Probing Cosmic Rays with VERITAS: Observations of M 31 and the Positron Fraction

Ralph Alan Bird, MSc MSci MA

11277246

The thesis is submitted to University College Dublin in fulfilment of the requirements for the degree of Doctor of Philosophy

School of Physics

Head of School: Prof. Padraig Dunne
Supervisor: Dr. John Quinn

August 2015

Contents

List of Figures

List of Tables

Abstract

In this thesis I present two attempts to measure cosmic rays with VERITAS, a very-high-energy (VHE, E>100 GeV) imaging atmospheric Cherenkov telescope array. The first is a study into measuring the positron fraction ($e^+/(e^+ + e^-)$) and the antiproton ratio ($\overline{p}/p$) using the Earth-Moon ion spectrometer. Using the cosmic ray shadow of the Moon after it has been deflected by the Earth's magnetic field as an ion spectrometer and applying analysis cuts to discriminate between leptonic and hadronic initiated showers it is possible to measure the relative abundances of cosmic ray types. This requires observations close to the Moon (2-3°) using special observing modes, reduced high voltage to the photomultiplier tubes and UV bandpass filters; the development and verification of these modes is also presented. This feasibility study showed that it will be very difficult to conduct this observation with VERITAS, taking tens of hours to detect the electron shadow and hundreds of hours for the positron shadow. Observations during a Lunar eclipse did allow for detection of the cosmic-ray shadow but with insufficient statistics to draw any detailed conclusions.

The second approach is to observe the diffuse γ-ray emission from M 31, the Andromeda galaxy. Produced through the interaction of cosmic rays with the interstellar medium, this emission presents a unique opportunity to study the cosmic ray dynamics of an entire galaxy. An updated *Fermi*-LAT analysis is also conducted which shows marginal evidence for extended emission and has reduced the errors on the spectral fit in comparison to the original publication. An upper limit in the 100-200 MeV energy bin is suggestive of the emission being due to pion decay but further work is required to verify this. No evidence of emission was detected with VERITAS and upper limits were put on the flux at 2.2-6.6% of the Crab Nebula flux depending upon the method used. Predictions of the flux at VHE energies suggests that, with expected improvements in the analysis, a further 500 hours of observations will be required to detect the emission.

Statement of Original Authorship

I hereby certify that the submitted work is my own work, was completed while registered as a candidate for the degree stated on the Title Page, and I have not obtained a degree elsewhere on the basis of the research presented in this submitted work.

During this thesis I have worked as part of the VERITAS collaboration*. The work presented in this thesis uses the equipment and software developed by the collaboration (described in Chapter 3) but the analysis and the specific tools developed to conduct this work (described in the later chapters) has been my own work. The only exception to this has been the work conducted on the bright moonlight programme which has been conducted in conjunction with Sean Griffin (McGill University). The analysis presented has been my own but we have worked together extensively to produce and corroborate the results using two independent packages.

I acknowledge the financial support of the DGPP which is Funded under the Programme for Research in Third-Level Institutions and co-funded under the European Regional Development Fund (ERDF).

*http://veritas.sao.arizona.edu/

Acknowledgements

This thesis is born from a lifelong enjoyment of wanting to know how things work, with my love of physics drawing on this by addressing questions on how the universe works. Thus I want to thank everybody who has ever endured me asking that silly question and hasn't told me to be quiet. In particular I have to thank my parents, their encouragement from a young age that I should keep exploring, keep thinking and never just accept an answer is the main reason I have produced this work. I can never say thank you enough.

Though they will probably never get to read this I have also to thank the two inspirational teachers in my life, Mr Wicking and Mr Gough. Mr Gough taught me so many lessons that I will never forget; principally to keep trying and never settle for what I viewed as acceptable. Mr Wicking introduced me to the joys of real physics, the thrill of seeing what happens if I have a go and the excitement of finding something out when I didn't already know the answer.

To complete this Ph.D. I moved to Ireland where I have been lucky enough to be part of an amazing church in Holy Trinity, Rathmines. The community there has been the source of so much fun, encouragement and growth over these years, I thank you all. In particular though I want to thank Rob and Dilys, your constant stream of love and support has been invaluable. Keep being yourselves, you are a shining example of living out the life we are called too.

I wish to thank the School of Physics in UCD for all their help and support throughout my time here. Prof. Padraig Dunne and Prof. Lorraine Hanlon, thank you for all the help you have given me during your tenures as Head of School and Bairbre Fox and John Brennan, thank you for numerous times you have sorted something for me, often at the last moment. I also with to acknowledge my funding, the Dublin Graduate Physics Programme, and Stephanie Mathews for all the work she has put in managing that for us all.

My colleagues in UCD (Anna, Anna, Eddie, Yerbol, Ste and Elisa) and within VERITAS (Sean and Simon in particular), thank you for keeping me sane, answering all of

my stupid questions and sharing so much beer with me. I also want to thank all of the senior members of the collaboration who have taught me so much, in particular Pascal Fortin, Jamie Holder, Ken Regan and David Hanna. That you are all so keen to support and encourage others has made this collaboration a joy to work in. Most of all though I want to thank John Quinn, that you always had your office door open, were always willing to hear my question and humour my silly idea, and always asked the question that I least wanted you to ask, thank you. I have learnt so much from our discussions.

The person who I have to thank the most is the person who has been my rock through it all. From encouragement when I was down to celebrating every success; with lots of love, food, support and gin in between. Emma you have made this period so much more enjoyable than it would have been without you. I dedicate this thesis to you.

1. Introduction

1.1. Introduction to Cosmic Ray Astrophysics

Cosmic rays are the highest-energy particles known to mankind, with energies recorded up to 10^{20} eV (Bird et al., 1995), roughly the energy of Andy Murray's second serve. Since their discovery in 1912 by Victor Hess (Hess, 1912) the origin of these cosmic rays has remained a mystery. Though significant work has been done and many advances have been made, the origin these particles and how they are accelerated to such extraordinary energies remains one of the key questions in astrophysics, and is still the subject of detailed research after over 100 years of study.

One of the best windows into the origin of cosmic rays is to study their interaction products, in particular electromagnetic emission as it travels through the universe unaffected by the intervening magnetic fields. The most energetic electromagnetic emission detected is in the form of γ-rays, which have recorded energies of $>10^{14}$ eV. Through these observations several candidates have been identified for the acceleration of cosmic rays, both within our Galaxy and beyond. What links all of these potential sources is that they have extreme energy densities; they can be the remains of supernovae, supermassive black holes powering large jets, or the last death throes of merging compact objects (black holes or neutron stars).

Combined with observations at other wavelengths, very-high-energy astrophysics (VHE, $E > 100$ GeV) is adding key information about the properties of these objects. VHE γ-rays, which are only produced by the most energetic particles, provide key information about the highest-energy processes powering these objects. Combined with our growing understanding of the make up and energies of the cosmic rays that arrive at the Earth, we are also developing a picture of how cosmic rays are accelerated and diffuse through our galaxy and beyond.

In this thesis I am present two attempts to conduct measurements of cosmic rays using the VHE γ-ray observatory, VERITAS:

1. A feasibility study into and preliminary results on using the shadow in the cosmic ray flux caused by the Moon and its apparent deflection due to the Earth's magnetic field as an ion spectrometer. This could allow the make up of the local cosmic ray flux, specifically the positron fraction ($e^+/(e^+ + e^-)$) and the antiproton ratio ($\overline{p}/p$) to be determined.

2. I have produced the most detailed high-energy (HE, $100\,\mathrm{MeV} < E < 100\,\mathrm{GeV}$) and VHE γ-ray analysis of M 31, the closest spiral galaxy to the Milky Way (MW), to date. M 31 presents a unique opportunity to study an entire galaxy in a way that is not possible for the MW where our location within the galaxy obscures and confuses observations.

This work adds to the body of knowledge that has been built up over the past 100 years on the origins of cosmic rays and will help future work to pin down their origins.

1.2. Cosmic Ray Acceleration and Gamma Ray Emission

Since cosmic and γ-rays are of such high energy they require extreme environments to produce them, thus they provide a window into some of the most interesting places in the universe. Their energies are such that they challenge existing understanding of how particles are accelerated. Our current understanding is outlined in Chapter 2 where I work through a basic model of acceleration of cosmic rays by shock fronts and then the interaction of cosmic rays with with interstellar medium to produce γ-rays. This basic understanding addresses the key processes but still leaves many questions unanswered. More-detailed models exist which better match the data but we are still lacking the conclusive evidence that identifies the sites of hadronic cosmic ray acceleration and explains the full details of these objects. These answers will be found through detailed observations, constraining or disproving existing theories and allowing more detailed models to be tested.

1.3. Cosmic and Gamma Ray Observations

Cosmic and γ-rays share many similarities in their detection techniques. In its most simple form, these techniques are all based around having a detector volume within which energy is deposited and a way of measuring that energy. How that is conducted depends

primarily upon the energy of the particles being detected. If charge discrimination is required then a magnetic field is applied to the detector volume.

Lower energy particles are significantly more abundant and require a smaller mass to deposit all of their energy. Thus their detectors can be smaller. However, since the Earth's atmosphere will stop these particles and absorb all of the energy produced in the interaction, they have to be high up in/above it. Thus they are placed on satellites or carried on long duration balloon flights.

Higher-energy particles are less abundant and require a significantly larger mass within which to deposit their energy. However, by using the Earth's atmosphere as a part of the detector it is possible to conduct observations from the ground and make use of a significantly larger volume than could be carried by a satellite or a balloon (the interaction depth of the atmosphere is equivalent to about 1 m of lead). This larger volume both allows the interaction products to be better sampled (giving a better measure of properties of the incident particle) and increases the area for interactions to occur in, increasing the number of particles detected.

Summarised below are a few of the key detectors of various types that are applicable to this thesis. This is not an exhaustive list and several other detector types are used but these cover the energy ranges of interest in this thesis and are the experiments whose data is either directly or indirectly used in this work. The differential sensitivities of the different γ-ray observatories are shown in Figure 1.1, showing their complementary energy ranges and the performance upgrade coming from the proposed Cherenkov Telescope Array (CTA).

1.3.1. *Fermi*-LAT

The *Fermi* satellite was launched in June 2008 and has two instruments on board. The primary instrument is the Large Area Telescope (LAT, Atwood et al. (2009)) which was designed as an all sky survey telescope in the energy range $0.02 - 300\,\mathrm{GeV}$. It is an imaging telescope with a large field-of-view (FoV), covering $> 2\,\mathrm{sr}$ ($\sim$20% of the sky) and has produced the most detailed map of the γ-ray sky ever, significantly increasing the number of sources detected and greatly increasing our understanding of those sources. The second instrument is the Gamma-Ray Burst Monitor (GBM, Meegan et al. (2009)), its purpose is to detect rapid changes in flux at lower energies (a few keV to 30 MeV), principally γ-ray bursts, incredibly bright flashes of γ-rays most likely caused by the

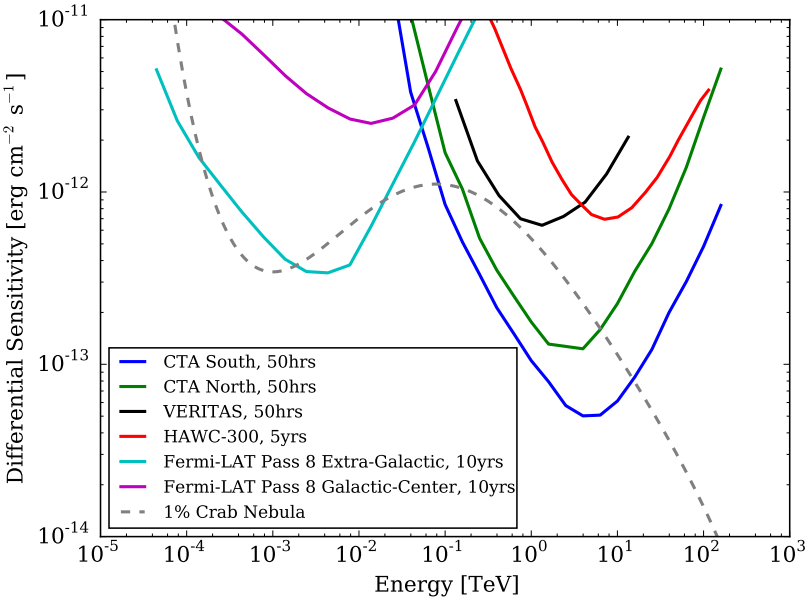

Figure 1.1.: The differential sensitivities of γ-ray observatories for point sources. The Crab Nebula spectrum is from (Meyer et al., 2010).

death throes of compact objects. Designed to detect, localise and provide spectroscopy on these flashes it has a near full-sky FoV and can rapidly localise these events to within a few degrees, allowing for follow-up observations. All data from the *Fermi* mission is made publicly available, along with the tools to analyse it.

In this thesis, since I am interested in imaging and spectral reconstruction of M 31 I will be working with data from the LAT. With a primary mission objective of γ-ray astronomy, the LAT is designed to reduce its background counts from charged cosmic rays (which are $10^5\times$ more abundant) as well as determine the trajectory and energy of the incident photons. The primary signature that is used for this is the pair conversion of the incoming photon into an electron-positron pair. By reconstructing the energies and directions of this pair, the energy and direction of the photons are determined. To reduce the background an anticoincidence shield surrounds the detector to detect incoming charged particles and act as a veto. Within the detector, multiple thin layers of high-Z conversion foils provide a target for the γ-rays, the resulting pairs are tracked (using silicon-strip detectors) before entering a CsI calorimeter, where they cause an

electromagentic shower and their energy is determined. The characteristic γ-ray signature in the LAT is therefore (1) no signal in the anticoincidence shield, (2) more than one track starting from the same location within the volume of the tracker, and (3) an electromagnetic shower in the calorimeter. These combined detectors give a peak effective area of $>8000\,\mathrm{cm}^2$ with an angular resolution that varies with energy (from $\sim 3.5°$ at 100 MeV to $<0.15°$ at >10 GeV) and an energy resolution that is less than 10%.

At low energies the LAT is background limited since differentiating γ-rays from the cosmic ray background is more difficult at these energies. Thus the significance of the signal increases as the square root of time. At higher energies the background rejection is much more efficient (due to the greater amount of energy deposited), however the flux of cosmic rays is much lower. This means that the detector is signal limited and the signal increases linearly with the time observed. Since the LAT is designed as a survey telescope it (primarily) operates in a mode where it repeatedly scans the whole sky, building up exposure on all sources and also giving it repeated sampling of transient and periodic sources.

For details of how *Fermi*-LAT data is analysed see Section 6.2.

1.3.2. AMS-02 & *PAMELA*

Fermi is designed to observe γ-rays, whereas AMS-02 (launched in 2011, Kounine (2012)) and *PAMELA* (launched in 2006, Boezio et al. (2009)) are two space based detectors designed for cosmic-ray studies. Both are magnetic spectrometers with multiple systems for determining the energy and direction of cosmic rays that pass through the detector. This is done through the use of a large permanent magnet and then series of trackers, calorimeters, time-of-flight detectors and anti-coincidence counters (each has slightly different layouts). Both are capable of discriminating between different cosmic ray species in the range of a few tens of MeV to a few hundred GeV, but AMS-02 has significantly more collection area and discrimination power.

The aim of the experiments is to study the composition and spectra of cosmic rays, in particular the antimatter component. Since *PAMELA*'s launch in 2006 (the first satellite-based experiment dedicated to cosmic rays), and especially following the deployment of AMS-02 on the international space station, our knowledge of the cosmic rays has increased dramatically. Of particular interest to this thesis is their measurement of the positron fraction, the fraction of positrons in the total electron plus positron flux. With

their detailed measurements they have detected an unexpected rise above 10 GeV which is contrary to traditional models. Significant work has gone into understanding the origin of this excess. Further details and attempts to measure it with VERITAS are presented in Chapter 5

1.3.3. Imaging Atmospheric Cherenkov Telescopes (VERITAS, MAGIC & H.E.S.S.)

As the energy rises the fluxes of both cosmic and γ-rays fall dramatically (as power laws with indexes in the range -2.5 to -3). Thus the detector area of satellites becomes too small to measure a sufficient number of interactions and the detector volume too small to accurately measure the interaction. Fortuitously, at about the energy that satellites become so limited, the particles have sufficient energy that when they interact within the Earth's atmosphere they produce a shower of highly-energetic particles (an air shower) that cause the atmosphere to emit Cherenkov radiation that is sufficiently bright for it to be measured from the ground. By imaging this radiation with large, sensitive telescopes it is possible to determine the energy and direction of the shower (and thus of the initiating particle) and discriminate between showers initiated by γ-rays and those initiated by cosmic rays (primarily hadrons). Thus observations are no longer constrained by what can be lifted above the atmosphere, rather the atmosphere itself is used as a part of the target. This switch-over energy is also used as a boundary in the nomenclature between high energy ($0.1 - 100$ GeV, HE) and very high energy ($0.1 - 100$ TeV, VHE) astrophysics, a convention that is used throughout this thesis.

With a background rate that dwarfs even the brightest γ-ray emitters, the key breakthrough for VHE γ-ray astronomy was the ability to image the showers and use that information to discriminate between the signal (γ-ray initiated showers) and background (cosmic ray, primarily hadronic, initiated showers). This technique was pioneered by the Whipple collaboration with a single telescope which detected the first VHE source, the Crab Nebula (Weekes et al., 1989). Since then significant improvements have been made to the technique by a number of collaborations around the world (primarily the use of arrays of telescopes to improve the direction reconstruction and reduce the background) with the total number of VHE sources rising to 183*. Currently there are three major collaborations that are operating imaging atmospheric Cherenkov telescopes (IACTs),

*June 2015, http://tevcat.uchicago.edu/

VERITAS (situated in Arizona, USA, (Holder et al., 2008)), MAGIC (La Palma, Spain, (Aleksić et al., 2012)) and H.E.S.S. (Khomas highland, Namibia, (Aharonian et al., 2006)). All three arrays are roughly similar in their operation and capabilities, providing coverage of both the northern and southern sky. Work is also under way on the successor telescope array, the Cherenkov Telescope Array (CTA, (Acharya et al., 2013)) which will increase sensitivities by about an order of magnitude over current arrays. Details of the IACT technique, in particular as applied by VERITAS are given in Chapter 3.

Of particular interest to this thesis are two experiments which adapted IACTs to attempt to measure the cosmic-ray Moon shadow.

ARTEMIS

ARTEMIS was an experiment that used the Whipple telescope to attempt to measure the antiproton ratio (Pomarède et al., 2001). To do this, it replaced the standard Whipple camera with one that is more sensitive to UV light (since air showers develop mainly below the ozone layer and thus the UV light reaches the telescope) and used an optical filter to limit the sensitivity to the UV light. With 36 hours of data they did not detect any deficit, simulations predicted that (provided that they could improve their systematic errors) they would require 140 hours to detect the deficit to 5σ.

CLUE

The CLUE experiment was located on the island of La Palma, Spain and comprised of nine telescopes. The aim of CLUE was to measure the ratio of matter to antimatter in the cosmic radiation at energies around 1 TeV by detecting the cosmic-ray Moon shadow. To achieve the desired UV sensitivity it used a multiwire proportional chamber filled with TMAE (TetrakisMethylAminoEthylene, $C_4H_{11}NO$). This substance has a particularly good photoconversion in the 180-240 nm range . With this it was able to detect a deficit in the cosmic-ray flux in the direction of the Moon (Dokoutchaeva et al., 2001) but the sensitivity was too limited to measure the antiproton flux (Lorenz & Wagner, 2012).

1.3.4. Water Cherenkov Telescopes (Milagro & HAWC)

At higher energies and/or higher altitudes, the particles in the air shower reach to the ground, allowing for them to be detected by their Cherenkov emission when they pass

through water tanks (or other means). This technique was used by Milagro (2001-2008, New Mexico, USA, (Abdo et al., 2012)) which had a single large tank with multiple detectors, and is used by HAWC (2014-, Puebla, Mexico, (Mostafa, 2013)) which has multiple smaller tanks to allow a larger area to be covered for a lower cost. By situating at a high altitude (2530 m & 4100 m respectively) they have been able to reduce their energy threshold to a few TeV. Since they are not relying on optical telescopes to detect the Cherenkov light from the showers, they operate in an all sky survey mode 24 hours a day. However, they are less sensitive than IACTs and have a larger PSF. IACTs and water Cherenkov telescopes are complementary to each other, and the current period of γ-ray astronomy, with a combination of IACTs, HAWC and the *Fermi*-LAT provides a golden opportunity to simultaneously study sources over many decades of energy.

1.3.5. Ultra-High-Energy Detectors (Pierre Auger & Telescope Array)

Though not directly used in this work, various Ultra-High-Energy (UHE, $> 10^{18}$ eV) experiments provide valuable information about cosmic rays at the highest energies. Both Pierre Auger (Abraham et al., 2004) & Telescope Array (Kawai et al., 2008) use two main methods in combination; ground detectors which detect the shower as it reaches the ground and fluorescence detectors which look for the fluorescence of nitrogen molecules caused by the showers. These UHE measurements have found significant kinks in the cosmic ray spectrum, variations in the composition with energy and anisotropy in the direction of the arrival directions (see Section 2.1 for further details).

1.3.6. Multi-Messenger Astrophysics (IceCube and LIGO)

We are currently experiencing the birth of multi-messenger astrophysics, using cosmic rays, gravitational waves and neutrinos along with photons to observe the universe. Of these, gravitational waves and neutrinos are the two that hold the most interest since, like photons, they are unaffected by the magnetic fields in the galaxy and beyond, and thus their source can be identified (or at least localised) for follow up by other instruments.

IceCube (Abbasi et al., 2009) and ANTARES (Ageron et al., 2011) are operational neutrino telescopes that use Cherenkov emission from muons produced by upwards travelling muon neutrinos that have travelled through the Earth. This enables them to be used as neutrino telescopes and look for signals of astrophyisical neutrino sources. The detection of such high energy neutrinos will, due to the emission interaction process of

VHE cosmic rays, provide conclusive evidence for the source being a site of hadronic cosmic ray acceleration (see Subsection 2.5.1). IceCube has observed the first astrophysical high-energy neutrino flux, with significance at the 5.7 sigma level after analyzing three years of data (2010-2013) (Aartsen et al., 2013).

Gravitational waves are predicted to be detectable from, amongst other things, the mergers of compact objects and supernovae. LIGO (which was operational from 2002 to 2010) failed to detect gravitational waves and the detectors are currently being upgraded with an estimated order of magnitude improvement in sensitivity and improved angular resolution.

Though no multi-messenger sources have yet been identified work is ongoing and constraints have already been placed using existing data. In the next few years, with upgrades to existing systems and new systems (for example CTA and KM3NeT (Margiotta, 2014)) which will improve both sensitivity and angular resolution, the first detection of a source in the multi-messenger regime is approaching.

1.4. The VHE Gamma-Ray Sky

VERITAS is sensitive to γ-rays in the energy range 85 GeV to >30 TeV. This enables it to detect and characterise a variety of VHE sources, which can be divided into six main categories. These are pulsar wind nebulae (PWN), (compact) binary sources, active galactic nuclei (AGN), supernova remnants (SNRs), starburst galaxies, and unidentified sources (sources without counterparts at other wavelengths). These categories cover a wide variety of astronomical objects, from the Geminga PWN* in our own galaxy at distance of ∼0.169 kpc to the gravitationally lensed AGN S3 0218+35 located at a redshift of z = 0.944†. The sources detected in the VHE sky are shown in Figure 1.2

The most similar sources to M 31 that are known to emit VHE γ-rays are the starburst galaxies, M82 (Acciari et al., 2009a) and NGC 253 (Abramowski et al., 2012). With their high star formation rates they have a high density of PWN, SNRs, and binaries which as well as emitting γ-rays in their own right are believed to be the sites of cosmic ray acceleration. These cosmic rays then interact with the interstellar medium (ISM) to produce diffuse γ-rays emission.

*VERITAS has not detected the Geminga PWN, it has been detected by Milagro (Abdo et al., 2009) and HAWC (Pretz, 2015). Work on a VERITAS detection is ongoing.

†http://tevcat.uchicago.edu/

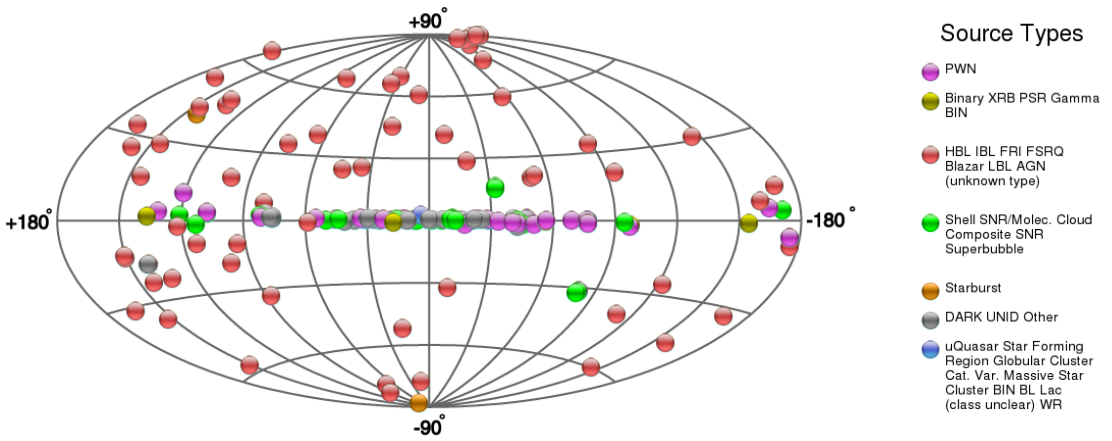

Figure 1.2.: The current VHE source catalogue *Image Credit:* *http://tevcat.uchicago.edu/*, *June 2015*

The single largest class of VHE emitters are the AGNs, primarily blazars. These are galaxies which have a supermassive black hole in the centre which generates a large jet that is kiloparsecs long. This jet is moving at relativistic speeds and, if pointed directly at the Earth, can be detected in VHE γ-rays. These objects, as well as being amongst the most powerful objects in the universe, can also be extremely variable, with variations recorded on the timescales of minutes.

The Galactic VHE sources (SNRs, PWN and binaries) are all potential candidates for cosmic-ray accelerators. By studying these objects at VHE energies we are able to develop an understanding of the mechanisms going on in them and their potential contribution to the cosmic ray flux.

1.4.1. Cosmic Ray Observations with VERITAS

In addition to its primary role studying γ-rays, VERITAS also detects cosmic rays and has been used to study their composition and spectra, principally the electron + positron spectrum (Stasak, 2015) and a search for heavy nuclei in the cosmic ray flux (Wissel, 2007). The work presented in this thesis builds upon VERITAS's reputation as a key tool in studying the origins of cosmic rays and their interactions within the Galaxy.

1.5. M 31, The Andromeda Galaxy

The Andromeda Galaxy or M 31 is a spiral galaxy that lies approximately 780 kpc from Earth. Visible with the naked eye and recorded in early star charts (often under the name *"Little Cloud"*), the first description of telescopic observations was given by Simon Marius in 1612 (Mayr (Marius), 1614) and it was recorded by Charles Messier in his first catalogue (Messier, 1774). After much debate, Hubble (1929) showed that M 31 lay outside the MW and thus it had to be a galaxy of its own. Since then M 31 has been one of the most studied objects in the sky, with detailed surveys in all wavelengths including the "most detailed picture of outer space"*. A known HE γ-ray emitter (Abdo et al., 2010), it is, as yet, undetected at higher energies, though upper limits have been published by HEGRA (Aharonian et al., 2003) and presented by HAWC (Brewer & Smith, 2015). It is of particular interest to the VHE community for exactly the same reasons as it is of interest at other wavelengths, as the closest spiral galaxy it gives us a unique opportunity to study the emission throughout the galaxy. Since it is so well studied at other wavelengths, it has the potential to significantly to our understanding of cosmic ray acceleration and diffuse on a galactic scale.

1.6. Outline of this Thesis

This thesis uses two ways to investigate galactic cosmic rays.

The first is a feasibility study into using VERITAS to measure the positron fraction $(e^+/(e^+ + e^-))$ in the $0.5 - 1$ TeV range. This is done by measuring the deficit in the cosmic ray flux caused by the Moon which appears deflected due to the Earth's magnetic field and thus acts as an ion spectrometer. In addition I examine the antiproton ratio $(\overline{p}/p)$ at higher energies, including observations conducted during a lunar eclipse where the first detection of the shadow of the Moon with a PMT based IACT is observed. See Chapter 5 for the details of this analysis. To conduct this analysis required adapting VERITAS for operation under bright moonlight, this is described in Section 4.3.

The second is the most detailed γ-ray analysis of M 31 to date. As the closest spiral galaxy to the MW it is the obvious target for the study of cosmic rays throughout an entire galaxy. By studying the γ-ray emission from M 31 I aim to increase our understanding of the cosmic ray population and its interactions, and thus add to the

*http://www.nasa.gov/sites/default/files/thumbnails/image/hs-2015-02-a-hires_jpg.jpg

knowledge that is gained by studying individual objects within our own galaxy. An updated *Fermi*-LAT analysis as well as an analysis of 54.6 hours of VERITAS data are presented. Chapter 6 describes the results of this work and the methods employed to overcome the analysis challenges for VERITAS for such an extended, optically-bright object.

In order to conduct both of these observations work had to be done to reduce the impact of optical brightness and variations in the night-sky-background on VERITAS observations. Chapter 4 describes the methods employed as well as techniques developed to conduct observations under moonlight. This includes two new techniques for VERITAS, reducing the high-voltage to the photomultiplier tubes to reduce the gain and the use of UV bandpass filters to reduce the amount of background light entering the cameras.

A description of the science in this thesis is presented in Chapter 2 and the VERITAS array, data collection and analysis is presented in Chapter 3. A summary of the work and the conclusions drawn is given in Chapter 7

In addition to the above work, the appendices contain a summary of work that I have undertaken to reduce the deadtime of the array (Appendix A) and a novel technique to create a extended calibration source by conducting a raster scan on the Crab Nebula (Appendix B).

2. Cosmic and Gamma-Ray Physics

Primary cosmic rays are charged particles which are accelerated (via a non-thermal process) to very high energies (VHE, E>100 GeV). These VHE cosmic rays then escape from the source and diffuse through interstellar space (and for the very highest energy particles, intergalactic space) until they interact. This interaction can either provide us with a direct detection or produce secondary cosmic rays, including γ-rays, which we subsequently detect. To understand the acceleration and subsequent interactions several parameters can be measured and used to constrain models of these processes. The primary requirement is that the putative source(s) must accelerate sufficient particles to sufficient energies to power the cosmic rays that we observe. Then, either directly or through secondary processes (interaction or decay) they must produce the required particle types and spectra (Figure 2.1) within the constraints of other measurements of properties within the galaxy - for example the galactic magnetic field and interstellar material (ISM) density and composition. By balancing multiple measurements of different cosmic ray properties it is possible to constrain the acceleration and interaction mechanisms.

Presented in this chapter is a brief discussion of the basic methods of acceleration, diffusion, interaction and detection. This is neither an exhaustive study of the state of the art models, nor does it seek to explain the full complexity of the modelling required. Rather, it aims to provide a summary of the basic mechanisms involved and present the ways in which the observations presented in this work can provide additional light on the subject matter. For a more detailed study see the references, in particular Gaisser (1990) for a summary of the physics underpinning the field and Blasi (2013) for a summary of the current theoretical basis.

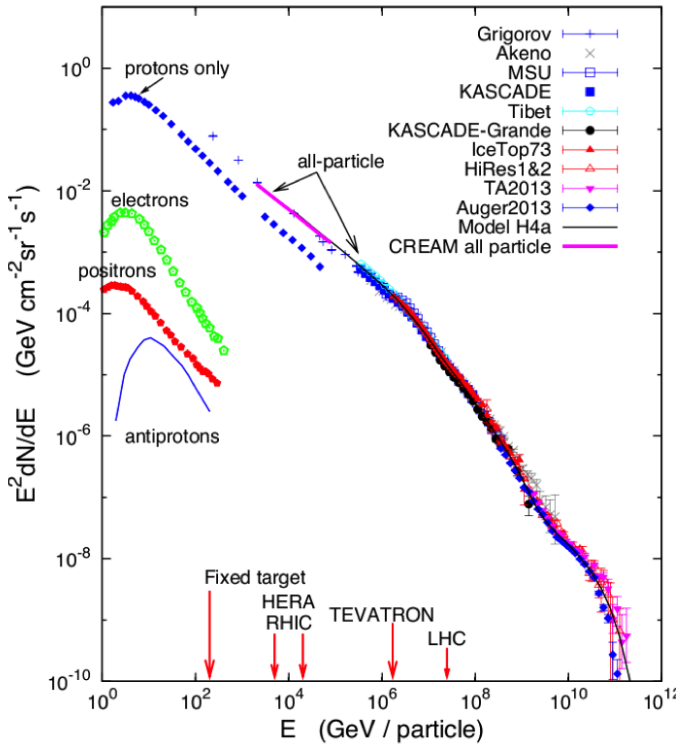

Figure 2.1.: The cosmic ray spectrum (from Blasi (2013)) as measured at the Earth. The subdominant contributions (electrons, positrons and antiprotons) are measured by *PAMELA*.

2.1. Source of Cosmic Rays

Given the observed cosmic-ray spectrum there must be at least one source that is capable of accelerating sufficient particles to populate the galaxy to that level (on the assumption that the measured spectrum is broadly similar to the spectrum throughout the Galaxy away from cosmic ray sources, there is no particular reason to believe that the Earth is located at a distinct part of the galaxy, though measurements do show differences to the local spectrum at e.g. the Galactic Center (Abramowski et al., 2014) where there are multiple potential cosmic-ray sources).

The differential intensity (I) of primary nucleons, over the energy range of a few GeV to somewhat beyond 100 TeV, is

$$I_N(E) \simeq 1.8 \times 10^4 \left(\frac{E}{1\,\mathrm{GeV}}\right)^{-\alpha} \text{ nucleons m}^{-2}\text{ s}^{-1}\text{ sr}^{-1}\text{ GeV}^{-1}. \tag{2.1}$$

where $\alpha(\equiv \gamma + 1) = 2.7$ is the differential spectral index of the cosmic ray flux, and γ is the integral spectral index (Olive, 2014). This gives a total energy density, ρ_E, of $\sim$1 eV cm^{-3} (Blasi, 2013).

If we assume that this is a typical energy density throughout the galaxy, then the power required to supply all of the galactic cosmic rays is

$$L_{CR} = \frac{V_D \, \rho_E}{\tau_R} \sim 5 \times 10^{40} \, \text{erg s}^{-1} \tag{2.2}$$

where the volume of the galactic disk, $V_D = \pi R^2 d \sim \pi (15 \, \text{kpc})^2 (200 \, \text{pc}) \sim 4 \times 10^{66} \, \text{cm}^3$ and τ_R is the residence time in the galactic disk which is derived from the boron to carbon ratio as $\sim 6 \times 10^6$ years (Gaisser, 1990).

