We shall then make two different measurements of the Mrk 421 spectral energy distribution. The first one will be a simple measurement where we assume that the Mrk 421 spectral index is $\alpha = 2.5$. We shall then do a more refined spectral analysis, where we make our own measurement of the spectral index and then use it to calculate the overall flux level. Finally, we shall discuss the implications of our Mrk 421 results.

10.1.1 Multi-wavelength Comparisons

Studies of BL Lac objects are most informative when done simultaneously at several different wavelengths. Multi-wavelength observations are necessary because modelling these sources requires information from across the electromagnetic spectrum. Simultaneous observations are necessary because BL Lac objects are highly variable.

Mrk 421 is a well studied object, as one would expect for the first extra-galactic source detected in VHE γ -rays. We therefore can make direct comparisons between STACEE observations and those of other instruments. In particular, we shall compare our measured final γ -ray rate to data from the Rossi X-ray Timing Explorer (RXTE) and the Whipple telescope. The Whipple telescope operates close to the STACEE energy range and we would therefore expect to see good correlation between the STACEE and Whipple rates. The correlation between X-ray and γ -ray fluxes is more complicated and the expected degree of correlation is model dependent. The generally favored Synchrotron Self-Compton (SSC) emission models lead us to expect a high degree of correlation between the X-ray and γ -ray fluxes, as was described in Section 2.2. Hadronic emission models lead to expectations of lower correlation.

For RXTE we use results from the All Sky Monitor (ASM) instrument, which detects X-rays in the range 2 - 10 keV. The ASM results are available online from the MIT archive.² We retrieved the ASM-measured Mrk 421 rate from each *dwell* (ie observation), then averaged the ASM measured rates for each day. Specifically, we calculated a weighted average rate using the dwells that occurred in the period [0,0.6] of the Julian day. We use only data from this portion of the Julian day since it brackets the period of STACEE observations. The result is that the mean time of the RXTE observations matches reasonably well with the mean time of the STACEE observations, as we shall show. The weights for the averaging are set to be $1/\sigma^2$, where σ is the error on the individual ASM rate measurement. After this averaging, we remove any days where the final error on the measured ASM rate was greater than 2.0 cts sec⁻¹ (following the procedure

²See <http://xte.mit.edu/asmlc/srcs/mkn421.html#data>.




Figure 10.3: RXTE, STACEE and Whipple light curves for Mrk 421 during the 2002-2004 seasons. The dotted lines marks January 1st of each year.

described in Cui (2004) [34]).

The Whipple results are presented in Blazejowski et al. (2005) [21], and were kindly provided by Wei Cui. The Whipple data consists of run-by-run measurements of the Mrk 421 γ -ray rate. We use the same averaging procedure to compute daily Whipple rate measurements. The final error on the daily Whipple measurements was always small, so we did not remove any nights from our data set.

Note that for all three instruments we are dealing with experimental trigger rates,




Figure 10.4: Difference in the mean time of observation for the 26 nights of contemporaneous observations. The left-side plot shows the time difference between STACEE and RXTE observations. The right-side plot shows the time difference between STACEE and Whipple observations.

rather than ‘proper’ source flux measurements. In general there will be a strong correlation between the measured trigger rate and the source flux, but this correlation will be moderated by changing detector sensitivity. For instance, the STACEE detector efficiency changes with the source elevation, a fact that our simple comparison of the γ -ray rate does not take into account. This caveat should be borne in mind.

The results of this procedure are shown in Figure 10.3, which shows the 2002-2004 Mrk 421 light curves as seen by these three instruments. It is clear from this figure that there is a great deal of variability in the Mrk 421 flux levels. It is also clear that the coverage of this source is very different for the three instruments. Both STACEE and Whipple can only make observations on clear moonless nights, which explains why there is far less data from those two experiments. Because of this, we must carefully choose the exact periods for which to make comparisons.

For the following analysis, we have chosen nights on which there was data from all three instruments; there are a total of 26 such nights. Figure 10.4 shows the distribution of the difference between the mean time of the STACEE observations and the mean time of the observations of the other two experiments. Overall, the STACEE data sets are temporally well matched to the other experiments; there is only an average of an hour difference between the STACEE observations and those of the other instruments. Even so, the temporal overlap with STACEE observations is not entirely satisfactory; significant variations in Mrk 421 fluxes are known to occur on hourly timescales [1,21]. Unfortunately, it is difficult to make comparisons for shorter timescales than a day, because stricter time overlap requirements would result in a much smaller set of STACEE data. The overall STACEE detection significance of the smaller data set would be very weak and would make com-




Figure 10.5: Comparison between RXTE and STACEE measured rates for Mrk 421. The data shown is for 26 nights of simultaneous data. The straight line fit is constrained to pass through the origin.

comparisons between STACEE and other experiments inconclusive. The imperfect temporal match is therefore a second caveat for the following comparisons.

We shall start with the X-ray/ γ -ray comparison. Figure 10.5 shows a scatter plot of the mean daily Mrk 421 rates for STACEE and RXTE. The plot shows the results for the 26 nights of contemporaneous observations. There are at least two ways in which we can characterize the degree of correlation. The first is to calculate the Pearson correlation coefficient, r . A description of the correlation coefficient is given in Appendix A. The crucial point is that r can range between -1.0 and 1.0, with a value of 1.0 indicating a perfect linear correlation between two quantities. For this set of data the correlation coefficient is $r_{rxte} = 0.58^{+0.12}_{-0.15}$. Taking the non-Gaussian errors on r_{rxte} into account, r_{rxte} is different from zero at the 3.3σ level [116]. There is therefore evidence for some correlation between the STACEE and RXTE rates, but the correlation is not 100%.

As detailed in Appendix A, at least part of the reason that r_{rxte} is lower than 1.0 can be attributed to the known errors on the measured rates. We need to be aware that our known measurement errors means that our calculated r_{rxte} is therefore less than the ‘true’ r_{rxte} that we could calculate if our measurement errors were zero. Despite this problem with our technique, it is interesting to note that this measurement of the STACEE/RXTE correlation coefficient is consistent with results on Mrk 421 by the MAGIC group. During a 2005 observation of Mrk 421, they found that the correlation coefficient between




Figure 10.6: Comparison between Whipple and STACEE measured rates for Mrk 421. The data shown is for 26 nights of simultaneous data. The straight line fit is constrained to pass through the origin.

measured MAGIC and ASM rates was $r = 0.57^{+0.12}_{-0.14}$ [6].

Because of this concern about the interpretation of r_{rate} , we also use a second procedure to quantify the degree of correlation between the STACEE and RXTE rates. We fit a straight line to the data in Figure 10.5; the line is constrained to pass through the origin. The χ^2 of this fit is related to the degree of correlation between the two quantities. Here, the χ^2 is 44 for 25 degrees of freedom. This has a chance probability of 1.0%, which means that if the straight line is a good description of the data would only expect a χ^2 of 44 or higher in one out of a hundred attempts. There is therefore a low probability that there is a perfect linear correlation between the rates of the two experiments. These results are consistent with there being only a partial correlation, as suggested by our calculated r .

Figure 10.6 shows a similar comparison between the Whipple and STACEE Mrk 421 rates. We perform the same two tests of the linear correlation as for the RXTE/STACEE comparison. We measure the correlation coefficient to be $r_{whip} = 0.78^{+0.07}_{-0.09}$; this value is different from zero at the 5.2σ level. It is interesting to compare this value of r to the measurement for the RXTE/STACEE comparison; the comparison shows that the Whipple/STACEE rates are marginally more correlated than the RXTE/STACEE rates. One might argue that such a comparison is meaningless, since, as noted, the correlation coefficient does not account for the errors on our measurements and the errors are different for

Table 10.1: Summary of Mrk 421 STACEE results for low, medium and high periods.

Period	Livetime (ks)	Final Rate (counts min ⁻¹)	Detection Significance
Low	31.6	0.92 ± 0.36	2.6σ
Medium	20.4	2.46 ± 0.38	6.5σ
High	17.0	3.63 ± 0.41	8.9σ

the RXTE and Whipple points. However, the dominant error for either comparison is the error on the STACEE measurement. Since this error is common to both comparisons, the difference between the two correlation coefficients should be a good measure of whether the STACEE rates are more correlated with Whipple than with RXTE.

We can also do the same straight line fit for the Whipple/STACEE scatter plot. In the case of this comparison, the χ^2 is 27 for 25 degrees of freedom, which has a chance probability of 38%. These results are therefore consistent with there being a greater degree of linear correlation between the STACEE and Whipple rates than between the STACEE and RXTE rates. This is as expected, since STACEE and Whipple cover almost the same range of the electromagnetic spectrum.

Whether or not there is *more* correlation between the Whipple and STACEE than RXTE and STACEE, there is at least *some* correlation in both cases. We can demonstrate this more clearly by averaging our nightly rates over longer timescales. In Blazejowski et al. (2005) [21] the 2002-2004 data set was divided into low-, medium- and high-state periods based on PCA rates.³ For the purposes of the present work these different periods are defined as covering December 2002 - May 2003 (low state), January 2004 - March 2004 (medium state) and April 2004 (high state). This doesn't correspond exactly to periods described in Blazejowski et al. (2005) because that work used stricter conditions in defining periods of comparison. Of the 26 simultaneous nights defined previously, 11 nights are during the low period, 7 nights are during the medium period and 8 nights are during the high period. We calculate the average of the nightly averages for each of the periods for STACEE, RXTE and Whipple data. The average STACEE rates for these periods are summarized in Table 10.1.

The comparisons between these average rates are shown in Figure 10.7. As expected, the plots show that the average STACEE rates do indeed increase with the Whipple and RXTE rates from the low to medium to high state periods. Given Mrk 421's known variability, this method of averaging may be questioned. It is unclear what physical significance to place on a multi-month average of a sparse series of daily rate averages of a

³The PCA (Proportional Counter Array) is another instrument on RXTE.




Figure 10.7: Coarse rate comparison for Mrk 421. The left-side plot shows the average rates for the low, medium and high periods for STACEE and RXTE. The right-side plot shows the same, except for STACEE and Whipple.

source that varies on an hourly timescale. Figure 10.7 cannot therefore be used to make any quantitative statements about correlation between the different instruments. Nevertheless, the plots are a compelling demonstration that correlations do exist.

10.1.2 Simple Flux Measurement

Our next goal is to convert our measured γ -ray rate into a proper source flux measurement. We shall present two different methods for doing this. Both methods use STACEE's Mrk 421 response curve to convert our final observed γ -ray rate into a flux normalization. The difference is in what spectral index to use in this procedure. In our initial, *simple* flux measurement we shall simply assume that the spectral index is $\alpha = 2.5$. In our second, *refined* flux measurement, we shall first calculate the spectral index using the method outlined in Chapter 9 and then use that to calculate the flux normalization. It should be noted that both methods assume that the underlying form of the SED is well described by a power law.

It would seem clear that it is preferable to use a measured α rather than an assumed α . Why therefore do we perform the simple flux measurement at all? There are several reasons:

- The spectral index measurement cannot be done for the 2002-2003 Mrk 421 data, since the data set used the monocanted heliostat scheme. The simple flux measurement is therefore the only flux measurement that will be done for the 2002-2003




Figure 10.8: HA-weighted final effective area curve for Mrk 421 2003-2004 after all event cuts.

data.

- The assumption that $\alpha = 2.5$ is the same assumption that we will use in calculating upper limits for 3C 66A and OJ 287. Providing a flux measurement using the same assumption makes comparison between the results for the three different AGN more direct.
- The simple method allows us to discuss subtle aspects of the flux normalization, such as systematic energy scale errors, without the additional complication of spectral index measurements.

We start, therefore, with our simple flux measurements. To perform them we need our final effective area curve for Mrk 421. Figure 10.8 shows this curve for the 2003-2004 season (also shown in Figure 8.14). We convolve the effective area with an assumed form for the source γ -ray flux; the resulting distribution is called the *response curve*. As noted, we assume that the source flux is well described as having an $\alpha = 2.5$ power law energy distribution, ie

$$\frac{dN}{dE} = N_0 (E/E_{thr})^{-2.5}, \quad (10.1)$$

where N_0 is the normalization constant. The integral of the response curve gives the

expected STACEE γ -ray trigger rate, R . That is to say

$$\begin{aligned} R &= \int_0^\infty \frac{dN}{dE} A_{eff}(E) dE \\ &= N_0 \int_0^\infty (E/E_{thr})^{-2.5} A_{eff}(E) dE, \end{aligned} \quad (10.2)$$

where A_{eff} is the final effective area. We know the final γ -ray rate as well as the effective area and have made an assumption about the power law form of the source spectrum. We can therefore calculate the normalization constant N_0 .