Baade & Zwicky (1934) and Ginzburg & Syrovatskii (1964) were the first to identify that this was suggestive of supernovae as the source, since the luminosity of supernovae within the galaxy ($L_{SNR} \sim 3 \times 10^{42} \, \text{erg s}^{-1}$) is sufficient to energise all of the galactic cosmic rays if they can transmit their energy at an efficiency of only a few percent (a level recorded at lower energies e.g. during solar flares (Lin & Hudson, 1976)). Since then significant additional but not yet conclusive evidence has built up (see e.g. Blasi (2013) for a summary). Alternative theories have also been proposed, for example through evaporation from neutron stars (Ahlen & Tarlé, 2014), micro-quasars (Guo et al., 2013) or young stellar winds (Hörandel, 2008).

The shape of the observed spectrum (Figure 2.1) provides clues to the origin. Below 100 MeV the flux is affected by the solar wind and varies depending upon the solar cycle. Above the energies affected by solar modulation and up to $\sim 1 \times 10^6$ GeV/particle the spectrum is (nearly) constant and it is thought that the source of these particles are galactic accelerators. At about $\sim 1 \times 10^6$ GeV/particle there is a feature known as the "knee" where the spectrum softens, this likely due to changes in the composition of the particles or as the origin changes from galactic to extragalactic cosmic rays (Gaisser, 2010). At about 1×10^9 GeV/particle the spectrum hardens again at a feature known as the ankle, this harder spectrum is then cut off at about 1×10^{11} GeV/particle. The origin of this feature is debated between either a change in the composition (with progressively heavier elements cutting off at higher energies) or the Greisen-Zatsepin-Kuzmin limit (Greisen, 1966; Zatsepin & Kuz'min, 1966) where, over astronomical distances, cosmic rays with energies above this limit interact with CMB photons and thus have another loss mechanism.

2.2. Fermi Acceleration

In order for supernovae and their remnants to be cosmic ray accelerators there needs to be a mechanism for the transfer of energy from the explosive shocks to the cosmic rays. A common mechanism that is used to explain this was first derived by Enrico Fermi (Fermi, 1949) and provides a basis on which much work has been done to explain the details of the spectra and fit to specific observations and data.

Fermi acceleration is a way of explaining the transfer of macroscopic kinetic energy of a moving, magnetised plasma to individual charged particles. This increases the energy per particle to many times its initial value, creating the non-thermal energy distribution characteristic of particle acceleration. In this thesis I present a derivation taken from Gaisser (1990).

Consider a test particle of energy, E, that increases its energy by an amount proportional to its energy with each "encounter", thus $\Delta E = \xi E$ per encounter. Thus, after n encounters, an injected particle of initial energy E_0 has an energy of

$$E_n = E_0(1+\xi)^n \tag{2.3}$$

The number of encounters required to reach an energy E is therefore

$$n = \frac{ln\left(\frac{E}{E_0}\right)}{ln(1+\xi)}. \tag{2.4}$$

For n encounters to occur, the particle must not escape during that time. If the probability of escape is P_{esc} per encounter then, after n encounters the probability of remaining is $(1-P_{esc})^n$. Thus the proportion of particles reaching an energy E is

$$N(\geq E) \propto \sum_{m=n}^{\infty}(1-P_{esc})^m = \frac{(1-P_{esc})^n}{P_{esc}}. \tag{2.5}$$

Substituting Equation 2.4 into Equation 2.5 gives

$$N(\geq E) \propto \frac{1}{P_{esc}}\left(\frac{E}{E_0}\right)^{-\gamma} \tag{2.6}$$

where

$$\gamma = \frac{ln\left(\frac{1}{1-P_{esc}}\right)}{ln(1+\xi)} \approx \frac{P_{esc}}{\xi}, \quad (\xi, P_{esc} \ll 1) \tag{2.7}$$

Since Equation 2.6 is the number of particles that have attained an energy of at least E and some fraction of them will have continued to higher energies, this leads to a spectrum of the form

$$N(E)\,\mathrm{d}E = \text{constant} \times E^{-1-\gamma}\,\mathrm{d}E. \tag{2.8}$$

Thus a power-law energy distribution is recovered which depends upon the probability of escape P_{esc} (which is related to the characteristic time for escape from the acceleration region) and the fractional energy gained per encounter ξ.

2.2.1. Fermi Acceleration Mechanisms

Acceleration can be considered in two physical situations, when a charged particle "scatters" off a planar shock front (Figure 2.2b), or when it encounters moving clouds of plasma and "scatters" on irregularities in the magnetic field (Figure 2.2a). (In both cases scattering is off the magnetic field rather than off particles as that would result in collisional energy losses.)

Notation

For the rest of this chapter the following notation is used:

- **Primed** ($'$) quantities are in the **rest frame** of the shock/cloud.
- **Unprimed** quantities are in the **lab frame**.
- Energies with subscript **0** refer to the energy **before** interaction with the shock/cloud.
- Energies with subscript **1** refer to the energy **after** interaction with the shock/cloud.

In both cases the ambient gas is assumed to be at rest. For the gas cloud it is moving at a velocity $\vec{V}$ in the laboratory frame. For the planar shock, in the laboratory frame the shock is moving at a velocity $-\vec{u}_1$. Behind the shock, the shocked gas is moving away from the shock at a velocity of $\vec{u}_2$ relative to the shock front, thus in the the laboratory frame it is moving at a velocity $\vec{V} = \vec{u}_2 - \vec{u}_1$.

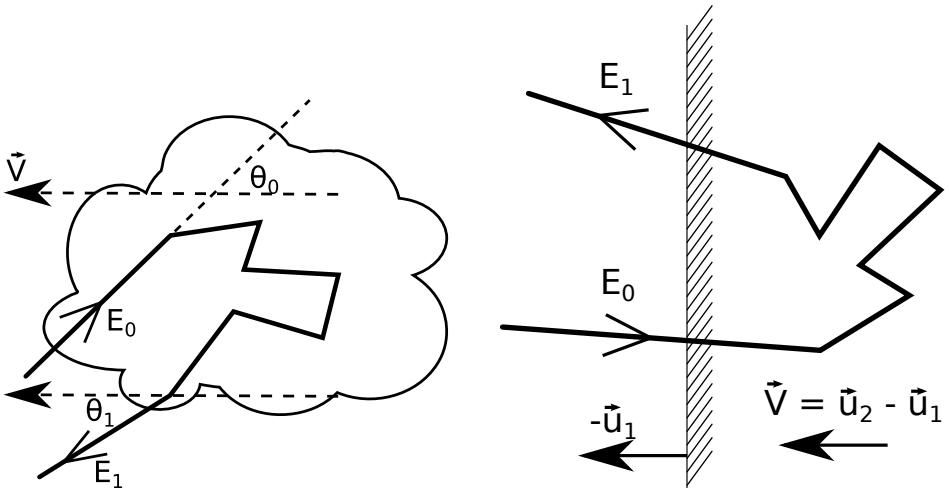

(a) The laboratory view of a charged particle accelerated from E_0 to E_1 by a moving, partly ionised gas cloud travelling at a velocity $\vec{V}$.

(b) The laboratory view of a charged particle accelerated from E_0 to E_1 by a planar shock front moving at velocity $-\vec{u}_1$. Behind the shock, the shocked gas moves away from the shock at a velocity of $\vec{u}_2$ relative to the shock front, thus in the the laboratory frame it is moving at a velocity $\vec{V} = \vec{u}_2 - \vec{u}_1$

Figure 2.2.: The two forms of Fermi acceleration, off a cloud of moving plasma (a) and off a planar shock (b).

Acceleration off a moving cloud of moving plasma

The first case presented is that of a particle accelerated by a moving, partly ionised gas cloud. The acceleration occurs when the particle enters the cloud and then, due to elastic scattering of the magnetic field lines, has the same average motion as the cloud.

This can be expressed through two Lorentz transformations the first into the rest frame of the cloud (and for a particle that is sufficiently relativistic so that $E \approx pc$)

$$E_0' = \gamma E_0(1 - \beta \cos\theta_0), \tag{2.9}$$

this can be transferred back into the laboratory frame when the particle exists the cloud via

$$E_1 = \gamma E_1'(1 + \beta \cos\theta_1'). \tag{2.10}$$

Now, using the fact that within the rest frame of the gas cloud, all collisions are with the magnetic field and thus elastic, i.e. $E_1' = E_0'$, it is possible to characterise the fractional change in energy due to the cloud ($^{\Delta E}/_{E_0} = {}^{(E_1 - E_0)}/_{E_0}$) as a function of θ_0 and θ_1'

$$\frac{\Delta E}{E_0} = \frac{1 - \beta\cos\theta_0 + \beta\cos\theta_1' - \beta^2\cos\theta_0\cos\theta_1'}{1 - \beta^2} - 1 \tag{2.11}$$

The angular average can be taken to obtain the average fraction energy gain per encounter, $\xi = \langle {}^{\Delta E}/_{E_0} \rangle$,

$$\frac{\mathrm{d}n}{\mathrm{d}\cos\theta_1'} = \text{constant}, \; -1 \leq \cos\theta_1' \leq 1, \tag{2.12}$$

$$\therefore \langle \cos\theta_1' \rangle = 0. \tag{2.13}$$

Applying this to Equation 2.11, gives,

$$\frac{\langle \Delta E \rangle}{E_0} = \frac{1 - \beta\cos\theta_0}{1 - \beta^2} - 1 \tag{2.14}$$

Since the probability of collision is proportional to the relative velocity between the cloud and the particle,

$$\frac{\mathrm{d}n}{\mathrm{d}\cos\theta_0} = \frac{c - V\cos\theta_0}{2c}, -1 \leq \cos\theta_0 \leq 1, \tag{2.15}$$

$$\therefore \langle \cos\theta_0 \rangle = \frac{-V}{3c}. \tag{2.16}$$

Thus,

$$\xi = \frac{1 + \frac{1}{3}\beta^2}{1 - \beta^2} - 1 \sim \frac{4}{3}\beta^2. \tag{2.17}$$

Acceleration off a planar shock

The other case is acceleration off a planar shock. With a shock moving at velocity $-\vec{u}_1$ and with the shocked gas moving away from the shock with a velocity $\vec{u}_2$, Equation 2.11 still applies but with $\beta = {}^{V}/_{c} = \frac{|\vec{u}_2 - \vec{u}_1|}{c}$. This means that, when the angular average is

taken to obtain the average fraction energy gain per encounter, ξ, this is the projection of an isotropic flux onto a plane

$$\frac{\mathrm{d}n}{\mathrm{d}\cos\theta_1'} = 2\cos\theta_1', \;\; 0 \leq \cos\theta_1' \leq 1, \tag{2.18}$$

$$\therefore \langle\cos\theta_1'\rangle = 2/3, \tag{2.19}$$

Applying this to Equation 2.11 gives,

$$\frac{\langle\Delta E\rangle}{E_0} = \frac{1 - \beta\cos\theta_0 + \frac{2}{3}\beta - \frac{2}{3}\beta^2\cos\theta_0}{1-\beta^2} - 1 \tag{2.20}$$

The average of $\cos\theta_0$ is again the projection of an isotropic flux onto a plane. However, this time over the range $-1 \leq \cos\theta_0 \leq 0$ so that $\langle\cos\theta_0\rangle = -\frac{2}{3}$. Thus,

$$\xi = \frac{1 + \frac{4}{3}\beta + \frac{4}{9}\beta^2}{1-\beta^2} - 1 \sim \frac{4}{3}\beta. \tag{2.21}$$

Comparison

Both scattering off shocks and clouds is capable of accelerating particles to high energies. However scattering off clouds only gives a fractional energy gain that is proportional to β^2 (hence it is commonly referred to as second-order Fermi acceleration) whereas shocks can accelerate at an efficiency proportional to β (hence it is referred to as first-order Fermi Acceleration). Since first-order acceleration is so much more efficient at accelerating it is expected that it will be the dominant acceleration mechanism so second-order acceleration is not considered any futher.

2.2.2. First-Order Fermi Acceleration Spectrum

In order to find the value of the exponent in Equation 2.8 we require a value for ξ (given in Equation 2.21) and P_{esc}. Assuming a large, planar shock, the rate of encounter is the projection of an isotropic cosmic ray flux onto the shock front. For cosmic rays of density ρ_{CR}, the rate of encounters is given by

$$\int_0^1 \mathrm{d}\cos\theta \int_0^{2\pi} \mathrm{d}\phi \frac{c\rho_{CR}}{4\pi}\cos\theta = \frac{c\rho_{CR}}{4} \tag{2.22}$$

The escape of particles is by convection downstream, away from the shock is $\rho_{CR} \times u_2$, thus

$$P_{esc} = \frac{\text{rate of escape}}{\text{rate of encounter}} = \frac{\rho_{CR}\, u_2}{c\,\rho_{CR}/4} = \frac{4u_2}{c}. \tag{2.23}$$

Substituting Equation 2.21 and Equation 2.23 into Equation 2.7,

$$\gamma = \frac{P_{esc}}{\xi} = \frac{3}{u_1/u_2 - 1}. \tag{2.24}$$

Figure 2.2b shows the laboratory view of the shock, moving to the shock rest frame. Relative to the shock, particles upstream of the shock are moving at a velocity $\vec{u}_1$ whereas downstream they are moving at a velocity $\vec{u}_2$. The equation of continuity requires that

$$\rho_1 u_1 = \rho_2 u_2 \tag{2.25}$$

(ρ is the gas density) and since this is in the region of strong shocks, and with γ_{HCR} as the ratio of their heat capacities (Longair, 2011)

$$\frac{\vec{u}_1}{\vec{u}_2} = \frac{\vec{\rho}_2}{\vec{\rho}_1} = \frac{\gamma_{HCR} + 1}{\gamma_{HCR} - 1}, \tag{2.26}$$

For a fully-ionised plasma, $\gamma_{HCR} = \frac{5}{3}$, thus the ratio of these velocities, $\vec{u}_1/\vec{u}_2$, is 4. This gives a value of $\gamma = 1$, and a spectrum (from Equation 2.8) of

$$N(E)\,\mathrm{d}E = \text{constant} \times E^{-2}\,\mathrm{d}E. \tag{2.27}$$

Thus a power law is derived with a harder spectrum than that observed in cosmic rays. This problem cannot be neglected, however, more detailed analyses, for example including the feedback into the magnetic field of the cosmic rays, somewhat soften the index. Also, there is evidence for spectral softening as cosmic rays diffuse away from putative cosmic ray sources. Though not the final solution, this method clearly shows that it is possible for charged particles to be efficiently accelerated to high energies in planar shocks and to produce a power-law spectrum. For a more up to date and detailed treatment see Blasi (2013).

2.2.3. Maximum Energy of Fermi Acceleration

With a mechanism for accelerating particles and potential sources which could accelerate particles with sufficient luminosity, the final constraint is that the particles must be accelerated to the observed energies. There are a number of constraints on the maximum

energy that can be reached during the acceleration of cosmic rays. Following the presentation of Drury (2012), the first constraint is that, since the particles are magnetically confined whilst accelerating, the gyroradius (r_g) of the particles has to be less than the size of the system. That is, for a particle of energy $E = cp$, in an accelerator region of size R with a magnetic field of strength B

$$r_g = \frac{p}{eB} = \frac{E}{ecB} < R \implies E < ecBR \tag{2.28}$$

This is an absolute upper bound and the actual energy obtained will be much lower. In the case of diffusive shock acceleration, the requirement can be tightened such that the diffuse length-scale of the particles is small compared to the shock radius,

$$\kappa \ll RV \tag{2.29}$$

where κ is the diffusion coefficient at momentum p and $V(= \dot{R})$ is the shock speed of a shock of radius R (assuming a spherical shock expanding into a stationary medium). Taking the assumption that the lower limit on the diffusion coefficient is the Bohm limit (that is, the mean free path is comparable to the gyroradius) then the limit is

$$\frac{1}{3} r_g c < \kappa \implies r_g \ll \frac{3VR}{c} \tag{2.30}$$

dropping the factor of three for simplicity and writing $V/c = \beta$ gives

$$E < \beta ecBR \tag{2.31}$$

This is the original Hillas limit that was first shown in the Hillas plot (Figure 2.3) where various astrophysical systems and the limit are shown on the B, R plane (Hillas, 1984).

In addition to the limits on the size of the accelerating region, there is also a limit on the time for the acceleration, that is the accelerating time (t_{acc}) must be less than the dynamical time of the system

$$t_{acc} < \frac{R}{V} \tag{2.32}$$

Since in a linear shock the time scale is given by

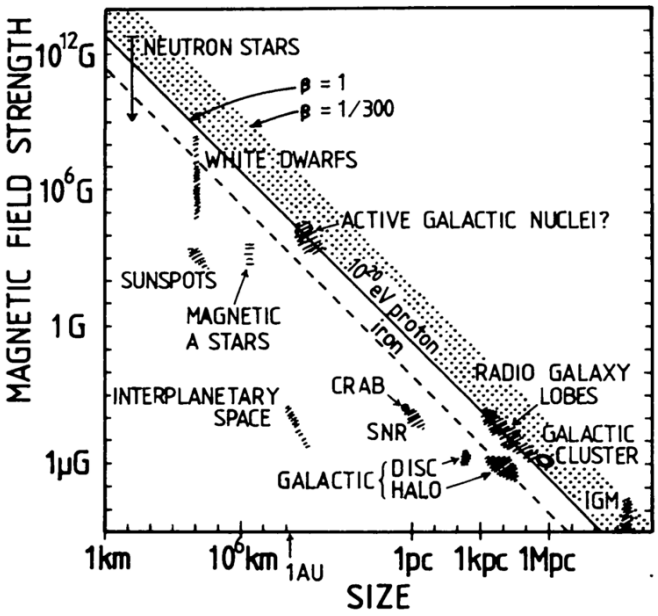

Figure 2.3.: The original Hillas plot (Hillas, 1984) showing various astrophysical systems on the B, R plane and the limit required to accelerate particles to 10^{20} eV, only those objects above the line are capable of doing so given the simple assumptions made.

$$t_{acc} = \frac{3}{u_1 - u_2}\left(\frac{\kappa_1}{u_1} + \frac{\kappa_2}{u_2}\right) \tag{2.33}$$

$$\approx 10\frac{\kappa}{V^2} \tag{2.34}$$

substituting Equation 2.32 into Equation 2.34 gives

$$\kappa < \frac{1}{10}RV \tag{2.35}$$

where v_1, v_2 are the upstream and downstream flow velocities and κ_1, κ_2 the corresponding diffusion coefficients (Drury, 1991). This is somewhat stricter than the Hillas limit (Equation 2.28) when the numerical factors are included.

There are also other, competing energy loss processes that are important (in particular for electrons) which can place tighter constraints on the acceleration time and maximum energy.

Given these constraints particle acceleration within the galaxy can reach all but the highest energies of the particles detected (Figure 2.1) and the mechanisms outlined above are sufficient to explain all of the measurements taken for this thesis. An explanation of the acceleration of the highest-energy particles is beyond the scope of this work, for further details see the extensive literature on this subject including Gaisser & Stanev (2008) for a review.

2.3. Dark Matter as a Source of Cosmic and Gamma Rays

Dark matter is a hypothetical type of matter that accounts for 26.8% of the energy in the universe (with baryonic matter making up 4.9% and dark energy the rest (Hinshaw et al., 2013)). Because it neither emits or absorbs electromagnetic radiation its existence has only been inferred from its gravitational impact on large astronomical bodies, such as galaxies and galaxy clusters. One of the prime theories for dark matter is that it is made up of WIMPs (weakly interacting massive particles) and that it is cold (nonrelativistic at the time of structure formation, Porter et al. (2011)). These WIMPs could annihilate or decay to produce standard model particles (that may themselves subsequently decay) that can be detected. The detection of these decay products would provide strong evidence for dark matter, but not conclusively prove it since astrophysical backgrounds could not be ruled out. Rather, confirmation will only be achieved in combination with direct detection or collider experiments detecting particles that show similar properties to those that produce the astrophysical signature.

A full treatment of the annihilation and decay of potential dark matter particles is beyond the scope of this thesis (see e.g. Porter et al. (2011) or Heo & Kim (2013)). What is relevant is that end products of these processes can include γ-rays, electrons, positrons, protons or antiprotons, thus it could act as a cosmic-/γ-ray source in a wide variety of ways. If it does appear as a major source it will likely appear as a line or with quite a steep cut-off in energy, which will be its most prominent signature. With γ-ray emission, though it may first be detected in other regions such as the galactic center, the best evidence will come from observations of targets with weak and/or well constrained astrophysical backgrounds such as dwarf-spheroidal galaxies, allowing for better characterisation of the γ-ray signal.

2.4. Cosmic Ray Diffusion

Cosmic rays diffuse throughout the galaxy (and beyond) interacting with the ISM. Two approaches are commonly taken depending upon whether the interest is in the spectra and interactions of the cosmic rays for comparison with direct measurements, when you treat them as particles, or whether the focus is on the stability and heating of the ISM when they are treated as a weightless, collisionless, relativistic gas. In this thesis, since I am interested in the spectra and interactions of cosmic rays, I will stick to the former treatment.

Cosmic rays propagation can be explained with the following equation (Strong et al., 2007) which describes an individual species,

$$\begin{aligned}\frac{\partial\psi(\vec{r},p,t)}{\partial t} &= q(\vec{r},p,t) + \vec{\nabla}\cdot(D_{xx}\vec{\nabla}\psi - \vec{V}\psi) \\ &+ \frac{\partial}{\partial p}\left(p^2 D_{pp}\right)\frac{\partial}{\partial p}\left(\frac{1}{p^2}\psi\right) \\ &- \frac{\partial}{\partial p}\left[\dot{p}\psi - \frac{p}{3}(\vec{\nabla}\cdot\vec{V})\psi\right] \\ &- \frac{1}{\tau_f}\psi - \frac{1}{\tau_r}\psi\end{aligned} \tag{2.36}$$

where

$\psi(\vec{r},p,t)$ is the cosmic ray density per unit of total particle momentum p at a position $\vec{r}$

$\psi(p)\,\mathrm{d}p = 4\pi p^2 f(\vec{p})\,\mathrm{d}p$ in terms of the phase space density $f(\vec{p})$

$q(\vec{r},p,t)$ is the source term including primary, spallation and decay contributions

D_{xx} is the spatial diffusion coefficient

$\vec{V}$ is the convection velocity

D_{pp} is the momentum diffusion coefficient for diffusive reacceleration

$\dot{p} \equiv \frac{\mathrm{d}p}{\mathrm{d}t}$ is the momentum gain/loss rate

τ_f/τ_r are the timescales for loss by fragmentation/decay

Though complex it can be broken down into a few general concepts. Since all the situations treated in this thesis assume that the object is in steady state, $\frac{\partial\psi(\vec{r},p,t)}{\partial t}$ (the rate of change of cosmic ray density) is equal to zero.

Sources are described by $q(\vec{r},p,t)$, that is objects that are accelerating cosmic rays (such as SNRs) and also cosmic rays that originate from the decay or interactions of other cosmic rays, e.g. proton-proton collisions producing pions which decay to produce electrons. The spallation and decay parts depend upon the progenitor species, both cosmic rays and (if required) target material. An additional potential source term is dark matter, as discussed in Section 2.3.

The change in ψ due to the bulk motion of cosmic rays through diffusion and convection is described by D_{xx} and $\vec{V}$. D_{xx} depends upon the particle species, energy and the magnetic field, it may be isotropic but, more realistically, it is anisotropic and influenced by the cosmic rays. $\vec{V}$ depends upon the nature of the galactic wind and the term $\vec{\nabla}\cdot\vec{V}$ describes the adiabatic momentum gain or loss due to such winds which have inhomogeneous, frozen-in magnetic fields whose inhomogeneities scatter cosmic rays.

D_{pp} describes diffusion in momentum, which is a continuous change in contrast to the two loss terms (τ_f,τ_r) which are catastrophic and remove the cosmic ray entirely.

$\psi=0$ is often chosen as a boundary condition at the edge of the Galactic Halo where particles escape into intergalactic space, this is obviously just an approximation since the intergalactic flux is not zero so this boundary condition can be relaxed to $\psi>0$ (at the cost of additional complexity in computation).

To solve it, this equation can be approximated and simplified in a myriad of ways. The most powerful for tool for modelling the formation and diffuse of cosmic rays and their products in the MW is GALPROP* (Section 6.1.2). This uses observations of a myriad of parameters to predict the cosmic and γ-ray spectra for different models of the MW cosmic ray acceleration and diffusion parameters.

2.4.1. Antiproton Ratio Predictions

Antiprotons are principally formed through the interaction of cosmic-ray protons with interstellar material. Thus, the source term in the diffusion equation for antiprotons is

$$q_{\bar{p}}(\vec{r},p,t)=\int_{E_p^0}^{+\infty}\frac{\mathrm{d}\sigma_{pH\to\bar{p}}}{\mathrm{d}E_{\bar{p}}}\left\{E_p\to E_{\bar{p}}\right\}n_H v_p\psi_p(\vec{r},p,t,E_p)\,\mathrm{d}E_p \tag{2.37}$$

*http://galprop.stanford.edu/

where n_H is the hydrogen density, v_p is the proton velocity, and $\psi_p(\vec{r}, p, t, E_p)$ is the cosmic ray proton density at a position, $\vec{r}$, time, t, with a momentum, p, and an energy, E_p. Using this source term (and additional terms for heavier element interactions which take the same form) in the diffusion equation allows for predictions as in Bottino et al. (1998) and M. Boudaud et al. (2015). These predictions of pure secondary production show a decrease in the antiproton ratio $\bar{p}/p$ with increasing energy as shown by the solid line (Donato et al., 2001) in Figure 2.4. Alternative production mechanisms that have additional components have also been postulated, such as from heavy dark matter (blue, short dashed line (Cirelli et al., 2009)), extragalactic components (black, long dashed line (Stephens, 1985)) and with an additional $\bar{p}$ component accelerated at cosmic-ray sources (black, dot-dash line (Blasi & Serpico, 2009)). An attempt by VERITAS to measure the antiproton ratio is described in Chapter 5.

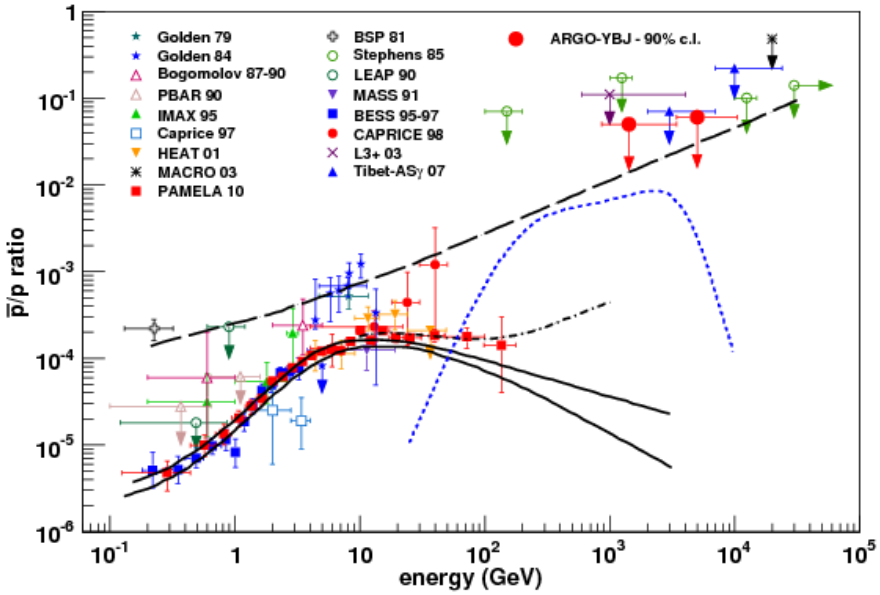

Figure 2.4.: Measurements and upper limits of the antiproton ratio (Bartoli et al., 2012) along with predictions (see text for explanations of the lines plotted). At low energies the turnover is an effect of the threshold for the formation of antiprotons from secondary particles and from the effects of solar modulation.

2.4.2. Positron Fraction Predictions

As with antiprotons, if positrons are only produced as secondary particles from cosmic-ray protons interacting with the interstellar medium then a similar decrease in the positron fraction $e^+/(e^+ + e^-)$ is expected, with the exact solution varying slightly due to the reduced interaction length of leptons. Early measurements of the positron fraction by balloons in the 1970's suggested an excess in the positron fraction above a few GeV, a hint that was reinforced by the HEAT-e$^{\pm}$/HEAT-pbar results from 1994 through 2004 in the 10 GeV to 50 GeV energy range. These hints were confirmed by PAMELA (Adriani et al., 2009) and observations by the *Fermi*-LAT (Ackermann et al., 2012b) and AMS-02 (Accardo et al., 2014) have reduced the uncertainty and increased the energy range (Figure 2.5). An attempt by VERITAS to measure the positron fraction is described in Chapter 5.

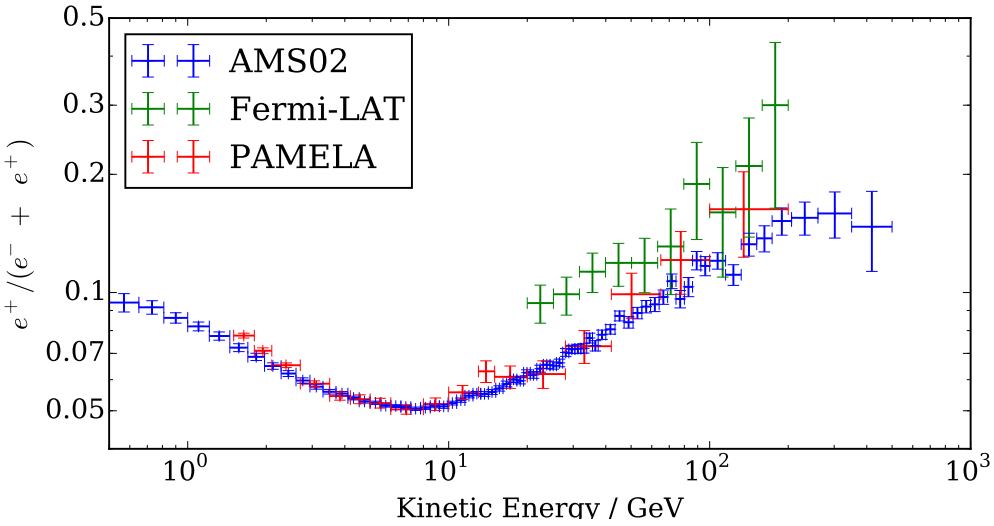

Figure 2.5.: The positron fraction as measured by PAMELA, *Fermi*-LAT and AMS-02 (from the Cosmic Rays Database (http://lpsc.in2p3.fr/crdb/), error bars are statistical and systematic combined). The rise above ∼10 GeV indicates that there must be a local source of positrons, the behaviour of the decrease in the fraction above ∼300 GeV will provide evidence on the source type.

A full and detailed summary of the models generated to match this prediction is beyond the scope of this thesis (see Linden & Profumo (2013) for a good summary).

However, the basics of the different classes of models is presented. The first class of models that quickly garnered interest were those that introduced additional source terms through the dark matter decay (see Section 2.3 for further details), essentially that dark matter decays to produce positrons. These models have come under severe pressure due to constraints from other observations, such as diffuse γ-ray emission and the antiproton ratio. The currently favoured solutions are based upon pulsars, with the cause being either a single pulsar that is local to the Earth's vicinity (such as the Geminga pulsar, (Boezio et al., 2009)) or an effect caused by the Earth's position relative to the spiral arms of the MW where the majority of pulsars are located (Blasi & Amato, 2012). Many additional models also exist, in order to discriminate between them detail of the behaviour of the fraction at higher energies is required, primarily, the energy at which the positron fraction starts to decrease again and its behaviour above that energy.

2.5. γ-ray Production from Cosmic Ray Interactions

γ-rays are produced by cosmic rays through four main cosmic-ray interactions:

hadron-hadron inelastic collisions produce intermediary products that decay to give γ-rays ("hadronic" emission, Figure 2.6a),

inverse-Compton scattering occurs when very-high-energy leptons inverse-Compton scatter low energy photons to γ-ray energies ("leptonic" emission, Figure 2.6b),

bremsstrahlung radiation occurs when leptons (primarily electrons) experience a sharp retardation through interactions with hadrons or leptons to produce γ-rays (Figure 2.6c), and

lepton-lepton annihilation (primarily electron-positron) is when a particle and its antiparticle collide and self annihilate to produce low-energy γ-rays ($\ll$100 MeV, since this is below the energy ranges considered in this thesis this is discussed no more).

Though not producing γ-rays, **synchrotron radiation** (Figure 2.6d) is also of significance for cosmic-ray energy losses and can be used to constrain cosmic ray populations and spectra.

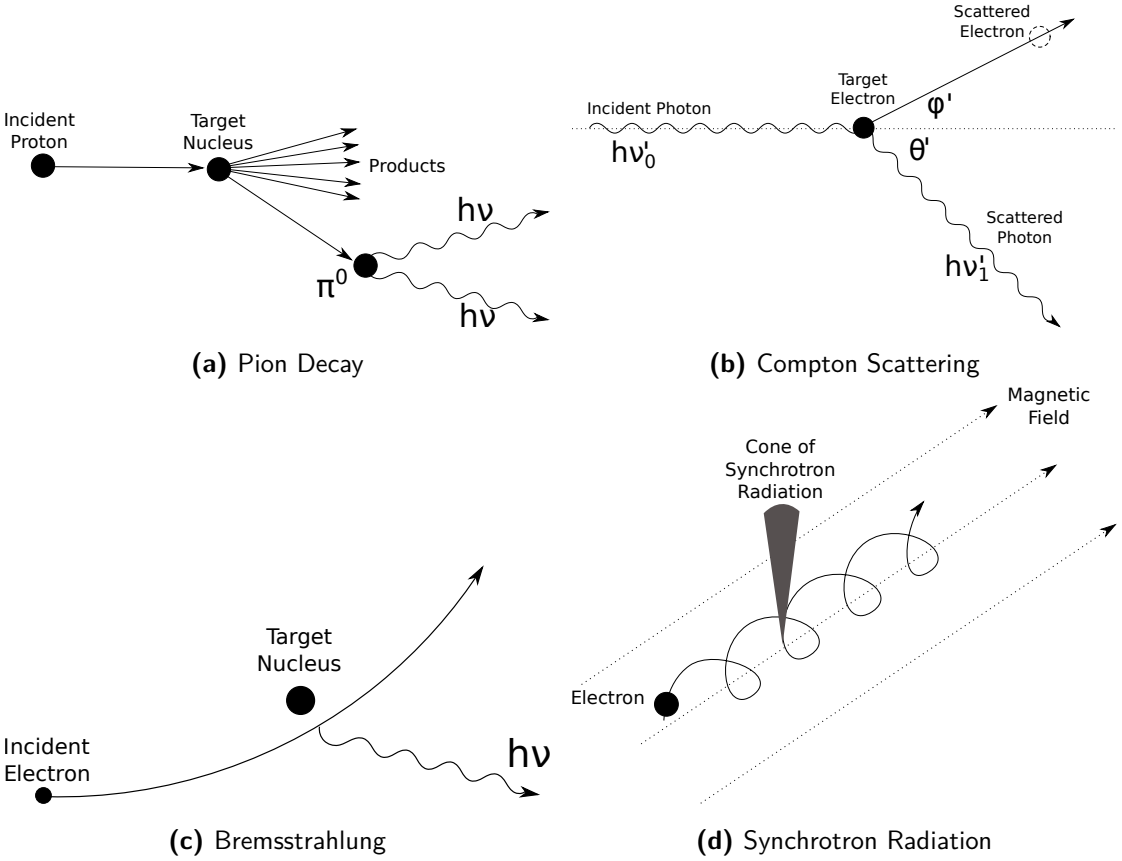

(a) Pion Decay **(b)** Compton Scattering

(c) Bremsstrahlung **(d)** Synchrotron Radiation

Figure 2.6.: The four energy loss mechanisms for cosmic rays that produce radiation that is clearly detectable and provide information about the cosmic ray spectra. The first three produce γ-rays, whereas synchrotron emits in the radio to X-ray bands.

2.5.1. Hadronic Emission

Hadronic emission occurs when hadronic cosmic rays (primarily protons, but also helium and heavier elements) collide with the interstellar medium (again, primarily protons) to produce particles that decay to produce γ-rays. The key reactions are:

$$\text{CR} + \text{ISM} \rightarrow \begin{cases} \pi^0(+\text{anything}) & \rightarrow \gamma\gamma \\ \pi^+ & \rightarrow \mu^+ + \nu_\mu \\ \pi^- & \rightarrow \mu^- + \bar{\nu}_\mu \\ \text{other products} & \end{cases} \tag{2.38}$$

The main interaction is between protons and hydrogen atoms/molecules/ions to produce pions (and other products including η-mesons but since pions are the primary product it is frequently referred to as *pion decay*), so that is what is considered in detail. Heavier elements are added using an enhancement factor that takes in their relative abundances and cross sections. In addition to producing γ-rays, hadronic emission also produces neutrinos which, as part of the multimessenger approach to astrophysics, could provide significant evidence on the emission processes (Subsection 1.3.6).

The spectra of γ-rays produced from pion decay from an incident proton spectrum $J_p(E_p)$ is onto a target gas of density n_H can be calculated using

$$\Phi_\gamma(E_\gamma) \equiv \frac{\mathrm{d}N_\gamma}{\mathrm{d}E_\gamma} = c\, n_H \int_{E_\gamma}^{\infty} \sigma_{inel}(E_p) J_p(E_p) F_\gamma\left(\frac{E_\gamma}{E_p}, E_p\right) \frac{\mathrm{d}E_p}{E_p}. \tag{2.39}$$

where c is the speed of light, E_p, E_γ is the energy of the proton, γ-ray, σ_{inel} is the inelastic cross-section and $F_\gamma\left(\frac{E_\gamma}{E_p}, E_p\right)$ is the energy spectrum of γ-rays produced from pp interactions (Kelner et al., 2006). These parameters are all either measurable in the laboratory (e.g. σ_{inel}), can be determined from observations (n_H) or can be assumed based upon local measurements ($J_p(E_p)$, though this can also be determined by working backwards from a measurement of the diffuse γ-ray spectrum). Working with measured values requires significant computational overhead, instead in this thesis I have used the formulation of Kafexhiu et al. (2014) applied in Zabalza (2015) which uses analytical functions fitted to each of the solutions. The principal aspect to note is that the rate of γ-ray production from pion decay is proportional to the amount of target material and the density of hadronic cosmic rays. Thus dense gas regions can be used as proton calorimeters, allowing the hadronic cosmic ray spectrum to be calculated for those regions since the γ-ray index is the same as the injected cosmic ray index.

At low energies, due to the threshold energy for pion formation, a *pion bump* is observed at $\sim$70 MeV below which the production of γ-rays is heavily suppressed. This

bump is a characteristic of pion decay and is difficult to explain with any other processes, thus is considered a smoking gun for the presence of hadronic cosmic rays.

2.5.2. Inverse-Compton Scattering

Inverse-Compton scattering occurs when a low energy photon collides with a high energy electron and is upscattered, this is physically the same as standard Compton scattering where the reverse occurs.

Cross Section

The likelihood of a collision occurring depends upon the cross section, which depends upon the energy of the incident photon in the rest frame of the electron. The full (quantum mechanical) calculation gives the Klein-Hishina cross section (using $x = h\nu_0'$)

$$\sigma_{KN} = \frac{3}{4}\sigma_T \left\{ \frac{(1+x)}{x^3} \left[\frac{2x(1+x)}{1+2x} - \ln(1+2x) \right] + \frac{1}{2x}\ln(1+2x) - \frac{1+3x}{(1+2x)^2} \right\} \text{cm}^2 \tag{2.40}$$

where the classical electron radius, $r_o = {}^{e^2}/_{m_e c^2}$, and

$$\sigma_T = \frac{8\pi r_0^2}{3} = 0.665\,\text{barn} \tag{2.41}$$

is the Thompson cross section.

In the low energy limit ($x \ll m_e c^2$, the Thompson regime)

$$\sigma_{KN} \sim \sigma_T \left(1 - 2x + \frac{26x^2}{5} \right) \tag{2.42}$$

whereas for $x \gg m_e c^2$ (the Klein-Nishina regime)

$$\sigma_{KN} \sim \frac{3\sigma_T}{8x} \left[\ln(2x) + \frac{1}{2} \right] \tag{2.43}$$

Energy Change of a Single Scatter

Inverse Compton scattering for a single scatter can be described by analysing Compton scattering in the rest frame of the electron. Working in the rest frame of the electron and using the geometry in Figure 2.6b conservation of energy gives

$$m_e c^2 + h\nu_0' = E_1 + h\nu_1' \tag{2.44}$$

where the prime ($'$) indicates that this is in the rest frame of the electron, ν_0' is the frequency before collision, ν_1' is the frequency after and E_1 is the total energy of the electron. Conservation of momentum also gives

$$\frac{h\nu_0'}{c} = \frac{h\nu_1'}{c}\cos\theta' + \gamma m_e \cos\phi' \tag{2.45}$$

$$0 = \frac{h\nu_1'}{c}\sin\theta' - \gamma m_e \sin\phi' \tag{2.46}$$

These equations reduce to

$$h\nu_1' = \frac{h\nu_0'}{h\nu_0'(1-\cos\theta') + m_e c^2} \tag{2.47}$$

The results of this collision then needs to be transformed back into the lab frame, with the result that the scattered photons are boosted to high energies. Assuming an incident angle of α_0 between the photon and the electron (in the lab frame), and a post collision angle of α_1

$$h\nu_0' = h\nu_0[\gamma(1-\beta\cos\alpha_0)] \tag{2.48}$$

$$h\nu_1 = h\nu_1'[\gamma(1+\beta\cos\alpha_1)] \tag{2.49}$$

the solution to this is messy but insightful observations can be made in the two extremes, $h\nu_0' \ll m_e c^2$ and $h\nu_0' \gg m_e c^2$.

If $h\nu_0' \ll m_e c^2$ (which is frequently called the Thompson regime) then

$$h\nu_1 \approx \gamma^2 h\nu_0(1+\beta\cos\theta')(1-\beta\cos\alpha_0) \tag{2.50}$$

this can then be averaged over all incident and scattered angles so that

$$h\nu_1 < 4\gamma^2 h\nu_0 \tag{2.51}$$

with an average energy, $\langle h\nu_1 \rangle$

$$\langle h\nu_1 \rangle = \frac{4}{3}\gamma^2 h\nu_0 \tag{2.52}$$

This means that the spectrum is very strongly peaked towards the maximum energy, and thus inverse Compton scattering is a very efficient process for transferring energy from electrons to γ-rays.

If $h\nu_0' \gg m_e c^2$ (the Klein-Nishina regime) then, Equation 2.47, after transformation into the lab frame, becomes

$$h\nu_1 \sim \gamma h\nu_0 \tag{2.53}$$

It can be summarised that the resulting γ-rays from inverse Compton scattering will have characteristic energies of $\gamma^2 h\nu_0$ in the Thompson regime or $\gamma h\nu_0$ in the Klein-Nishina regime, with a decreasing cross section with increasing energy. For an input differential electron spectrum that is proportional to $E_e^{-\Gamma_e}$, the resultant differential γ-ray spectrum is then proportional to $E_\gamma^{(-\Gamma_e+1)/2}$.

Power Radiated

Using the energy radiated per scatter and the cross section it is possible to calculate the power radiated through inverse Compton scattering. Working in the rest frame of the electron, the power radiated, P', depends upon the energy density of the incident radiation, $U'_{rad,0}$ as

$$P' = \sigma_T \, c \, U'_{rad,0} \tag{2.54}$$

and since power is Lorentz invariant,

$$P = P' = \sigma_T \, c \, U'_{rad,0} \tag{2.55}$$

Also,

$$U_{rad} = N h\nu \tag{2.56}$$

where N is the photon number density in the lab frame. Thus

$$U'_{rad} = N' h\nu' \tag{2.57}$$

$$= \frac{4\gamma^2 - 1}{3} U_{rad} \tag{2.58}$$

Substituting Equation 2.57 into Equation 2.55 and removing the initial power in the radiation field (in the lab frame) gives the energy loss rate for the electron population.

$$P = \sigma_T \, c \, (U'_{rad,1} - U'_{rad,0}) \tag{2.59}$$

$$= \sigma_T \, c \, \frac{4\gamma^2 - 1}{3} U_{rad,0} - \sigma_T \, c \, U_{rad,0} \tag{2.60}$$

$$= \frac{4}{3} \, \sigma_T \, c \, \beta^2 \gamma^2 U_{rad,0} \tag{2.61}$$

n.b. $\gamma^2 - 1 = \beta^2 \gamma^2$

Usage in this thesis

The full treatment of inverse Compton emission, in particular the approximations to the Klein-Nishina formula, are beyond the scope of this thesis, rather this analysis used the approximations described in Khangulyan et al. (2014), implemented in the code by Zabalza (2015). The important aspect to note is that this emission depends upon the target photon field density and the electron spectrum. Since electrons lose energy through other processes more readily than protons, this emission will largely be localised to regions close to cosmic ray acceleration sites where there is a higher density of electrons, and in regions containing large numbers of photons, such as star formation regions.

2.5.3. Bremsstrahlung Radiation

Bremsstrahlung, a German word meaning "braking radiation", is the name given to the radiation emitted when charged particles (primarily electrons) are decelerated in the Coulomb field of an atomic nucleus or electron. In the case of relativistic electrons, the emitted photon energy can match that of the accelerated electron. Thus, in the case of a power-law distribution of electron energies, the bremsstrahlung-emitted photons will have the same spectral index.

For the astrophysical objects studied in this work, at the energies I am concerned with, bremsstrahlung emission is significantly less than that of other processes and the models used to predict it are complex, see e.g. Strong et al. (2000) and Baring et al. (1999) for a detailed treatment. This thesis used the approximations of Baring et al. (1999), implemented in the code by Zabalza (2015).

What is important is that the cross section is proportional to Z^2 where Z is the number of protons in the nucleus and that the total emission depends on the gas density and the electron energy distribution. Thus the bulk of the emission will happen in dense, metal rich regions close to sources of electrons, i.e. in regions where there will be significant emission from other sources.

Bremsstrahlung, however, plays a very important role in the development of air showers, as described in Subsection 2.7.1 and Subsection 2.7.2.

2.5.4. Synchrotron Radiation

Though not emitted at γ-ray energies, synchrotron emission is an important emission process for VHE cosmic rays and it can provide a lot of information about the cosmic-ray population in the source. When a relativistic particle is travelling through a magnetic field it follows a helical path around the lines of field. This emission is beamed into a cone of angle $\theta \approx {}^{m_e c^2}/_{E}$ and emitted in a continuous spectrum (Figure 2.7) peaked at

$$\omega_c = \left(\frac{3eH}{2mc}\right)\gamma^2 \sin\phi, \tag{2.62}$$

where ϕ is the pitch angle between the direction of the magnetic field and the electron and H is intensity of the field normal to the velocity vector of the electron. For typical locations in the interstellar medium, H is measured in microgauss and E in GeV, $\omega_c \approx 100 H E^2 \sin\phi$ MHz, thus emission is in the radio waveband. Whereas, within objects such as pulsar wind nebulae, H can be many orders of magnitude higher and synchrotron emission can reach X-ray energies.

The average power radiated from an isotropic population of high energy electrons is

$$P = \frac{4}{3} c \sigma_T U_{mag} \beta^2 \gamma^2 \tag{2.63}$$

If it can be shown that the same population of electrons is responsible for both the synchrotron and inverse Compton emission (for example, through comparing their morphology of the emission region), the ratio of the relative loss rates (η) from Equation 2.63 and Equation 2.61 is

$$\eta = \frac{\left({}^{\mathrm{d}E}/_{\mathrm{d}t}\right)_{IC}}{\left({}^{\mathrm{d}E}/_{\mathrm{d}t}\right)_{sync}} = \frac{U_{rad,0}}{{}^{B^2}/_{2\mu_0}} \tag{2.64}$$

This can be used to determine the magnetic field in a region.

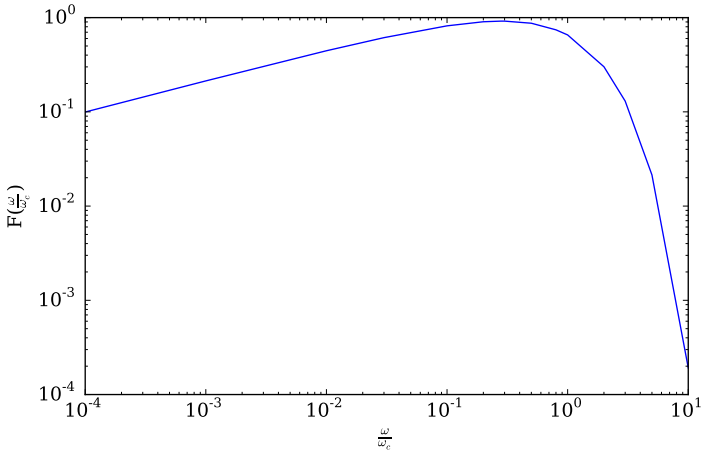

Figure 2.7.: The spectrum of synchrotron emission of a single electron in terms of ω/ω_c.

Synchrotron-self-Compton is the process whereby the radiation emitted by the same population of particles both emits the low energy radiation and upscatters it to higher energies. If η is greater than 1 then low energy synchrotron photons are upscattered by the same electron population to X-ray energies and then again up to γ-ray energies and the electrons rapidly loose energy at the highest energies and the source will be a powerful X-/γ-ray emitter. This is a process known as the synchro-Compton catastrophe.

2.6. Astrophysical Evidence for Cosmic Ray Acceleration

With a possible mechanism for cosmic ray acceleration in Fermi acceleration and an understanding of the physics behind cosmic ray interactions, observations of astrophysical objects can be used to confirm whether our understanding is correct.

Breaking down the cosmic rays into leptonic and hadronic components, the key signature for leptons is synchrotron radiation. This is clearly visible in a number of objects, for example Tycho's SNR (Figure 2.8a) which shows a shell of electrons moving ahead of the hot gas heated by the shock waves from the SNR. These electrons have been accelerated by the forward shock and are now emitting synchrotron radiation in the hard X-ray band showing that they must be at extremely high energies.

Evidence for hadronic acceleration can be found in the spectral energy distributions (SEDs) of IC443 and W44 (Figure 2.8b) where the pion bump, the sharp cutoff at

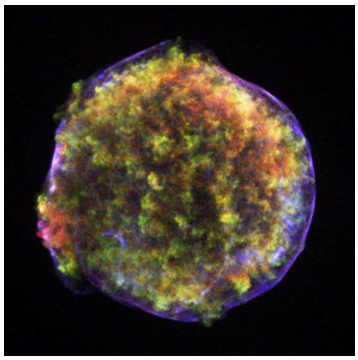

(a) A composite Chandra image of Tycho's SNR, the hot gas (red and green) is surrounded by a shell of electrons (blue) emitting synchrotron radiation in the hard X-ray band showing that they are at extremely high energies in a strong magnetic field (*Picture credit: NASA/CXC/Rutgers/J.Warren & J.Hughes et al.*)

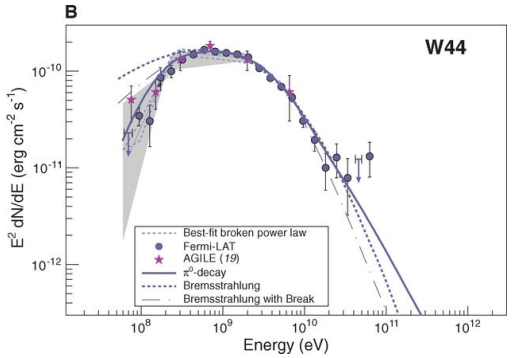

(b) The SED of W44 (Ackermann et al., 2013) showing the characteristic pion bump at low energies, along with a similar SED for IC 443 they show that hadrons must be being accelerated in SNRs

Figure 2.8.: Evidence of cosmic ray acceleration, synchrotron radiation showing the presence of VHE leptons (a) and the pion bump (b).

$\sim$100 MeV can best be explained using a hadronic model.

VHE γ-ray observations of these and similar objects (including pulsar wind nebulae, pulsars and binary systems) show that particles are being accelerated to tens if not hundreds of TeV. As exposures deepen and statistics increase at higher energies it will be possible to determine what happens at the highest energies and whether such objects are capable of accelerating particles up to the energies of the cosmic ray knee.

Observations of diffuse emission, especially from gas clouds close to known TeV accelerators, have the potential to show that particles are escaping from these sources and interacting. This will show that, as well as being capable of accelerating particles, they are able to power the observed cosmic ray population. Currently, VHE diffuse emission has been observed from the Galactic Centre by H.E.S.S. (Abramowski et al., 2014), VERITAS (Smith, 2015) and MAGIC (Fruck, 2015), and from the Galactic Plane by HAWC (Pretz, 2015) and, in particular from the Cygnus region, by Milagro (Abdo et al., 2008) which shows that there must be variations in the cosmic ray density within the

galaxy and since these regions contain multiple VHE sources they suggest that cosmic ray acceleration is occurring there and that the particles are able to escape into the ISM.

2.7. Observations of Cosmic and Gamma Rays

Cosmic ray and γ-rayastronomy has two major complications: they cannot be focussed to increase the signal-to-noise and they interact with the Earth's atmosphere. The result is that they either have to be detected above the atmosphere from space (or from very high altitude balloons) or, for energies above a few 10's of GeV, use their interaction with the atmosphere as a part of the detector.

The inability to focus them means that the number of particles detected scales directly with the size of the detector. For energies below 1 TeV the flux of particles is sufficient that they can be detected using a volume that can be flown on a satellite. These detectors work through the particles (or photons) interacting with the detector to deposit energy via various mechanisms, usually in the presence of a magnetic field to allow for charge discrimination. A variety of different techniques have been used based upon detecting the products of interactions within the detector volume, the methods used by satellites whose results are used in this thesis are presented in Section 1.3.

At higher energies the fluxes are lower and the products travel further, meaning that detector volumes would have to increase to a size that is not feasible for a satellite. Fortunately, there is a cross over, above a few 10s of GeV the products of the interactions between cosmic and γ-rays are sufficiently energetic that they create air showers, a cascade of particles moving through the atmosphere faster than the speed of light in the atmosphere, that can be detected from the ground. These particles can be detected either indirectly through their emission of Cherenkov light in the atmosphere or, if the energy is high enough, through the shower particles interaction in ground based detectors (for example in detectors like HAWC or, at extremely high energies, Pierre Auger) or the floresence light they induce (for example by Pierre Auger).

In this thesis I am principally using the VERITAS array which works in the detection range from a few 10s of GeV to $\sim$100 TeV, where it is the Cherenkov light from air showers that produces the best detection method.

2.7.1. Electromagnetic Air Showers

Electromagnetic air showers are those initiated by γ-rays or leptons, the only difference is in the first interaction. The first interaction will typically occur after the initiating particle has traversed one radiation length* of the atmosphere†, at an altitude of about 20 km. For γ-rays this is pair production, the production of an electron/positron pair which, on average, then goes on to interact in exactly the same manner as lepton induced showers, but one radiation length further down. After another radiation length, these VHE leptons will emit γ-rays through the process of bremsstrahlung (Subsection 2.5.3). This process of γ-rays pair producing and the resultant leptons emitting γ-rays continues down through the atmosphere and with the emission angle $\propto {}^{m_e c}/_{E}$ rad, the result is a remarkably tightly bunched cascade as can be seen in Figure 2.9a.

After a number of radiation lengths, the average energy drops to a point where ionization energy and radiation losses become equal. At this point the cascade has reached *shower maximum* and the number of particles in the cascade gradually diminishes and the cascade dies away. Depending on the energy of the initiating particle, this can occur well before it reaches sea level, thus, in order to observe the lowest-energy events observations must be conducted at altitude.

During the progress of the shower, if the leptons are above the threshold energy for Cherenkov emission (Subsection 2.7.3) then they will stimulate the atmosphere to radiate in the forward direction ($\theta \approx 1.3°$ at sea level) producing a shower of Cherenkov photons on the ground ($300\,\text{nm} \lesssim \lambda \lesssim 450\,\text{nm}$).

2.7.2. Hadronic Air Showers

In contrast with electromagnetic air showers, the initiating process in hadronic air showers is pion decay (Subsection 2.5.1), with much higher mass products than electromagnetic showers resulting in larger emission angles. Within a hadronic air shower some of the energy goes into electromagnetic sub-showers. The fraction of energy that goes into electromagnetic sub-showers, which are efficient at producing Cherenkov radiation, is quite small and thus the typical light yield from a shower is 2 to 3 times less than an electromagnetic shower. The hadronic interactions also result in larger emission angles

*$37.1\,\text{g cm}^{-2}$, the mean distance over which a high-energy electron looses all by $1/e$ of it energy by bremsstrahlung and $7/9$ of the mean free path for pair production

†the total thickness is about $1030\,\text{g cm}^{-2}$

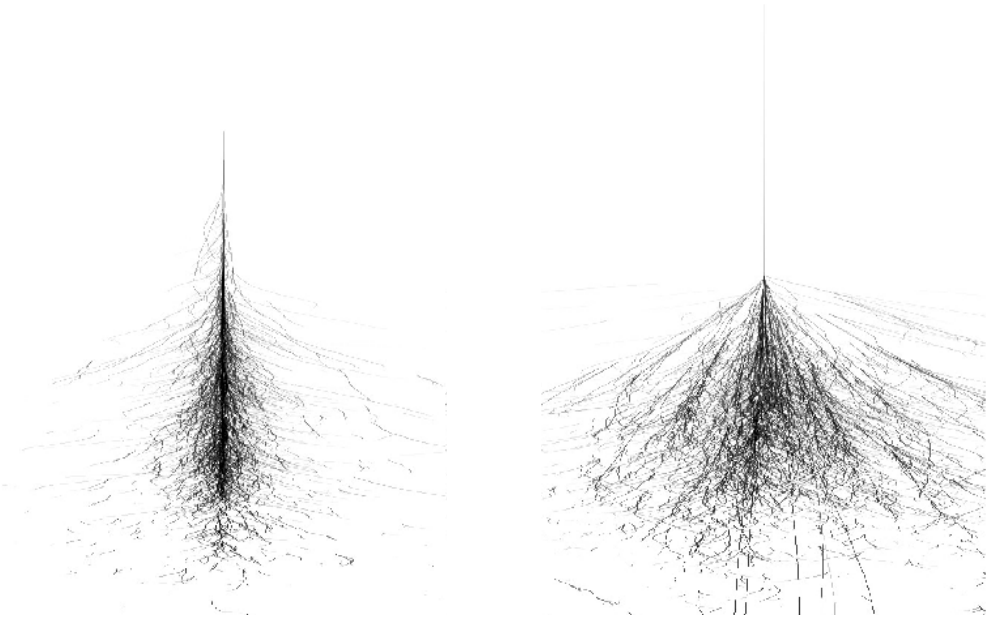

(a) An electromagnetic cascade **(b)** A hadronic cascade

Figure 2.9.: Simulations of γ-ray and proton showers: the γ-ray shower is more slender and smoother whereas the hadronic shower is broader and has substructure (Völk & Bernlöhr, 2009)

between products (and thus between the electromagnetic sub-showers), resulting in a less regular and geometrically larger shower (it is commonly described as "clumpy"), as is visible in Figure 2.9b.

2.7.3. Cherenkov Radiation

If a particle is travelling (at a velocity v) faster than the speed of light in a dielectric medium (c/n, where c is the speed of light in a vacuum and n is the refractive index) then it emits Cherenkov radiation. This occurs because the particle induces a polarization in the molecule as it passes, this change in polarization causes the molecule to radiate. If the particle is moving slower than the speed of light in the medium then the disturbance is symmetrical around the particle and there is no residual electrical field - therefore no detectable radiation. However, if $v > c/n$ then along the trajectory of the particle there will be a resultant dipole in the electric field which can produce a detectable effect. As the particle traverses the dielectric medium, each finite element radiates

a brief electromagnetic pulse. In general, these wavelets will destructively interfere, however, in the forward direction they will constructively interfere, as in Figure 2.10 where $\theta = \cos^{-1}(c/nv)$, producing a cone of light.

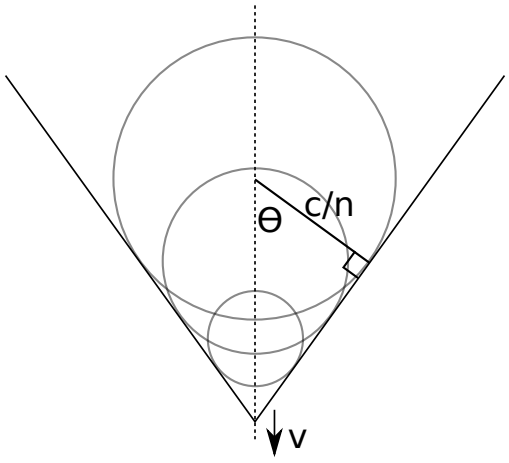

Figure 2.10.: Cherenkov radiation is emitted when a particle is travelling at a velocity, v, that is faster than the phase velocity of light, $= c/n$, then light is coherently emitted at an angle $\theta = \cos^{-1}(c/nv)$.

For electrons moving through the air at ground level ($n = 1.000029 \rightarrow \theta_{max} = 1.3°$, n decreases with altitude, reducing θ) the threshold energy is 21 MeV and the light yield in the visible range is about 30 photons m^{-2} (for a few hundred GeV γ-ray in the central 120m radius region from the shower axis) in the blue-UV part of the spectrum (Figure 2.11).

2.7.4. Observational Properties of Electromagnetic Air Showers

Since Cherenkov light is emitted in a cone around the high energy charged particles and electromagnetic air showers are tightly bunched, these cones align to produce a pool of light on the ground. The refractive index decreases with altitude, thus θ decreases, this causes the pool to form a ring structure (of radius ∼125 m centred around the shower axis (Figure 2.12)). Depending on the energy of the incident particle, the observation conditions and the elevation of the observation the density of Cherenkov photons hitting the mirror varies from ∼1 to a hundreds of photons per m^2. IACTs are positioned and designed based around the energy range they are targeting, VERITAS, which is aiming for the core energy range of IACT observations is at a moderate height to increase

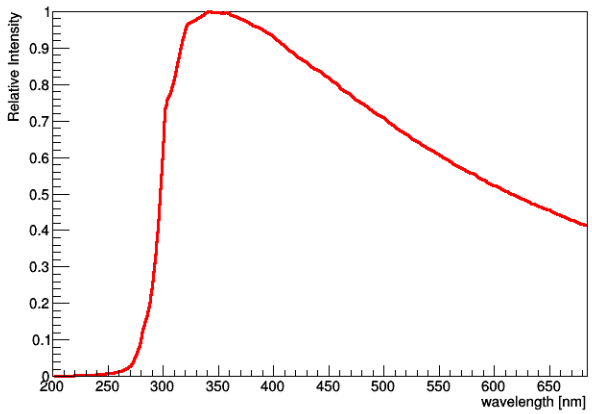

Figure 2.11.: The average Cherenkov spectrum from showers initated by a 500 GeV γ-ray. The spectrum continues to longer wavelengths but is not plotted since that is above the sensitivity range of the PMTs used in VERITAS.

the effective area at the cost of energy threshold. Each telescope has a mirror area of $140\,\mathrm{m}^2$ which is a compromise between the aim of collecting the faintest (lowest energy, $\sim$80 GeV) showers where the photon density is less than $1\,\mathrm{m}^{-2}$ (giving $\mathcal{O}(100\,\text{photons})$) with the cost of building larger telescopes.

2.8. Summary

This chapter has addressed the basic physics underlying the production, diffusion, interaction and detection of cosmic and γ-rays. These processes allow us to develop an understanding of how particles can be accelerated to such extreme energies, and thus provide insights into some of the most extreme places in the universe. By using different measurements of the different cosmic ray types and interactions we are able to constrain different aspects of the physics underlying the processes, this allows us to build up the complete picture of what is going on within the galaxy. The physics outlined above provides the basis for the detection method used by the VERITAS array and the physics studied in the Moon shadow (Chapter 5) and M 31 (Chapter 6) .

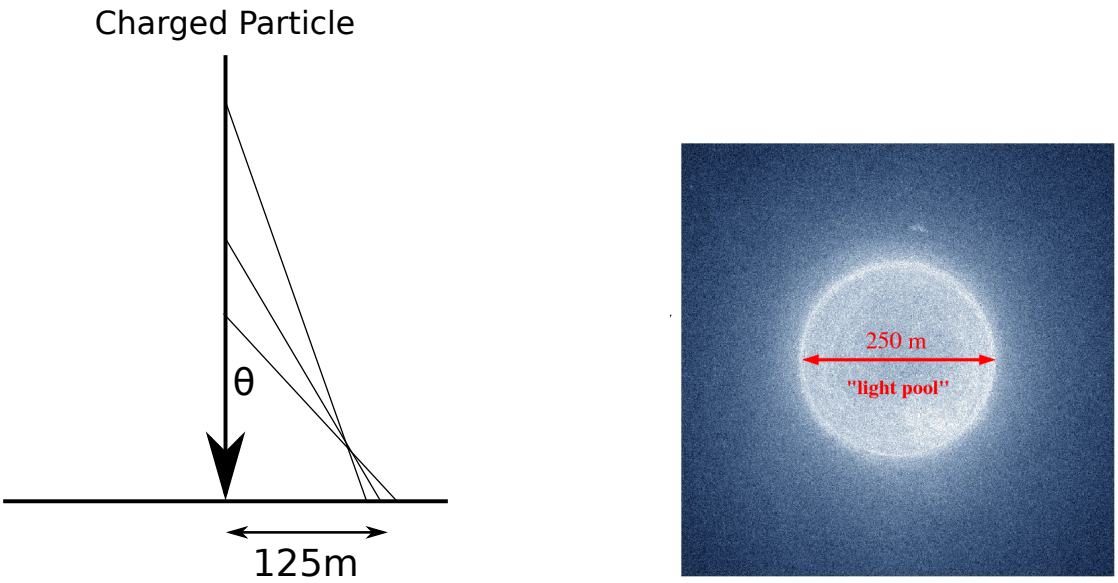

(a) Side on view

(b) A simulated Cherenkov light pool on the ground (Völk & Bernlöhr, 2009)

Figure 2.12.: With decreasing air density at higher altitudes, n decreases, thus the emission angle for Cherenkov emission from a single charged particle also decreases and it produces a ring of emission around the particle. With multiple particles these rings overlap to produce a light pool on the ground of radius $\sim 125\,\mathrm{m}$.

3. The VERITAS Array and Data Analysis

The Very Energetic Radiation Imaging Telescope Array System (VERITAS) is an array of four 12 m telescopes is located at the basecamp of the Fred Lawrence Whipple Observatory in southern Arizona, USA ($31°40'30.21''$, $-110°57'7.77''$) at an elevation of 1270 m (Figure 3.1). VERITAS started operating as a single telescope in early 2005 and was completed as an array in 2007. It has a field-of-view (FoV) of $\sim$3.5° with an angular resolution of $\sim$0.1° at 1 TeV and an energy range extending from $\sim$85 GeV to over 30 TeV (Holder et al., 2008). Since completion of the array it has undergone two major hardware upgrades and a number of smaller hardware and software improvements. The first major upgrade was conducted in the summer of 2009 when telescope 1 was moved to make the array more symmetric, increasing the overall sensitivity by $\sim$30% (Perkins & Maier, 2009). A second major upgrade was conducted in 2012 with an improvement in the trigger system and a replacement of all of the PMTs with high-quantum-efficiency bialkali devices (Kieda, 2011). It can currently detect a 1% Crab source in $\sim$25 hours (Park, 2015).

Since September 2012 VERITAS has pursued a programme of bright moonlight observations, allowing observations up to $\sim$80% Moon illumination through reducing the high voltages applied to the photomultiplier tubes or the use of optical filters. This has provided an additional $\sim$300 hours of observing time (an increase of $\sim$30%, (Griffin, 2015)).

3.1. The VERITAS Array

VERITAS consists of four identical telescopes, each one is comprised of four major units, the optical support structure and mirror, the camera, the trigger system and the data acquisition system.