The final response curve for Mrk 421 2003-2004 is shown in Figure 10.9. As noted, our energy threshold is defined to be the energy at which the differential trigger rate is maximized. For our purposes, we will define our measured energy threshold as being the midpoint between the highest two points in Figure 10.9. This definition is straightforward and adequate given the large systematic error in energy threshold that will be discussed in the next section. In the case of the 2003-2004 Markarian 421 data we obtain an energy threshold of 185 GeV. With the energy threshold we can solve for the normalization constant and calculate that the flux at E_{thr} is therefore

$$\frac{dN^{Mrk421}}{dE}_{03-04}(185 \text{ GeV}) = (3.17 \pm 0.24) \times 10^{-8} \text{ GeV}^{-1} \text{ m}^{-2} \text{ s}^{-1}$$

The error on the flux is statistical only and is based on the error on the measured rate.

The more conventional way of describing flux is as $E^2 dN/dE$, which is the same quantity as νF_ν . In addition, since we know the spectral index, we may write the full form of the spectral energy distribution. For the 2003-2004 Mrk 421 it is

$$E^2 \frac{dN^{Mrk421}}{dE}_{03-04} = (10.8 \pm 1.0) \times 10^{-4} \left(\frac{E}{185 \text{ GeV}} \right)^{-0.5} \text{ GeV m}^{-2} \text{ s}^{-1}.$$

An identical procedure is applied to our Mrk 421 results for 2002-2003. In that case we find that the energy threshold is again 185 GeV and that the spectral energy distribution is

$$E^2 \frac{dN^{Mrk421}}{dE}_{02-03} = (3.9 \pm 1.0) \times 10^{-4} \left(\frac{E}{185 \text{ GeV}} \right)^{-0.5} \text{ GeV m}^{-2} \text{ s}^{-1}.$$

This is our simple measurement of the Mrk 421 flux in 2002-2003 and 2003-2004.

Systematic Error on Energy Threshold and Flux Normalization

This is an appropriate point to analyze the systematic errors on our flux measurements. In this analysis we largely follow the work of Oser (2000); the key point of that work was




Figure 10.9: Mrk 421 2003-2004 Response Curve for $\alpha = 2.5$ spectrum. The line is meant only to guide the eye.

that the calorimetric nature of EAS simplifies the discussion of systematic errors and that the systematic errors on our flux measurements could be simply related to errors in our energy scale. Our energy scale defines how much charge we expect to be produced in our detector by a γ -ray shower with a given energy [111].

For instance, given the calorimetric nature of EAS, a 200 GeV γ -ray will be identical to a 180 GeV γ -ray, except that the 180 GeV shower will produce a 10% lower yield of Cherenkov photons on the ground and hence a 10% lower charge in our detector. Suppose, however, that the throughput of our detector was 10% lower than expected; for instance, suppose the actual PMT gains were 10% lower than the measured PMT gains. In that case a 200 GeV γ -ray would produce a STACEE event that appeared exactly as we would expect a 180 GeV γ -ray to appear. This scaling will be true for any other element of our throughput, such as atmospheric attenuation or mirror reflectivity.

Let us therefore define β to be the ratio of our true throughput to the nominal throughput. So, as noted, if $\beta = 0.9$, then a 200 GeV γ -ray event would look like a 180 GeV γ -ray event. The true effective area would then be related to the nominal effective area by

$$A_{eff}^{true}(E) = A_{eff}^{nom}(\beta E) = A_{eff}^{nom}((1 \pm \Delta\beta)E), \quad (10.3)$$

Cherenkov production	5%
Atmospheric attenuation	15%
Heliostat reflectivity	10%
Heliostat pointing	5%
Secondary reflectivity	3%
Secondary pointing	5%
PMT Q.E and gains	14%
Total ($\Delta\beta =$)	25%

Table 10.2: List of principal sources of systematic error on the energy scale.

where $\Delta\beta$ defines our uncertainty in the energy scale. Of even more importance is how β affects our energy threshold. It is easy to show that for a power law spectrum,

$$E_{thr}^{true} = E_{thr}^{nom} (1 \pm \Delta\beta). \quad (10.4)$$

Our measurement of flux also has a factor of energy and thus our measurement of the flux at the energy threshold will also pick up a systematic error that goes like $(1 \pm \Delta\beta)$. The true flux measurement is therefore related to our nominal flux measurement as

$$E^2 \frac{dN^{true}}{dE} = E^2 \frac{dN^{nom}}{dE} (1 \pm \Delta\beta). \quad (10.5)$$

In order to understand our systematic error on the energy threshold and flux, we therefore sum in quadrature all the individual uncertainties in our throughput in order to calculate $\Delta\beta$. There are many sources of these uncertainties. The following are the factors which we believe contribute significantly to the overall systematic error:

1. **Cherenkov production:** Our simulation of the expected Cherenkov photon production may not be correct. In order to estimate this source of error, we compared two different EAS simulation: CORSIKA and MOCCA. The Cherenkov production differed by 5% [59]. The estimated error is thus 5%.
2. **Atmospheric Attenuation:** As noted in Section 6.1.1 there is a 15% difference in the expected atmospheric attenuation between summer and winter. Since we do not account for this we conservatively assign a 15% error.
3. **Heliostat Reflectivity:** Measurements were made of the reflectivity of the heliostat mirrors. The systematic error on this measurement is estimated to be 10% [111].
4. **Heliostat Pointing:** Comparison of the expected and actual size and position of heliostat images of the sun on the solar tower allow us to estimate the error on

heliostat pointing to be 5% [60].

5. **Secondary Reflectivity:** The secondary reflectivities are measured monthly; the error is estimated to be 3%.
6. **Secondary Pointing:** The pointing of the secondary mirrors is periodically checked with *cameraspots*; these tests allow us to estimate the error from secondary pointing to be 5%.
7. **PMT Q.E. and Gains:** The error on PMT quantum efficiency is estimated to be 10% [111]. The PMT gains are monitored extensively, as noted in Section 4.6.2; the error is estimated to be 10%. The total error is therefore 14%.

These factors are all summarized in Table 10.2. Assuming that the factors are uncorrelated, we calculate that our overall energy scale uncertainty, $\Delta\beta$, is 25%. An additional check of this estimate is possible using tests from our *drift scans*. The drift scans allow us to calculate the PMT currents that result from a bright star passing through our field of view. It is also quite easy to calculate the predicted current increase. Tests of this type have found that the real and predicted currents differ on average by 20% [83]. This is quite similar to the estimate of the throughput error calculated from each individual element. The drift scan test does not probe all the sources of systematic error listed above. Nevertheless, this test does suggest that the scale of the problem calculated above is approximately correct. We shall set $\Delta\beta = 25\%$ as our uncertainty in the STACEE energy scale.

Our measurement of the spectral energy density at the energy threshold therefore has both statistical and systematic errors. A summary of our simple flux measurements for both years of Mrk 421 data is shown in Table 10.3. It is worth emphasizing that the large systematic error on the energy scale does not affect the statistical significance of our detection. So while the measurement of the spectral energy density for Mrk 421 during 2003-2004 includes the 25% systematic error, our measurement of the final rate ($2.88 \pm 0.27_{stat}$ counts min^{-1}) does not. We believe that the systematic error on the final rate is much smaller than the statistical error, so we are still correct to say that we have detected Mrk 421 with a significance of 10.9σ .

10.1.3 Refined Flux Measurement

We can now turn to our refined flux measurement. This measurement has two steps. The first is to estimate the Mrk 421 spectral index for this data set using the spectral reconstruction techniques devised and tested in Chapter 9. The second step is to calculate the

Season	Final Rate (counts min ⁻¹)	E_{thr} (GeV)	Spectral Energy Density at E_{thr} (10 ⁻⁴ GeV m ⁻² s ⁻¹)
2002-2003	$1.01 \pm 0.26_{stat}$	$185 \pm 46_{sys}$	$3.9 \pm 1.0_{stat} \pm 1.0_{sys}$
2003-2004	$2.88 \pm 0.27_{stat}$	$185 \pm 46_{sys}$	$10.8 \pm 1.0_{stat} \pm 2.7_{sys}$

Table 10.3: Mrk 421 2002-2004 simple flux measurements.

flux normalization using this spectral index. The flux normalization and the spectral index together will give us our refined measurement of the Mrk 421 spectral energy distribution. As with our analysis of the Crab, the energy reconstruction will only be applied to the data set that had paracanted heliostats, namely the data taken in 2003-2004.

Spectral Measurement

The Template Fitting part of the spectral reconstruction was performed in the same manner as described in Chapter 9. A different template file was generated for Mrk 421, since the source is at a different declination than the Crab. The predicted energy resolution was very similar to what was found for the Crab. As before, our γ -ray simulations were originally done using an $\alpha_{org} = 2.4$ power law distribution. We then created synthetic distributions of E_{rec} for other spectral indices by re-weighting the original γ -ray simulation.

Figure 10.10 shows the difference between the ON and OFF distributions of E_{rec} . Given the relatively large current difference between the ON and OFF Mrk 421 fields, it is important that the distributions have been corrected for the differences between ON and OFF livetimes. As before, the essence of the Forward Folding is comparing this distribution to the distributions from simulated γ -rays. Figure 10.10 also shows the simulations whose spectral index results in the best match with the real data; in this case the best spectral index is near $\alpha = 2.2$. Figure 10.11 shows the distribution of reduced χ^2 from the real/simulation comparisons as a function of the spectral index of the synthetic γ -ray simulation. We fit a parabolic curve to the region near the minimum. This analysis results in a measurement of $\alpha_{mrk} = 2.32 \pm 0.13$ for Mrk 421 during the 2003-2004 season. The minimum χ^2 is 9.9 for 6 degrees of freedom, which has only a 13% chance probability. This indicates that the fit of real and simulated data is not entirely satisfactory.

Systematic Error on Spectral Index

Like our measurement of the energy threshold and flux normalization, our measurement of the spectral index also has a systematic error. Unlike the energy threshold and flux normalization, however, the spectral index should not depend strongly on systematic errors in




Figure 10.10: Mrk 421 2003-2004 real and simulated γ -ray reconstructed Energy Distribution. The plot shows the difference between the ON and OFF distributions of E_{rec} , which is the result of Mrk 421 γ -rays, as well as the distribution for the $\alpha = 2.2$ synthetic simulation.




Figure 10.11: Reduced χ^2 for comparisons between real and simulated Mrk 421 2003-2004 distributions of E_{rec} . The reduced χ^2 are plotted versus the spectral index of the γ -ray simulations. The dotted line indicates the χ^2 value at which the error of α is calculated.




Figure 10.12: Alternate spectral reconstruction for Mrk 421. The distribution of χ^2 versus spectral index are shown for two sets of γ -ray simulations; one where the original spectral index was 2.4 (our original set of simulations), one where it was 1.2. The two sets of simulations give slightly different measurements of α_{mrk} .

our energy scale. Given the calorimetric nature of γ -ray EAS, systematic errors in the energy scale only affect the overall expected brightness (number of photoelectrons) of each shower. Suppose, for instance, that what we believe to be a ‘200’ GeV γ -ray may actually be a 180 GeV γ -ray. This mistake only affects the energy threshold and flux normalization. The spectral index depends on the ratio of the number of ‘200’ GeV and ‘300’ GeV γ -rays we detect. This ratio will be the same even if the γ -rays in question were actually 180 and 270 GeV respectively. Our measurement of the spectral index should therefore be insensitive to errors in our energy scale.

Our measurement of the spectral index will be more susceptible to subtle systematic differences between the simulated and real data. For instance, differences between real and simulated pulse shapes and hence saturated pulse fitting might result in different distributions of measured charges for higher energy events. This could in turn lead to a systematic error on the spectral index. Quantifying these differences is, however, difficult.

One obvious test is to redo the spectral reconstruction using a different set of γ -ray simulations. In particular, we shall redo the analysis using a set of simulations with a spectral index of 1.2. Our original simulations used a spectral index of 2.4, which is suspiciously similar to the spectral index that we measured. This raises the concern that our process of creating synthetic distributions is flawed in some manner and simply returns

the spectral index of the original simulation. Redoing the analysis using a $\alpha_{org} = 1.2$ simulation allows us to test this concern.

This alternate set of γ -ray simulations were used to make a second measurement of the spectral index for the full 2003-2004 Mrk 421 data set. The resulting distribution of χ^2 as a function of spectral index is shown in Figure 10.12. Remember, $\alpha_{org} = 1.2$ and $\alpha_{org} = 2.4$ refer to the spectral index of the original simulations, whereas the values on the x-axis refer to the spectral index of the *synthetic* spectrums after re-weighting. We would hope that the reconstructed α value would be the same, no matter what the α_{org} of the original simulations.

Figure 10.12 shows that this is not quite the case. Using the $\alpha_{org} = 1.2$ simulation, we find that the final reconstructed spectral index is $\alpha_{mrk}^{alt} = 2.22 \pm 0.15$. This new measurement of the Mrk 421 spectral index differs from the original one at the 0.1 level. We shall therefore estimate that our measurement of the spectral index has a ± 0.1 systematic error. *Our full measurement of the 2003-2004 Mrk 421 spectral index is therefore* $\alpha_{mrk} = 2.32 \pm 0.13_{stat} \pm 0.10_{sys}$.

As noted, there is a previous STACEE measurement of the Mrk 421 spectral index during 2003-2004, which found that $\alpha_{mrk}^{alt} = 1.83 \pm 0.33_{stat}$ [29]. The approximate agreement between the two measurements is encouraging.

Medium/High States

Although the measurement of the average Mrk 421 spectral index is in itself an achievement, the spectral index of Mrk 421 is known to vary with the flux state (see Section 2.3.2). Both Whipple and HEGRA observed that the spectrum became harder as the overall flux level increased [1, 77]. Indeed this same effect was seen in the 2002-2004 Mrk 421 data by Whipple, albeit at a lower significance [21].

Given that the spectral index of Mrk 421 is known to vary, we should attempt to measure it on the shortest possible timescale. The STACEE sensitivity precludes us from doing the measurement on a daily timescale, but it is possible to do the spectral index measurements for the medium and high state periods. We shall use the same 7 and 8 nights of data (for medium and high periods) that was used in the discussion of correlations with other experiments. We cannot do spectral reconstruction for the low period data, because it was taken in 2002-2003 with monocanted heliostats.

The spectral reconstructions for these two periods are shown in Figure 10.13. For the medium state period the result is a measured spectral index of $\alpha_{medium} = 2.57 \pm 0.25_{stat} \pm 0.10_{sys}$. For the high state period the result is a measured spectral index of $\alpha_{high} = 2.16 \pm 0.17_{stat} \pm 0.10_{sys}$. The minimum χ^2 is found to be 5.9 and 6.8 for the

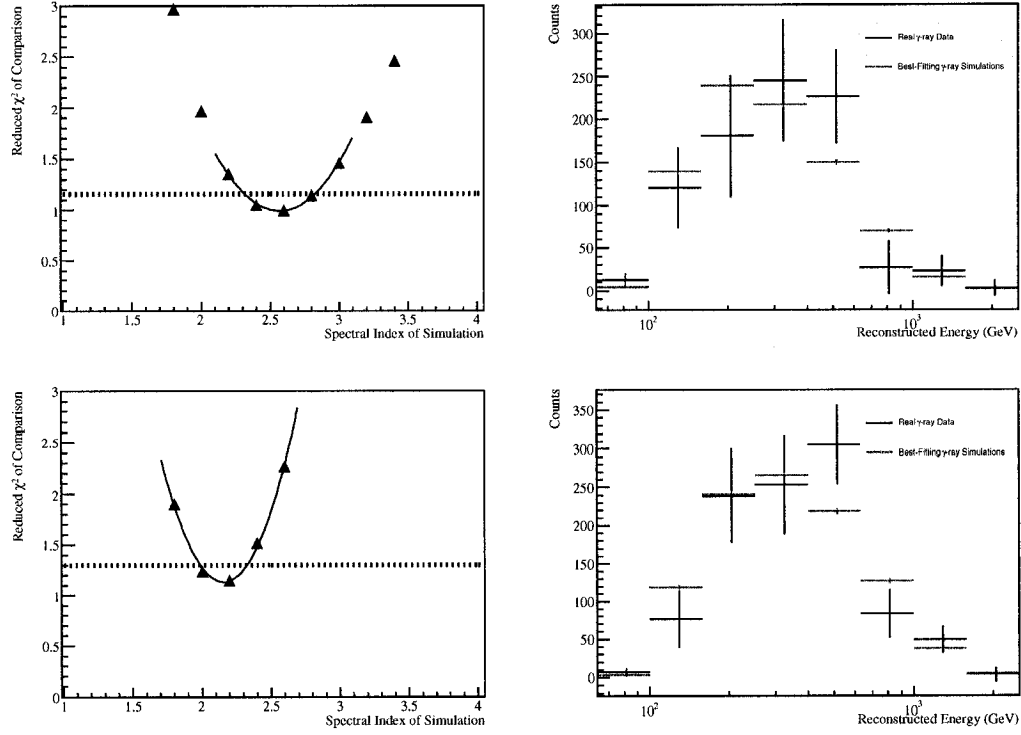


Figure 10.13: Spectral reconstruction for Mrk 421 during medium/high state periods. The left-side plots show the distribution of reduced χ^2 of comparisons as a function of the spectral index of the γ -ray simulations. The right-side plots show the difference between the ON and OFF E_{rec} distributions, as well as the best fitting γ -ray distributions. Upper plots are for medium state period; lower plots are for high state period.

medium and high states, respectively, indicating good fits in both cases.

The STACEE measurements of the spectral index are consistent with the Whipple measurements from the same periods. The Whipple measurements from the medium and high states were $\alpha_{med}^{whip} = 2.40 \pm 0.18$ and $\alpha_{high}^{whip} = 2.11 \pm 0.14$ respectively [21]. The γ -ray SED as measured by STACEE therefore seems to be simply a continuation of the Whipple spectrum. Note that these Whipple results are based on fits that assume an exponential cutoff at 4.9 TeV; this form of the SED was favored by the earlier Whipple work [77]. The STACEE spectral measurement does not use an exponential cutoff, but this has a minimal effect on our results, since we have very limited sensitivity above 1 TeV.

The difference between the STACEE measurements of the medium and high period spectral index is 0.41 ± 0.30 . These results are therefore *consistent* with the phenomenon of spectral hardening; STACEE's measurements of the spectral index do not, however,

constitute *proof* of the phenomenon.

Now that we have our own measurement of the spectral index during 2003-2004 we can compute our own refined flux measurements. We use an identical procedure to Section 10.1.2 for calculating the flux normalization, except using different spectral indices. Our final result for the Mrk 421 flux measurement during the medium state period is therefore

$$E^2 \frac{dN}{dE} = (9.5 \pm 1.5_{stat} \pm 2.4_{sys}) \times 10^{-4} \left(\frac{E}{185 \text{ GeV}} \right)^{-0.57 \pm 0.25_{stat} \pm 0.10_{sys}} \text{ GeV m}^{-2} \text{ s}^{-1}. \quad (10.6)$$

This spectral measurement is for the STACEE energy range, which we define to be approximately 100 - 1000 GeV. The systematic error on the normalization is defined entirely by the 25% systematic error on our energy scale and the systematic error on the spectral index is simply the ± 0.10 determined earlier. In reality, there may be some mixing of the two sources of systematic error, but the difference from what is present in Equation 10.6 should be fairly small.

Likewise, the final result for the Mrk 421 flux during the high state period is

$$E^2 \frac{dN}{dE} = (11.0 \pm 1.2_{stat} \pm 2.8_{sys}) \times 10^{-4} \left(\frac{E}{185 \text{ GeV}} \right)^{-0.16 \pm 0.17_{stat} \pm 0.10_{sys}} \text{ GeV m}^{-2} \text{ s}^{-1}. \quad (10.7)$$

10.1.4 Discussion

The final medium and high state flux measurements are shown in Figure 10.14 in the context of other high energy γ -ray results. The STACEE results are presented as ‘butterfly’ plots. The butterfly plot is defined by the four curves generated by taking Equations 10.6 and 10.7 and varying both the normalization and the spectral index by their 1σ errors. It should be noted that we use only the statistical errors when generating the butterfly plots.

Figure 10.14 also shows the contemporaneous Whipple flux measurements of Mrk 421, as well as archival EGRET results from 1992-1995 [21, 55]. Given the variability of this source, the EGRET points serve only to give a general sense of the inverse Compton peak and should not be used to do exact extrapolations. The comparison between the contemporaneous Whipple and STACEE results is more interesting. The Whipple and STACEE results for the medium state period agree very well; the STACEE SED seem to be a simple extrapolation of the Whipple SED. This suggests both that our analysis method is robust and that the relative energy calibration of STACEE and Whipple is well understood. The medium state results are therefore very encouraging.




Figure 10.14: Mrk 421 spectral energy distribution. The STACEE butterflies are shown for the medium and high periods, as well as contemporaneous Whipple data. The EGRET data is archival and is only provided to give a general sense of the Compton peak.

For the high state period, however, the Whipple fluxes are significantly higher than the STACEE fluxes. A possible explanation for the difference in high state fluxes is that the measurements were not exactly contemporaneous. The Whipple observations were consistently a few hours after the STACEE observations. In addition, the choice of which nights were included in the Whipple high state SED measurement was different from STACEE. For STACEE we simply used data from any night in April 2004 where there were observations from STACEE, Whipple and the ASM. Whipple used a more restrictive condition, only using nights where the PCA rates were above a certain level. This requirement may have resulted in the Whipple high state measurements being taken during a period when Mrk 421 was systematically in a higher state than the STACEE high state period, despite both measurements being contained within the same month. We should therefore not be too concerned by the apparent discrepancy of the high state measurements.

Overall, the good correlation between the Whipple and STACEE measured rates, as




Figure 10.15: 3C 66A 2003-2004 ξ_{cor} distributions. The left-side plot shows the distributions of ξ_{cor} for the ON and OFF data; the distributions are corrected for the small live-time differences between the ON and OFF data. The right-side plot shows the difference between the ON and OFF distributions. The dashed red line shows our cut value.

well as the agreement between the measured spectral indices, strongly suggests that the Mrk 421 SED in the STACEE energy range is a continuation of the SED in the Whipple energy range. The lack of change in the SED suggests that the peak of Mrk 421's Compton distribution is below the STACEE energy range. *The STACEE results are therefore consistent with the 2003-2004 Mrk 421 Compton peak being at approximately 100 GeV.*

10.2 3C 66A Observations and Results

STACEE observed 3C 66A from September to December 2003; the full data set comprised 87 ON/OFF pairs. We used the standard data quality cuts to remove periods where the detector malfunctioned or the weather was poor. 31% of the data was removed, resulting in a final ON source livetime of 83 ks. We used monocanted heliostats for the duration of this dataset. The STACEE observations were a part of a multi-wavelength 3C 66A observation campaign, which included instruments from radio to VHE γ -rays. Many of the results of this campaign have already been summarized in Böttcher et al. (2005) [24] and initial STACEE measurements using this data set were presented in Bramel et al. (2005) [28].

Figure 10.15 shows the distributions of ξ_{cor} for the 3C 66A data set, after the padding cut has been applied. There is no evidence for an ON-source excess. Our standard cut of $\xi_{cor} < 0.35$ results in a final γ -ray rate of -0.35 ± 0.22 counts min^{-1} , with a final




Figure 10.16: Distribution of significances for each 3C 66A pair. The significance for each pair is defined in terms of number of σ . Each point is weighted by the livetime; so a pair that has a full 28 minutes of data has twice the weighting of a pair with 14 minutes of data. A Gaussian function has been fit to the data.

significance of -1.6σ . We shall therefore be interpreting our result in the context of an upper limit.

As an additional test of the data quality, Figure 10.16 shows a histogram of the detection significance for each of the pairs of 3C 66A data. The distribution of significances is well described as a Gaussian centered at 0.0 with a width of 1.0. This is the expected distribution of significances for a non-detection for a well-behaved detector where the data quality issues have been correctly accounted for. The figure also demonstrates that there are no single pairs that are significantly affecting the mean value.⁴

10.2.1 3C 66A Rate Upper Limit

Our final 3C 66A rate measurement is consistent with zero. We would therefore like to calculate an upper limit on the final rate and hence a flux upper limit. To do this we use the “bounded upper limit method” [13]. The procedure is as follows. We assume that our probability density distribution is Gaussian; but we bound the distribution to the range $[0, \infty]$, since negative fluxes are unphysical. We must therefore re-normalize our Gaussian distribution such that

⁴We did not present a similar plot for the Mrk 421 since a source with a detectable and varying flux is not expected to have a Gaussian pair-wise significance distribution with a width of 1.0.




Figure 10.17: The left-side shows the HA-weighted final effective area curve for 3C 66A after all event cuts. The right-side shows the response curve, after convolution with a $\alpha = 2.5$ spectrum.

$$\int_0^{\infty} G(x; R, \sigma_R) dx = 1, \quad (10.8)$$

where G is our Gaussian distribution for our final rate measurement and rate error, R and σ_R . Our rate upper limit, R_{UL} , is then defined as the value at which

$$C = \frac{\int_0^{R_{UL}} G(x; R, \sigma_R) dx}{\int_0^{\infty} G(x; R, \sigma_R) dx}, \quad (10.9)$$

where C is equal to the desired confidence level. In our case we wish to calculate 99% CL upper limits, so we set $C = 0.99$. We use our measured rate of -0.35 ± 0.22 counts min^{-1} to solve Equation 10.9 and find that our 99% CL final rate upper limit is 0.37 counts min^{-1} .