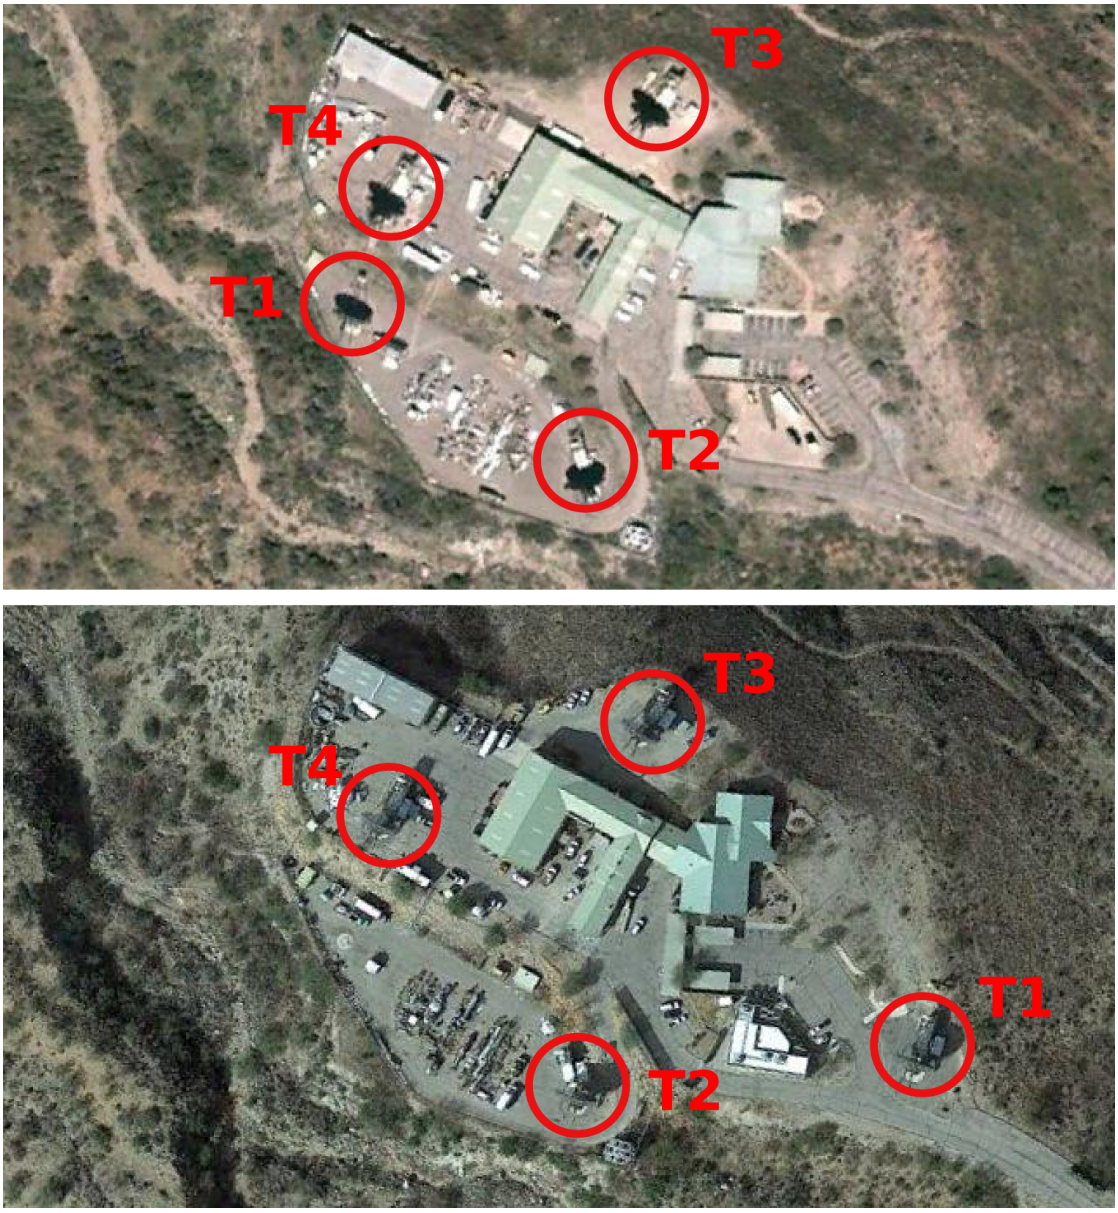

Figure 3.1.: The VERITAS telescope array, top is the old array before the move of telescope 1 (T1), bottom is the array in the more symmetrical arrangement after the T1 move. The telescopes are stowed pointing north to prevent damage to the cameras and reflections from the mirrors causing fires. *Picture credit: Gooogle Earth*

3.1.1. Optical Support Structure & Mirror

The telescope design consists of an alt-azimuth mount supporting a tubular steel optical support structure (OSS). Onto the OSS 345 identical hexagonal mirrors are mounted following the segmented Davis and Cotton design (Figure 3.2) to generate a 12 m mirror with an total reflector area of about 140 m^2. At the focus of these mirrors the camera is supported by a quadrapod arm structure which uses a mechanical bypass to direct the camera load onto a set of counterweights at the back of the OSS. Due to the load on the OSS it distorts with variations in zenith angle pointing, adding the potential for pointing inaccuracies and increasing the PSF. This distortion is regularly measured and (partially) corrected for in the telescopes' pointing.

Figure 3.2.: A section of the hexagonal mirrors on a VERITAS telescope. Each mirror is mounted on three custom mounts to allow for accurate alignment. Also visible is part of the OSS that holds the camera.

3.1.2. Camera

Each camera is located at the focus of the reflector and consists of 499 PMTs. The PMTs are UV-sensitive and chosen for their fast rise time and high quantum efficiency. They were upgraded in 2012 and the properties are outlined in Table 3.1. The main

improvements in the upgraded PMTs are an increase in the quantum efficiency and a narrowing of the pulse width. The pixel spacing is 0.15° giving a FoV of 3.5°. In front of the PMTs are light concentrators (modified Winston Cones (Winston, 1970)) that both reduce the dead space between the pixels, increasing the collecting area, and reduce the amount of background light entering the PMTs (Figure 3.3).

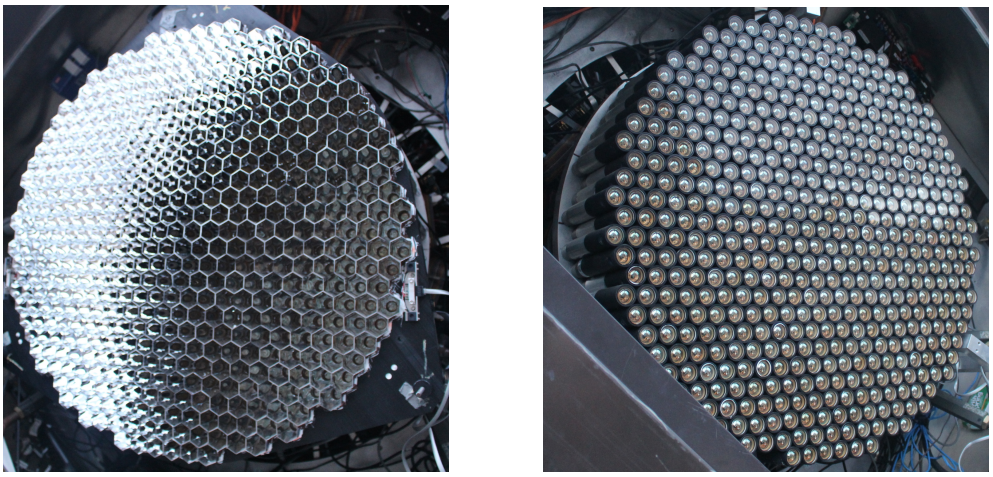

(a) With lightcones installed **(b)** Without lightcones installed

Figure 3.3.: The VERITAS camera with (a) and without (b) lightcones. Without the lightcones the 499 PMTs are visible. The lightcones, with their hexagonal entrances remove the gaps between the PMTs increasing their effective collection area (by ∼60%) and reducing the amount of background light that enters the PMTs.

The PMTs are powered with a multi-channel power supply that is capable of controlling each PMT individually. By varying the high-voltage (HV) supplied to each channel the gain in each PMT can be adjusted to flat-field the camera (see Subsection 3.2.2 for details on flat-fielding).

Both before and after the upgrade the pixels operate with a nominal gain of 2×10^5 (Otte, 2011), though for the same voltage, the post-upgrade PMTs have a gain that is ∼1.5× the pre-upgrade gain.

A high-bandwidth pre-amplifier is integrated into the base of the PMTs to provide an extra gain of 6.6. The pre-amplifier also provides a direct DC output for anode-current monitoring purposes (primarily for the protection of the PMTs from excess currents).

Table 3.1.: Properties of the PMTs used in the VERITAS cameras before and after the upgrade (Otte, 2011).

	Type	Q.E.	Single Photo Electron FWHM
Pre-Upgrade	Photonis XP2970/02	20% @ 320 nm	6.8 ns
Post-Upgrade	Hamamatsu R10560-100-20 MOD	37% @ 330 nm	4.2 ns

3.1.3. Trigger System

The trigger system receives the signals from the pre-amplifiers and is designed to trigger only on light from potential air showers, rather than from fluctuations in the night-sky background (NSB) or from local muons. The trigger has three hierarchical conditions as explained in detail in Weinstein (2008) and laid out in Figure 3.4.

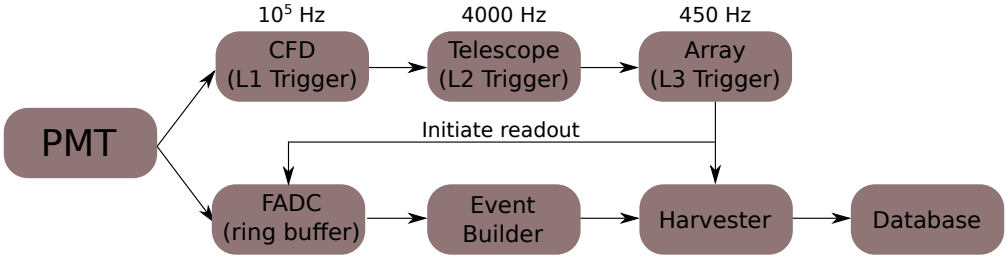

Figure 3.4.: A schematic showing the trigger and data acquisition systems used in VERITAS and the rates at each level of the trigger. The signal from the PMT is split, one part passes through the trigger chain, at an array level trigger the second part, which is stored in a ring bugger, is read out and combined with the signals from the other PMTs and telescopes in the data acquisition system.

The first level of the trigger is on the single-pixel level (L1) and combines a programmable constant fraction discriminator (CFD) and a rate feedback (RFB) circuit that reduces timing jitter and stabilizes the CFD trigger rate in the presence of changing NSB levels. The L1 trigger produces a 10 ns logic output pulse if the sum of the voltages from the PMT pulse and a time-delayed copy crosses a set threshold. The RFB system effectively raises the trigger threshold with increasing NSB (60 mV per MHz),

this has many advantages but also has potential issues with analysing data affected by the raising of the L1 trigger threshold for parts of the camera, especially when trying to detect a signal in an area of high NSB (Hall et al., 2003)

The second trigger is at a camera level (L2) and is often called the pattern selection trigger. The L2 trigger criterion is satisfied if 3 adjacent pixels surpass the L1 trigger within a coincidence window (typically 6 ns, (Zitzer, 2013)). The aim of this trigger is to reject single-pixel fluctuations due to NSB and PMT afterpulsing. Since the 2012 upgrade individual problem pixels can be removed from the trigger by the operator using a GUI (e.g. if a malfunction in the pixel causes it to repeatedly cross the L1 trigger threshold).

The final trigger level is an array level trigger (L3), the main purpose of which is to provide additional NSB trigger suppression and to reject local muons which, by nature of their limited light pool size, are only able to trigger a single telescope (since the move of telescope 1). It requires that two or more telescopes produce an L2 trigger within 50 ns and ensures a stereoscopic view of the event. This coincidence window is a compromise between the desire to have as narrow a window as possible to reduce accidental triggers yet maintaining sufficient width to allow for the shower front to pass across all four telescopes which is particularly wide at low elevations due to geometric effects. When an L3 trigger occurs, a logic signal (which has been appropriately delayed) is sent to all four telescopes to initiate the readout of the buffer by the DAQ system (see Subsection 3.1.4). During this readout time, no further triggers can be accepted, introducing a deadtime into the system. This deadtime is of the order of 10% at a rate of 400 Hz, see Appendix A for discussion of methods investigated to reduce this.

In order to set the trigger levels, bias curves are taken (Figure 3.5) where a run is conducted scanning across different trigger thresholds. If the trigger threshold is too low then the triggers are dominated by night-sky-background (NSB) this increases the deadtime of the array, reducing its sensitivity. As the trigger threshold level rises, the rate of NSB triggers decreases dramatically. Above a certain threshold the triggers are dominated by cosmic ray triggers and the trigger rate decreases at a lower more slowly since the cosmic ray triggers are less sensitive to the trigger level. The impact of the L3 trigger is also particularly apparent, with both the trigger rates and the break threshold lower in the array trigger than for the individual telescopes. The ideal is to have a trigger level that is as low as possible, to maximise the number of cosmic ray triggers, without

lowering the sensitivity of the array by swamping it with NSB triggers and increasing the deadtime. Since the NSB rate depends upon the brightness of the FoV, the trigger threshold is set for the expected NSB levels during those observations. For some sources which will have a known lower NSB and where a low energy threshold is particularly desirable a lower CFD threshold can be applied.

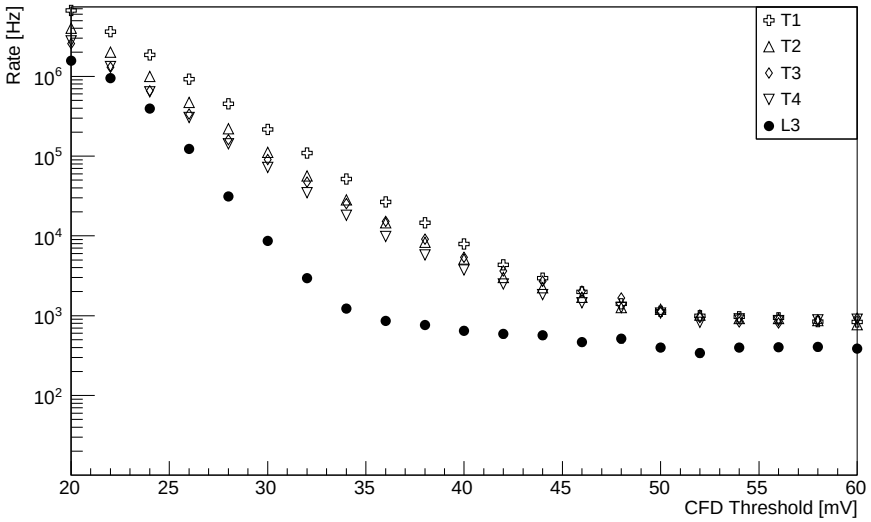

Figure 3.5.: The results of a bias curve following the upgrade of the PMTs. The CFD threshold is varied between 20 and 60 mV and the trigger rate varies, at low threshold levels the triggers are dominated by NSB, at about 35 mV the decrease in the L3 rate (black dots) changes slope as the triggers are now dominated by cosmic rays rather than NSB. As the threshold increases above this value the rate decreases as increasing numbers of cosmic rays fail to trigger the array. The L2 rates follow a similar pattern, but without the array trigger the break point and rates are higher.

3.1.4. Data Acquisition System

The data acquisition (DAQ) system employs 500 megasample-per-second flash analogue-to-digital converters (FADCs) (Hays, 2007). These FADCs continuously digitise the analogue PMT signal at a rate of 2 ns/sample which is then stored in a ring buffer for 32 µs while awaiting an L3 trigger signal. The FADC system has an 8-bit range giving

a dynamic range of 256 digital counts (dc). If this is saturated by the signal then an analogue switch connects the FADC to a delayed low-gain path extending the dynamic range to 1500 dc.

When an L3 trigger is received, the standard operation is that the DAQ reads out a section of the buffer (typically 16 (post-upgrade) or 20 (pre-upgrade) samples) for each PMT signal (see Appendix A for an explanation of alternative readout methods). This is then passed to a telescope-level event builder which assembles complete telescope events, tests them and then passes them on to an array-level data harvester. The harvester combines the telescope-level event information with array level information such as event numbers and times, before saving into a custom VERITAS data format for storing on disk.

Additional information, such as the trigger settings, the high-voltage values and the weather conditions (Subsection 3.2.4) are recorded to a database, along with the target, observation mode and any comments from the observers. This database is accessed during the early stages of the offline data analysis to provide relevant details and to assist in data quality monitoring (DQM, Subsection 3.4.1) and data selection.

3.2. Ancillary Hardware

In addition to the main data acquisition systems VERITAS employs a number of ancillary systems to monitor the telescopes performance, calibrate the other systems and assess and monitor the climactic conditions. A full summary of these systems and techniques is beyond the scope of this thesis but those which are particularly pertinent to the data selection or techniques that have been used or adapted for work in this thesis are outlined below.

3.2.1. Mirror Alignment System

The current mirror alignment method employs a raster scan technique whereby a CCD camera is placed at the focal point of the reflector to take a series of images as the telescope performs a raster scan about a star. Properly aligned mirror facets appear bright when the telescope is pointing at the star, if a facet appears brightest when pointing away from the star then the angle between the star and the facet is twice the misalignment angle. The mirrors, which are mounted on three alignment screws

(Figure 3.2) can then be adjusted to correct the alignment (McCann et al., 2010).

3.2.2. Laser & Flasher Systems

Inter-pixel calibration has been conducted using two methods during the lifetime of the array. Both rely on taking short (2 to 5 min), dedicated, calibration runs pointing the array at a dark patch of the sky The first method (before 2009) used a laser to uniformly illuminate the entire camera with a single flash and the telescope is externally triggered to read out. The second method (since 2009) uses an LED system (Figure 3.6) to uniformly illuminate the telescopes with a ramp of light pulses produced by 0 to 7 LEDs, again with an external trigger (Figure 3.6). Both techniques provide the same basic information, but the latter technique, with a ramp of illumination levels (using cheaper and more robust hardware) provides additional information and is thus the preferred method.

Figure 3.6.: The VPM (white box) that monitors the positions of the LEDs in the camera box to allow for correction due to deflection in the OSS. Besides the VPM camera is the flasher assembly contained inside a modified torch.

Three key calibrations can be conducted with the flasher system:

The relative timing information between pixels

Due to differences in cable lengths and electronic delays, timing differences between channels naturally arise. With a known arrival time at all channels, the measured

arrival time can be corrected for these delays to ensure that all channels within a telescope have the same relative timing.

The relative gain between pixels

Different pixels will have different gains, either due to the inherent properties of the PMTs (which will vary with pixel age) or pre-amps, or due to the applied voltage. By measuring the relative response between channels to a known illumination it is possible to correct for these differences by adjusting the gain, this is done monthly by adjusting the HV supplied to each pixel, a process known as flat fielding. A flasher run is taken every night for each observing mode, the relative gains are calculated and then used to correct each channel for that night's operation.

Absolute gain

The absolute gain of the pixels **from the first dynode of the PMT** can be measured using the response of the PMTs to different light levels. This is an important parameter in the calculation of the energy of the shower and thus is regularly monitored. As the PMTs age, the absolute gain will decrease, thus it is normal to increase the voltage applied to all PMTs at regular intervals to keep the absolute gains within a given range (this is normally done at the start of each observing season).

The full details of the method of analysing the information from a flasher run is provided in Hanna et al. (2010).

3.2.3. VERITAS Pointing Monitor (VPM)

VERITAS uses a pair of CCD cameras on each telescope to track the pointing of the telescopes. One camera is mounted by the mirrors and tracks a selection of bright stars within its FoV (close to the target), the other is mounted beside the flasher, half way up the camera support arms (Figure 3.6) and tracks four LEDs mounted on the light cone plate (one pair is just visible in Figure 3.3 on the right hand side of the plate, just above the mounting screw). By comparing the relative positions of the LEDs to the stars it is possible to correct the telescopes pointing, improving the PSF.

3.2.4. Atmospheric Monitoring

Since the atmosphere is an integral part of the detector, Cherenkov telescopes are particularly sensitive to weather, especially the presence of atmospheric particles and clouds,

and thus require monitoring to ensure that data being analysed is not being adversely affected. At VERITAS this is done in a variety of ways:

Observers Records

The traditional method is still the most important, the observers keep an eye on the weather and record the sky conditions. These observations are recorded in the database for each run and form the first pass at DQM.

Weather Station

A weather station records the wind speed, rainfall, temperature and humidity and this is recorded in the VERITAS database. This is primarily used by the observers to determine if it is safe to observe as there are limits on the safe wind speed and humidity for operations, obviously operations cannot be conducted whilst it is raining.

Far-Infrared (FIR) Pyrometer

VERITAS employs three FIR Pyrometers (FIRs) to detect cloud cover, one is mounted pointing vertically upwards, another two are mounted on telescopes two and three, pointing parallel to the optical axis. The temperature of a cloud is typically warmer ($-5\,^{\circ}$C to $-10\,^{\circ}$C) than the clear sky ($\sim -50\,^{\circ}$C). This means that when thin cloud is passing overhead/through the FoV of the telescopes the temperature is raised. Variations in the temperature are indicative of bands of cloud passing overhead. More difficult to detect, but just as detrimental to data analysis, in particular for absolute rate calculations, is thin cloud which is uniform but has raised the temperature of the sky.

Plotting the FIR against the L3 rate of the array clearly shows the correlation between the two and the FIR is useful for selecting periods of a run that are not affected by cloud and can thus be analysed, enabling the extraction of useful data which would otherwise be discarded.

Ceilometer

VERITAS uses a ceilometer (a form of LIDAR; Vaisala CL51*) to detect the height and thickness of cloud vertically above the array. The ceilometer is particularly useful

*http://www.vaisala.com/en/products/ceilometers/Pages/cl51.aspx

for detecting uniform thin cloud that is hard to detect using the FIRs, it also provides detailed quantitative information on the thickness and types of cloud and on the presence of dust in the lower atmosphere. Work is currently under way to determine if it is possible to use this information to "correct" the observations for the presence of thin cloud/dust.

All-Sky Monitor and VPM cameras

An all-sky monitor provides a view of the complete sky and the VPM cameras (Subsection 3.2.3) provide views of the night sky where the telescopes are pointing. On both devices the integration time and other settings can be adjusted to allow the observers to monitor cloud in greater detail than is possible with the naked eye.

3.3. Online Data Analysis

VERITAS has a simplified analysis program (quicklook) which provides a real-time analysis of the target and all observations conducted that night (Wakely et al., 2003). This analysis is useful for monitoring the quality of the data being taken and for determining whether a source is "*flaring*" and if so whether it satisfies the trigger criteria for additional observations.

3.4. Offline Data Analysis

The aim of the analysis is to work backwards from the detected signal to make an estimate of the spectrum and morphology from a point in the sky, and from that the source morphology (if the source is extended). To do this it needs to invert all of the processes that affect the γ-rays as they are converted into the recorded traces, for example, the initial interaction height, the volume of atmosphere being sampled, the shower generation and transmission through the atmosphere, the reflectivity of mirrors and quantum efficiency of the cameras, losses in the electronics, and so on. Since there are so many variables, many of which are impossible to measure on their own, it is impossible to directly solve an equation for this, especially as there are many a priori unknowns in the generation of the shower, e.g. shape, height, position and energy. Therefore, instead the showers must be compared with simulations to make an estimate of the input properties that result in the measured output. That there are no calibration

TeV sources to validate these simulations against is a caveat that must be remembered with all TeV results.

Compounding this challenge is the significant background that IACTs face, for faint sources a run might contain $\mathcal{O}(10^6)$ events, of which only a handful are due to γ-rays, the remainder being triggers caused by protons, electrons and fluctuations in the NSB. Selection criteria have been developed (see Subsection 3.4.6) that are very efficient at reducing this background and significant work is ongoing to improve this differentiation.

Within the VERITAS collaboration, two different analysis packages are employed, Eventdisplay (ED) and VEGAS (VEritas Gamma-ray Analysis Suite). Since I have worked entirely with VEGAS for the duration of my thesis I will explain the VEGAS analysis method. Eventdisplay is largely similar but, for the sake of clarity, the differences will not be explored within this work. Results are generally verified with both analysis packages prior to publication.

3.4.1. Data Quality Selection

All data that is taken by VERITAS is passed through an early-stage analysis which is tuned to produce plots to allow for the data quality to be assessed and the results are saved to a database. Before conducting a detailed analysis of the data these plots are checked to ensure that the data meets suitable quality standards, for example periods affected by clouds are identified, the tracking is checked and any problems with the camera are flagged. A next-day look at the DQM plots is conducted by a member of the collaboration to provide a preliminary selection and to check that the telescopes are operating correctly. In addition, before any data was analysed for this thesis I conducted my own DQM to ensure that all the data was assessed to the same level and that data that failed to meet the required standard was removed.

3.4.2. Camera Calibration and Gain Correction

The first stage of the analysis is to download the target and tracking information from the database, along with records of the HV and weather throughout the runs. A check is made to identify any bad channels in the camera, for example channels that are always or never triggering due to a fault, these channels are removed from the analysis. The daily flasher run is then analysed to determine the gain and timing parameters for each channel. These values are then used to flat-field the camera and correct the timing in

the data runs (Subsection 3.2.2, see also Hanna et al. (2010)).

The PMTs are AC coupled and the signal is measured as a negative fluctuation from the mean of the trace. Noise, primarily NSB but also electronic noise, causes fluctuations around the mean of this trace. In order to prevent the signal fluctuating positively it is offset from zero with a pedestal (nominally −16 dc). A measure of the magnitude of the noise is the RMS of the difference of each sample from the pedestal, this is called the pedvar.

During a data run the telescopes are subject to an artificial trigger at a rate of 1 Hz to record the NSB, this allows for determination of the pedestal and the pedvar. A first pass at the integration of the charge is also conducted, this is done by taking the average trace of the event in a camera and using that to determine where to place the seven sample readout window to maximise the integrated signal (Figure 3.7). To do this the *TZero* value is calculated, this is the 50% point of the falling edge of the trace, calculated by linear interpolation. The window is then placed with 30% of the readout window before *TZero*, this ensures that the entire trace is read out, the signal in each channel is then integrated using the same readout window position.

3.4.3. Image Cleaning

The event image is now processed, the raw image (Figure 3.8a) is gain corrected and the pedestal is removed. The result is a noisy image of the shower (if one is present) as in Figure 3.8b. The brightest pixels are clearly visible on the left hand side but it is difficult to determine whether the signal in the weaker pixels is from the shower or whether the pixels are only recording NSB fluctuations, this is especially true for smaller images which have fewer bright pixels. To reduce the impact of these NSB fluctuations image cleaning (Reynolds et al., 1993) is employed, which is the process whereby a decision is made as to whether a pixel is recording a part of an image of a shower or just the NSB. This is done in a variety of ways, but the basic process is the same for all.

The first pass is to select pixels that have almost certainly been triggered by the shower, that is pixels that have a large signal. This is done by selecting pixels that have a signal-to-noise (S2N) ratio that is at least 5, that is, their integrated charge within the integration window is at least 5 times the pedvar for that channel. Cleaning thresholds are set as a multiple of the pedvar to work with the lowest possible threshold to ensure that the smallest images are recovered (increasing the sensitivity and reducing the energy

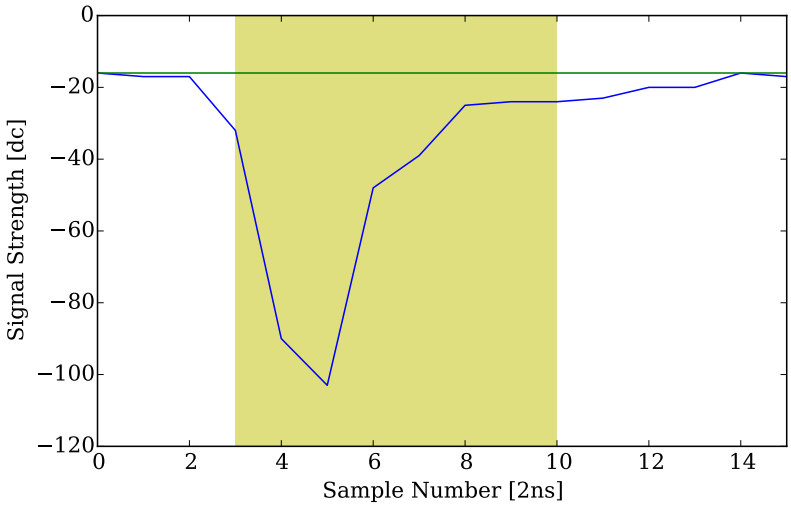

Figure 3.7.: A trace from a single channel that has received a significant signal. The FADC boards record a negative pulse but it is customary to work with absolute values. The 7 sample (1 sample = 2 ns) window has been placed to integrate the bulk of the trace whilst excluding the background. The trace is offset from zero by a pedestal of -16 dc.

threshold). If the threshold is set to a fixed value then it would either be discarding some of the observed showers (if the NSB was low) or it would be increasing our false trigger rate (if the NSB was high). These pixels are now identified as *picture* pixels and are retained.

With the *picture* pixels identified, their immediate neighbours are now examined. Since they are adjacent to a *picture* pixels they are likely to be a part of the shower, thus they are subjected to a lower cleaning threshold, a S2N of 2.5. These pixels are now identified as *boundary* pixels and are retained. After both passes have been conducted the image is said to have been cleaned (Figure 3.8c), with the retained pixels referred to as *image* pixels.

This basic cleaning (standard cleaning) was used by the Whipple collaboration but has since been replaced by *EventDisplay* cleaning (so named as it was first implemented within Eventdisplay), wherein an additional check is also conducted. If a pixel has passed cleaning but has no neighbouring *image* pixels then (due to the shower physics) it is highly likely that the pixel is only recording noise and it is removed (Figure 3.8d).

This additional level of cleaning removes pixels that will likely distort the reconstruction of the shower, either reducing the quality of the fit or making it fail completely.

When a shower is imaged there will be a time gradient due to the different propagation times from different parts of the shower. The first pass at image cleaning uses an integration window that is fixed in time to the camera average, this will reduce the sensitivity to parts of the shower which are away from the camera mean. To correct for this, after the image has been cleaned, the time gradient is calculated and the integration windows are varied across the camera to allow for this, the cleaning process is then repeated (Holder, 2005).

Time Cleaning

In order to improve the image cleaning I have worked on adding an additional constraint based upon the timing of the signals between adjacent pixels. To do this an additional requirement was placed on the *boundary* pixels, their *TZero* has to be within a given time of the *TZero* of their adjacent *picture* pixel. This will be implemented as the default in VEGAS when the next round of simulations are analysed. For full details of this work see Section 4.1.

Fixed Threshold Cleaning

In regions with gradients in the NSB, as were frequently encountered during this work, the existing image cleaning methods have a significant drawback in that they will have a gradient in the cleaning threshold. The result of this is that parts of the camera with the higher threshold will have a lower sensitivity and a higher energy threshold. This is often seen in regions around stars where the starlight raises the NSB (Mohanty et al., 1998). To reduce this a fixed cleaning threshold across the camera can be used, thus, instead of using a S2N threshold the *picture* and *boundary* thresholds are given in digital counts for the integrated signals. Since the values were fixed for the entire analysis (reproducing all of the simulations required for multiple thresholds is significantly time consuming) optimal values had to be chosen depending on the observed pedvars in order to maximise the signal but minimise the impact of the gradient. See Section 4.2 where this was employed for the selection of the thresholds in each mode of operation and further details on the sensitivity resulting from those selections.

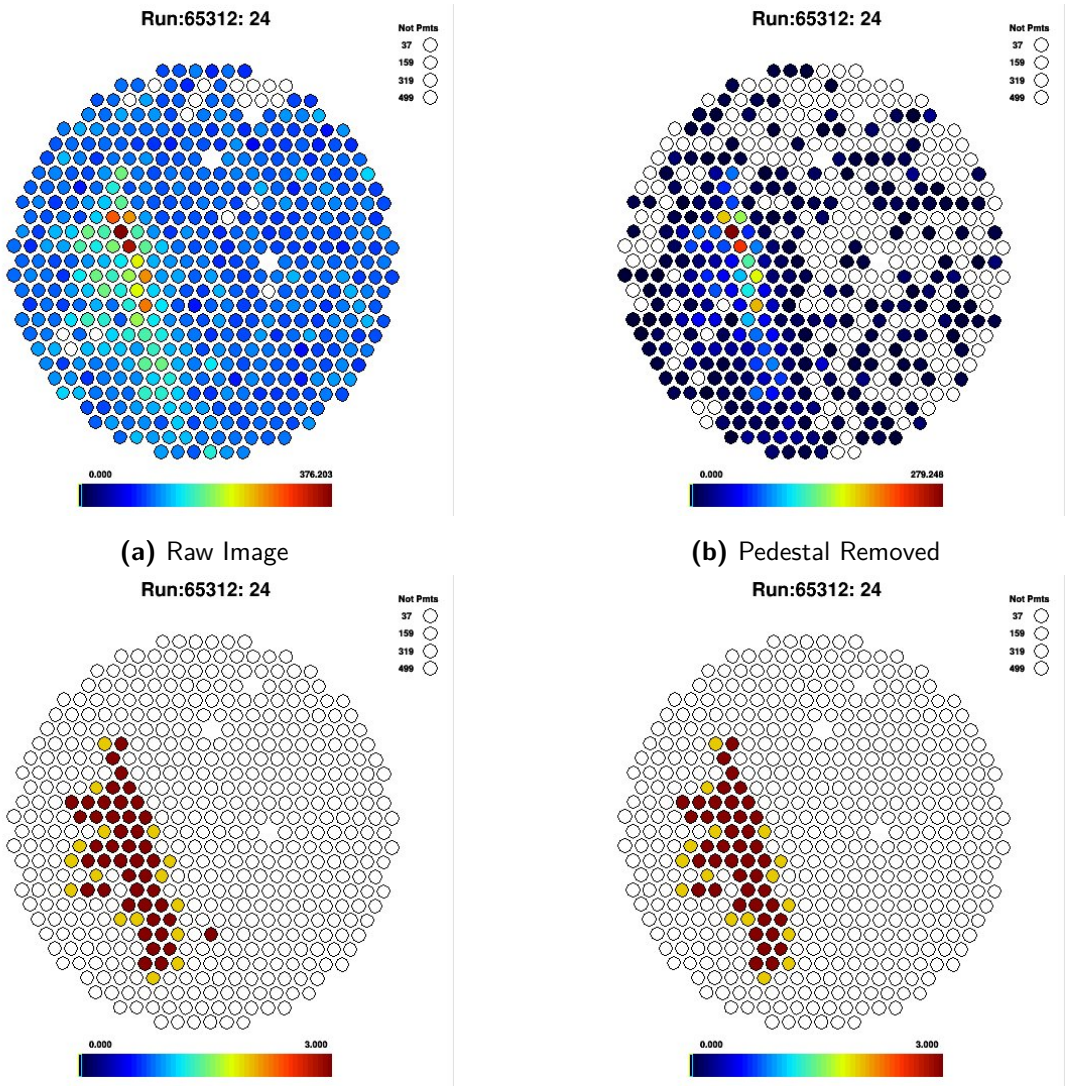

(a) Raw Image

(b) Pedestal Removed

(c) Standard Image Cleaning

(d) EventDisplay Image Cleaning

Figure 3.8.: It is possible to identify this large shower in the raw image in the camera (after gain correction and bad channel removal). However, this is not normally the case for smaller showers. After removal of the pedestal the situation is clearer. When image cleaning is applied the *picture* (brown) and *boundary* (yellow) pixels are clearly identified. The single isolated pixel in the standard image cleaning is unlikely to overly adversely affect this large shower but it is clearly an outlier and the analysis benefits from its removal using *EventDisplay* image cleaning.