10.2.2 3C 66A Energy Threshold and Flux Upper Limit

Our process for calculating the flux upper limit for 3C 66A is identical to the procedure described in Section 10.1.2, except that in this case we have a final rate upper limit of 0.37 counts min^{-1} , instead of an actual rate measurement. The final effective area for the 3C 66A data set is shown on the left-side of Figure 10.17. As before, we convolve this

Table 10.4: Summary of 3C 66A 2003-2004 final results.

ON-source livetime	83.2 ks
Final γ -ray rate	-0.35 ± 0.22 counts min $^{-1}$
Detection significance	-1.6σ
Rate Upper limit (99% CL)	0.37 counts min $^{-1}$
Energy Threshold	$185 \pm 46_{stat}$ GeV
Flux Upper limit (99% CL) at E_{thr}	1.1×10^{-4} GeV m $^{-2}$ s $^{-1}$
C&G Flux Prediction at E_{thr}	0.13×10^{-4} GeV m $^{-2}$ s $^{-1}$

curve with a power law of the form

$$dN/dE = N_0(E/E_{thr})^{-2.5}. \quad (10.10)$$

The result is our response curve, which is shown in the right-side plot of Figure 10.17. We find that our energy threshold is 185 GeV for this data set.

For 3C 66A the integral of the response curve must match our final rate upper limit; this condition allows us to solve for the normalization constant N_0 . We thereby calculate that the 99% CL upper limit on the flux at E_{thr} is therefore

$$E^2 \frac{dN}{dE}^{3C\ 66A} (185 \pm 46_{sys} \text{ GeV}) < 1.1 \times 10^{-4} \text{ GeV m}^{-2} \text{ s}^{-1}.$$

Alternately, our 3C 66A integral flux upper limit is 15% of the Crab Nebula integral flux above the same energy, using previous STACEE measurements of the Crab [44]. A summary of 3C 66A results is shown in Table 10.4.

10.2.3 3C 66A Discussion

Figure 10.18 shows the spectral energy distribution of 3C 66A in the γ -ray regime, including the new STACEE 2003 flux upper limit. The plot also shows the original STACEE upper limit for the same data set [27, 28], as well as contemporaneous results from the Whipple telescope [24]. In addition, the plot shows archival results from EGRET, Whipple and HEGRA [2, 55, 62]. With the exception of the EGRET points, all the results are upper limits. It should also be remembered that the EGRET data is actually for the source 3EG 0222+4253, which is probably the result of both 3C 66A and a pulsar. The true HE γ -ray flux from 3C 66A is probably weaker, but also harder.

The new STACEE upper limit is noticeably more constraining than the original STACEE measurement. This is expected, since the original work did not have the added cosmic-ray




Figure 10.18: Spectral Energy Distribution for 3C 66A. The plot shows STACEE 99% CL upper limit, including the 25% systematic error on the energy threshold. The plot also shows upper limits from other VHE experiments, as well as archival EGRET data. The dashed line marks the level of the C&G prediction.

suppression of the Grid Alignment technique. The new STACEE result is less constraining than the IACT measurements for the purpose of evaluating a simple extrapolation of the EGRET curve. Given, however, the high redshift of this source there may be a large amount of EBL absorption for the higher energy γ -rays that Whipple and HEGRA observed. The effect of EBL absorption will be less for the STACEE measurements with their lower energy threshold. Our measurements should therefore be useful for constraining certain models of 3C 66A emission.

As an example, we can compare our measurements to the predictions made in Costamante and Ghisellini (C&G; 2002) [32]. As noted in Chapter 2, STACEE chose to observe 3C 66A and OJ 287 partly because they appeared on the C&G list of potential VHE emitters; the paper included predictions for flux levels of VHE γ -rays. In particular, we use the predictions based on a phenomenological model of blazar emission presented in Fossati et al. (1998) [45], which was modified for C&G. It should be noted that this model does




Figure 10.19: OJ 287 2003-2004 ξ_{cor} distributions. The left-side plot shows the distributions of ξ_{cor} for the ON and OFF data; the distributions are corrected for the small live-time differences between the ON and OFF data. The right-side plot shows the difference between the ON and OFF distributions. Distributions are after the padding cut has been applied.

not account for absorption from the EBL, since part of the goal of the model was to allow measurement of the EBL.

The predictions in the C&G paper are presented in terms of integral fluxes above 40, 300 and 1000 GeV. In order to convert these predictions into differential flux measurements at 185 GeV we use a power law interpolation between the 40 GeV and the 300 GeV predictions. The flux prediction at 185 GeV for 3C 66A is

$$E^2 \frac{dN^{pred}}{dE} (185 \text{ GeV}) = 0.13 \times 10^{-4} \text{ GeV m}^{-1} \text{ s}^{-1}.$$

The predicted flux is significantly below our measured upper limit, meaning that we can neither confirm nor reject this particular model. This is disappointing because including EBL absorption in the models would decrease the predicted flux even further. The C&G prediction is shown in Figure 10.18. It can be seen that the predicted flux level is below the upper limits of the other VHE experiments as well. As always, however, predictions of BL Lac emission are somewhat tentative; even if the prediction is a correct description on average, the variability of a BL Lac object means that the instantaneous flux could be much larger.




Figure 10.20: Distribution of significances for each OJ 287 pair. The significance for each pair is defined in terms of number of σ . Each point is weighted by the livetime; so a pair that has a full 28 minutes of data has twice the weighting of a pair with 14 minutes of data. A Gaussian function has been fit to the data.

10.3 OJ 287 Observations and Results

STACEE observed OJ 287 from December 2003 to February 2004; the full data set comprised 28 ON/OFF pairs. We used the standard data quality cuts to remove periods where the detector malfunctioned or the weather was poor. 48% of the data was removed, resulting in a final ON source livetime of 21 ks. This is the smallest of the AGN data sets presented in this work. We used paracanted heliostats for the duration of this data set.

Figure 10.19 shows the distributions of ξ_{cor} for the OJ 287 data set. There is no evidence of an ON-source excess. Our standard cut of ξ_{cor} results in a final γ -ray rate of 0.35 ± 0.39 counts min^{-1} , with a final significance of 0.9σ . Using the same procedure as before we calculate a 99% CL upper limit on the final γ -ray rate as being 1.29 counts min^{-1} . As with the 3C 66A results, Figure 10.20 shows the distributions of significances for each pair. Once again, the distribution is consistent with the expectations for a non-detection. This provides additional confidence in the quality of this data set.




Figure 10.21: The left-side shows the HA-weighted final effective area curve for OJ 287 after all event cuts. The right-side shows the response curve, after convolution with a $\alpha = 2.5$ spectrum.

We use an identical process to that of the 3C 66A analysis to determine the energy threshold and flux upper limit. The final effective area and response curve are shown in Figure 10.21. With the assumption of a $\alpha = 2.5$ spectrum, we find that our energy threshold is 145 GeV. We compute the 99% CL upper limit on the flux at the energy threshold as being

$$E^2 \frac{dN}{dE}^{OJ287} (145 \pm 36_{sys} \text{ GeV}) < 4.0 \times 10^{-4} \text{ GeV m}^{-2} \text{ s}^{-1}.$$

As before, we can express this in terms of the Crab; our OJ 287 integral flux upper limit is 52% of the Crab Nebula integral flux above the same energy. The OJ 287 results are summarized in Table 10.5.

Table 10.5: Summary of OJ 287 2003-2004 final results.

ON-source livetime	21.1 ks
Final γ -ray rate	$0.35 \pm 0.39 \text{ counts min}^{-1}$
Detection significance	0.9σ
Rate Upper limit (99% CL)	$1.29 \text{ counts min}^{-1}$
Energy Threshold	$145 \pm 36_{stat} \text{ GeV}$
Flux Upper limit (99% CL) at E_{thr}	$4.0 \times 10^{-4} \text{ GeV m}^{-2} \text{ s}^{-1}$
C&G Flux Prediction at E_{thr}	$0.49 \times 10^{-4} \text{ GeV m}^{-2} \text{ s}^{-1}$




Figure 10.22: Spectral Energy Distribution for OJ 287. The plot shows STACEE 99% CL upper limit, including the 25% systematic error on the energy threshold. The solid lines denote archival EGRET data and dashed lines denote the level of C&G prediction.

10.3.1 Discussion

As with 3C 66A, we can compare our flux upper limit to the flux predictions in C&G, again doing a power-law interpolation between the 40 GeV and 300 GeV integrated fluxes. The flux prediction for OJ 287 at 145 GeV is

$$E^2 \frac{dN^{pred}}{dE} (145 \text{ GeV}) = 0.49 \times 10^{-4} \text{ GeV m}^{-1} \text{ s}^{-1}.$$

As before, this is noticeably below our upper limit. We cannot therefore say anything meaningful about this particular model. In any case, C&G note that this model appears to be a poor match in the case of OJ 287; rejecting the model would therefore not be very meaningful.

Figure 10.22 shows the spectral energy distribution of OJ 287 in the γ -ray regime. The plot shows the EGRET results, which are not contemporaneous, as well as the C&G predicted flux at 145 GeV [32, 55]. The STACEE upper limit is the only measurement of

this source in the VHE γ -ray regime and may therefore be of significant importance to future modelling of this source. This upper limit may also be of importance if the apparent 12 year optical periodicity of OJ 287 is found to be of relevance in the γ -ray regime. In that case, long term observations, including the STACEE upper limit, may be useful for constraining periodicity models.

Chapter 11

Conclusion

This completes our study of BL Lac objects with the STACEE γ -ray detector. As is natural with a new and complicated field, we have spent a considerable time describing the physics of extensive air showers and the details of the STACEE detector. We have also described in detail the STACEE analysis technique, emphasizing in particular the improvements that have resulted from the installation of our FADCs. The first improvement was the development of the Library Padding technique to correct for promotion, which was a systematic bias of our measured γ -ray rate. The second improvement was the development of the Grid Alignment technique. This technique allowed for a significant suppression of our cosmic-ray background; the result was an increase in sensitivity of approximately 70%. The final improvement was the implementation of a spectral reconstruction technique. All these improvements were extensively tested on data taken on the Crab Nebula, as well as a series of star data sets. We therefore have confidence in all three techniques. The results are the culmination of a considerable amount of work by the STACEE collaboration.

These techniques have been applied to the data taken on three BL Lac objects during the 2002-2004 observing seasons. The first of these objects was Mrk 421, which is normally the brightest AGN in the γ -ray sky. We detect the source at the 3.9σ level in 2002-2003 and 10.9σ level in 2003-2004. The later detection is the strongest ever detection of a source by STACEE during a single year. This detection allowed for more interesting investigation of Mrk 421 behavior. Comparisons of STACEE's measured Mrk 421 rates with those measured by the Whipple experiment showed a strong degree of correlation between the results of both experiments. A similar result was found when comparing STACEE results to those of RXTE, though the correlation was less pronounced. In addition, we measured the spectral index of Mrk 421 to be $\alpha_{mrk} = 2.32 \pm 0.13_{stat} \pm 0.10_{sys}$ in the 100 - 1000 GeV energy range for the 2003-2004 season. Our measurement is con-

sistent with an extrapolation of Whipple's Mrk 421 spectrum for the same season. Our results suggest that the Compton peak was located slightly below the STACEE energy range, probably near 100 GeV.

The other two sources, 3C 66A and OJ 287, were promising LBL candidates for VHE γ -ray emission. Unfortunately, we did not detect either object. Our upper limits may, however, be of considerable interest in future modelling of these sources, particularly for OJ 287. Given OJ 287's possible 12 year periodicity, extensive observations of the VHE emission of this object are important. Indeed, preliminary evidence indicates that the current optical outburst has begun earlier than expected, in November 2005 [145]. Given the uncertainty regarding the binary supermassive blackhole model for this source, it is important that as many VHE observations as possible be made of OJ 287, particularly over the next two years.

We expect many more interesting VHE γ -ray discoveries over the following years. The STACEE experiment will continue observations for at least another year, concentrating on the study of gamma-ray bursts. The future of VHE astronomy, however, is the next generation of imaging atmospheric Cherenkov telescopes, such as HESS, VERITAS and MAGIC. These experiments use arrays of Whipple-like telescopes to take stereoscopic images of the extensive air showers. The resulting increase in sensitivity has already profoundly affected VHE astronomy. The list of sources detected at VHE energies has been increasing rapidly since HESS started observing three years ago; the number of different *types* of sources has also been increasing. In the area of VHE AGN studies, the new detectors have resulted in a doubling of the number of detected BL Lac objects. Even more importantly, detections have been made of BL Lacs with higher and higher redshifts. Already these results have led to interesting work on EBL absorption [5]. More discoveries are being made every day.