3.4.4. Hillas Parametrisation

After the event has been cleaned, if a shower image remains, it is parameterised for further analysis using the method of Michael Hillas (Hillas, 1985), hence their name *Hillas* parameters. *Hillas* parameters are calculated from a moment analysis of the shower (and some related quantities) and have proven to be very useful in characterising showers, distinguishing between different source particle and determining the energy and source position. The main *Hillas* parameters for a single telescope (that are still used in stereoscopic arrays) are (Fegan, 1997):

Size
: The total integrated light content of the shower, i.e. the sum of the channels that pass image cleaning.

nTubes
: The number of *image* pixels.

Width
: The rms spread of light along the minor axis of the image, a measure of the lateral development of the cascade.

Length
: The rms spread of light along the major axis of the image, a measure of the vertical development of the cascade.

Distance
: The distance from the centroid of the image to the centre of the field-of-view of the camera.

3.4.5. Stereoscopic Event Reconstruction

With all of the individual telescope images cleaned and the *Hillas* parameters determined, the images can now be combined to reconstruct the shower. With two of more telescopes imaging the same shower*, stereo reconstruction produces a new set of parameters that are far more powerful than the single-telescope *Hillas* parameters, improving

*The number of telescopes imaging the shower is referred to as $nTels$.

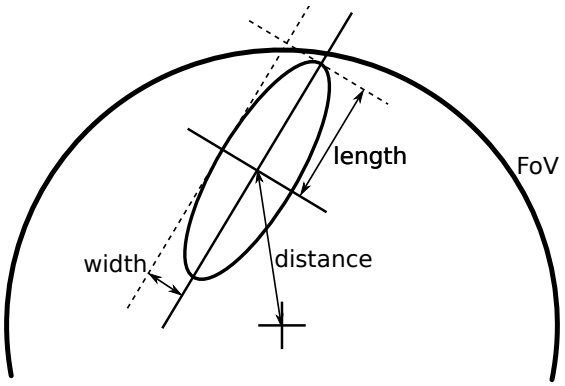

Figure 3.9.: The main *Hillas* parameters for a single single telescope.

the γ/hadron separation. The stereoscopic view of the showers also enables the reconstruction of the trajectories of the showers and the elimination of triggers caused by muons which will likely fail such reconstruction.

Before reconstruction is undertaken, a set of quality cuts are applied (Table 3.2). These cuts remove images that will be problematic to reconstruct, either because they are too small or because they are too close to the edge of the camera and liable to be truncated.

Table 3.2.: The standard quality cuts used in VEGAS, the size cut varies depending upon the analysis, with lower *Size* cuts reducing the energy threshold whereas a higher *Size* cut improves the PSF and gamma/hadron separation.

Cut	Value
$nTubes$	>5
$Distance$	$<1.43°$
$nTels$	≥ 2
$Size$	variable

Shower Source and Impact Location

Showers develop along the original direction of the initiating particle. With at least two good images of the shower it is possible to conduct a geometric reconstruction of the shower to find the source location (the point on the sky where the particle came from) and the impact location (the point where the center of the shower hits the ground). The

fitted ellipse has a major axis that points back to the source location and down to the impact location.

Since the shower is much closer to the telescopes than the source its image will not appear in the same position in each of the cameras. If the multiple images of the shower are overlaid in a single camera then the intersection point of the major axes is the source position of the shower (Figure 3.10a). If there are more than two images of the shower it is unlikely that the major axes will overlap at a single point, rather an estimate must be made of the source position given the information in the showers. This is done by minimising the square of perpendicular distances from each major axis to the estimated source position, with each distance weighted by the size of each image and the angles between the major axes (to reduce the weight for pairs of nearly parallel images that could introduce a large error).

A commonly-used parameter in stereoscopic analysis is *thetaSq*, which is the square of the distance between the reconstructed source position and the point of interest. In an ideal case, with a correctly-handled background and acceptance (Subsection 3.4.9) and no sources, then, if the plot of the acceptance-weighted number of showers against *thetaSq* should be flat. For this reason it is common to refer to areas on the sky by their *thetaSq* rather than their radius or their true area (including the factor of π).

A similar approach is used to find the impact position only instead of superimposing the images in a single camera they are positioned at the telescope locations and the reconstruction is done in the "ground plane" (Figure 3.10b).

Stereo Parameter Calculation

Similar to the *Hillas* parameters which parametrise the view of a shower from a single telescope, a set of stereo parameters can be determined that combine the information from all of the imaging telescopes. The parameters are based upon the single telescope parameters, however a simple averaging process will not suffice as the distances from the shower to telescopes are different, giving different weights to different telescopes.

First proposed in Daum et al. (1997) the solution is to compare the observed shower with simulated γ-ray showers that produce an image of the same size in each of the telescopes. This requires that the simulated showers have the same "observing" conditions (distance, angle, telescope configuration, etc.), as the observed shower. A goodness of fit for the width and length is then calculated which enables selection of γ-rays (which

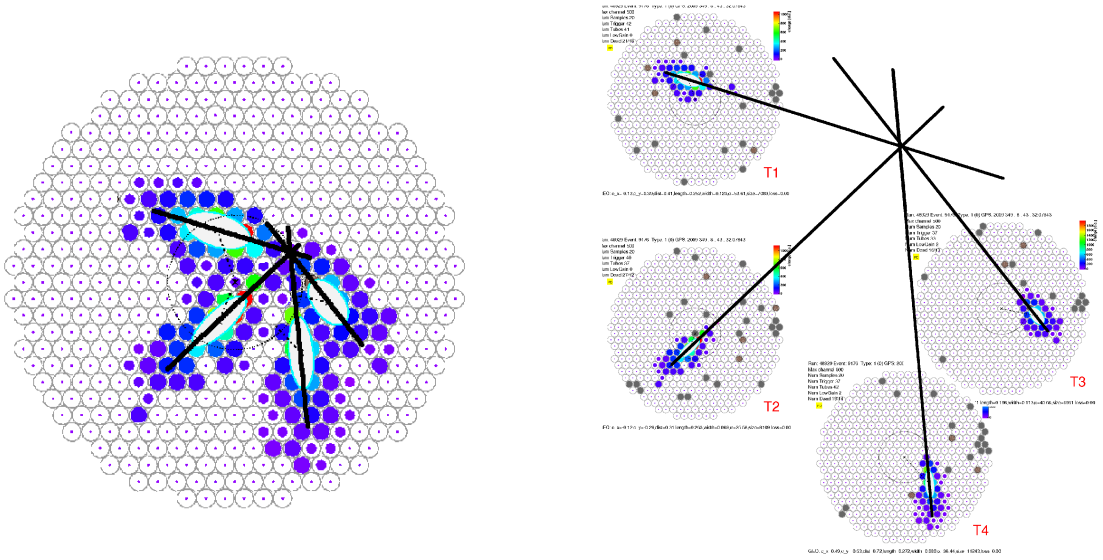

(a) Source Position

(b) Impact Position

Figure 3.10.: Using the stereoscopic imaging of a shower the source (left) and impact (right) positions can be calculated. The source position is reconstructed by overlaying the images in a single camera and the source location is at the intersection of the major axes. The impact location is found using a similar method, only the showers are reconstructed in the "ground plane" with the images positioned at the telescope positions. *Image Credit: A. McCann*

will be well fit) from cosmic rays (which won't). In VEGAS this is done by calculating the mean scaled width and mean scaled length (MSW, MSL) parameters, where

$$MSW = \frac{1}{N_{tel}} \sum_{k=1}^{N_{tel}} \frac{w_k}{\langle w_{sim}(\theta_k, b_k, size_k) \rangle} \tag{3.1}$$

where k is the telescope number and $\langle w_{sim}(\theta_k, b_k, size_k) \rangle$ is the prediction of the parameter in that telescope given the telescopes pointing (both zenith and azimuth affect the shower), θ_k, the impact distance (the distance to the core location), b_k and the shower size, $size_k$ (MSL uses the same formulation but for the length rather than the width). The summation is conducted over all telescopes that have images.

The predicted parameter is an average over many simulated γ-ray-induced showers to allow for fluctuations between showers. These simulations are as accurate-as-possible reconstructions of everything that a real shower will experience, including the atmosphere,

NSB, telescope optics and electronics and analysis chain. The results of these simulations are stored in lookup tables which enable an estimate of the stereo parameters to be made for any given observing conditions (or, more often, by interpolating between lookup tables of given sets of observing conditions).

At this stage it is also possible to determine the height of peak emission (shower maximum, a useful parameter for discrimination at lower energies) using the centroids of the shower images, the impact and source locations, and the energy of the shower (which is recovered during the calculation of the MSW/L since the simulations which are used to derive them are of known energy).

3.4.6. Event Selection

The selection of γ-ray events from the background cosmic-ray events is done using a simple box-cut analysis, that is parameters are cut with no consideration of the other parameters and a box in phase space is selected. This is primarily done using the parameters MSW and MSL, both of these will peak at 1 for γ-ray events, thus by cutting at values close to this it is possible to find an optimum values that maximise the number of γ-ray whilst minimising the number of background events (Figure 3.11). The exact values chosen depend upon a variety of factors, primarily the source strength and spectral index, but also whether the principal interest is maximising detection significance or reducing the PSF. Standard cuts are generated for the analysis of the majority of sources, based upon the the source strength and spectral index, they are soft (low *Size* cut, optimised for a soft spectral index or a strong source), hard (high *Size* cut, hard spectral index or a weak source) and medium (a compromise between the other two sets of cuts and used for most sources). They are optimised on datasets other than the target of interest to give the maximum signal, typical values are MSW < 1.1, MSL < 1.3.

3.4.7. Source and Background Estimation

Having selected the events and reconstructed their source positions and energies, the excess number of events from the source direction (N_{excess}) is determined. To do this a source (*on*) region is defined (the region that the analysis is interested in, usually the (potential) source), this takes the form of a circle of a given radius chosen by consideration of the source's extension and the PSF. The number of showers that came from within this region is called N_{on}. Similarly, background (*off*) regions are defined, away

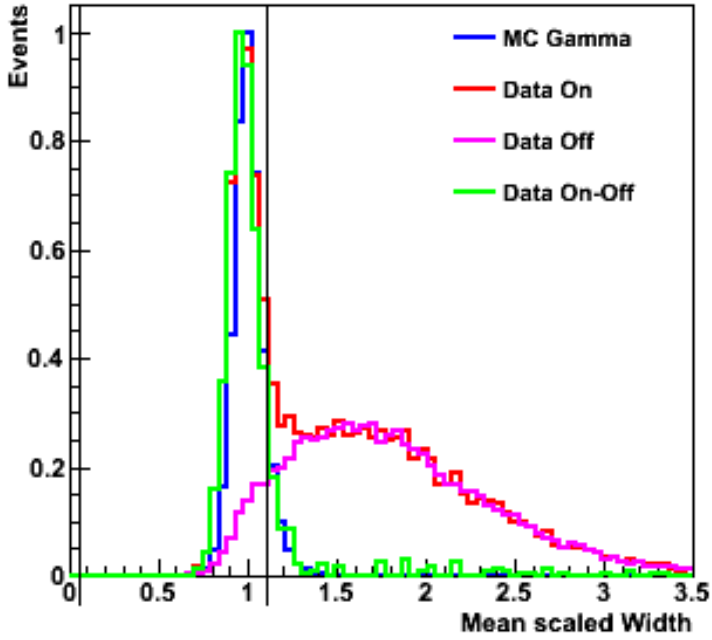

Figure 3.11.: The MSW distribution for simulated γ-rays and data. The MSW of the γ-rays are peaked at 1, the MSW of the cosmic rays at over 2, with a much wider distribution, by cutting at a MSW of 1.1 the majority of the cosmic ray events are removed, leaving the γ-ray events.

from known or potential sources, and from bright stars (see Section 4.3), the number of events from this region is called N_{off}. The relative acceptance* of the two regions ($\alpha = {}^{Acceptance_{on}}/_{Acceptance_{off}}$). For a point source the optimum radius is calculated by finding the radius that gives the maximum detection significance on a test source, this radius will vary with the cuts used as the PSF varies, and is optimised in the same way as the cuts.

Four different background methods are used with VERITAS to model the background depending on the observing technique used and the priorities of the analysis being conducted.

*Acceptance is the likelihood that a signal is to be measured. Within IACTs it depends upon two main factors (provided the observations are matched in zenith, azimuth and noise), the length of the observation time and the radial distance of the position from the centre of the camera.

On/Off Observations and Analysis

On/Off was the original method used by the early telescopes and still used for extended sources since the entire FoV can be used as the on source region (Cawley & Weekes, 1995). An *on* run is taken (normally ∼20 min long) and immediately after the run has been completed the telescope is moved to target a position the length of the run behind the source's position in RA (in order to match the *on* run's azimuth and elevation) and an *off* run is taken. This second run is then used as the background estimator with the hope that there has been minimal variation in the NSB between runs and thus the background will be the same (small differences can be corrected for using padding but that is beyond the scope of this work (Cawley & Weekes, 1995)). The number of *off* events is then subtracted from the number of *on* events (after correction for NSB/run length differences if required) and the excess calculated ($N_{excess} = N_{on} - N_{off}$). The major disadvantage of this method is that it requires a significant commitment of time to observing background regions, at the expense of observing sources, and has an α of 1.

Wobble Observations and Reflected Regions Analysis

To overcome the loss of observing time with On/Off observations wobble mode displaces the source within the camera's FoV (Figure 3.13a) allowing for the estimation of the background from the same run (Fomin et al., 1991). This has the advantage that no *off* runs are required at the penalty of the slight decrease in the efficiency of the telescopes by displacing the source from the centre of the FoV. By displacing the source a fixed distance from the center of the FoV in each of the four cardinal directions for different runs it is possible to remove a number of potential systematic errors from this observing mode. The distance displaced is a compromise between the decrease in collection efficiency at larger distances with the increase in the background region. For point sources a distance of 0.5° is often used, whereas for extended sources a larger distance (0.7° or even 1.0°) is used to allow for sufficient background estimation and to ensure that the source does not overlap the center of the FoV to allow for the radial acceptance function (Figure 3.12) to be determined.

In reflected-region analysis (Figure 3.13a) the background is estimated using regions the same size as the source region and at the same radial distance from the camera center, but rotated away from the source. This enables multiple, non-overlapping regions to be

used to make the estimate of the background with a low value of α. Since the acceptance depends on the radius and energy (there is a zenith angle effect as well, but that is much weaker) this is the preferred method for spectral analysis due to the constant radial distance and thus the constant acceptance. However, it cannot be used close to the center of the camera (since there is no space to calculate background regions) or with multiple/extended sources (difficult to find background regions that are not contaminated).

In this observing mode $\alpha = 1/(\text{number of } off \text{ regions})$, N_{off} is the sum of the counts in all the *off* regions and $N_{excess} = N_{on} - \alpha N_{off}$.

Ring Background Model Analysis

An alternative method that is commonly used is the ring background model (RBM, Figure 3.13b, Rowell (2003)). RBM uses a ring on the sky around the source region as the background region. This enables the entire FoV to be studied since it should always be possible to find a region of the camera to use as the background (unless the source fills the entire camera, in which case an alternative method must be used). The major disadvantage is that since the background region samples areas with different radial distances from the center of the camera the acceptances will vary. Therefore, it is necessary to weight the showers based upon their distance from the centre of the camera, this is done through the use of a radial acceptance function (Figure 3.12), the density of showers at a given radial distance. The radial acceptance is calculated using γ-ray like events that are outside of the excluded regions, their number as a function of distance is calculated either from the runs being analysed or from a dedicated set of observations that match in zenith and NSB. The radial acceptance is also a function of energy and this makes spectral analysis using the RBM method unreliable.

Crescent Background Analysis

The crescent background method (Figure 3.13c, Zitzer (2015b)) is an implementation designed to be a compromise between the two other methods. A ring is constructed at the same radius from the center of the FoV as the source, but excluding the source region. The background is then calculated from this region after it has corrected for the radial acceptance, since the circular source region will be sampling the radial acceptance function differently to the background region. This has the advantage of increasing the

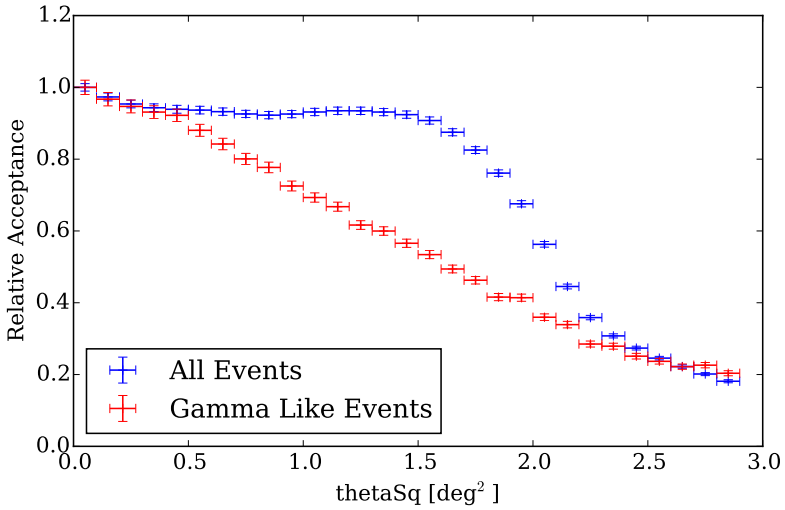

Figure 3.12.: The (smoothed) radial acceptance (number of counts per unit area) varies across the FoV, decreasing with increasing distance from the center. In VEGAS the radial acceptance function is calculated from γ-ray like events that lie outside of the source and exclusion regions and is the inverse of the density of showers at a given radial distance. If the statistics are too low then a radial acceptance function can be generated using other observations, though care must be taken to match the zenith and azimuth of the observations.

background statistics in comparison to the reflected region analysis, but with reduced dependence on the radial acceptance function than the RBM analysis.

3.4.8. Detection Significance

With N_{on}, N_{off} and α known then the detection significance is determined using the formula of Li & Ma (1983). The significance, S, for a single observation is calculated using

$$S = \sqrt{2}\left\{N_{on}\, ln\left[\frac{1+\alpha}{\alpha}\left(\frac{N_{on}}{N_{on}+N_{off}}\right)\right] + N_{off}\, ln\left[(1+\alpha)\left(\frac{N_{off}}{N_{on}+N_{off}}\right)\right]\right\}^{1/2} \tag{3.2}$$

If multiple observations of the same pointing are to be combined then α is constant, so

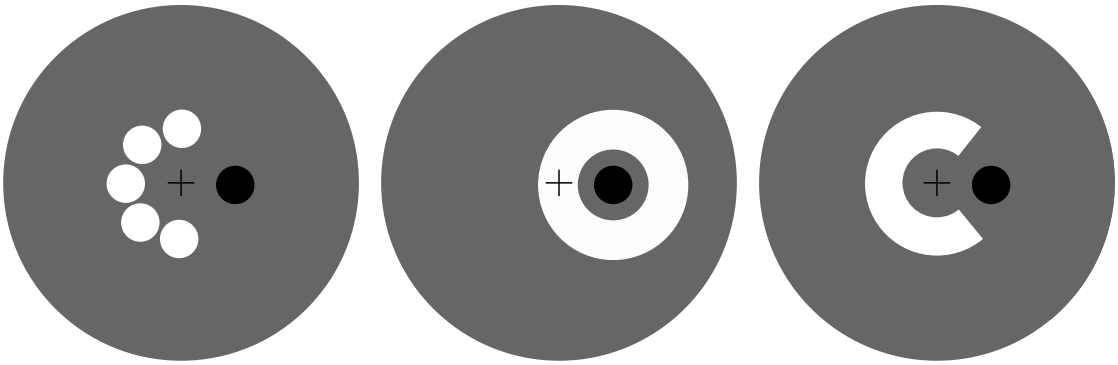

(a) Reflected Regions **(b)** Ring Background **(c)** Crescent Background

Figure 3.13.: The three background analysis methods, the source position (black circle) is offset from the center (black cross) of the FoV (grey circle) by an amount (normally $0.5°$). The white regions are used to determine the background, they are positioned to avoid the potential source and any other (potential) sources or bright stars that could affect the flux in the region, distorting the result.

it is simply a case of summing the N_{on} and N_{off} for all the runs and apply Equation 3.2. However, when combining different pointings care must be taken as α may vary due to changes in the exclusion regions for different wobble directions.

Bright stars increase the noise in that part of the camera, increasing the threshold and altering the number of showers detected around it. To counteract this, areas impacted by stars are excluded from the background, similarly any known or expected sources are also excluded. Since these exclusions will change with the array's pointing direction, the value of α will also vary also vary, this is particularly pertinent for wobble observations (Section 3.4.7) where the telescopes are pointed a fixed distance away from the source in the four cardinal directions.

To overcome this the generalised Li & Ma equation (Equation 3.3) is used (Aharonian et al., 2004).

$$S = \sqrt{2}\left\{\sum_i N_{on}^{(i)}\, ln\left[\left(\frac{\sum_i N_{on}^{(i)}}{\sum_i \frac{\alpha_i}{1+\alpha_i}\left(N_{on}^{(i)} + N_{off}^{(i)}\right)}\right)\right] + \sum_i N_{off}^{(i)}\, ln\left[\left(\frac{\sum_i N_{off}^{(i)}}{\sum_i \frac{1}{1+\alpha_i}\left(N_{on}^{(i)} + N_{off}^{(i)}\right)}\right)\right]\right\}^{1/2} \tag{3.3}$$

This is the standard method used within VEGAS for the reflected region analysis. With RBM analysis a value for α is calculated based upon the observation time and acceptance of the entire background region on a bin by bin basis and then summed, thus Equation 3.2 is used.

3.4.9. Skymap Generation

Using the RBM method it is possible to calculate the excess and significance for an observation for all of the points within the FoV of the observation. This enables a skymap to be produced showing the excesses and significances of all points, which is particularly useful for determining source morphology and for checking that an analysis has been conducted correctly without errors, for example gradients, or additional sources/regions of high NSB that need exclusion. To confirm that all regions that should be excluded from the analysis (known or potential sources and bright stars) are excluded and that the remaining variation is statistical, a histogram of the significances in each bin is produced (known as a significance distribution), this should produce a Gaussian of mean 0, standard deviation 1 (Figure 3.14).

3.4.10. Rate & Flux Calculation

The rate of γ-rays (R_γ) detected by the array is simply

$$R_\gamma = \frac{N_{on} - \alpha N_{off}}{\tau} \tag{3.4}$$

where τ is the livetime of the array, that is the the duration of the observations less time removed by time cuts (for bad weather, etc.) and the time lost due to deadtime when the data acquisition system was busy and unable to record events. Since this rate is specific to a set of observation conditions (elevation, weather etc.) it is not useful for absolute measurements. However, it is often used as a first-pass comparison to see if

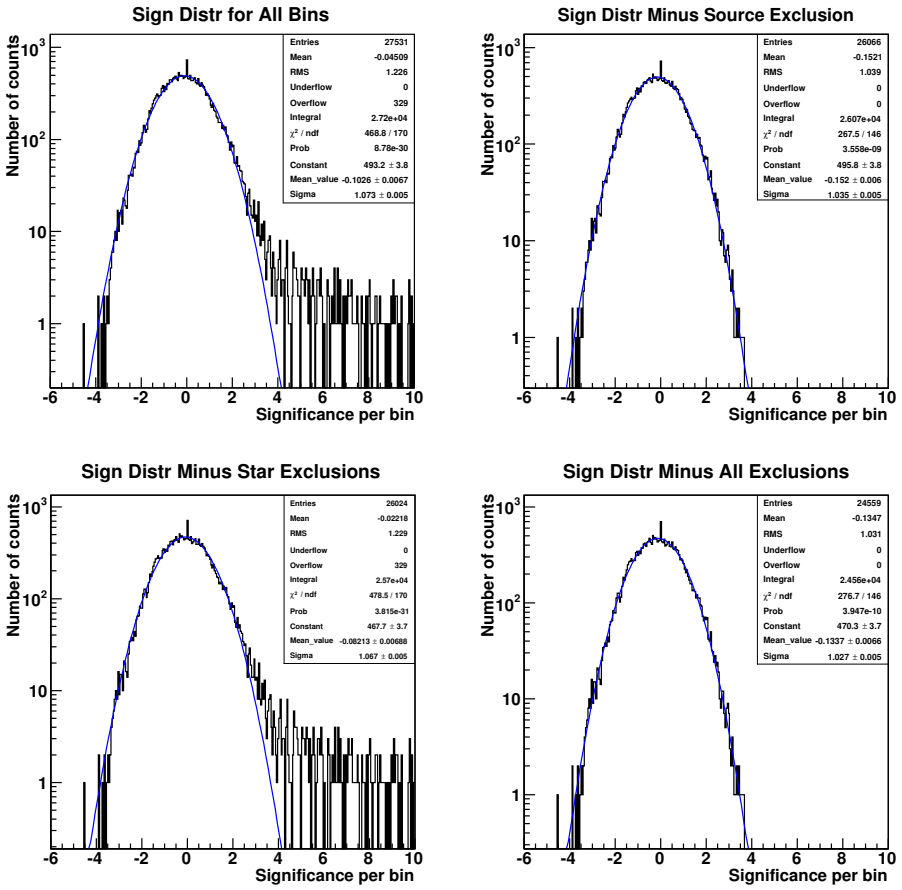

Figure 3.14.: Significance Distributions -The significances of the bins in a skymap can be be put into a histogram. If the correct exclusions have been applied and there is no untoward pathology then the histogram should be well fit with a Gaussian of mean 0 and sigma 1 as in the bottom right plot above. This distribution may not show up gradients in the skymaps which must be checked for by eye.

the array is behaving as expected (if the rate is vastly different from what is expected, especially from a "steady" source then it suggests there is a fault) or if a source is flaring and showing a different rate from that expected given the observation conditions.

To determine the flux I need to know how "big" our detector is, the "bigger" the detector, the more showers will be detected, the higher the rate for a given flux. This is calculated using the effective area (A_{eff}), the area around the array which, if a γ-ray enters, the resulting shower will be reconstructed. This is primarily a geometric effect, with a larger zenith angle exposing the telescopes to a larger area of atmosphere (increasing the effective area) but also increasing the depth of atmosphere (reducing the amount of light transmitted to the array, reducing the effective area at low energies and raising the energy threshold). However, other observing conditions also play a part, such as NSB, weather, pointing offset and azimuth angle.

In a similar way to the shower parameter reconstruction (Section 3.4.5) the easiest way to determine the effective area is with simulations. To determine A_{eff} a number of showers (N_{thrown}) are thrown over a region of area A_0 that is centred on the array in a plane perpendicular to the optical axis and is sufficiently large that if additional showers were thrown outside this area they would not trigger the array. These simulated showers are then analysed using exactly the same method as the data and the number of showers after all the cuts have been applied is calculated (N_{recon}). For a given energy (E),

$$A_{eff} = A_0 \frac{N_{recon}(E)}{N_{thrown}(E)}. \tag{3.5}$$

The effective area can then be used to determine the flux (the number of showers per unit area, per unit time). It is impossible to generate an effective area for every energy, zenith and pointing offset therefore it needs to be calculated in bins. Converting from a (binned) rate to a (binned) flux using a (binned) effective area is not trivial and it is not simply a case of doing the naive $^{\text{rate}}/_{\text{effective area}}$ = flux. The main reason for this is the finite precision associated with binning and with the energy resolution of the array, coupled with sources that are typically described with power-law energy spectra meaning that slight errors in reconstructed energy can result in major differences in the reconstructed spectra. Correcting for these effects requires careful and complex consideration, beyond the scope of this thesis, a detailed explanation can be found in Albert et al. (2007) and Mohanty et al. (1998).

3.4.11. Energy Bias and Resolution

Errors in the analysis chain and variation in the propagation of showers results in errors in the reconstructed energies of showers. By analysing simulations and working out the error in reconstructed energy ($\log_{10}(E_{rec}) - \log_{10}(E_{true})$) and plotting that against $\log_{10}(E_{true})$ the error can be split into two components, the mean error at a given energy (the energy bias) and the standard deviation about that energy (the energy resolution). The energy bias is corrected for by shifting all reconstructed showers by the appropriate difference. The resolution dictates the minimum width of the spectral bins, and thus the minimum width of any spectral features that can be detected.

3.4.12. Spectral Reconstruction

Once the flux in each energy bin is calculated then the result can be fitted to give the spectrum of the detected source. This is done by iteratively adjusting the model spectrum $\left(\frac{\mathrm{d}F}{\mathrm{d}x}\right)$ until

$$\frac{\mathrm{d}N_{rec}}{\mathrm{d}x'} = \tau \int \frac{\mathrm{d}R(x';x}{\mathrm{d}x'} A_{eff}(x) \frac{\mathrm{d}F}{\mathrm{d}x}\,\mathrm{d}x \tag{3.6}$$

where $x = \log(E)$, $\frac{\mathrm{d}N_{rec}}{\mathrm{d}x'}$ is the reconstructed spectrum (calculated using the reconstructed energies E'), τ is the livetime (exposure), $A_{eff}(x)$ is the effective area and $\frac{\mathrm{d}R(x';x}{\mathrm{d}x'}$ is the energy resolution.

3.4.13. Upper-Limit Calculations

If no significant detection is made (typically the detection criterion is 5σ **post trials**) then an upper limit on the flux is calculated, to do this a upper limit on the number of excess events detected is compared to the number of excess events that would be expected from an input spectrum and the spectrum scaled accordingly. The upper limit on the number of excess counts can be determined in a number of ways, in this thesis the method of Rolke et al. (2005) is used. This uses the frequentist properties of confidence intervals computed by the profile likelihood method. In particular it is suitable for dealing with situations where there are additional nuisance parameters with statistical or systematic errors. In VEGAS, the class TRolke, a part of the CERN ROOT framework (Brun & Rademakers, 1997), is used, taking as inputs N_{on}, N_{off}, α and the confidence interval (typically 0.95) and returns an upper limit on $N_{excess,UL}$.

This is compared with the excess from an input spectrum that is assumed to match that of the potential source (i.e. of the same type and spectral index) and a given normalisation (a_{guess}) is multiplied by the effective area of the observations and the livetime. This will give a value of $N_{excess,guess}$ for that spectrum, which can be used to determine the upper limit on the normalisation of the input spectrum

$$a_{UL} = a_{guess} \frac{N_{excess,UL}}{N_{excess,guess}}. \tag{3.7}$$

3.5. Summary

VERITAS is the basis of the majority of the work presented in this thesis. In this chapter I have outlined the how VERITAS works and the standard analysis methods used, from detection of the Cherenkov light through to production of spectra and upper limits. Later chapters present the details of the methods used to overcome specific problems when analysing the targets in this thesis, primarily associated with high levels of NSB (Chapter 4) and significantly extended sources (Chapter 6).

4. Reducing the Impact of the Night Sky Background: Image Cleaning Improvements and Bright Moonlight Observations

In addition to light emitted from air showers, IACTs detect background light from various sources and are also affected by electronic noise within the detector. The background light, generally referred to as the night-sky-background (NSB) is made up of many different constituents, primarily stars, zodiacal light, light from human activities and, if it is above the horizon, light from the Moon. This light presents the primary limit to the operation of IACTs since the NSB both reduces their sensitivity to the light emitted from air showers and, if sufficiently bright, damages the detectors.

In order to reduce the impact of the NSB on the sensitivity of the array, image cleaning forms part of the standard VERITAS analysis (some more advanced analyses do not use image cleaning) retaining only those channels that image the Cherenkov signal from the shower whilst removing the channels that only contain noise. The current method for doing this is through a simple, two-stage, signal-to-noise cut, with a lower threshold for channels adjacent to pixels which pass a higher threshold in Subsection 3.4.3. Two aspects of image cleaning are explored here, a new cleaning method to help discriminate the signal from the background using the pixel timing is presented (Section 4.1) and the efficiency of different image cleaning approaches are compared for their ability to reduce the impact of significant variations in the NSB in Section 4.2.

Standard observations are conducted with less than ∼30% Moon illumination. At higher levels of Moon illumination the NSB level is sufficiently high to cause damage to the photomultiplier tubes which puts limits to the amount of observing time available*.

*Photomultiplier tubes are damaged in normal operation by nature of how they work but in this case

By reducing the high voltage and using UV bandpass filters it is possible to reduce this damage and conduct observations at higher levels of Moon illumination, extending the range from ∼30 to >65% illumination. The development of these techniques and their verification are discussed in Section 4.3.

4.1. Image Cleaning Improvements - Time Cleaning

4.1.1. Existing Image Cleaning Methods

Image cleaning is the process whereby the noisy image recorded in the camera is processed to aid the fitting of the shower and the calculation of the *Hillas* parameters. This is done so that shower parameterisation is only conducted using those pixels which are imaging the shower. Good image cleaning improves the quality of the *Hillas* parameterisation, reducing the PSF, lowering the energy threshold, increasing the sensitivity and improving the signal-to-background of the observation. The aim is to select as many pixels that are imaging the shower as possible whilst minimising the number of background pixels which will affect the reconstruction of the shower.

The basic method of image cleaning (*CleanUpStd* within VEGAS) performs two levels of cleaning based upon the pixel's signal-to-noise (S2N) ratio. The S2N of a pixel is calculated for a single event by summing the digital counts within the 7 sample (14 ns) integration window (a subset of the readout window, Figure 3.7) and dividing it by the pedestal variance. If the S2N is greater than 5.0 (the default value) then the pixel is labelled as a *picture* pixel and retained. This high threshold selects pixels that, provided there are no hardware issues, are almost certainly imaging the shower. However, it misses a lot of smaller-signal pixels that are important in characterising the edge of the shower, which are particularly useful for gamma/hadron separation, and to lower the energy threshold. To catch these pixels a second pass over the camera is conducted and all pixels that have a S2N greater than 2.5 (again, the default) *and* that are adjacent to a *picture* pixel are labelled as a *boundary* pixels and retained. The combined list of *picture* and *boundary* pixels are the *image* pixels of the event and used for the calculation of the *Hillas* parameters.

This simple image cleaning has been improved by the removal of single isolated pixels

we are looking at photocathode currents (and thus damage rates) which are at least an order of magnitude higher, if not more.

(confusingly called *EventDisplayCleanUp* (ED cleaning) within VEGAS). In this cleaning, any single image pixels (i.e. those without any *picture* or *boundary* neighbours) are identified and removed since, due to the shape of γ-ray showers, there are not expected to be any isolated pixels. The aim of this cleaning is to remove any pixels that have been spuriously triggered by noise (either sky noise or, due to faulty electronics, electronic noise) and are distorting the reconstructed image.

4.1.2. A Simple Time-Cleaning Method

The addition of timing information of the pulses, made possible by using FADCs, to the image-cleaning algorithm adds another parameter for discrimination between signal and noise channels, this will reduce the impact of areas of higher NSB (e.g. stars) in the field of view. Though there have been previous attempts at including timing in the image cleaning method, they have been rather complex and required the tuning of several parameters to optimise the cleaning and thus have not been adopted. Due to this, a method was proposed whereby only a single additional parameter was introduced: that *boundary* pixels, in addition to passing a S2N cut, must also have a *TZero* (the half fall time of the pulse) that is within a certain time, *deltaT*, of the *TZero* of the *picture* pixel. This cleaning was motivated by examining individual shower images where a number of noise pixels that pass image cleaning were clearly identified with peak times that are significantly different to the shower pixels they neighbour. It was felt that this cleaning method would add an additional constraint to remove these pixels without adding an undue load to the optimisation and analysis speed. This method has been implemented within VEGAS 2.5.4 and is available as *TimeCleanUp2* (TC2 cleaning, *TimeCleanUp* is an existing, more complex cleaning algorithm that has not been tested as a part of this work) and is the subject of this section.

During the implementation and testing I also examined the impact of introducing a time cut on neighbouring *picture* pixels, but this reduced the sensitivity gain. I also considered whether using the trace peak would be better than the *TZero*, but this was rejected as there was no substantive reason to make the change and it would require additional coding.

In short, TC2 = ED cleaning (the current standard) + additional time constraint between *picture* and *boundary* pixels.

4.1.3. *deltaT* Cut Motivation

The motivation for this additional cut comes from plotting the *deltaT* between adjacent pixels. The *deltaT* distribution is plotted (Figure 4.1) with and without standard ED cleaning applied.

Data

Figure 4.1 shows the result of this plotting for a data run (65777) both for the full time window and zoomed in about the central peak. Overlaid over the data (blue) are two Gaussians (black) and their sum (red) which was fitted to the data (for the zoomed region the fit was done using the zoomed region only). The twin Gaussian structure is explained by the larger, wider Gaussian being a result of the time difference between random noise pixels. The narrower central Gaussian is due to the time difference between pixels that are imaging the shower, since a shower front is narrow and neighbouring pixels are neighbouring parts of the front the narrower distribution is to be expected. The effect of ED cleaning is clearly visible, the majority of the background pixels are removed though there is still a clear background. The aim of this new cleaning method is to remove pixels that are triggered by random noise, i.e. those pixels that pass ED cleaning but are not caused by the shower. The simplistic time cut of $|\text{deltaT}| < 2.6\,\text{ns}$ is appropriate since this is about the time difference where a pixel is more likely to be triggered from random noise than from a shower (as visible by the background Gaussian crossing the shower Gaussian). It is not optimal to retain all of the signal pixels since this will end up retaining a significant number of background pixels and thus reduce the effectiveness of the cleaning.

Simulations

Since data is dominated by cosmic rays, to confirm whether the tighter cut is also beneficial for γ-rays a similar analysis was conducted on CARE* detector simulations (a simulation package used by VERITAS and CTA) using simulated showers from CORSIKA† and optical ray tracing by GrOPTICS‡. When the code was naively applied to the

*http://otte.gatech.edu/care/
†https://www.ikp.kit.edu/corsika/
‡http://www.physics.utah.edu/gammaray/GrISU/

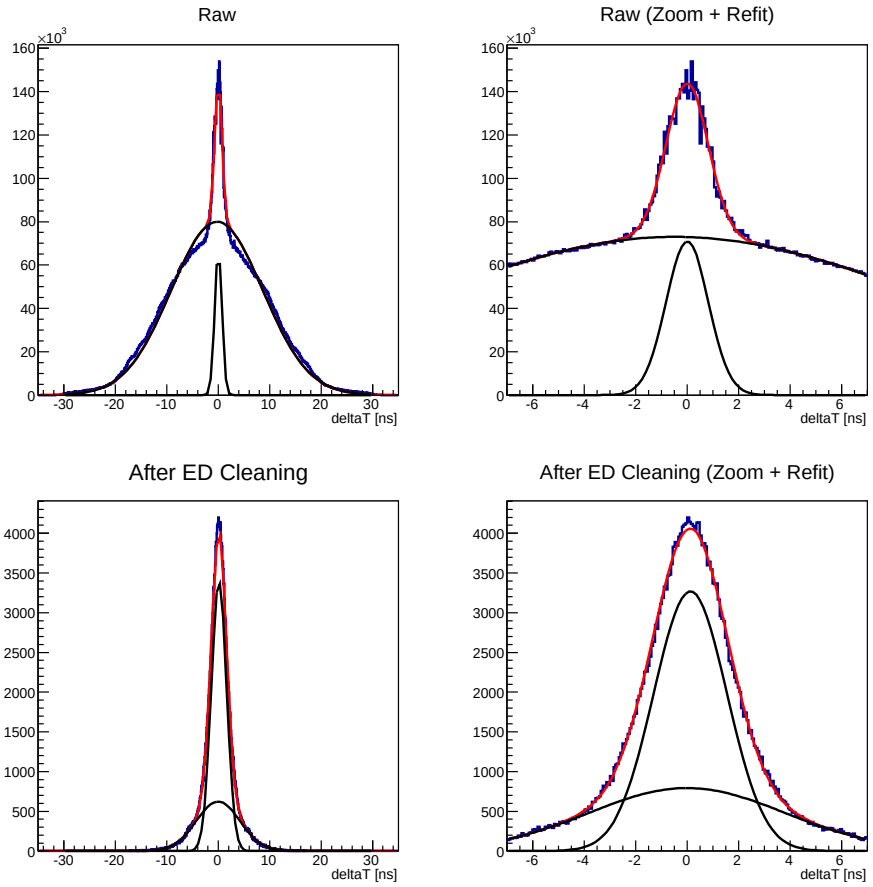

Figure 4.1.: Data: The *deltaT* between neighbouring pixels before (raw, top row) and after ED cleaning (*picture* S2N = 5, *boundary* S2N = 2.5, bottom row) for all time differences (left column) and zoomed to the central peak (right column). The blue line is the raw data, the black lines are the individual Gaussians from the summed Gaussian fit (red line). Each fit was conducted on the data presented in that histogram. The vertical axis is the number of pixel pairs at that *deltaT*. 5000 events using all four telescopes and all 499 channels were used from run 65777.

simulations no central peak was visible in the raw cleaning plots and only one Gaussian was visible in the post-ED cleaning plots (Figure 4.2).

The reason for this is that in the data there will, inevitably, be small timing differences between neighbouring pixels due to differences in delays, etc. This is not the case in the simulations where all of the channels have the same delay and thus the readout window times are identical. The result of this difference is that, instead of generating a smooth distribution of peak times (and thus the smooth Gaussian structure) only a finite number of integer time differences are possible corresponding to the time between the samples. These spikes dwarfed the shower peak. To correct for this a random time shift drawn from a Gaussian distribution of mean 0 ns standard deviation 0.4 ns was added to the *deltaT* (0.4 ns was chosen since it was the smallest time that smoothed out the peaks sufficiently to allow identification of the shower peak). This had the desired effect of smoothing out the peaks, allowing for the shower Gaussian distribution to be clearly visible (Figure 4.3). The crossing time was similar (though slightly broader than for the data), this was in part due to the relative heights of the two Gaussian fits and the additional width from the Gaussian smoothing. It was chosen to stick the with original 2.6 ns as identified in the data.

4.1.4. Testing

The main challenge in testing is to determine what parameter to optimise on, ideally it would be on the detection significance (over a given energy grange) of a typical source (for example, the Crab Nebula scaled to 5% of its normal flux). However, to perform such a test requires analysing data through the whole analysis chain *and* processing the simulations to conduct the analysis *with the same cleaning*. This is clearly not feasible for multiple different cleaning algorithms and thresholds. Many different methods which do not require a full analysis have been tried previously with differing levels of success.

Care must be taken though when using a partial analysis to optimise the thresholds as improvements in one parameter may result in an overall reduction to the sensitivity or energy threshold. For example, if optimising the containment radius it is preferable to only retain the center of larger showers since this reduces the direction reconstruction, thus smaller showers are cleaned out and cosmic rays are cleaned so aggressively that they look more like γ-rays. The result is that the energy threshold is increased and the overall sensitivity goes down. Thus, in this work it was decided to test the whole way

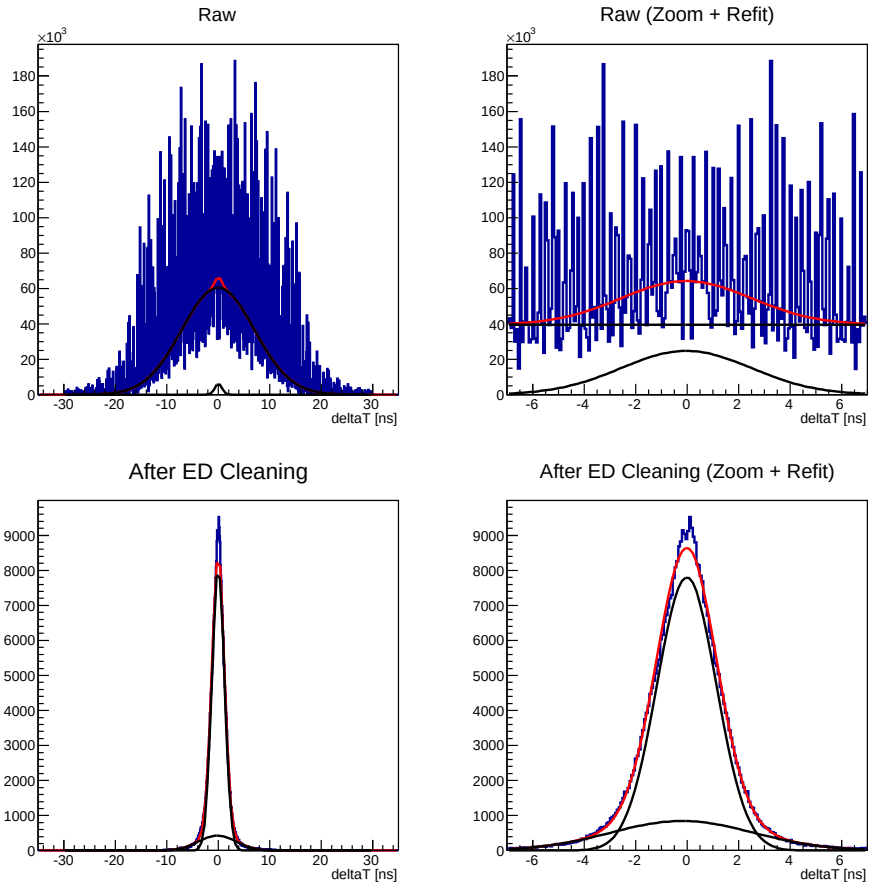

Figure 4.2.: Raw Simulation: The *deltaT* between neighbouring pixels before (raw, top row) and after ED cleaning (*picture* S2N = 5, *boundary* S2N = 2.5, bottom row) for all time differences (left column) and zoomed to the central peak (right column). The blue line is the raw data, the black lines are the individual Gaussians from the summed Gaussian fit (red line). Each fit was conducted on the data presented in that histogram. The vertical axis is the number of pixel pairs at that *deltaT*. 5000 events using all four telescopes and all 499 channels were used from a simulation at 0 deg zenith and with a pedvar of 6.5. Without smoothing there are very pronounced spikes due to binning which swamp the interesting features without ED cleaning. With ED cleaning the peak due to shower pixels is very pronounced.

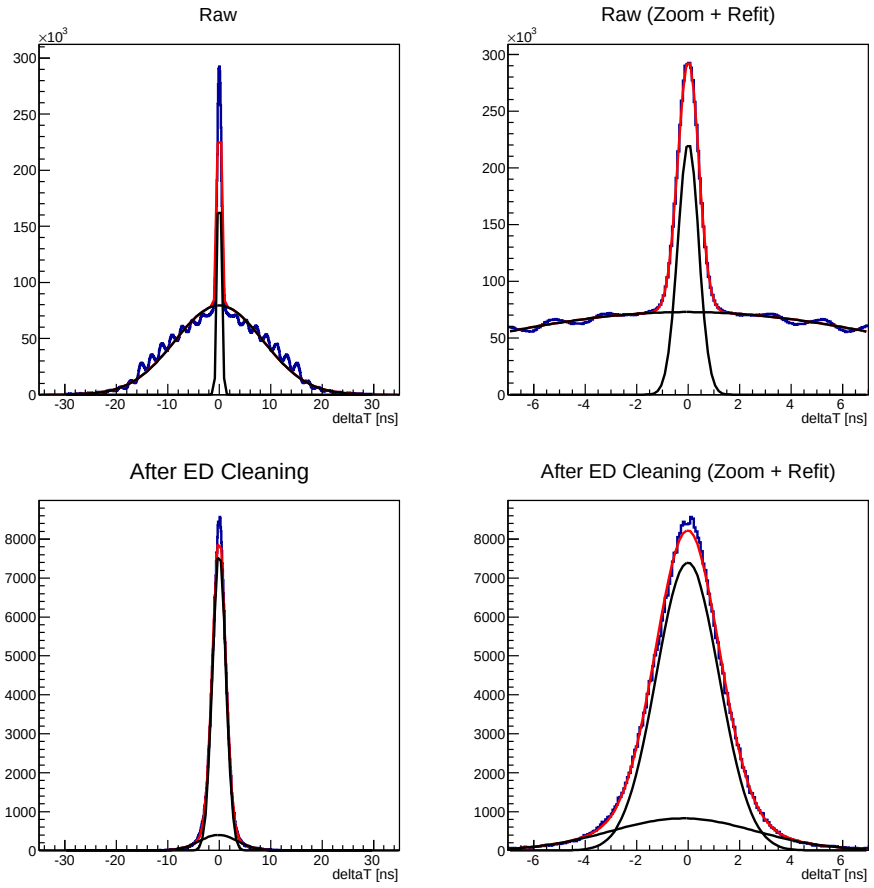

Figure 4.3.: Smoothed Simulation: The *deltaT* (smoothed with a Gaussian of width 0.4) between neighbouring pixels before (raw, top row) and after ED cleaning (*picture* S2N = 5, *boundary* S2N = 2.5, bottom row) for all time differences (left column) and zoomed to the central peak (right column). The blue line is the raw data, the black lines are the individual Gaussians from the summed Gaussian fit (red line). Each fit was conducted on the data presented in that histogram. The vertical axis is the number of pixel pairs at that *deltaT*. 5000 events using all four telescopes and all 499 channels were used from a simulation at 0 deg zenith and with a pedvar of 6.5. Despite the smoothing the integer value peaks are still clearly visible, further smoothing was not desirable though as it would have further broadened the peaks.

through an analysis for a small number of runs, with a limited noise, zenith and azimuth range but using existing lookup tables and effective areas. Once a threshold is optimised the newly optimised values were used to generate a full set of simulations and a full test was conducted to verify that the results are valid.

Initial Test

An initial test was conducted on 16 Crab runs, 4 at each wobble offset, for a total exposure time of 406.4 minutes. This runlist was analysed using ED and TC2 cleaning with S2N thresholds of *picture* = 5.0, *boundary* = 2.5 and for TC2, $deltaT$ = 2.6 ns (Table 4.2). Both cleaning methods were analysed using newly-generated lookup tables and effective areas that spanned the noise and zenith range of the observations using standard VEGAS 2.5 soft and medium cuts (Table 4.1).

Table 4.1.: Cuts Used During Analysis

Cuts	Quality Cuts			Selection Cuts			
	Distance	Size	nTubes	MSW	MSL	Height	thetaSq
Soft	<1.43°	>400 dc	>5	0.05-1.1	0.05 - 1.3	>7 km	<0.03
Medium		>700 dc					<0.01

There is a clear increase in the significance (∼2-5% for soft and medium cuts) arising from an increase in both the number of gamma and background events with TC2. This is due to an improvement in the direction reconstruction of events, there is a cut within VEGAS that rejects events whose direction reconstruction fails, the improvement in reconstruction direction reduces the number of events that fail this cut. To check that this increase in the background events does not result in a reduction in sensitivity with weaker sources, the RBM results were used to scale the Crab signal to a 1% of its nominal flux (Table 4.3) which show a similar gain. This analysis shows a slightly larger increase in the sensitivity for soft cuts over medium cuts and no change in the increase with source strength. In addition, the source is fit with a symmetrical Gaussian function to measure the PSF, it was found that the width decreased by 4-5% for both sets of cuts (Only 100% Crab was examined).

Examining the skymaps by eye shows no obvious difference in their form, including in the impact of the bright star, $\zeta-$Tauri (for a detailed examination of this see Section 4.2).

Table 4.2.: Results for the full Crab runlist of 16 Runs analysed using the reflected regions method (n.b. the significance is calculated using Equation 3.3 with α varying from run to run)

	Soft		Medium	
	ED	TC2	ED	TC2
On	8979	10296	3817	4367
Off	5633	6926	1076	1423
Significance	109.0	113.9	110.4	115.96
Difference	-	+4.50%	-	+5.04%
Rate Gamma [min^{-1}]	18.69	21.15	9.08	10.34
Rate Background [min^{-1}]	3.40	4.19	0.31	0.41

Table 4.3.: Results for the full Crab runlist of 16 Runs analysed using the ring background method

	Soft		Medium	
	ED	TC2	ED	TC2
On	9133	10478	3859	4399
Off	11553	14188	3200	4142
Effective Alpha	0.084	0.084	0.027	0.027
100 % Crab σ	141.89	148.95	135.8	142.84
Difference	-	+4.98%	-	+5.18%
PSF [deg]	0.0665 $\pm$0.0003	0.0638 $\pm$0.0003	0.0577 $\pm$0.0002	0.0550 $\pm$0.0002
Difference	-	-4.1%	-	-4.75%
1 % Crab σ	2.48	2.55	3.73	3.77
Delta	-	+2.8%	-	+1.1%

The reconstructed spectra are almost identical (ED cleaning: $N_0 = 3.65 \pm 0.07 \times 10^{-11}\text{cm}^{-2}\ \text{s}^{-1}\ \text{TeV}^{-1}$, $\alpha = 2.42 \pm 0.01$; TC2 cleaning: $N_0 = 3.53 \pm 7 \times 10^{-11}\text{cm}^{-2}\ \text{s}^{-1}\ \text{TeV}^{-1}$, $\alpha = 2.42 \pm 0.01$; Pivot Energy = 1 TeV). A comparison of the effective areas (Figure 4.4) shows that there is a slight increase across almost the entire energy range of about $10 - 20\%$, with the increase larger at smaller energies as one would expect since image cleaning is particularly important for smaller showers.

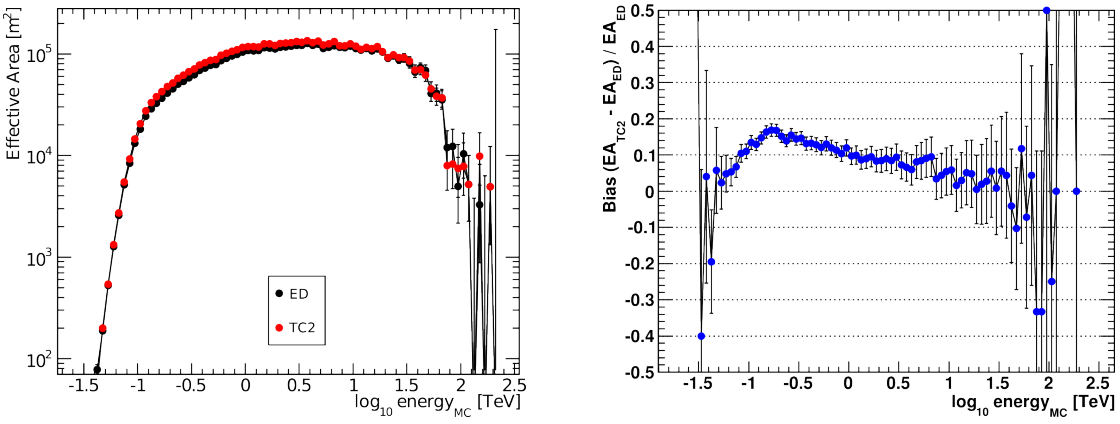

(a) Effective Areas for ED (black) and TC2 (red) cleaning, across the "plateau" of the effective area, TC2 cleaning has a slightly higher effective area.

(b) The improvement ((TC2-ED)/ED) is in the range 10 - 20% at energies between 65 GeV and 5 TeV.

Figure 4.4.: A comparison between two Effective Areas, both effective areas are at zenith = 0° and noise = 4.26, a typical low noise level.

deltaT and *Boundary* threshold Optimisation

In order to optimise on different values of *deltaT* and S2N, a scan on *deltaT* and the *boundary* S2N threshold was conducted (No scan was conducted on the *picture* threshold which remained constant at 5.0 throughout). No new tables were generated for this scan, rather those tables generated for the initial tests were used (*boundary* = 2.5, *deltaT* = 2.6 ns). Though this will introduce a slight bias against values that are significantly different to the values used in the generation of the lookup tables, this is likely to be small and enable an optimisation on the cleaning thresholds.

The results of the optimisation are shown in Figure 4.5a which shows that the optimal values are slightly higher than those initially tested. Due to the unexpected structure in

this scan with its double peak structure, the scan was repeated with a different data set. The results (presented in Figure 4.5b) confirm the general structure that increasing the *boundary* threshold increases the significance with a smaller variation with the time cut and a peak around $boundary = 2.5$, $deltaT = 2.9$ ns but does not have the double peaked structure. Depending upon the analysis a balance has to be struck between energy threshold and sensitivity, in order not to push up the energy threshold too much it is desirable to not increase the cleaning thresholds too high. Thus, based upon both scans and this requirement, it was decided to generate a new set of tables with $boundary = 2.7$, $deltaT = 2.7$ ns.

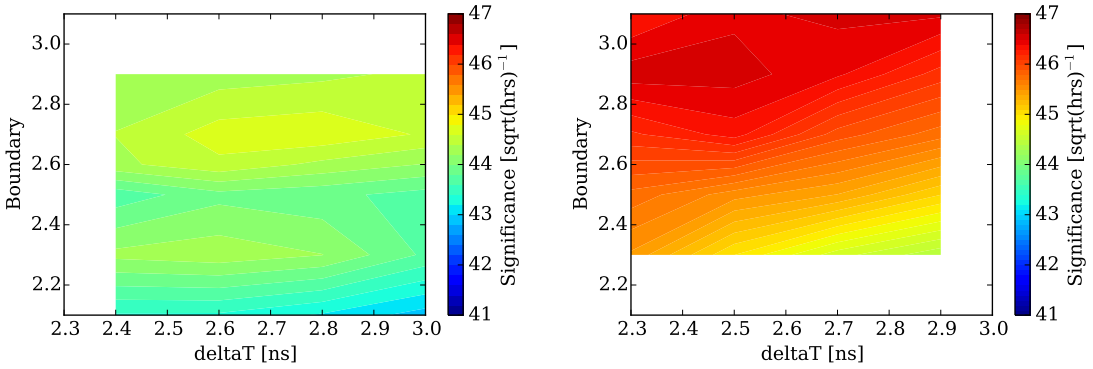

(a) Crab runs 65777-9 & 66112, exposure = 101.236 min). For comparison, when the same four runs are analysed with ED cleaning the significance/sqrt(hours) is 42.46.

(b) Crab runs 66002-3, 66113 & 67044.

Figure 4.5.: Optimisation of the *boundary* and *deltaT* thresholds (*picture* fixed at 5) using soft cuts. All runs were analysed using the same cuts (standard soft) and lookup tables. The plots share a similarity in that there is a general increase with increasing *boundary* threshold and decreasing *deltaT*, both of which reduce the sensitivity to the smallest showers. There are differences in the structure, with the limited test dataset size this is not unexpected and I am more interested in the general trends. Ideally ∼50 runs would be used to do such an optimisation but computing resources did not permit this.

An independent Crab dataset was analysed using new lookup tables, no increase was observed in this dataset in comparison with the original TC2 analysis with *bound-*

ary = 2.5, *deltaT* = 2.7 ns. With no measurable increase in sensitivity and a slight increase in the energy threshold due to the harder cleaning, the original TC2 cleaning thresholds were retained.

MSW/MSL Cut Optimisation

After optimising the cleaning thresholds, a cut optimisation was attempted for medium cuts using data on the Crab Nebula but no improvement upon the standard values was found. With the reduction in the PSF it is possible to reduce the *thetaSq* cut a small amount, but within the variation in values determined during the cut optimisation no improved value was found. Additional work on improving the PSF is being done by collaborators which, when combined with this work, will hopefully enable a reduction in the *thetaSq* cut in the near future.

4.1.5. Test with 1ES 1011+496

In order to test the performance over a wide variety of fluxes for a soft-spectrum ($\alpha \simeq 3$) source it was decided to analyse the 2013-14 observing season winter atmosphere data on *1ES 1011+496*, which typically has a flux of a few percent of the Crab Nebula flux but experienced a flare to ∼75% of the Crab Nebula flux during these observations. Using ED cleanup the cumulative significance was 45.1 (N_{On} = 3767, N_{Off} = 10104), there was a slight (∼2%) increase in the significance to 46.2 (N_{On} = 3913, N_{Off} = 10456) using TC2 cleanup with *picture* = 5.0, *boundary* = 2.5 and *deltaT* = 2.6 ns.

A comparison of the light curve and time-averaged spectrum is shown in Figure 4.6 and Figure 4.7. Looking at these comparisons, there is a small difference in both the index and the normalisation largely driven by differences at the highest energies where TC2 has reconstructed more events. The light curves show good agreement.

4.1.6. Time Cleaning 2 Conclusions and Outcomes

- The new cleaning method has been shown to work, producing a 2-5% increase in the sensitivity. Though not a significant improvement on its own these small improvements add up to a significant improvement over time. A 2% improvement is equivalent to 20 hours observing time, about the average amount of time spent observing a typical source each year.

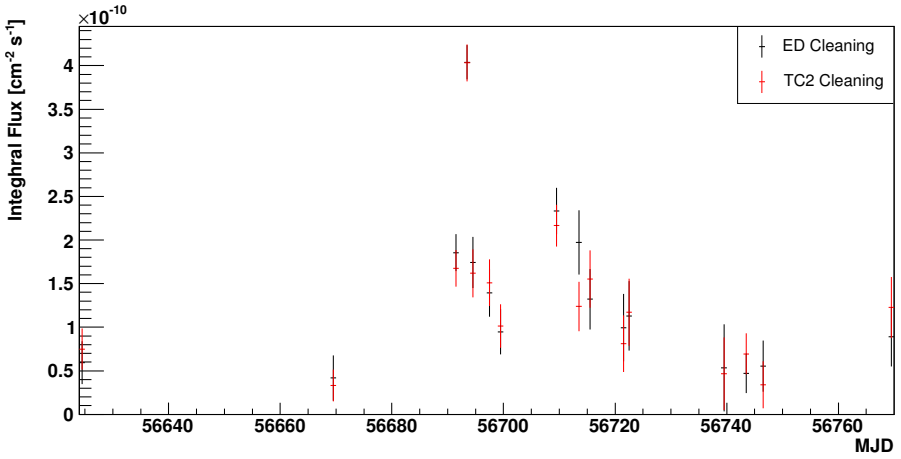

Figure 4.6.: *1ES 1011+496* lightcurve for the 2013-14 observing season, winter atmosphere data analysed with soft cuts. The black points are the standard (ED) image cleaning, the red points are the new (TC2) image cleaning. Both sets of data were analysed using lookup tables and effective areas that were specially generated using identical methods apart from the image cleaning.

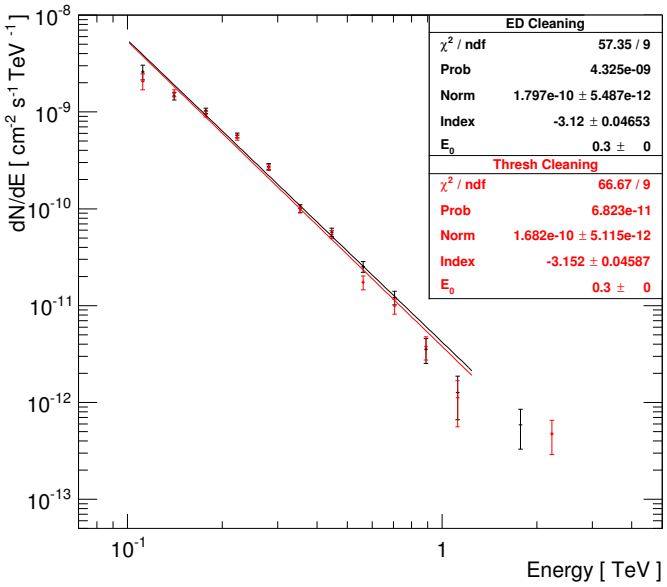

Figure 4.7.: *1ES 1011+496* time average spectra for standard (ED) cleaning (black) and the new (TC2) cleaning (red).

- There is no discernible additional overhead in computation time.
- The spectrum and flux for the Crab Nebula and *1ES 1011+496* have been reproduced, with only small differences between the results for different cleaning thresholds.
- Cleaning thresholds have been scanned and suitable values identified ($picture = 5$, $boundary = 2.5$, $deltaT = 2.6$ ns).
- Implementing this cleaning method as the default requires the reanalysis of all simulations. For a 2-5% gain this was not warranted but switching for the next round of simulations was recommended and will be implemented.
- TC2 has not been adopted for the remaining work in this thesis since that work was addressing sources with significant gradients in the NSB and thus different methods have been employed to reduce its impact (Section 4.2).

4.2. Image Cleaning Improvements - Reducing the Impact of Variations in the NSB

Variations in the night-sky-background (NSB), for example around stars, have the potential to impact IACTs in two ways:

- The raised NSB raises the trigger rate, this causes an increase in the RFB (Subsection 3.1.3). Since this is not correctable in software after the data has been taken, it is discussed no further.
- The raised NSB in a pixel reduces the S2N for a given signal size, reducing the chances of it passing image cleaning.

The result of both of these factors is that the part of the shower closest to the star is removed and when the direction is reconstructed it introduces an error and directs the showers away from the star. This results in a hole in the skymap at the position of a bright star, for example $\zeta-$Tauri (Figure 4.8), and those positions are normally excluded from the analysis.

This work was conducted to reduce the impact of significant variations in the NSB across the FoV at the analysis level, in particular through image-cleaning methods. Two

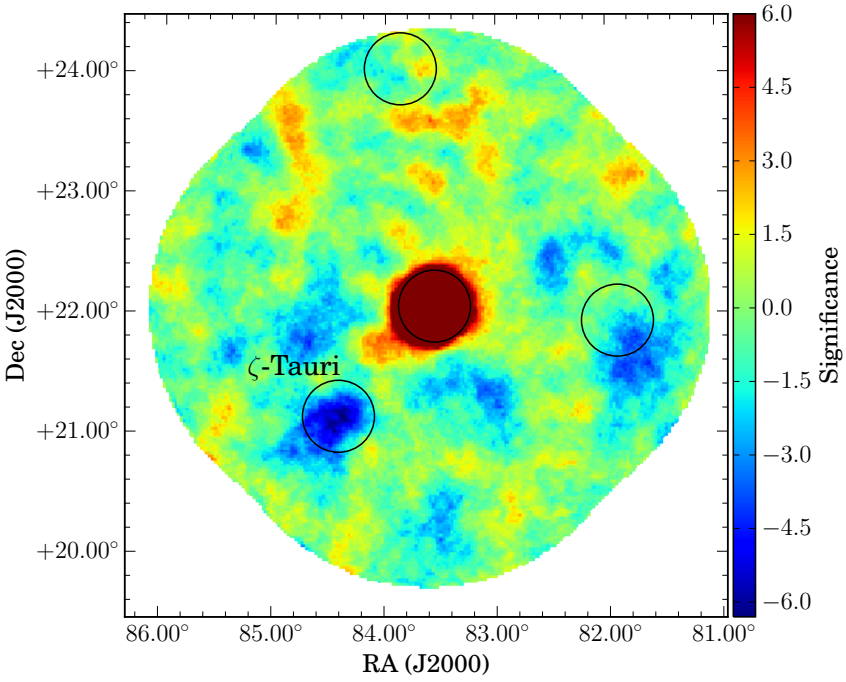

Figure 4.8.: Skymap from 16 Crab Nebula runs, soft cuts (*Size* $>$ 200 dc) showing the impact of the optical brightness from $\zeta-$Tauri with a significant deficit due to it. The scale has been limited to 3σ as the Crab Nebula is at $\sim 110\sigma$ in this dataset, the black circles correspond to the exclusion regions (bright stars).

methods were explored, the first used the timing of the signal (Section 4.1), the second used the same fixed cleaning threshold across the camera for all runs, rather than using a cleaning level based upon individual pixels that varied from run to run (*Threshold* cleaning, this removes the impact of NSB variations in the cleaning threshold). During this study the thresholds used were *picture* = 60 dc, *boundary* = 30 dc*. These values were chosen based upon a pedvar of 12 dc (a level above all but the brightest pixels in the FoV of M 31, the primary focus of this thesis) and using the ED cleaning S2N thresholds of 5 and 2.5 respectively.

*n.b. no *deltaT* cut was used in this work since no improvement was found with its application in a small test.

To conduct this work 16 "perfect" Crab Nebula observations were analysed, 4 at each (0.5°) wobble direction. The effectiveness of the different cleaning is calculated through the following parameters:

Significance

The ring background model (RBM) significance map is fit with a two dimensional asymmetrical Gaussian function (to allow for fitting to $\zeta-$Tauri which, due to its location at the edge of the FoV, is often elongated). The significance is taken to be the amplitude of the fit rather than the peak of the skymap.

FWHM

Using the same asymmetrical Gaussian, the two standard deviations are combined using $\sigma_{tot} = \sqrt{\sigma_x^2 + \sigma_y^2}$. The FWHM is calculated from FWHM $= 2\sqrt{2\ln 2}\sigma_{tot}$.

Comparison

To compare the impact on both the Crab Nebula and $\zeta-$Tauri a comparison is generated using the two fits. The comparison statistic is $-1 \times (\int \text{ Crab Fit})/(\int \zeta\text{−Tauri Fit})$ (the integral rather than the amplitude is used to correct for any changes in the PSF from the different cleaning methods).

The aim of this study is to investigate different image cleaning techniques to maximise the signal from the Crab Nebula whilst minimising the impact of $\zeta-$Tauri by maximising the *Comparison* parameter.

These parameters are calculated for each of the cleaning methods using standard medium cuts (*Size* $>$ 700 dc, Table 4.4). In addition a further analysis using hard cuts (*Size* $>$ 1200 dc, Table 4.4) for the standard (*EventDisplay*) image cleaning is conducted, this is the standard way of reducing the impact of NSB variations since the increased *Size* cut also reduces the impact of the variation of the NSB.