In addition to advances in ground-based γ -ray astronomy, we will also soon see the launch of the GLAST telescope. Assembly of the LAT has been completed; the instrument is currently undergoing stress testing at the Naval Research Laboratory. Things are on schedule for a launch late in 2007. GLAST, with its larger area than EGRET and superior direction reconstruction, is expected to detect thousands of new γ -ray sources. With the launch of GLAST the gap between ground-based and space-based γ -ray detectors should finally be closed. This is important for AGN research, since it will result in unambiguous measurements of the Compton peak for sources such as Markarian 421. In addition, the overlap will also lead to a significant reduction in the systematic error in the energy scale of ground-based γ -ray telescopes. The energy scale of the LAT instrument is well understood, since it has undergone extensive calibrated beam testing [33].

More exotic is the possibility of future detection of neutrinos from AGN. The production of large neutrino fluxes is generally predicted for the hadronic class of BL Lac emission models. Conclusive detection of neutrinos from AGN would therefore be interesting, as well as being extremely complementary to VHE studies. There is the possibility that the current generation of high energy neutrino telescopes might already be sensitive enough to detect the neutrino flux from BL Lacs [50]. If this turns out not to be the case, then the next generation of km-scale neutrino detectors, such as ICECUBE, will have an even better chance of detecting AGN neutrinos [10].

In addition to new instruments, other types of astronomical objects also hold the potential for improving our understanding of AGN. In the past fifteen years, it has been found that relativistic jets also occur from the smaller, stellar mass blackholes that occur in our own galaxy. These systems, which have been dubbed ‘microquasars’, are believed to be scaled down versions of AGN. They contain the blackhole, the accretion disk and the relativistic jet. One of these microquasars, LS 5039, has also been discovered to be an emitter of VHE γ -rays by the HESS collaboration [3]; another has been detected by the MAGIC collaboration [7]. The detection of VHE γ -rays from microquasars suggests that similar types of processes are occurring on systems that differ in scale by greater than six orders of magnitude. Microquasars therefore may provide a different way of studying the processes that occur in AGN, in objects that are far closer to us.

For all these reasons the next ten years of AGN research promise to be very exciting. New instruments and longer observations of AGN and BL Lacs will lead to the answering of old questions and to the discovery of new questions.

Appendix A

Correlation Coefficient

The goal of calculating the correlation coefficient is to measure the degree of the linear correlation between two quantities, x and y . We will be using the *Pearson correlation coefficient*, r , which is numerically defined as

$$r = \frac{\sum_i^N (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_i^N (x_i - \bar{x})^2 \sum_i^N (y_i - \bar{y})^2}}, \quad (\text{A.1})$$

where N is the number of samples in the data set and \bar{x} and \bar{y} are the averages of the x and y values. The value of this quantity ranges between -1 and 1. $r = 1$ indicates a perfect positive correlation, $r = -1$ indicates a perfect negative correlation and $r = 0$ indicates a perfect lack of correlation. Examples of the correlation coefficient for different distributions is shown in Figure A.1. The figure shows the range of different behaviour from no correlation to perfect positive correlation.

A.1 Calculation of Error on r

The calculation of the correlation coefficient is straight-forward. A more difficult problem is determining the error on r , as well as determining the statistical significance of the correlation coefficient. The correlation between two quantities can be highly statistically significant, even if the degree of correlation is relatively low.

One method of calculating the error and significance is to use *Fisher's z-transformation*. This method assumes that the x and y distribution jointly form a *binormal* or *two-dimensional Gaussian* distribution around their mean values. The method can be used if the actual x and y distribution is only approximately binormal, which is probably the case for the distribution analyzed in this work. Nevertheless, if the actual distribution is significantly different from binormal, then the following treatment will be incorrect.

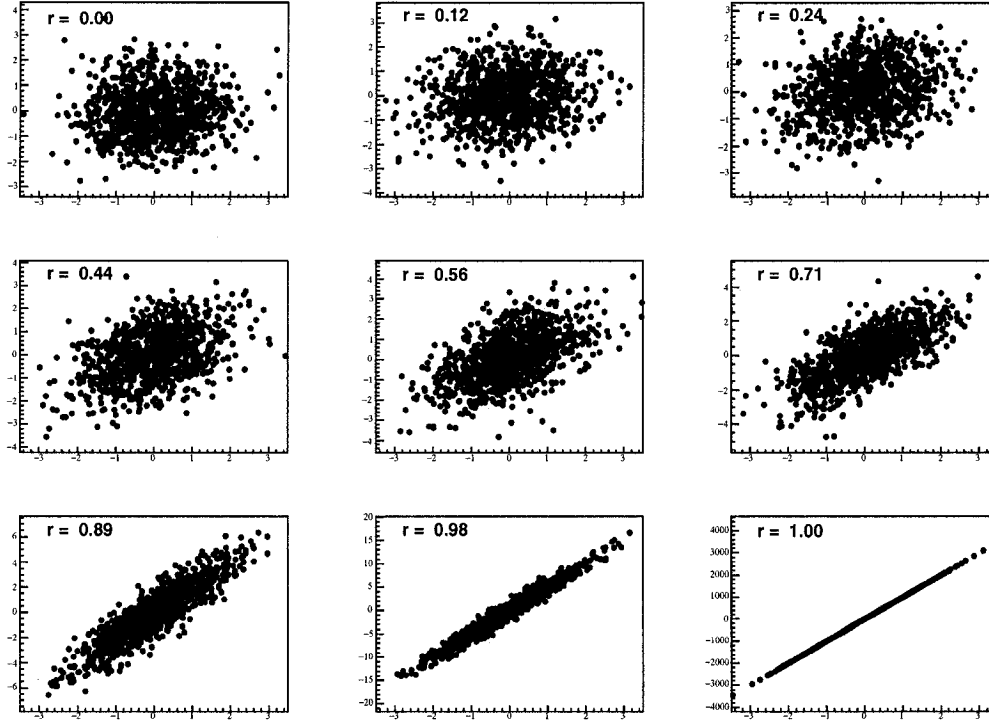


Figure A.1: Series of 1000 point distributions. The correlation coefficient, r , is shown for each distribution; the values of r range from 0 to 1.

If the x and y distribution is binormal, then we can associate our measured r with a value z such that

$$z = \frac{1}{2} \ln \left(\frac{1+r}{1-r} \right). \quad (\text{A.2})$$

The quantity z has a Gaussian distribution, with a standard deviation of

$$\sigma_z = \frac{1}{\sqrt{N-3}}. \quad (\text{A.3})$$

where N is the number of samples in the data set. The error on z allows us to calculate the error on r by using the inverse transformation to Equation A.2. It also allows us to calculate the statistical significance of our measured r from zero correlation as being simply z/σ_z . Consider, for instance, the case of the lower-left distribution in Figure A.1; the measured r is 0.89 for 1000 points. A z -transformation therefore shows that $z = 1.42 \pm 0.03$. The full value of the correlation coefficient is therefore $r = 0.89^{+0.006}_{-0.006}$ and the statistical difference from no correlation is 45σ [116, 74].

A.2 Deficiency of Correlation Coefficient

The Pearson correlation coefficient is a measure of how well the variation of one quantity explains the variation of the other quantity. One major drawback of the correlation coefficient, however, is that r is calculated using only the data points and not their errors. So at least part of the reason that the r we calculated in Chapter 10 is lower than 1.0 can be attributed to the known errors on the measured rates. We need to be aware that our known measurement errors mean that our calculated r_{rate} is therefore less than the ‘true’ r_{rate} that we could calculate if our measurement errors were zero. Another way of stating this is that our measured correlation coefficient does not fully use all the information we have available, since it does not account for the known rate errors. This problem also means that the calculation of the correlation coefficient would not treat a point with a large measurement error any differently from a point with a small measurement error.

Glossary

Azimuth (Az):

Angle along the horizon between North and the bearing of the source.

Blazar:

A term uniting two types of γ -ray bright objects: BL Lacs and FSRQs.

Breakpoint:

The nominal discriminator threshold above which our L2 trigger rate is dominated by Cherenkov events.

Cameraspots:

A test of the secondary mirror alignment. The test is done by projecting an image of the full moon on the camera plane. The image position is compared to the actual PMT position. See Section 4.4.2.

Canting point:

The canting point is a point at a fixed number of radiation lengths into the atmosphere where the shower maximum would occur for a 100 GeV γ -ray that landed at the center of the heliostat field. The canting point is also the point that the majority or all of the heliostats are tracking.

Compton Peak:

The high energy peak of the spectral energy distribution of BL Lacs.

Core Position:

The point on the ground where the incident particle would have landed had the atmosphere not intervened.

Corsika:

Package for Monte Carlo simulation of EAS. See Section 6.1.1.

Declination (Dec):

Celestial coordinate of a source. Corresponds to latitude projected on the sky.

Drift Scan:

A test of the heliostat alignment. The test by allowing a bright star to drift through the field of view of the heliostats; monitoring the PMT currents allows us to check the heliostat alignment. See Section 4.4.1.

Elec:

Package for Monte Carlo simulation of the electronic part of STACEE detector. See Section 6.1.3.

Elevation (El):

The angle from the horizon to a source.

Grid Alignment:

The Grid Alignment technique involves realigning the FADC traces based on a grid of different assumed core positions. The method searches for the position that maximizes the height of the summed FADC trace. The method allows us to calculate a quantity called ξ_{cor} that can be used to distinguish between γ -rays and cosmic-rays. See Chapter 8.

Hour Angle (HA):

Angle between the observer's meridian and the hour circle on which the source lies. Like right ascension, but fixed with respect to the observer.

Library Padding:

Technique used by STACEE to account for the systematic bias caused by promotion. Library Padding involves adding additional noise to FADC traces and re-imposing a trigger condition at a higher threshold. See Chapter 7.

Monocanting:

Canting scheme where all 64 heliostats are pointing at the canting point. See Section 9.1.2.

Paracanting:

Canting scheme where 48 heliostats are pointing at the canting point and 16 heliostats are pointed directly in the source direction. See Section 9.1.2.

Parsec (pc):

Standard astronomical measure of distance. Equals 3.086×10^{16} m.

Pass0:

Program used for performing calibration and basic analysis of STACEE data. See Section 5.2.1.

Pass1:

Program used for performing sophisticated analyses of STACEE data. Elements of Pass1 include Library Padding, Grid Alignment and Template Fitting schemes.

Promotion:

Promotion is when a sub-threshold Cherenkov event triggers the experiment because of the additional fluctuations caused by NSB noise. Promotion causes a systematic bias of the STACEE γ -ray rate. See Chapter 7.

Right Ascension (RA):

Celestial coordinate of a source. Corresponds to longitude projected on the sky.

Sandfield:

Package for ray-tracing simulation of the optical part of STACEE detector. See Section 6.1.2.

Shower Maximum:

Point at which number of particles in Extensive Air Shower is maximized.

Synchrotron Peak:

The low energy peak of the spectral energy distribution of BL Lacs.

Template Fitting:

Template Fitting is the method that STACEE uses to determine the energy and core position for each event. The method is based on a χ^2 comparison between real and predicted charges measured using the FADCs. The predicted charges are defined by a set of templates. See Chapter 9.

Transit Point:

The point in the sky where the elevation of a source is maximized.

Zenith Point:

The point in the sky directly above the observer.