The different cleaning methods were used to analyse simulations and generate lookup tables and effective areas. These simulations were either thrown at a single offset of 0.5° (Table 4.5) or at multiple offsets in the range 0.0° to 2.0° in steps of 0.25° (so called "All offset" simulations, Table 4.6). The reason for using both sets of lookup tables is that All offset tables, with 9 different offsets to simulate, take significantly more time to produce and, since the majority of observations are at 0.5° wobble they are not as well verified (and not available for TC2). However, they produce better reconstruction at the edge of the FoV and thus will give a more accurate result for the fit to $\zeta-$Tauri

Both sets of analysis show roughly similar results for the Crab Nebula, but with increased significance for $\zeta-$Tauri with the All offset tables (this is due to the better reconstruction away from an offset of 0.5°). If the sole aim is the reduction of the deficit due to $\zeta-$Tauri then there is a significant benefit for analysing the data using *Threshold* cleaning, since this almost halved the significance of $\zeta-$Tauri (though made it slightly wider) in comparison to the other cleaning methods. The significance of the Crab Nebula also increased for *Threshold* cleaning with respect to both ED and TC2 cleaning (which also showed a gain over ED cleaning). This increase was unexpected and had not been seen in earlier studies. Further work showed that, if the noise is perfectly matched (as these simulations turned out to be) *Threshold* cleaning is a better method for image cleaning than ED. This optimum range is small though and it is better to generate simulations for analysis of all sources rather than have to generate new lookup tables and effective areas for each analysis.

Since the aim of this study is to maximise the *Comparison* parameter, the result is that *Threshold* cleaning is clearly the best (even if the Crab signal is reduced to that of ED cleaning then the reduction in $\zeta-$Tauri still leads to a marked improvement). The use of hard cuts does not provide as large a benefit as expected from previous studies, though an improvement was more noticeable with the All offset tables.

Table 4.4.: Cuts Used During Analysis

Cuts	Quality Cuts				Selection Cuts		
	Distance	Size	nTubes	MSW	MSL	Height	thetaSq
Medium	<1.43°	>700 dc	>5	0.05-1.1	0.05 - 1.3	>7 km	0.01
Hard		>1200 dc			0.05 - 1.4	-	

Looking at the significance maps (Figure 4.9 and Figure 4.11) verifies these results by eye. The impact of hard cuts is slightly more pronounced here, especially with the All offset tables, with the deficit from $\zeta-$Tauri more localised and smaller. The significance distributions (histograms of the significances in each of the bins across the skymap, Figure 4.10 and Figure 4.12) both show a marked reduction in the low significance bins for the medium cut *Threshold* cleaning and hard cut ED cleaning in the 050 offset analysis. In the All offset analysis the hard cut improvement is less marked but the *Threshold* cleaning clearly shows an improvement. A check of the spectra produced showed that they all had almost identical indices and normalisations, well within statistical errors.

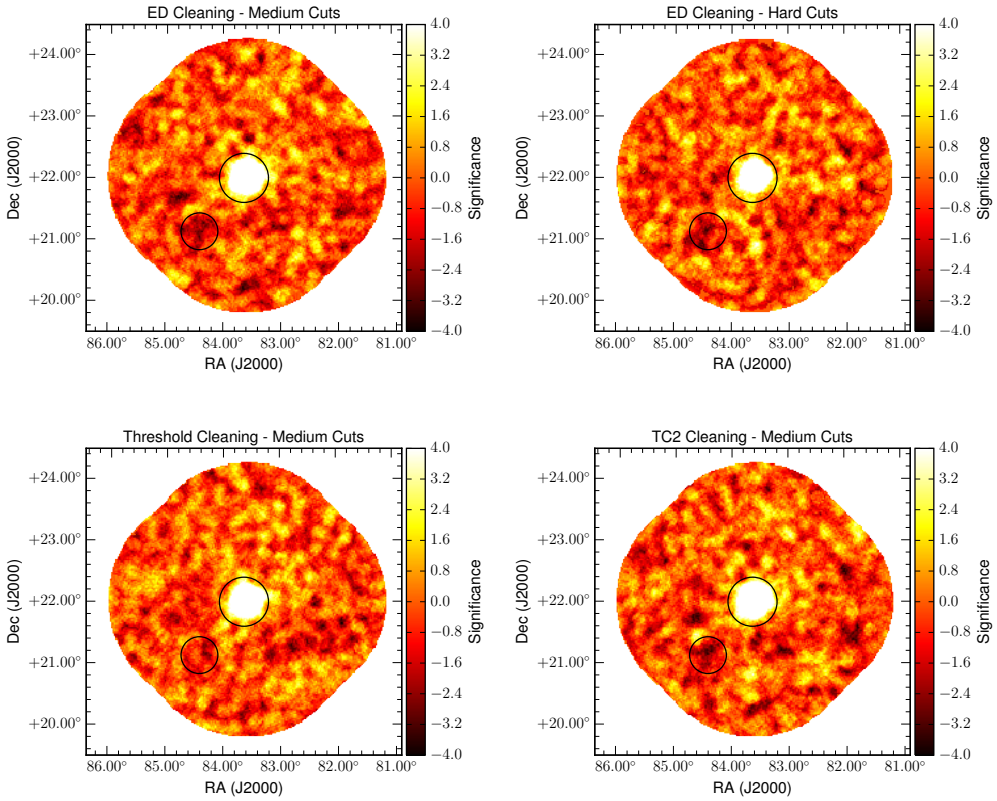

Figure 4.9.: Significance maps (with the scale limited from $-4\,\sigma$ to $4\,\sigma$) showing the impact of $\zeta-$Tauri (bottom left circle) for each of the cleaning and cuts combinations analysed using the single offset tables. The Crab is located in the central circle and has a peak at about 120σ.

Table 4.5.: Using single offset ($0.5°$) lookup tables, a comparison of different image cleaning methods for their impact on both the Crab Nebula and $\zeta-$Tauri

Cleaning	Cuts	Crab		$\zeta-$Tauri		Comparison
		Sig	FWHM [°]	Sig	FWHM [°]	
EventDisplay	Med	126.7	0.345	-2.54	0.629	16.21
EventDisplay	Hard	117.7	0.327	-2.71	0.565	16.33
Threshold	Med	133.3	0.331	-1.56	0.727	22.33
TC2	Med	131.8	0.343	-2.56	0.753	15.29

Table 4.6.: Using All offset (generated in the range $0.0°$ to $2.0°$ in steps of $0.25°$) lookup tables, a comparison of different image cleaning methods for their impact on both the Crab Nebula and $\zeta-$Tauri

Cleaning	Cuts	Crab		$\zeta-$Tauri		Comparison
		Sig	FWHM [°]	Sig	FWHM [°]	
EventDisplay	Med	127.9	0.345	-3.17	0.614	13.95
EventDisplay	Hard	118.8	0.327	-3.46	0.540	15.55
Threshold	Med	133.7	0.330	-1.87	0.644	25.00
TC2	Med	No simulations available				

4.2.1. Reducing the Impact of Variations in the NSB - Conclusions and Outcomes

Following the results of this study in analyses that are significantly affected by, or have backgrounds with significant variation in, the night-sky-background there are clear advantages to working with *Threshold* cleaning. The use of hard cuts did not present as much of a benefit but it did produce some positive gains, in particular with the All offset tables and provides a useful cross check the the *Threshold* cleaning using a standard analysis method.

4.3. Adapting VERITAS for Operation in Bright Moonlight

High levels of NSB due to the Moon place a limit on the amount of time the array is able to conduct observations during a month and it is desirable to conduct observations

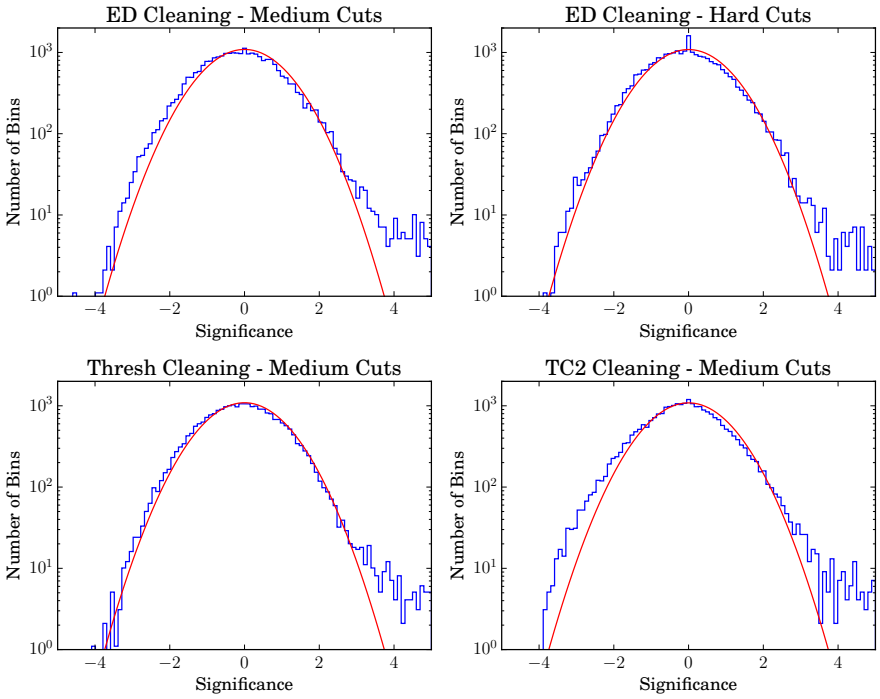

Figure 4.10.: Significance distributions (histograms of the significances in each bin in the significance maps) for the 050 offset analysis. The impact of bright stars is an excess of negative significance above a Gaussian distribution of mean 0, sigma 1 (red line). The maximum significance is clipped at 5 to allow comparison, the highest bin is about 120σ.

during this time which is traditionally not used. I have been leading work that has been undertaken to develop methods to conduct observations during periods when then levels of moonlight are such that normal dark time observations are not possible.

There are two main factors that limit the NSB level that VERITAS can operate under:

1. Reduced Sensitivity.

   The raised NSB increases the accidental trigger rate, increasing the deadtime of the DAQ system. It also reduces the sensitivity of the array to the Cherenkov emission from smaller showers which are swamped by the increased background noise.

2. Damage to the PMTs.

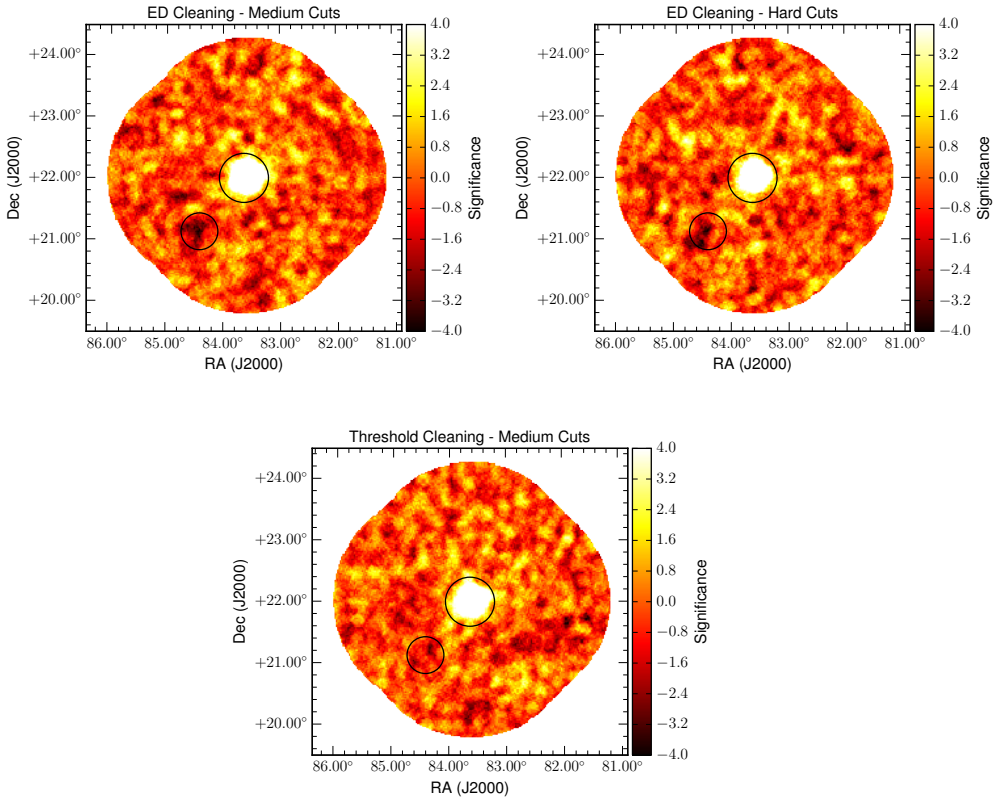

Figure 4.11.: Significance maps (with the scale limited from $-4\,\sigma$ to $4\,\sigma$) showing the impact of $\zeta-$Tauri (bottom left circle) for each of the cleaning and cuts combinations analysed using the All offset tables. The Crab is located in the central circle and has a peak at about 120σ.

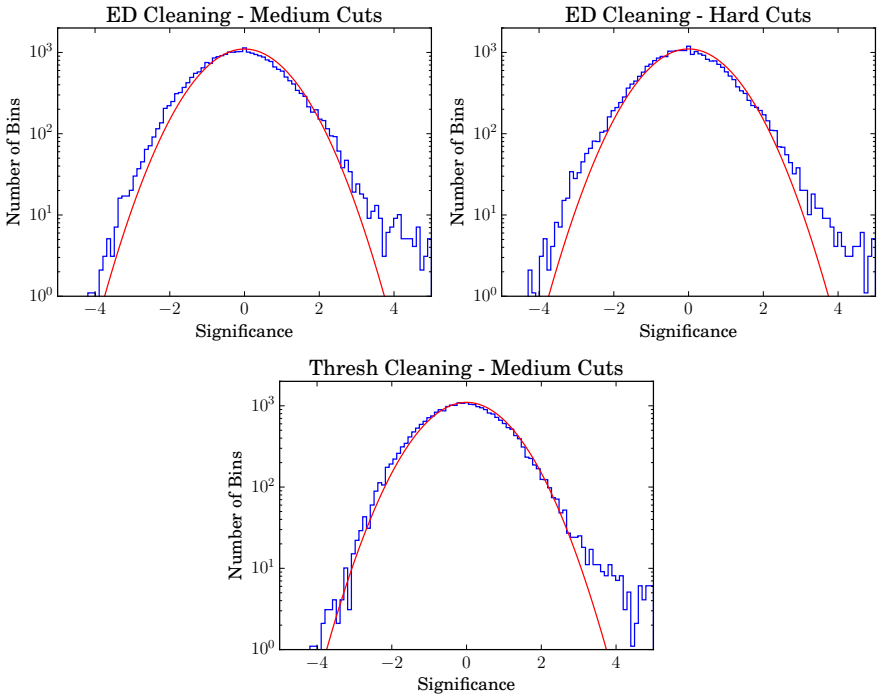

Figure 4.12.: Significance distributions (histograms of the significances in each bin in the significance maps) for the All offset analysis. The impact of bright stars is an excess of negative significance above a Gaussian distribution of mean 0, sigma 1 (red line). The maximum significance is clipped at 5 to allow comparison, the highest bin is about 120σ.

The higher NSB increases the currents passing through the PMT, this can cause damage in two ways.

a) Damage to the last dynode.

   There is a limit to the current that can pass through the last dynode of the PMT without damaging it. VERITAS employs a safety cut-out (30 µAmp), triggered by the current monitors (Subsection 3.1.2) to reduce the rate of damage to the last dynode. If the current through the PMT gets too high then the HV supplied to the PMT is turned off.

b) Damage to the photocathode.

   There is a limit to the number of photoelectrons that the photocathode can

supply. Exposure to bright light, especially in the presence of a high voltage can prematurely age the photocathode, reducing its ability to produce photoelectrons.

Three methods have been developed to allow operation for all but the brightest nights (Figure 4.13). Each mode is able to cope with brighter levels of moonlight (and the modes can be combined) but at the cost of reduced sensitivity.

Raised CFD Thresholds

Raising the CFD thresholds reduces the number of NSB triggers, reducing the rate and the deadtime, but provides no additional protection for the PMTs.

Reduced High Voltage

The high voltage provided to the PMT is reduced, this reduces the gain and thus the current through the last dynode. This protects the last dynode but not the photocathode and thus is used in moderate levels of moonlight.

UV Bandpass Filters

Since the Cherenkov spectrum is peaked in the UV and the atmosphere absorbs the UV light reflected from the Moon, a UV bandpass filter both reduces the amount of light incident on the PMT (protecting the photocathode) and improves the signal-to-noise.

The actual NSB that the cameras are subjected to depends upon many factors in addition to the Moon illumination. The primary factor is the distance between the target and the Moon (the PMT currents reach a minimum at a separation of about 90°, beyond that the camera is pointing back past/through the mirrors at the Moon), as well as the Moon's elevation and atmospheric conditions (clouds reflect the light from the Moon and terrestrial sources). As a result, changes between modes are only approximate and are at the discretion of the observers, with a bias that it is better to sacrifice possible observing time for the mode with lower sensitivity to maximise the time with the better sensitivity.

4.3.1. Raised CFD Thresholds

When the currents reach a predefined threshold that is above the standard operating levels (i.e. the level for which the CFD thresholds are set), but before the risk of damage

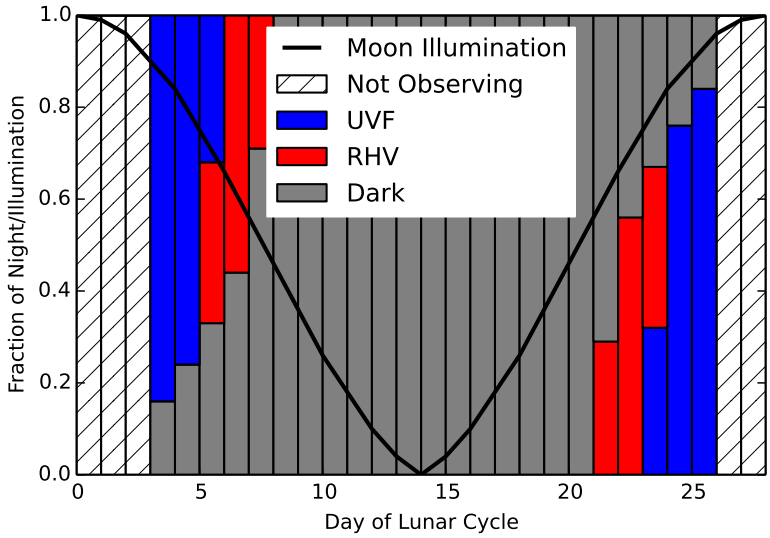

Figure 4.13.: A typical observing Month. The RHV mode provides a 13% boost in yearly exposure and the UVF mode provides a 16% boost in yearly exposure.

to the PMTs, the CFD thresholds are raised to reduce the number of accidental triggers, reducing the L3 rate and thus the deadtime of the array. This typically allows for operation up to ∼30% Moon illumination. This method has been used for many years and was first used by VERITAS in the detection of W Comae (Acciari et al., 2009b).

4.3.2. Reduced High Voltage (RHV) Observations

By reducing the HV (from the nominal which is around 1000 V) applied to the PMTs, the PMT gain is decreased, lowering the current through the last dynode and protecting the PMT. The larger the reduction in the HV the larger the reduction in the gain, thus the lower the current. However, this gain reduction comes at a cost, the signal size is also reduced, increasing the energy threshold and reducing the effective area of the array. It is therefore necessary to choose a level of reduction that provides a suitable compromise between the requirements, reducing the gain sufficiently so that the camera is protected whilst also maintaining the instruments response

In order to test this a series of flasher runs (Subsection 3.2.2) were conducted to measure the gain at different HV settings (Figure 4.14). These show that the behaviour of the gain with the HV is predictable with no step changes or other unwanted effects.

A fit to this data shows Gain $\propto$ HV$^{5.85}$. It was decided that a gain should be reduced to $\sim$30% of the nominal, a reduction in the HV to 81% of the nominal values, allowing for observations when the moon is <65% illuminated. After reducing the HV the camera was then flat fielded to ensure a uniform gain across the camera. The CFD settings are adjusted according to the NSB level and the target object, and a nightly flasher run is taken for gain correction.

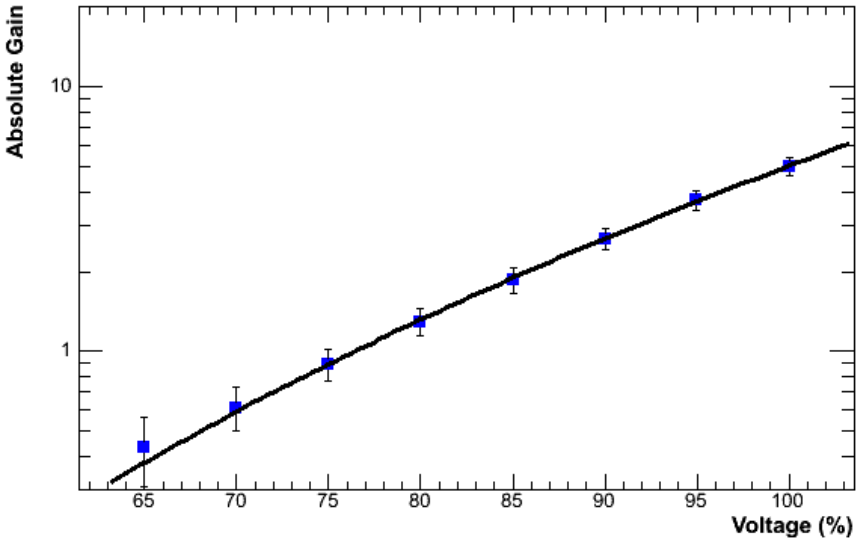

Figure 4.14.: The variation in PMT gain with change in the high voltage from nominal ($\sim$1000 V) overlaid with the fit, Gain $\propto$ HV$^{5.85}$. *Image Credit: A. Bouvier*

To continue observing above 65% illumination a further reduction in the HV would be required. This was disfavoured because lowing the HV does not reduce the photocathode current which prematurely ages the PMTs.

4.3.3. UV Filter (UVF) Observations

Ultraviolet filter (UVF) observations are conducted with special filters constructed from Schott's UG11 filter glass* placed between the light cones and the PMTs. When choosing the filter material several conditions had to be balanced:

- Maximise the transmission of light from a γ-ray shower.

*http://www.schott.com/advanced_optics/english/syn/advanced_optics/products/optical-components/optical-filters/optical-filter-glass/index.html

- Minimise the NSB.
- Allow continued operation of the existing flasher system with a wavelength of 380 nm
- Ease of use and installation, including robustness. Previous filters used by the ARTEMIS experiment on the Whipple 10 m (Pomarède et al., 2001) used a liquid filter between glass plates. That worked for installation on a single telescope for a whole night's operation but would not be practical for the planned observation strategies with VERITAS due to the difficulties in installing and removing the filters.
- Cost: the filter must be affordable.

Schott's UG11* satisfied these requirements, as can be seen for the the first three points in Figure 4.15.

The filter plate, shown in Figure 4.16, is comprised of three sheets of plastic and 499 3 mm thick filters. The central plate has holes the diameter of the filters into which the filters are placed, this is then sealed between two thin sheets with holes the diameter of the exit of Winston Cone, trapping the filters in place. The plate is then sandwiched between the VERITAS camera PMTs and the light cone plate. Installation of the filters onto all four telescopes takes ∼30 min for a team of four working as two pairs. After installation a flasher run is taken for gain correction (as for normal observations) but no additional changes to the array settings are required.

The filters also have additional benefits in bright star fields where standard observations are not possible due to the high background and lack of dark regions to use as background regions. Using the filters it is possible to observe these regions with a significant reduction in the impact of the stars. With reduced impact from the stars, the number of *off* regions can be increased, in part compensating for the reduced sensitivity. With the filters installed observations are conducted at Moon illumination levels >65%, with regular observations being conducted on all nights except the full moon and the night either side (though special observations have been conducted through the full Moon).

*http://www.schott.com/advanced_optics/english/syn/advanced_optics/products/optical-components/optical-filters/optical-filter-glass/index.html

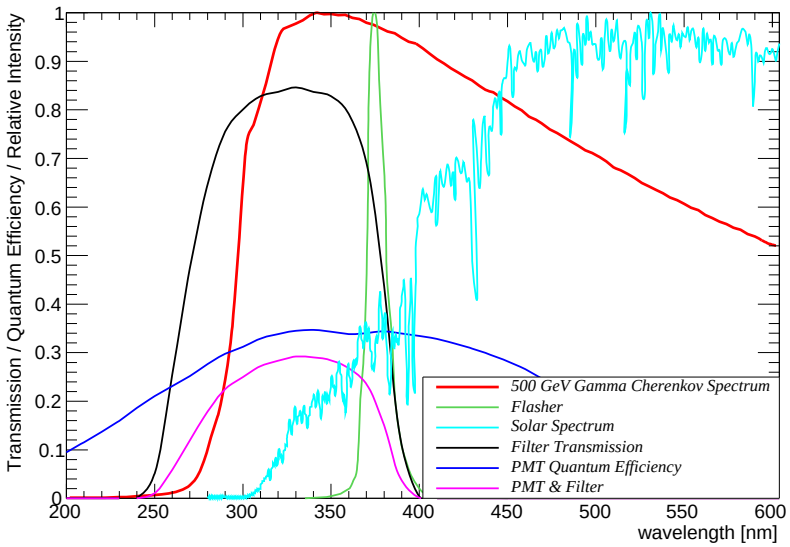

Figure 4.15.: UG11 was chosen as the filter as its transmission (black) was optimal for allowing the light emitted by an air shower ($500\,\mathrm{GeV}$ γ-ray Cherenkov spectrum shown in red) to pass whilst blocking the NSB (the solar spectrum at ground level which will be similar to that reflected from the Moon is shown in cyan). It also allows a significant fraction of the flasher light (green) to be transmitted. For comparison, the QE of the post-upgrade PMT is shown (blue) the effective QE of the combined PMT and filter (pink).

4.3.4. Summary of Observation Modes

Table 4.7 summarises the observing modes that are currently employed by VERITAS. It should be noted that there is quite a bit of overlap between modes, in particular depending upon the elevation of the Moon, the weather conditions and the source's location, both relative to the Moon and to any terrestrial light sources, it is quite common to have multiple observing modes in a single night, even if the Moon is above the horizon throughout observations.

Table 4.7.: VERITAS observing modes

Mode	Moon Illum. [%]	Description	CFD Setting	Other Conditions
Dark Time	0	Sun and Moon below horizon, average currents depend on source background, weather and position relative to terrestrial sources	45 mV	-
Low Moon	0 to 30	Moon is above (but often close to) the horizon and only slightly illuminated, currents slightly elevated but no effect on observations	45 mV	-
Moderate Moon	20 to 60	Moon partially illuminated and often close to the horizon, currents elevated and affecting L3 rate, no risk of camera damage	60 mV	-
RHV	35 to 65	Moon partially illuminated and well above the horizon, currents elevated across the camera to a level that will damage the last dynode under normal operation	25 mV or 35 mV depending on target	81% HV
UVF	≥ 65	Moon significantly illuminated and well above the horizon. Light level will damage the photocathode and the last dynode even with reduced HV	45 mV	UV filters installed

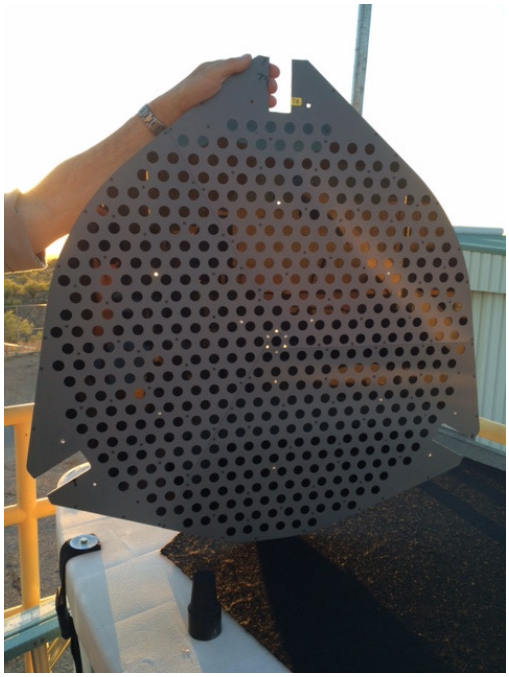

Figure 4.16.: The VERITAS filters. The dark circles are the filters sandwiched between the the grey plastic mounting plates. The bright spots are due to holes left over from the construction process and do not allow light to pass into the PMTs. *Image Credit: J. Holder*

4.3.5. Sensitivity and Spectral Comparison

With these two new modes in use, simulations have been produced for these observing modes to allow for reconstruction of events, event selection and spectral reconstruction. These take the standard VERITAS simulations and adjust the camera model accordingly, for the RHV the gain was reduced by a constant factor of 2.87* which produced good agreement between the data and Monte-Carlo (MC)†. For the UVF the situation is more complex, the first stage was to scale the quantum efficiency of the PMT by the effective transmission of the filters (including a correction for reflection of the filter faces). When a data-MC comparison was made the agreement was poor and further adjustments were

*Calculated from actual measurements on the change in the gain for the four telescopes.

†Simulation verification is based upon matching parameters measured in both the simulations and the data, for example the width and length distributions within individual telescopes and the MSW and MSL distributions for the whole array. In all, over 30 parameter distributions are compared in 6 energy bins for a typical data/MC comparison.

required. After extensive work it was discovered that the filter transmission is calculated for perpendicular transmission, whereas on the telescopes over half of the light exiting the Winston cones has an angle greater than 45°. This was identified during lab tests and confirmed by a colleague with ray-tracing which predicted a loss of about 14% in comparison to normal incidence light (Griffin, 2015). In order to determine the impact on the data a small amount of data that was taken prior to the 2012 upgrade (∼5 hours of observations of the Crab Nebula) were used to find a correction factor. This correction factor was found to be around 15% (very similar to the 14% effect expected from ray tracing) and thus a 15% gain reduction was applied to the post-2012 upgrade simulations to produce the results presented below. The data-MC comparison for UVF is now comparable with that for normal data given the statistics that are available. With further data improvements in the base model and further observations providing greater statistics I expect further improvements in the future.

Sensitivity

A standard measure of sensitivity for IACTs is to conduct observations on the Crab Nebula and to quote a significance per sqrt(hours). In this comparison (and all comparisons in this section) all of the good quality observations that have been taken on the Crab Nebula for RHV and UVF (7.29 and 6.80 hours respectively) are compared with a limited normal dark time runlist (NOM, 5.34 hours) which has been matched to the RHV and UVF data as closely as possible. The cuts used in this analysis are the standard VERITAS cuts optimised for sensitivity on the Crab Nebula scaled to 5% of its normal flux. For the UVF and RHV analysis, in order to reflect the reduction in the signal strength, the *Size* cut has been reduced from the standard 700 dc to 200 dc. These cuts are not optimal and work is under way to improve them and provide a suite of cuts optimised for different sources as exists for NOM observations.

With these cuts the following sensitivites are achieved:

NOM $43.95\sigma/\sqrt{hour}$

RHV $42.06\sigma/\sqrt{hour}$

UVF $21.36\sigma/\sqrt{hour}$

The sensitivity for RHV observations is almost identical to NOM observations but it should be noted that with such a low *Size* cut this analysis is close to the trigger energy

threshold (170 GeV) whereas the NOM observations (which with these cuts have an energy threshold of about 150 GeV) are well above their trigger energy energy threshold of ∼85 GeV. The UVF sensitivity is about a factor of two lower with a significantly higher energy threshold (ranging from 250 to 450 GeV, depending upon the level of the NSB). Most of this loss in sensitivity is as low energies, at higher energies, where showers are larger and brighter, the sensitivity is almost identical.

Effective Areas

Comparing the effective areas (Figure 4.17) shows the reduction in sensitivity at low energies that I would expect given the above results. The UVF EA is lower across all energies and is particularly low below 1 TeV.

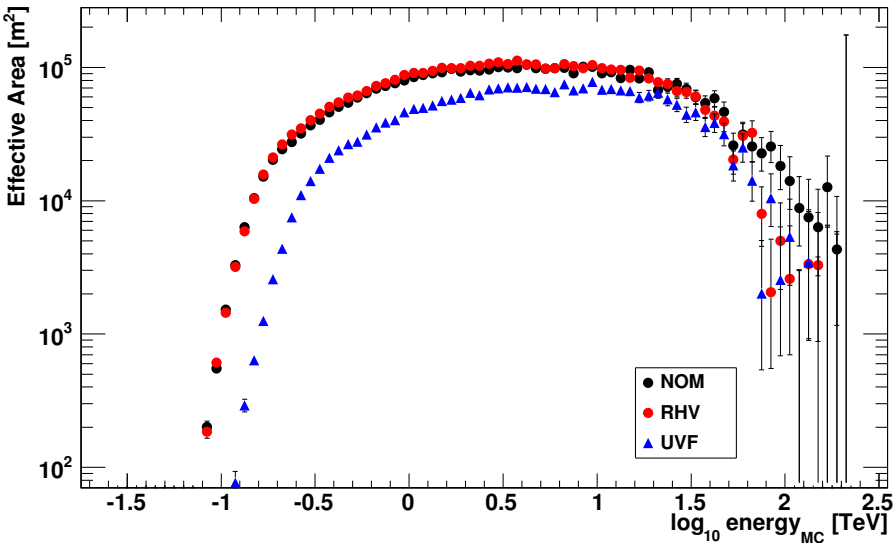

Figure 4.17.: A comparison of the effective areas in the three observing modes.

Spectral Reconstruction

The reconstructed Crab Nebula spectra all agree within errors (Figure 4.18), again the higher energy threshold of the UVF simulations is apparent. The spectra, fitted with a power law of the form

$$\frac{\mathrm{d}N}{\mathrm{d}E} = N_0 \left(\frac{E}{1\,\mathrm{TeV}}\right)^{-\alpha} \mathrm{TeV}^{-1}\ \mathrm{cm}^{-2}\ \mathrm{s}^{-1} \tag{4.1}$$

are

NOM $N_0 = (3.26 \pm 0.10) \times 10^{-11}$, $\alpha = 2.50 \pm 0.03$

RHV $N_0 = (3.26 \pm 0.08) \times 10^{-11}$, $\alpha = 2.51 \pm 0.02$

UVF $N_0 = (2.91 \pm 0.09) \times 10^{-11}$, $\alpha = 2.70 \pm 0.04$

The RHV and NOM spectral fits agree within statistical errors and the data points show no obvious discrepancy. Comparing the image parameters between the data and Monte-Carlo for both shows good agreement as well. The UVF has a softer spectrum than the other two and also a slightly lower normalisation, this is largely driven by the two highest-energy data points which significantly diverge from the UVF and NOM results, the others agree well. Examining these data points in more detail they both have very low counts in their *on* and *off* regions (7 *on*, 7 *off* and 6 *on*, 5 *off* respectively) and thus the Gaussian approximation to the Poisson distribution (which is relied upon for the standard analysis) is poor. It is therefore likely that these two points are statistical fluctuations and they are adversely affecting the fit, excluding them would both harden the index and increase the normalisation. However, rather than tuning simulations to fix spectra or placing artificial constraints on the spectral ranges, the aim instead is always to match other parameters, such as the distribution in the mean scaled width and the *Size* of the showers between the data and the Monte-Carlo, with this agreement looking good there is no justification to adjust the model. It should also be noted that the differences observed here are also comparable to differences between the different "versions" of the VERITAS array (prior to T1 move, after the T1 move but before the camera upgrade and after the camera upgrade and within a single observing season due to changes in the atmosphere). With additional data due to be taken and work ongoing to improve the base simulations model, it is expected that a better agreement will be achieved in future.