Abbreviations

3EG: Third EGRET Catalogue (of γ -ray sources)

AGN: Active Galactic Nuclei

ASM: All Sky Monitor (on RXTE)

CGRO: Compton Gamma-Ray Observatory

DAQ: Data AcQuisition System

DTIRC: Dielectric Total Internal Reflection Concentrator

EAS: Extensive Air Shower

EBL: Extragalactic Background Light

EGRET: Energetic Gamma-Ray Experiment Telescope (on CGRO)

FADC: Flash Analog-to-Digital Converter

FOV: Field of View

FSRQ: Flat Spectrum Radio Quasar

FWHM: Full Width at Half-Maximum

GLAST: Gamma-ray Large Area Space Telescope

GTC: Global Trigger Condition

HBL: High frequency peaked BL Lac

HE: High Energy

IAC: Imaging Atmospheric Cherenkov Telescope

IR: InfraRed (light)

LAT: Large Area Telescope (on GLAST)

LBL: Low frequency peaked BL Lac

LTC: Local Trigger Condition

MADDOG: McGill Asynchronous Digital Delays for Observations of Gammas

NSB: Night Sky Background

NSTTF: National Solar Thermal Test Facility

PCA: Proportional Counter Array (on RXTE)

PMT: Photomultiplier Tubes

RTV: Room Temperature Vulcanization

RXTE: Rossi X-ray Timing Explorer

SEC: Synchrotron External Compton (emission model)

SED: Spectral Energy Distribution

SSC: Synchrotron Self-Compton (emission model)

STACEE: Solar Tower Atmospheric Cherenkov Effect Experiment

TOF: Time-of-Flight

UV: UltraViolet (light)

VHE: Very High Energy

VLBI: Very Long Baseline Interferometry

Bibliography

- [1] F. Aharonian et al. Variations of the TeV energy spectrum at different flux levels of Mkn 421 observed with the HEGRA system of Cherenkov telescopes. *Astronomy and Astrophysics*, 393:89–99, October 2002.
- [2] F. Aharonian et al. Observations of 54 Active Galactic Nuclei with the HEGRA system of Cherenkov telescopes. *Astronomy and Astrophysics*, 421:529–537, July 2004.
- [3] F. Aharonian et al. Discovery of Very High Energy Gamma Rays Associated with an X-ray Binary. *Science*, 309:746–749, July 2005.
- [4] F. Aharonian et al. Observations of selected AGN with HESS. *Astronomy and Astrophysics*, 441:465–472, October 2005.
- [5] F. Aharonian et al. A low level of extragalactic background light as revealed by γ -rays from blazars. *Nature*, 440:1018–1021, April 2006.
- [6] J. Albert et al. Observations of Mkn 421 with the MAGIC Telescope . *Submitted to Astrophysical Journal*, 2006.
- [7] J. Albert et al. Variable Very-High-Energy Gamma-Ray Emission from the Microquasar LS I +61 303. *Science*, 312:1771–1773, June 2006.
- [8] J. R. P. Angel and H. S. Stockman. Optical and infrared polarization of active extragalactic objects. *Annual Review of Astronomy and Astrophysics*, 18:321–361, 1980.
- [9] F. Arqueros, J. Ballestrin, M. Berenguel, D. M. Borque, E. F. Camacho, M. Diaz, H.-J. Gebauer, R. Enriquez, and R. Plaga. Very high-energy γ -ray observations of the Crab nebula and other potential sources with the GRAAL experiment. *Astroparticle Physics*, 17:293–318, June 2002.

- [10] A. M. Atoyan and C. D. Dermer. Neutrinos and γ -rays of hadronic origin from AGN jets. *New Astronomy Review*, 48:381–386, April 2004.
- [11] W. B. Atwood. Prospects for observing dark-matter remnants with GLAST. *Advances in Space Research*, 37:1862–1867, 2006.
- [12] P. Auger, P. Ehrenfest, R. Maze, J. Daudin, and R. A. Fréon. Extensive Cosmic-Ray Showers. *Reviews of Modern Physics*, 11:288–291, July 1939.
- [13] R. M. Barnett et al. Review of Particle Physics. *Physical Review D*, 54:1–708, July 1996.
- [14] M. Beilicke et al. Observations of the Giant Radio Galaxy M87 at TeV energies with H.E.S.S. In *XXII Texas Symposium on Relativistic Astrophysics*, 2005.
- [15] K. Bernlohr. Impact of atmospheric parameters on the atmospheric Cherenkov technique. *Astroparticle Physics*, 12:255–268, January 2000.
- [16] G. M. Blake. Observations of Extragalactic Radio Sources Having Unusual Spectra. *Astrophysical Letters*, 6:201–, July 1970.
- [17] R. D. Blandford. Relativistic Accretion. In A. J. Sellwood and J. Goodman, editors, *ASP Conf. Ser. 160: Astrophysical Discs - an EC Summer School*, pages 265–, April 1999.
- [18] R. D. Blandford and D. G. Payne. Hydromagnetic flows from accretion discs and the production of radio jets. *Royal Astronomical Society, Monthly Notices*, 199:883–903, June 1982.
- [19] R. D. Blandford and M. J. Rees. Some comments on radiation mechanisms in Lacertids. In A. M. Wolfe, editor, *Pittsburgh Conference on BL Lac Objects, Pittsburgh, Pa., April 24-26, 1978, Proceedings. (A79-30026 11-90) Pittsburgh, Pa., University of Pittsburgh, 1978, p. 328-341; Discussion, p. 341-347. NATO-supported research*, pages 328–341, 1978.
- [20] R. D. Blandford and R. L. Znajek. Electromagnetic extraction of energy from Kerr black holes. *Royal Astronomical Society, Monthly Notices*, 179:433–456, May 1977.
- [21] M. Błażejowski et al. A Multiwavelength View of the TeV Blazar Markarian 421: Correlated Variability, Flaring, and Spectral Evolution. *Astrophysical Journal*, 630:130–141, September 2005.

-
- [22] L. Boone. *Observations of Markarian 421 with the Solar Tower Atmospheric Cherenkov Effect Experiment*. PhD thesis, University of California Santa Cruz, 2002.
- [23] L. M. Boone, J. A. Hinton, D. Bramel, E. Chae, C. E. Covault, P. Fortin, D. M. Gingrich, D. S. Hanna, R. Mukherjee, C. Mueller, R. A. Ong, K. Ragan, R. A. Scalzo, D. R. Schuette, C. G. Théoret, and D. A. Williams. STACEE Observations of Markarian 421 during an Extended Gamma-Ray Outburst. *Astrophysical Journal Letters*, 579:L5–L8, November 2002.
- [24] M. Böttcher. A Hadronic Synchrotron Mirror Model for the “Orphan” TeV Flare in 1ES 1959+650. *Astrophysical Journal*, 621:176–180, March 2005.
- [25] M. Böttcher, R. Mukherjee, and A. Reimer. Predictions of the High-Energy Emission from BL Lacertae Objects: The Case of W Comae. *Astrophysical Journal*, 581:143–154, December 2002.
- [26] D. A. Bramel. Analysis Chain with stoff_pair. STACEE internal note 04-03, 2004.
- [27] D. A. Bramel. *Observations of the BL Lac Object 3C66A with the Solar Tower Atmospheric Cherenkov Effect Experiment*. PhD thesis, Columbia University, January 2005.
- [28] D. A. Bramel, J. Carson, C. E. Covault, P. Fortin, D. M. Gingrich, D. S. Hanna, A. Jarvis, J. Kildea, T. Lindner, R. Mukherjee, C. Mueller, R. A. Ong, K. Ragan, R. A. Scalzo, D. A. Williams, and J. Zweerink. Observations of the BL Lacertae Object 3C 66A with STACEE. *Astrophysical Journal*, 629:108–114, August 2005.
- [29] J. E. Carson. *STACEE Observations of Markarian 421 above 100 GeV and a New Method for High-Energy Spectral Analysis*. PhD thesis, UCLA, 2005.
- [30] M. F. Cawley. The Application of Noise Padding to the Cherenkov Imaging Technique. In R. C. Lamb, editor, *Towards a Major Atmospheric Cherenkov Detector – II for TeV Astro/Particle Physics*, pages 176–+, 1993.
- [31] A. Celotti and R. D. Blandford. On the Formation of Jets. In L. Kaper, E. P. J. van den Heuvel, and P. A. Woudt, editors, *Black Holes in Binaries and Galactic Nuclei*, pages 206–+, 2001.
- [32] L. Costamante and G. Ghisellini. TeV candidate BL Lac objects. *Astronomy and Astrophysics*, 384:56–71, March 2002.

- [33] E. d. Couto e Silva et al. Results from the beam test of the engineering model of the GLAST large area telescope. *Nuclear Instruments and Methods in Physics Research A*, 474:19–37, November 2001.
- [34] W. Cui. X-Ray Flaring Activity of Markarian 421. *Astrophysical Journal*, 605:662–669, April 2004.
- [35] W. Cui, M. Blazejowski, M. Aller, H. Aller, H. Teräsanta, B. Mochejska, P. Boltwood, A. Sadun, M. Böttcher, A. Reimer, and VERITAS Collaboration. News from a Multi-Wavelength Monitoring Campaign on Mrk 421. In F. A. Aharonian, H. J. Völk, and D. Horns, editors, *AIP Conf. Proc. 745: High Energy Gamma-Ray Astronomy*, pages 455–461, February 2005.
- [36] S. Danaher, D. J. Fegan, N. A. Porter, T. C. Weekes, and T. Cole. Possible applications of large solar arrays in astronomy and astrophysics. *Solar Energy*, 28:335–343, 1982.
- [37] C. D. Dermer and R. Schlickeiser. Model for the High-Energy Emission from Blazars. *Astrophysical Journal*, 416:458–+, October 1993.
- [38] R. S. Dixon and J. D. Kraus. A High-Sensitivity 1415 MHz Survey at North Declinations between 19 and 37 degrees. *Astronomical Journal*, 73:381–407, August 1968.
- [39] D. M. Eardley and W. H. Press. Astrophysical processes near black holes. *Annual Review of Astronomy and Astrophysics*, 13:381–422, 1975.
- [40] D. O. Edge, J. R. Shakeshaft, W. B. McAdam, J. E. Baldwin, and S. Archer. A survey of radio sources at a frequency of 159 Mc/s. *MmRAS*, 68:37–60, 1959.
- [41] A. Fassò, A. Ferrari, and P. R. Sala. Electron-Photon Transport in FLUKA: Status. In A. Kling, F. Barão, M. Nakagawa, L. Távora, and P. Vaz, editors, *Advanced Monte Carlo for Radiation Physics, Particle Transport Simulation and Applications. Proceedings of the Monte Carlo 2000 Conference held in Lisbon, 23-26 October, 2000.*, page 159, 2001.
- [42] A. Fassò, A. Ferrari, P. R. Sala, and J. Ranft. FLUKA: Status and Prospects for Hadronic Applications. In A. Kling, F. Barão, M. Nakagawa, L. Távora, and P. Vaz, editors, *Advanced Monte Carlo for Radiation Physics, Particle Transport Simulation and Applications. Proceedings of the Monte Carlo 2000 Conference held in Lisbon, 23-26 October, 2000.*, page 955, 2001.

-
- [43] C. E. Fichtel, R. C. Hartman, D. A. Kniffen, D. J. Thompson, H. Ogelman, M. E. Ozel, T. Tümer, and G. F. Bignami. High-energy gamma-ray results from the second small astronomy satellite. *Astrophysical Journal*, 198:163–182, May 1975.
- [44] P. Fortin. *Observations of the Crab Nebula and Pulsar in the Optical and gamma-ray bands with STACEE*. PhD thesis, McGill University, January 2005.
- [45] G. Fossati, L. Maraschi, A. Celotti, A. Comastri, and G. Ghisellini. A unifying view of the spectral energy distributions of blazars. *Monthly Notices of the Royal Astronomical Society*, 299:433–448, September 1998.
- [46] W. Galbraith and Jelley J. V. Light Pulses from the Night Sky associated with Cosmic Rays. *Nature*, 171:349–350, 1953.
- [47] N. Gehrels and P. Michelson. GLAST: the next-generation high energy gamma-ray astronomy mission. *Astroparticle Physics*, 11:277–282, June 1999.
- [48] A. M. Ghez, S. Salim, S. D. Hornstein, A. Tanner, J. R. Lu, M. Morris, E. E. Becklin, and G. Duchêne. Stellar Orbits around the Galactic Center Black Hole. *Astrophysical Journal*, 620:744–757, February 2005.
- [49] D.E. Groom and et al. Review of Particle Physics. *The European Physical Journal*, C15:1+, 2000.
- [50] F. Halzen and D. Hooper. High energy neutrinos from the TeV Blazar 1ES 1959 + 650. *Astroparticle Physics*, 23:537–542, July 2005.
- [51] D. S. Hanna. Gain Drifts seen with Laser Runs. STACEE internal note 02-02, 2002.
- [52] D. S. Hanna. A New Way to Monitor Phototube Gains in STACEE. STACEE internal note 03-04, 2003.
- [53] D. S. Hanna, D. Bhattacharya, L. M. Boone, M. C. Chantell, Z. Conner, C. E. Covault, M. Dragovan, P. Fortin, D. T. Gregorich, J. A. Hinton, R. Mukherjee, R. A. Ong, S. Oser, K. Ragan, R. A. Scalzo, D. R. Schuette, C. G. Théoret, T. O. Tümer, D. A. Williams, and J. A. Zweerink. The STACEE-32 ground based gamma-ray detector. *Nuclear Instruments and Methods in Physics Research A*, 491:126–151, September 2002.

- [54] D. S. Hanna and R. Mukherjee. The laser calibration system for the STACEE ground-based gamma ray detector. *Nuclear Instruments and Methods in Physics Research A*, 482:271–280, April 2002.
- [55] R. C. Hartman, D. L. Bertsch, S. D. Bloom, A. W. Chen, P. Deines-Jones, J. A. Esposito, C. E. Fichtel, D. P. Friedlander, S. D. Hunter, L. M. McDonald, P. Sreekumar, D. J. Thompson, B. B. Jones, Y. C. Lin, P. F. Michelson, P. L. Nolan, W. F. Tompkins, G. Kanbach, H. A. Mayer-Hasselwander, A. Mücke, M. Pohl, O. Reimer, D. A. Kniffen, E. J. Schneid, C. von Montigny, R. Mukherjee, and B. L. Dingus. The Third EGRET Catalog of High-Energy Gamma-Ray Sources. *Astrophysical Journal Supplement Series*, 123:79–202, July 1999.
- [56] R. C. Hartman et al. Multiepoch Multiwavelength Spectra and Models for Blazar 3C 279. *Astrophysical Journal*, 553:683–694, June 2001.
- [57] D. Heck, G. Schatz, T. Thouw, J. Knapp, and J. N. Capdevielle. CORSIKA: A Monte Carlo Code to Simulate Extensive Air Showers. *Forschungszentrum Karlsruhe - Technik und Umwelt*. FZKA-6019.
- [58] A. M. Hillas, C. W. Akerlof, S. D. Biller, J. H. Buckley, D. A. Carter-Lewis, M. Catanese, M. F. Cawley, D. J. Fegan, J. P. Finley, J. A. Gaidos, F. Krennrich, R. C. Lamb, M. J. Lang, G. Mohanty, M. Punch, P. T. Reynolds, A. J. Rodgers, H. J. Rose, A. C. Rovero, M. S. Schubnell, G. H. Sembroski, G. Vacanti, T. C. Weekes, M. West, and J. Zweerink. The Spectrum of TeV Gamma Rays from the Crab Nebula. *Astrophysical Journal*, 503:744–+, August 1998.
- [59] J. Hinton. A comparison of the MOCCA and CORSIKA air-shower simulation codes. STACEE internal note 00-10, 2000.
- [60] J. Hinton, R. Dib, and R. Scalzo. Measurements of Heliostat Bias Drift and Light Collection using CCD Sunspot Images. STACEE internal note 00-14, 2000.
- [61] P. F. Hopkins, L. Hernquist, P. Martini, T. J. Cox, B. Robertson, T. Di Matteo, and V. Springel. A Physical Model for the Origin of Quasar Lifetimes. *Astrophysical Journal Letters*, 625:L71–L74, June 2005.
- [62] D. Horan et al. Constraints on the Very High Energy Emission from BL Lacertae Objects. *Astrophysical Journal*, 603:51–61, March 2004.
- [63] J. B. Hutchings and S. G. Neff. Optical imaging of QSOs with 0.5 arcsec resolution. *Astronomical Journal*, 104:1–14, July 1992.

-
- [64] J. V. Jelley. *Cerenkov Radiation and its Applications*. Pergamon Press, 1958.
- [65] J. V. Jelley and N. A. Porter. Cherenkov Radiation from the Night Sky, and its Application to Gamma-ray Astronomy. *Quarterly Journal of the Royal Astronomical Society*, 4, 1963.
- [66] N. N. Kalmykov, S. S. Ostapchenko, and A. I. Pavlov. Quark-Gluon-String Model and EAS Simulation Problems at Ultra-High Energies. *Nuclear Physics B Proceedings Supplements*, 52:17–28, February 1997.
- [67] G. Kanbach, D. L. Bertsch, C. E. Fichtel, R. C. Hartman, S. D. Hunter, D. A. Kniffen, B. W. Hughlock, A. Favale, R. Hofstadter, and E. B. Hughes. The project EGRET (Energetic Gamma-Ray Experiment Telescope) on NASA's Gamma-Ray Observatory (GRO). *Space Science Reviews*, 49:69–84, 1988.
- [68] A. K. Kembhavi and J. V. Narlikar. *Quasars and active galactic nuclei : an introduction*. Quasars and active galactic nuclei : an introduction /Ajit K. Kembhavi, Jayant V. Narlikar. Cambridge, U.K. : Cambridge University Press, c1999. ISBN 0521474779., April 1999.
- [69] M. R. Kidger. The 11 Year Period in OJ 287 Revisited: Is It a True Long-Enduring Period? *Astronomical Journal*, 119:2053–2059, May 2000.
- [70] J. Kildea. *Studies of the Crab Nebula and Pulsar at TeV Energies*. PhD thesis, University College Dublin, January 2002.
- [71] J. Kildea. Observations of the Crab Nebula and Pulsar with STACEE. In *29th International Cosmic Ray Conference*, 2005.
- [72] J. Kildea. Shower Reconstruction Techniques for STACEE. In *29th International Cosmic Ray Conference*, 2005.
- [73] A. Konopelko, A. Mastichiadis, J. Kirk, O. C. de Jager, and F. W. Stecker. Modeling the TeV Gamma-Ray Spectra of Two Low-Redshift Active Galactic Nuclei: Markarian 501 and Markarian 421. *Astrophysical Journal*, 597:851–859, November 2003.
- [74] G. A. Korn and M. S. Korn. *Mathematical Handbook for Scientists and Engineers*. McGraw-Hill, 1961.

- [75] H. Krawczynski, P. S. Coppi, and F. Aharonian. Time-dependent modelling of the Markarian 501 X-ray and TeV gamma-ray data taken during 1997 March and April. *Monthly Notices of the Royal Astronomical Society*, 336:721–735, November 2002.
- [76] H. Krawczynski, S. B. Hughes, D. Horan, F. Aharonian, M. F. Aller, H. Aller, P. Boltwood, J. Buckley, P. Coppi, G. Fossati, N. Götting, J. Holder, D. Horns, O. M. Kurtanidze, A. P. Marscher, M. Nikolashvili, R. A. Remillard, A. Sadun, and M. Schröder. Multiwavelength Observations of Strong Flares from the TeV Blazar 1ES 1959+650. *Astrophysical Journal*, 601:151–164, January 2004.
- [77] F. Krennrich et al. Discovery of Spectral Variability of Markarian 421 at TeV Energies. *Astrophysical Journal Letters*, 575:L9–L13, August 2002.
- [78] J. Krolik. *Active Galactic Nuclei*. 2003.
- [79] L. Kuiper, W. Hermsen, F. Verbunt, D. J. Thompson, I. H. Stairs, A. G. Lyne, M. S. Strickman, and G. Cusumano. The likely detection of pulsed high-energy gamma-ray emission from millisecond pulsar PSR J0218+4232. *Astronomy and Astrophysics*, 359:615–626, July 2000.
- [80] K. M. Lanzetta, D. A. Turnshek, and J. Sandoval. Ultraviolet spectra of QSOs, BL Lacertae objects, and Seyfert galaxies. *Astrophysical Journal Supplement Series*, 84:109–184, February 1993.
- [81] R. Le Gallou and the CELESTE Collaboration. 60 GeV observation of Mrk421 using the CELESTE experiment. In Societe Francaise d’Astronomie et d’Astrophysique, editor, *SF2A-2001: Semaine de l’Astrophysique Francaise*, pages 393–+, May 2001.
- [82] T.-P. Li and Y.-Q. Ma. Analysis methods for results in gamma-ray astronomy. *Astrophysical Journal*, 272:317–324, September 1983.
- [83] T. H. Lindner. Study of STACEE throughput using drift scans. STACEE internal note 04-05, 2004.
- [84] T. H. Lindner. Improvements and Calibration of Grid Alignment Scheme. STACEE internal note 06-01, 2006.
- [85] M. Livio, G. I. Ogilvie, and J. E. Pringle. Extracting Energy from Black Holes: The Relative Importance of the Blandford-Znajek Mechanism. *Astrophysical Journal*, 512:100–104, February 1999.

-
- [86] D. Lynden-Bell. Galactic Nuclei as Collapsed Old Quasars. *Nature*, 223:690–+, 1969.
- [87] M. Lyutikov. Role of reconnection in AGN jets. *New Astronomy Review*, 47:513–515, October 2003.
- [88] D. Maccagni, B. Garilli, R. Schild, and M. Tarengi. X-ray/optical brightness trends in 3C 66A. *Astronomy and Astrophysics*, 178:21–24, May 1987.
- [89] K. Mannheim. The proton blazar. *Astronomy and Astrophysics*, 269:67–76, March 1993.
- [90] K. Mannheim and P. L. Biermann. Photomeson production in active galactic nuclei. *Astronomy and Astrophysics*, 221:211–220, September 1989.
- [91] H. Manseri. *Astronomie Gamma au-dessus de 30 GeV. Une Nouvelle Methode D’Identification des Rayons Gamma Cosmiques a partir du Sol avec le Dectecteur CELESTE*. PhD thesis, École Polytechnique, 2004.
- [92] L. Maraschi, G. Ghisellini, and A. Celotti. A jet model for the gamma-ray emitting blazar 3C 279. *Astrophysical Journal Letters*, 397:L5–L9, September 1992.
- [93] B. E. Markarian. Galaxies with an ultraviolet continuum. *Astrofizika*, 3:24–38, 1967.
- [94] B. E. Markarian and V. A. Lipovetskij. Galaxies with ultraviolet continuum. V. *Astrofizika*, 8:89–99, 1972.
- [95] P. Marleau. Observations of Intermediate Energy Gamma-rays at CACTUS. In *Towards a Network of Atmospheric Cherenkov Detectors VII*, Palaiseau., 2005.
- [96] J. P. Martin and K. J. Ragan. A programmable nanosecond digital delay and trigger system. *IEEE Nuclear Science Symposium*, 2000.
- [97] P. F. Michelson. GLAST: a detector for high-energy gamma rays. In B. D. Ramsey and T. A. Parnell, editors, *Proc. SPIE Vol. 2806, p. 31-40, Gamma-Ray and Cosmic-Ray Detectors, Techniques, and Missions*, Brian D. Ramsey; Thomas A. Parnell; Eds., pages 31–40, October 1996.
- [98] J. S. Miller, H. B. French, and S. A. Hawley. Optical spectra of BL Lacertae objects. In A. M. Wolfe, editor, *Pittsburgh Conference on BL Lac Objects*, Pittsburgh, Pa., April 24-26, 1978, *Proceedings. (A79-30026 11-90)* Pittsburgh, Pa.,

- University of Pittsburgh, 1978, p. 176-187; Discussion, p. 187-191., pages 176–187, 1978.*
- [99] M. Miyoshi, J. Moran, J. Herrnstein, L. Greenhill, N. Nakai, P. Diamond, and M. Inoue. Evidence for a Black-Hole from High Rotation Velocities in a Sub-Parsec Region of NGC4258. *Nature*, 373:127–+, January 1995.
- [100] G. Mohanty, S. Biller, D. A. Carter-Lewis, D. J. Fegan, A. M. Hillas, R. C. Lamb, T. C. Weekes, M. West, and J. Zweerink. Measurement of TeV gamma-ray spectra with the Cherenkov imaging technique. *Astroparticle Physics*, 9:15–43, June 1998.
- [101] A. Mücke, R. J. Protheroe, R. Engel, J. P. Rachen, and T. Stanev. BL Lac objects in the synchrotron proton blazar model. *Astroparticle Physics*, 18:593–613, March 2003.
- [102] C. Mueller. CORSIKA for STACEE. STACEE internal note 05-02, 2005.
- [103] S. L. Mufson, D. J. Hutter, Y. Kondo, C. M. Urry, and W. Z. Wisniewski. An evolving relativistic jet model for the BL Lacertae object Markarian 421. *Astrophysical Journal*, 354:116–123, May 1990.
- [104] W. R. Nelson, H. Hirayama, and D. W. O. Rogers. The EGS4 Code System. SLAC-265, 1985.
- [105] X.. Ning, R. Winston, and J. O’Gallagher. Dielectric totally internally reflecting concentrators. *Applied Optics*, 26:300–305, 1987.
- [106] K. J. E. Northover. The radio galaxy 3C 66. *Monthly Notices of the Royal Astronomical Society*, 165:369–+, 1973.
- [107] R. A. Ong. Very high-energy gamma-ray astronomy. *Physics Reports*, 305:93–202, 1998.
- [108] R. A. Ong. ICRC 2005 - Rapporteur OG 1 - Gamma-ray Science and Techniques. In *International Cosmic Ray Conference*, 2005.
- [109] R. A. Ong, D. Bhattacharya, C. E. Covault, D. D. Dixon, D. T. Gregorich, D. S. Hanna, S. Oser, J. Québert, D. A. Smith, O. T. Tümer, and A. D. Zych. Detection of atmospheric Cherenkov radiation using solar heliostat mirrors. *Astroparticle Physics*, 5:353–365, October 1996.

-
- [110] S. Oser, D. Bhattacharya, L. M. Boone, M. C. Chantell, Z. Conner, C. E. Covault, M. Dragovan, P. Fortin, D. T. Gregorich, D. S. Hanna, R. Mukherjee, R. A. Ong, K. Ragan, R. A. Scalzo, D. R. Schuette, C. G. Théoret, T. O. Tümer, D. A. Williams, and J. A. Zweerink. High-Energy Gamma-Ray Observations of the Crab Nebula and Pulsar with the Solar Tower Atmospheric Cerenkov Effect Experiment. *Astrophysical Journal*, 547:949–958, February 2001.
- [111] S. M. Oser. *High Energy Gamma-Ray Observations of the Crab Nebula and Pulsar with the Solar Tower Atmospheric Cerenkov Effect Experiment*. PhD thesis, University of Chicago, August 2000.
- [112] P. Padovani and P. Giommi. The connection between x-ray- and radio-selected BL Lacertae objects. *Astrophysical Journal*, 444:567–581, May 1995.
- [113] E. Paré et al. CELESTE: an atmospheric Cherenkov telescope for high energy gamma astrophysics. *Nuclear Instruments and Methods in Physics Research A*, 490:71–89, September 2002.
- [114] D. Petry et al. Detection of VHE γ -rays from MKN 421 with the HEGRA Cherenkov Telescopes. *Astronomy and Astrophysics*, 311:L13–L16, July 1996.
- [115] F. Piron and The CAT Collaboration. Observation of the Gamma-ray Emission of the Blazar Markarian 421 Above 250 GeV with the CAT Cherenkov Telescope. In J. Paul, T. Montmerle, and E. Aubourg, editors, *Abstracts of the 19th Texas Symposium on Relativistic Astrophysics and Cosmology, held in Paris, France, Dec. 14-18, 1998*. Eds.: J. Paul, T. Montmerle, and E. Aubourg (CEA Saclay), meeting abstract., December 1998.
- [116] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C - The Art of Scientific Programming*, chapter 14-5. Cambridge University Press, Second edition, 1992.
- [117] J. R. Primack, R. S. Somerville, J. S. Bullock, and J. E. G. Devriendt. Probing Galaxy Formation with High-Energy Gamma Rays. In F. A. Aharonian and H. J. Völk, editors, *American Institute of Physics Conference Series*, pages 463–+, 2001.
- [118] M. Punch, C. W. Akerlof, M. F. Cawley, M. Chantell, D. J. Fegan, S. Fennell, J. A. Gaidos, J. Hagan, A. M. Hillas, Y. Jiang, A. D. Kerrick, R. C. Lamb, M. A. Lawrence, D. A. Lewis, D. I. Meyer, G. Mohanty, K. S. O’Flaherty, P. T. Reynolds,

- A. C. Rovero, M. S. Schubnell, G. Sembroski, T. C. Weekes, and C. Wilson. Detection of TeV photons from the active galaxy Markarian 421. *Nature*, 358:477–+, August 1992.
- [119] J. Quinn, C. W. Akerlof, S. Biller, J. Buckley, D. A. Carter-Lewis, M. F. Cawley, M. Catanese, V. Connaughton, D. J. Fegan, J. P. Finley, J. Gaidos, A. M. Hillas, R. C. Lamb, F. Krennrich, R. Lessard, J. E. McEnery, D. I. Meyer, G. Mohanty, A. J. Rodgers, H. J. Rose, G. Sembroski, M. S. Schubnell, T. C. Weekes, C. Wilson, and J. Zweerink. Detection of Gamma Rays with E greater 300 GeV from Markarian 501. *Astrophysical Journal Letters*, 456:L83+, January 1996.
- [120] K. J. Ragan. The Montreal digital delay system Specifications and Design. STACEE internal note 03-01, 2003.
- [121] N. Roos. Galaxy mergers and active galactic nuclei. *Astronomy and Astrophysics*, 104:218–228, December 1981.
- [122] E. E. Salpeter. Accretion of Interstellar Matter by Massive Objects. *Astrophysical Journal*, 140:796–800, August 1964.
- [123] R. A. Scalzo. Poisson Template Methods with Undercanting: Fitting Energies and Cores. STACEE internal note 02-05, 2001.
- [124] R. A. Scalzo. *Observations of the EGRET Blazar W Comae with the Solar Tower Atmospheric Cherenkov Effect Experiment*. PhD thesis, University of Chicago, 2004.
- [125] R. A. Scalzo, L. M. Boone, D. Bramel, J. Carson, C. E. Covault, P. Fortin, G. Gauthier, D. M. Gingrich, D. Hanna, A. Jarvis, J. Kildea, T. H. Lindner, C. Mueller, R. Mukherjee, R. A. Ong, K. J. Ragan, D. A. Williams, and J. Zweerink. High-Energy Gamma-Ray Observations of W Comae with the Solar Tower Atmospheric Cerenkov Effect Experiment (STACEE). *Astrophysical Journal*, 607:778–787, June 2004.
- [126] R. Scarpa, C. M. Urry, R. Falomo, J. E. Pesce, and A. Treves. The Hubble Space Telescope Survey of BL Lacertae Objects. I. Surface Brightness Profiles, Magnitudes, and Radii of Host Galaxies. *Astrophysical Journal*, 532:740–815, April 2000.
- [127] S. Schlenker, T. Lohse, C. Borgmeier, and C. Stegmann. Simulation of Night Sky Background Trigger Rate for the H.E.S.S. Camera. HESS internal note, 2001.

-
- [128] R. Schödel, T. Ott, R. Genzel, R. Hofmann, M. Lehnert, A. Eckart, N. Mouawad, T. Alexander, M. J. Reid, R. Lenzen, M. Hartung, F. Lacombe, D. Rouan, E. Gendron, G. Rousset, A.-M. Lagrange, W. Brandner, N. Ageorges, C. Lidman, A. F. M. Moorwood, J. Spyromilio, N. Hubin, and K. M. Menten. A star in a 15.2-year orbit around the supermassive black hole at the centre of the Milky Way. *Nature*, 419:694–696, October 2002.
- [129] A. Sillanpää, S. Haarala, M. J. Valtonen, B. Sundelius, and G. G. Byrd. OJ 287 - Binary pair of supermassive black holes. *Astrophysical Journal*, 325:628–634, February 1988.
- [130] A. Sillanpää, L. O. Takalo, T. Pursimo, K. Nilsson, P. Heinamaki, S. Katajainen, H. Pietila, M. Hanski, R. Rekola, M. Kidger, P. Boltwood, G. W. Turner, J. W. Robertson, R. K. Honeycut, Y. S. Efimov, N. Shakhovskoy, M. Fiorucci, G. Tosti, G. Ghisellini, C. M. Raiteri, M. Villata, G. de Francesco, L. Lanteri, M. Chiaberge, A. Peila, and J. Heidt. Double-peak structure in the cyclic optical outbursts of blazar OJ 287. *Astronomy and Astrophysics*, 315:L13–L16, November 1996.
- [131] M. L. Sitko and V. T. Junkkarinen. Continuum and line fluxes of OJ287 at minimum light. *Astronomical Society of the Pacific, Publications*, 97:1158–1162, December 1985.
- [132] M. Spada, G. Ghisellini, D. Lazzati, and A. Celotti. Internal shocks in the jets of radio-loud quasars. *Monthly Notices of the Royal Astronomical Society*, 325:1559–1570, August 2001.
- [133] T. Spreitzer. Stability of Discriminator Thresholds. STACEE internal note 03-10, 2003.
- [134] F. W. Stecker, O. C. de Jager, and M. H. Salamon. TeV gamma rays from 3C 279 - A possible probe of origin and intergalactic infrared radiation fields. *Astrophysical Journal Letters*, 390:L49–L52, May 1992.
- [135] A. A. Stepanyan, Y. I. Neshpor, N. A. Andreeva, O. R. Kalekin, N. A. Zhogolev, V. P. Fomin, and V. G. Shitov. Observations of the Flux of Very-High-Energy Gamma Rays from the Blazar 3C 66A. *Astronomy Reports*, 46:634–638, August 2002.
- [136] P. A. Strittmatter, K. Serkowski, R. Carswell, W. A. Stein, K. M. Merrill, and E. M. Burbidge. Compact Extragalactic Nonthermal Sources. *Astrophysical Journal*, 175:L7+, July 1972.

- [137] B. N. Swanenburg, K. Bennett, G. F. Bignami, R. Buccheri, P. Caraveo, W. Hermsen, G. Kanbach, G. G. Lichti, J. L. Masnou, H. A. Mayer-Hasselwander, J. A. Paul, B. Sacco, L. Scarsi, and R. D. Wills. Second COS B catalog of high-energy gamma-ray sources. *Astrophysical Journal Letters*, 243:L69–L73, January 1981.
- [138] T. Takahashi et al. Complex Spectral Variability from Intensive Multiwavelength Monitoring of Markarian 421 in 1998. *Astrophysical Journal Letters*, 542:L105–L109, October 2000.
- [139] D. J. Thompson, D. L. Bertsch, and R. H. O’Neal, Jr. The Highest-Energy Photons Seen by the Energetic Gamma Ray Experiment Telescope (EGRET) on the Compton Gamma Ray Observatory. *Astrophysical Journal Supplement Series*, 157:324–334, April 2005.
- [140] D. J. Thompson et al. Calibration of the Energetic Gamma-Ray Experiment Telescope (EGRET) for the Compton Gamma-Ray Observatory. *Astrophysical Journal Supplement Series*, 86:629–656, June 1993.
- [141] O. T. Tumer, J. T. O’Neill, A. D. Zych, and R. S. White. A promising large area VHE gamma ray detector with excellent hadron rejection capability. *Nuclear Physics B Proceedings Supplements*, 14:351–355, March 1990.
- [142] M.-H. Ulrich, T. D. Kinman, C. R. Lynds, G. H. Rieke, and R. D. Ekers. Non-thermal continuum radiation in three elliptical galaxies. *Astrophysical Journal*, 198:261–266, June 1975.
- [143] C. M. Urry and P. Padovani. Unified Schemes for Radio-Loud Active Galactic Nuclei. *Publications of the Astronomical Society of the Pacific*, 107:803–+, September 1995.
- [144] M. J. Valtonen, H. J. Lehto, and H. Pietilä. Probing the jet in the binary black hole model of the quasar OJ287. *Astronomy and Astrophysics*, 342:L29–L31, February 1999.
- [145] M. J. Valtonen, K. Nilsson, A. Sillanpää, L. O. Takalo, H. J. Lehto, W. C. Keel, S. Haque, D. Cornwall, and A. Mattingly. The 2005 November Outburst in OJ 287 and the Binary Black Hole Model. *Astrophysical Journal Letters*, 643:L9–L12, May 2006.

-
- [146] F. Vincent. Alternate method for determining the gain of a PMT. STACEE internal note 00-02, 2000.
- [147] R. M. Wagner and et al. Observations of the Crab Nebula with the MAGIC Telescope. In *29th International Cosmic Ray Conference*, 2005.
- [148] T. C. Weekes. Very high energy gamma-ray astronomy. *Physics Reports*, 160:1–2, March 1988.
- [149] T. C. Weekes. *Very high energy gamma-ray astronomy*. Very high energy gamma-ray astronomy, by Trevor C. Weekes. IoP Series in astronomy and astrophysics, ISBN 0750306580. Bristol, UK: The Institute of Physics Publishing, 2003, 2003.
- [150] T. C. Weekes, M. F. Cawley, D. J. Fegan, K. G. Gibbs, A. M. Hillas, P. W. Kowk, R. C. Lamb, D. A. Lewis, D. Macomb, N. A. Porter, P. T. Reynolds, and G. Vacanti. Observation of TeV gamma rays from the Crab nebula using the atmospheric Cerenkov imaging technique. *Astrophysical Journal*, 342:379–395, July 1989.
- [151] A. R. Whitney, I. I. Shapiro, A. E. E. Rogers, D. S. Robertson, C. A. Knight, T. A. Clark, R. M. Goldstein, G. E. Marandino, and N. R. Vandenberg. Quasars Revisited: Rapid Time Variations Observed Via Very-Long-Baseline Interferometry. *Science*, 173:225–230, July 1971.
- [152] B. J. Wills and D. Wills. C 66A: a Bright New Quasi-Stellar Object. *Astrophysical Journal Letters*, 190:L97+, June 1974.
- [153] Y. Xue and W. Cui. X-Ray Flares from Markarian 501. *Astrophysical Journal*, 622:160–167, March 2005.
- [154] Y. B. Zeldovich. Stellar Evolution and Release of Gravitational Energy in Accretion. *Doklady Acad. Nauk.*, 1656:67, 1964.
- [155] F. J. Zhang and L. B. Baath. The radio fine structure of the BL Lacertae object MK 421. *Astronomy and Astrophysics*, 236:47–52, September 1990.