4.3.6. Bright Moonlight Conclusions

The VERITAS bright moonlight programme which has been presented here has significantly increased the amount of data taken each year by the VERITAS array. Though operating at higher energy thresholds RHV observations have almost the same effective area above 200 GeV and have already been used in the publication of the detection of flaring activity from the blazar *1ES 1727+502* (Archambault et al., 2015). The UVF

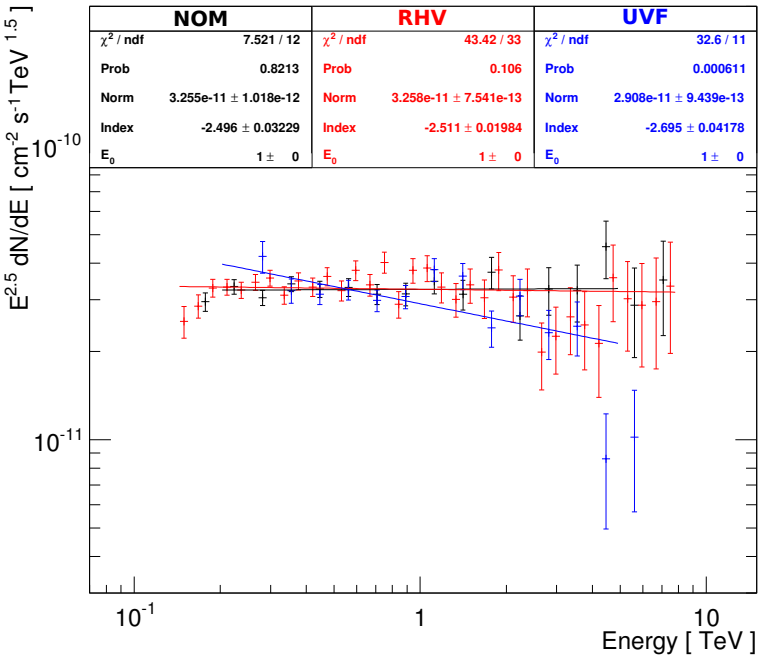

Figure 4.18.: A comparison of the Crab Nebula spectra in the three observing modes.

programme has been used to monitor hard spectrum blazars and also to conduct the novel measurements presented later in this thesis.

4.4. Conclusions

Work has been done to both reduce the impact of the NSB in normal observations and allow observations to be conducted in periods of bright moonlight when NSB levels would normally damage the PMTs. The work on image cleaning has resulted in a few percent increase for the standard observations (through the addition of a time constraint between adjacent pixels) and also shown that, by using *Threshold* cleaning, it is possible to reduce the impact of bright stars on the analysis. The bright moonlight programme has seen RHV and UVF observations conducted and the results verified. RHV observations are now regular observations with comparable sensitivity to standard observations but a higher energy threshold, adding 13% to the observing time by allowing observations up to ∼65% Moon illumination. UVF allows observations at even higher levels of Moon

illumination, though at reduced sensitivity and raised energy threshold, they add a further 16% to the annual observing time.

5. Moon Shadow Observations of the Positron Fraction and the Antiproton Ratio

Since their discovery over 100 years ago by Victor Hess (Hess, 1912) cosmic rays have been the subject of intense research, yet despite all this there is still no conclusive evidence for their origin. Key evidence to help solve this problem will come from their energy spectra and the relative proportions of the different cosmic ray types, in particular the ratio between matter and antimatter.

In this feasibility study (conducted prior to the release of the AMS-02 data) I look at the possibility of making two important measurements, the positron fraction ($e^+/(e^+ + e^-)$) in the ∼500 GeV to ∼1TeV range and the antiproton ratio ($\overline{p}/p$) in the ∼1.5 TeV to ∼3 TeV range (hadronic showers produce around one third of the Cherenkov light of electromagnetic showers (Subsection 2.7.2), thus the energy threshold is about a factor of three higher). Future results in this energy range will provide important extensions on these interesting measurements, constraining the existing theoretical models and improving our understanding of the origin of cosmic rays.

5.1. Summary of this Chapter

This chapter consists of a number of different aspects associated with measuring the positron fraction. The initial, introductary sections are:

Section 5.2 describes the Earth-Moon ion spectrometer technique used for these observations.

Section 5.3 and Section 5.4 describe the history and current status of positron fraction and antiproton ration observations and theoretical predictions.

Section 5.5 and Section 5.6 describe how the observations will be conducted with VERITAS.

The rest of the chapter is divided into three sections:

Section 5.7 describes the feasibility study conducted into these observations by looking at the energy threshold and sensitivity when observing with UVF and RHV. It also looks at the amount of time available and describes simulations of estimates to determine the amount of time required to conduct the observations.

Section 5.8 describes the observations that have been taken to date and analyses the data.

Section 5.9 describes observations that were taken during a lunar eclipse, where the cosmic ray shadow was detected, and analyses that data. It also uses the results to place an upper limit on the antiproton fraction given the observed shadow.

5.2. The Earth-Moon Ion Spectrometer

The Earth-Moon ion spectrometer (EMIS) is a technique to determine the relative fluxes of cosmic rays of different rigidities ($R = p/q$ where p is the particle's momentum, q is its charge). The basis of the technique is that the cosmic ray flux is uniform across the sky. The only exception is if there is an object obstructing this flux, creating a deficit or shadow, such as the Moon. For γ-rays (which don't have a charge) this shadow is (as for a traditional optical shadow) diametrically opposite the source. However, cosmic rays, which are charged, will be deflected by the Earth's magnetic field, with the amount of deflection depending upon the rigidity of the cosmic ray. Thus the deficit will also appear deflected away from the Moon. When viewing this shadow from the ground (Figure 5.1) one sees a band that passes across the Moon, with one charge deflected one way, one the other and with lower rigidity particles deflected further than higher rigidity particles. For a rising Moon in the East, the positive shadow lies above (higher elevation than) the Moon, the negative shadow lies below.

If this shadow is detected then the ratio between the positive and negative shadows can be measured to determine the ratio between positive and negative species. The energy of the "missing particles" can also be determined since the distance from the Moon depends upon that (thus providing a useful cross check to the simulations). By

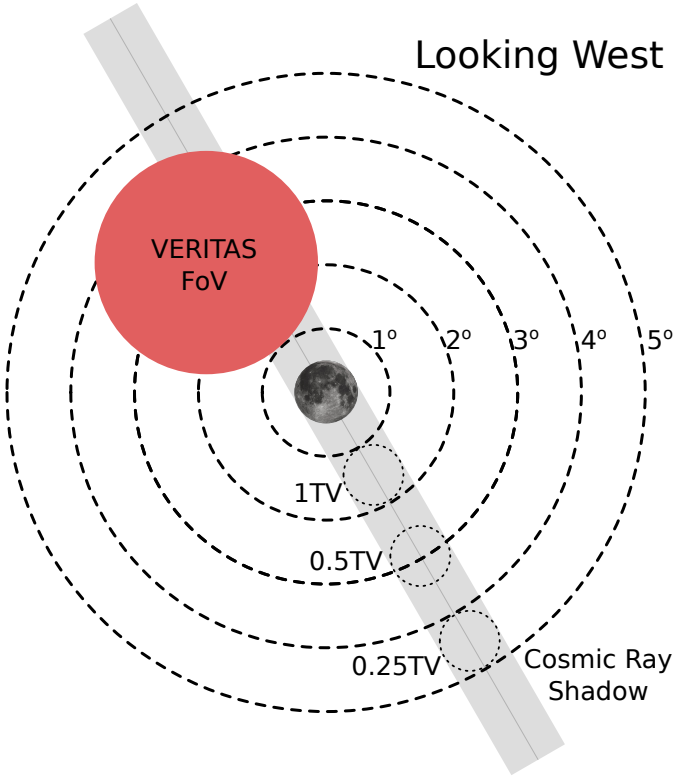

Figure 5.1.: A simplified diagram showing the cosmic ray shadow stretching away from the Moon, a 1 TV particle is deflected by about 1.5° depending on position in the sky. For a setting Moon in the west, the positive shadow is at a lower elevation.

applying analysis cuts hadronic and leptonic showers can be differentiated allowing for separate ratios to be determined for the two*.

5.3. The Positron Fraction

The positron fraction ($e^+/(e^+ + e^-)$) has recently become an area of great interest following measurements that show an increase above 10 GeV, with a myriad of explanations ranging from the distribution of sources to dark matter (see e.g. Linden & Profumo (2013)). However, it should be remembered that there have been indications of this rise for far longer (Figure 5.2).

*After the first interaction for γ-rays which produces an electron-positron pair leptonic showers develop in exactly the same way as electromagnetic showers (Subsection 2.7.1).

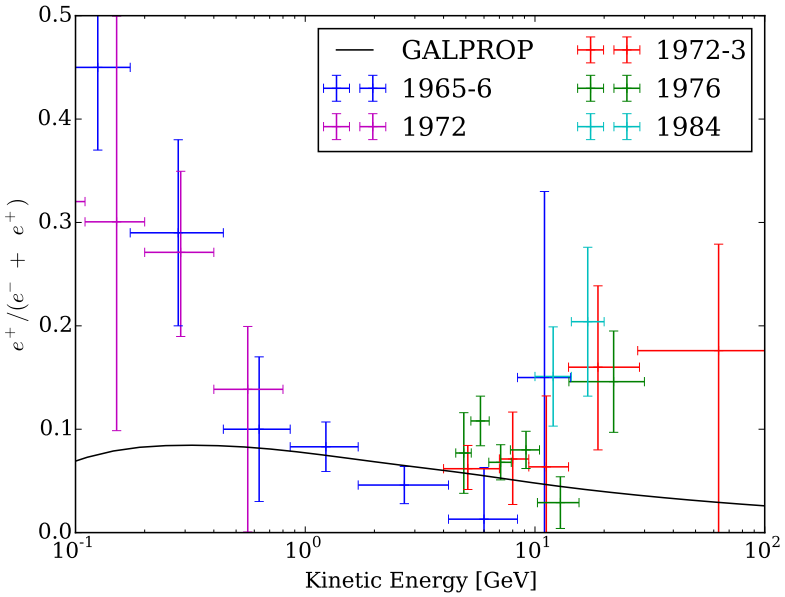

Figure 5.2.: The early balloon data showed hints of an increase in the positron fraction (in comparison to a standard GALPROP prediction, overlaid). References: 1965-6 (Fanselow et al., 1969), 1972 (Daugherty et al., 1975), 1972-3 (Buffington et al., 1975), 1976 (Golden et al., 1987), 1984 (Mueller & Tang, 1987). Data is from Maurin et al. (2014).

Though there was weak evidence from the the balloon flight HEAT (Barwick et al., 1997) and the space shuttle carried AMS-01 (Aguilar et al., 2007), the first conclusive result came from the *PAMELA* satellite (Adriani et al., 2009) which provided detailed measurements of the increase in the positron fraction. *Fermi*-LAT (Ackermann et al., 2012b) has since used an ingenious technique where, similar to the EMIS technique, the shadow in the cosmic-ray flux of one charge caused by the Earth is used to separate the two species. Though there is a noticeable difference in the measured values, it is clear that they both show a similar trend. AMS-02 has since dramatically improved our statistics for this measurement and increased the energy range up to a few hundred GeV (Aguilar et al., 2013) where there is a suggestion of a turnover (especially when the separate spectra are plotted rather than the ratio). The AMS-02 measurements are in good agreement with the *PAMELA* results. Plotting these results (Figure 5.3)

shows the strong agreement between them (with the *Fermi*-LAT fraction slightly higher at all energies suggesting a systematic difference which is not unexpected given the differences in how the measurement was conducted) and the clear divergence from a typical GALPROP* model of the expected fraction (for further details on GALPROP see Section 6.1.2).

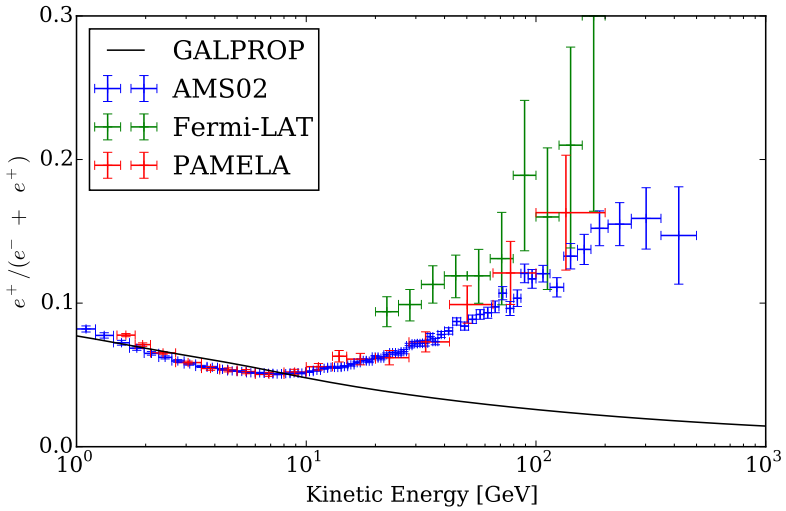

Figure 5.3.: Current measurements of the positron fraction from AMS-02 (Aguilar et al., 2013), *PAMELA* (Adriani et al., 2009) and the *Fermi*-LAT (Ackermann et al., 2012b). Overlaid is a typical GALPROP (Section 6.1.2) prediction of the positron fraction showing the divergence above a few 10's of GeV. Data is from (Maurin et al., 2014).

5.3.1. Current Theoretical Understanding

Predictions of what the positron fraction will be at higher energies come in a variety of forms depending upon the sources hypothesized (dark matter or astrophysical (either new sources or just using the known source population)) and the geometry (location of individual sources, Milky Way shape, general source distribution). A full summary of all of the models, their pros and cons (in particular in relation to other constraints such as the *Fermi*-LAT diffuse γ-ray emission, the antiproton ratio, etc.) is beyond the scope

*http://galprop.stanford.edu/

of this thesis (see e.g. Linden & Profumo (2013) for a summary of the models at that point in time).

Rather, what is important is how a measurement by VERITAS could aid these models. Figure 5.4 shows three typical predictions for the positron fraction:

A steep drop at an energy just above the current measurements

Often viewed as a smoking gun for a dark matter origin of the rise in the positron fraction (e.g. Cholis & Hooper (2013)) it can also be caused by astrophysical sources if distributed correctly within the Milky Way (e.g. Blasi (2013)).

A gradual drop over the next few hundred GeV

More typically associated with local pulsars (e.g. Linden & Profumo (2013)) a gradual curve over the next few hundred GeV. A detection with such a gradual drop-off would likely be constraining for dark matter models and favour an astrophysical solution.

The fraction remaining high (or rising higher) before dropping off

The least-common prediction and very hard to produce considering other constraints but not entirely ruled out and some models exist (e.g. Cowsik et al. (2014) which uses the temporal as well as the spatial distribution of sources).

The targeted energy band for VERITAS observations is also highlighted (yellow) showing that a measurement (or constraining upper limit) in this region could provide significant information for differentiating between these models.

5.3.2. MAGIC Observations

VERITAS is not the only IACT attempting to use the EMIS technique to measure the positron fraction, MAGIC is also attempting to conduct the same measurement (Colin et al., 2009, 2011). These observations are being conducted using a method largely similar to that used by VERITAS (but without the use of UV filters). They have yet to present any results, thus their work is not discussed further.

5.3.3. Previous IACT Observations

Previous observations have been conducted by ARTEMIS (Section 1.3.3) which failed to detect the shadow and CLUE (Section 1.3.3) which detected a deficit from the direction of the Moon.

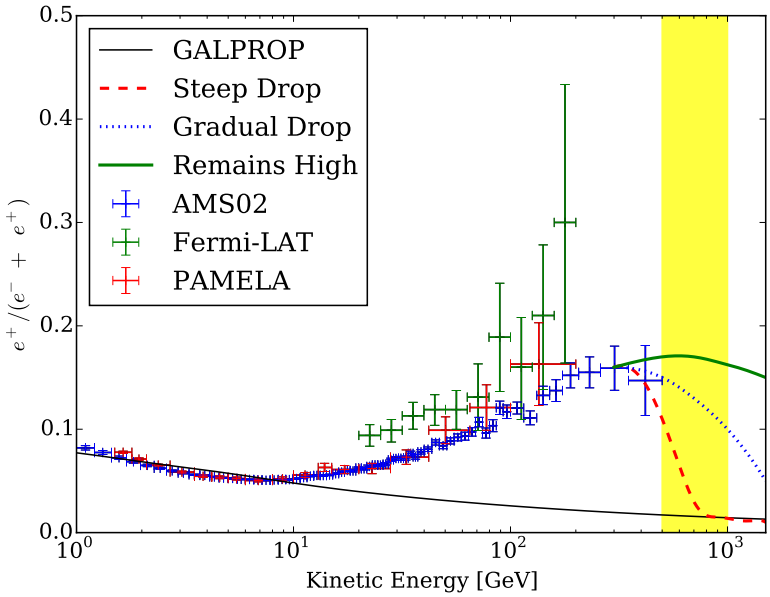

Figure 5.4.: Predictions of the behaviour of the positron fraction (see Figure 5.3 for details) at higher energies than the existing data, the yellow shaded region is the energy range of the potential VERITAS observations.

5.4. The Antiproton Ratio

The antiproton ratio ($\bar{p}/p$) also provides strong constraints on the origin of cosmic rays and their propogation. In particular, when taken into consideration with the positron fraction, it can provide limits on dark matter origins since it constrains how leptophylic dark matter has to be.

5.4.1. Sub-1TeV Satellite Observations - Measurements

Below 1 TeV the antiproton ratio is well measured by satellites, in particular *PAMELA* (Adriani et al., 2010) and the series of balloon flights launched by the BESS collaboration (shown is the data from the BESS-Polar (Abe et al., 2008), Figure 5.5). AMS-02 has not yet published its antiproton data but it has been presented at the AMS Days conference*. The data presented in Figure 5.5 is taken from a figure in the press release, it does not

*https://indico.cern.ch/event/381134/

cover the full range of the presented data and should be considered indicative only (no error bars were transcribed). The *PAMELA* results were entirely consistent with a secondary model of the antiproton ratio, the AMS-02 results present tension with typical GALPROP models that have been tuned to fit the *PAMELA* data but whether they are suggestive of another origin is still up for debate (e.g. Hamaguchi et al. (2015) and Giesen et al. (2015)). I await publication of the results and further tuning of the models to fit all of the data that is being published by AMS-02 before commenting further.

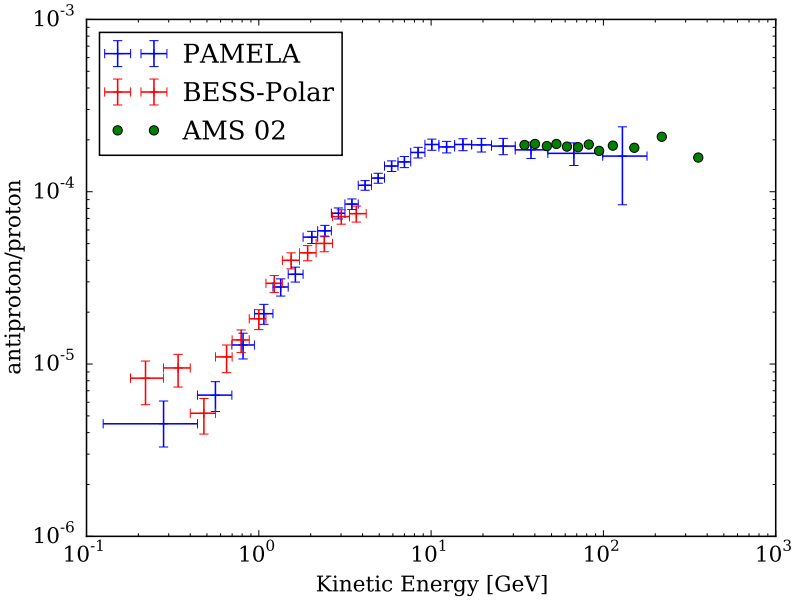

Figure 5.5.: The antiproton ratio below 1 TeV, from AMS-02 (AMS Days press release, https://indico.cern.ch/event/381134/, *PAMELA* (Adriani et al., 2010) and BESS-Polar (Abe et al., 2008). Data is from (Maurin et al., 2014)

5.4.2. Above 1TeV Ground Based Observations - Instrument Calibration and Upper Limits

Above 1 TeV the observations are all taken by ground-based detectors using the EMIS technique. This is primarily done to measure the displacement and the width of the shadow perpendicular to the magnetic field to determine their angular resolution and pointing accuracy. They also calculate upper limits based on the amount of deflection

and width of the shadow parallel to the shadow. As can be seen in Figure 5.6, at present none of these limits are constraining on typical models that fit the existing (and now updated) antiproton ratio measurements at lower energies.

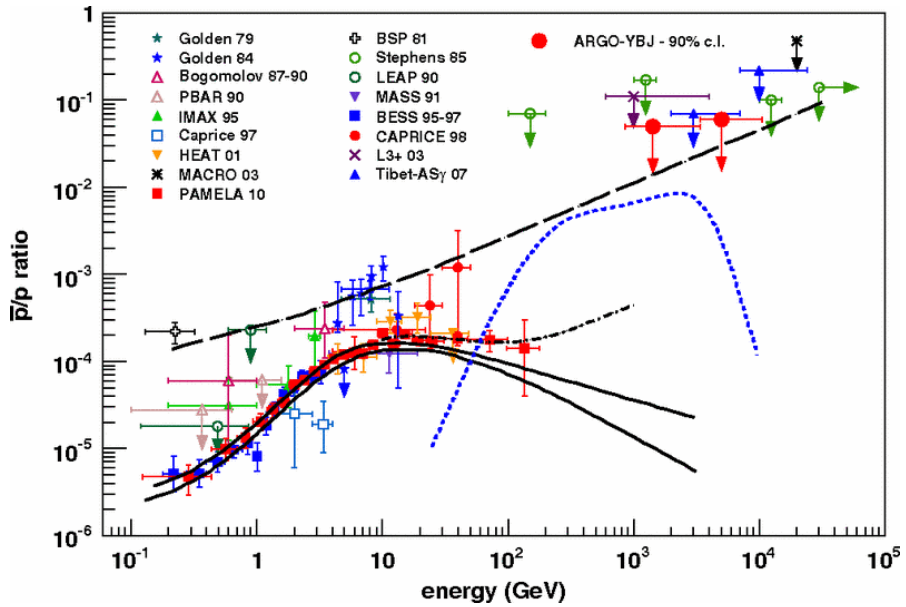

Figure 5.6.: The antiproton ratio (from Bartoli et al. (2012)) showing the upper limits from a variety of experiments above 1 TeV, the solid lines correspond to predictions based upon pure secondary production (Donato et al., 2009) and the dot-dash line includes contributions from additional sources $\bar{p}$ at cosmic ray accelerators (Blasi & Serpico, 2009).

5.5. Shadow Position

In order to conduct the observation I need to know the position of the Moon, the direction of the shadow and how stretched it is. The position of the Moon is a standard ephemeris calculation and in this work I use either the position calculated in the VERITAS telescope tracking software which uses the equations from Meeus (1991), which has been cross checked with the positions from PyEphem* (which is derived from XEphem†) and NOVAS‡, the Naval Observatory Vector Astrometry Software.

*http://rhodesmill.org/pyephem/
†http://www.clearskyinstitute.com/xephem/
‡http://aa.usno.navy.mil/software/novas/novas_info.php

The deflection of a particle is a function of the magnetic field and its rigidity. Since the shadow is a deficit of particles this deficit also appears shifted the same as a particle with the same properties passing through the position of the Moon would be. That is, the 500 GV shadow centroid will be at the position that a 500 GV particle that would have passed through the center of the Moon would arrive at. Thus, in order to determine the position of the deficit, rather than working out the positions of all of the particles not obstructed I can instead work out the positions of particles that have their origin at the position of the Moon.

In order to find this position I used the code TJ2010 (an updated version of TJI 1995[*] with the latest magnetic field, kindly provided by Don Smart). This code works by determining the origin of cosmic rays that arrive at a specific location by "lauching" an antiparticle (that is particle with the opposite charge) in that direction and tracking the path backwards through the Earth's magnetic field. The magnetic field model that is used is the IGRF2011 (Finlay et al., 2010) which is a series of mathematical models of the Earth's main field and its annual rate of change (secular variation). The IGRF specifies the numerical coefficients of a truncated spherical harmonic series: for dates until 2000 the truncation is at n=10, with 120 coefficients, but from 2000 the truncation is at n=13, with 195 coefficients. TJ2010 has been edited so that the test particle is a proton (at the energies I am concerned with the difference between the kinetic energies for a proton and an electron with the same total energy is negligible) with either a positive or negative charge as required. Also, the radius of the final position has been redefined as the radius of the orbit of the Moon (fixed at 60 Earth radii). No consideration has been taken of the impact of the solar wind, comparing results that have taken the solar wind into consideration has resulted in a difference that is well within the angular resolution of the VERITAS array and, since its impact cannot be predicted, the IGRF was deemed acceptable.

For a given arrival direction (as viewed from the VERITAS facility) TJ2010 provides the asymptotic position[†] of the particle when it reaches the cut-off radius, this is then converted to provide the azimuth and zenith of that position as viewed from the VERITAS facility. By iteratively updating the pointing direction until the finishing direction points at the Moon the position of the shadow can be determined for particles of a given rigidity. Plotting a selection of "Moon positions" (Zen = 25°, 35° and 45°, and Az = 90°,

[*] ftp://nssdcftp.gsfc.nasa.gov/models/cosmic_rays/cutoff_rigidity_sw/

[†] The latitude and longitude of the closest point on the surface of the Earth vertically below it.

180° and 270°, Figure 5.7) shows the variation in the shape and the amount of stretch at each position. This variation shows the necessity of a detailed model and the challenge when analysing the data.

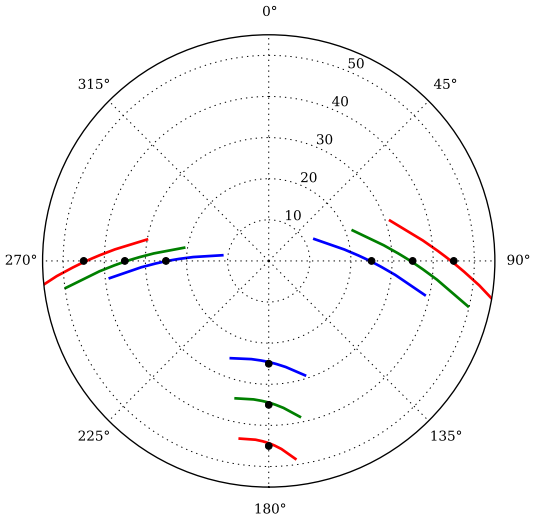

Figure 5.7.: The shape of the shadow for different positions of the Moon (black dots), the rigidities plotted range from $-100\,\mathrm{GV}$ to $100\,\mathrm{GV}$. *Red* = Moon 45° from zenith, *Green* = 35° *Blue* = 25°. Azimuth is measured from North

5.6. Observing with VERITAS

5.6.1. Observing Close to the Moon

To conduct these observations the telescopes must be pointed close to the Moon. However, the Moon provides a significant light source which will increase the background light, reducing the signal-to-noise ratio and potentially damaging the PMTs. To reduce this impact two techniques were employed:

- Reducing the high voltage supply to the PMTs (RHV)
- Employing filters to reduce the amount of light entering the PMTs (UVF)

These methods and their impact on the performance of the array are discussed in detail in Section 4.3. For the standard Moon-shadow observations a combined approach

with both RHV and UVF was used to enable pointing as close to the Moon as possible (2-3°). On 2014-04-15 the VERITAS site experienced a total Lunar eclipse during which observations were conducted. For these observations the Moon was sufficiently dark that it was possible to observe with only UVF (in part due to the fact that the light that is reflected off a totally-eclipsed Moon is that which is scattered onto the Moon through the Earth's atmosphere and thus is strongly peaked in the red part of the spectrum where the filters remove 100% of the the light).

5.6.2. Telescope Pointing

To analyse the data a version of wobble analysis was conducted. The telescopes are pointed 0.5° above and below the shadow axis as shown in Figure 5.8 as close to the Moon as the possible given the constraints on currents in the PMTs. The highest elevation shadow is always targeted as this has the lower energy threshold and is also the shadow that is further from the illuminated side of the Moon. To analyse the data in this way the tracking software for the array was edited to calculate the position of the shadow relative to the Moon (both direction and the amount of stretch). The shadow axis was defined using the Moon and the position of the ± 1 TV shadow, the wobble positions are then $\pm 0.5°$ perpendicular to this line (for a shadow of width 0.5° and an angular resolution of about 0.1° this allows sufficient distance for none of the shadow region to fall in the background side of the camera).

5.6.3. Analysis Setup

With a target that is moving across the sky and changing shape as it moves, a dedicated analysis is required since the standard analysis is set up to analyse objects that are fixed in RA/Dec. To do this the data are analysed using the standard analysis tools until the showers have been reconstructed and the direction, MSW, MSL and energy have been determined. Since, at the time of doing this analysis, I did not have simulations of protons with RHV and UVF, standard γ-ray lookup tables are used. This means that the MSW, MSL and energy are not correct and cannot be used to improve the analysis above the basic analysis presented here (work is ongoing at present with this and they will be completed in the near future). However, the direction reconstruction which does not require lookup tables, since it is calculated using the major axes of each *Hillas* ellipse, will be correct.

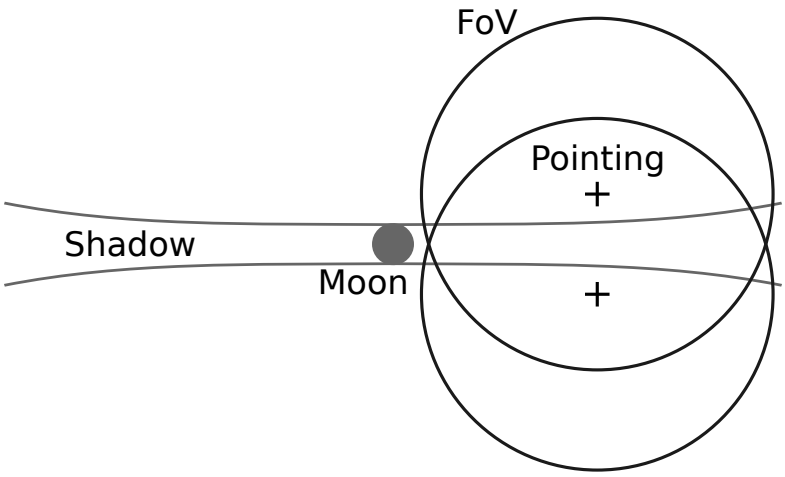

Figure 5.8.: The telescope pointings (crosses) either side of the shadow, perpendicular to the shadow axis. The telescopes are pointed so that the Moon lies outside the FoV. Notice how the shadow broadens with increasing distance from the Moon, this is due to the deteriorating angular resolution with reducing energies and this needs to be considered when choosing the wobble distance.

The data is then read into a purpose written code to conduct the rest of the analysis. The basis of the analysis is, for each event:

1. Determine the position of the Moon (this is done using pyephem*)
2. Determine the position of the desired ± 1 TV shadow (this depends on whether the positive of the negative shadow is desired and what energy range is of interest).
3. Calculate the origin of the shower in "camera coordinates" (see Section 5.6.3 for explanation of camera coordinates).
4. Weight the shower by its radial distance from the camera center by the using the radial acceptance function (calculated using other observations).
5. Rotate the shower position into "shadow coordinates" (see Section 5.6.3 for explanation of shadow coordinates).
6. Stretch the distance parallel to the shadow axis to correct for the stretch in the distance of the ± 1 TV shadow from the Moon as the Moon moves across the sky.

*http://rhodesmill.org/pyephem/, these positions have been cross checked with http://ssd.jpl.nasa.gov/horizons.cgi

A skymap is then generated using the weighted events in the shadow coordinate system. This skymap can then be analysed using a versions of the reflected region/ring background model as described in Subsection 5.7.5 and Section 5.9.

A note on coordinates

In this analysis I use two atypical coordinate systems that are defined here for clarity and shown in Figure 5.9.

The first set are "camera coordinates", they are centred on the Moon and used for calculating the angular distances between two points and aligned with the az/alt coordinate system of the telescopes. They are used as they correct for the fact that, in the azimuthal direction, with increasing altitude, a given angular separation corresponds to an increasing number of degrees in azimuth. To correct for this, the azimuthal direction is scaled by $cos((alt_1 + alt_2)/2)$*.

The second set are shadow coordinates which are based upon camera coordinates but they are rotated so that the azimuthal direction in camera coordinates is aligned with the shadow axis (the line from the Moon to the shadow test position used, with the positive direction defined as being from the Moon to the positive shadow). The zenith direction is rotated by the same amount and thus it now lies perpendicular to the shadow axis. The two axes are referred to as *Parallel* (parallel to the shadow axis) and *Perpendicular* (perpendicular to the shadow axis).

The conversion between coordinates is done thus. For a shower (red star in Figure 5.9) at a position (az/alt =) Sh_{Az}/Sh_{Alt} and the Moon at a position Mn_{Az}/Mn_{Alt}. First, convert the shadow position into "camera coordinates":

$$Sh_{AltC} = Sh_{Alt} - Mn_{Alt} \tag{5.1}$$

$$Sh_{AzC} = (Sh_{Az} - Mn_{Az}) \times \cos\left(\frac{Sh_{Alt} + Mn_{Alt}}{2}\right) \tag{5.2}$$

These are then rotated by the angle θ using the rotation matrix for coordinates

$$\begin{bmatrix} Sh_{Par} \\ Sh_{Perp} \end{bmatrix} = \begin{bmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{bmatrix} \begin{bmatrix} Sh_{AzC} \\ Sh_{AltC} \end{bmatrix} \tag{5.3}$$

*Since I are only working with small differences a simple correction factor based upon the average zenith angle is sufficient rather than the full calculation as in (Meeus, 1991)

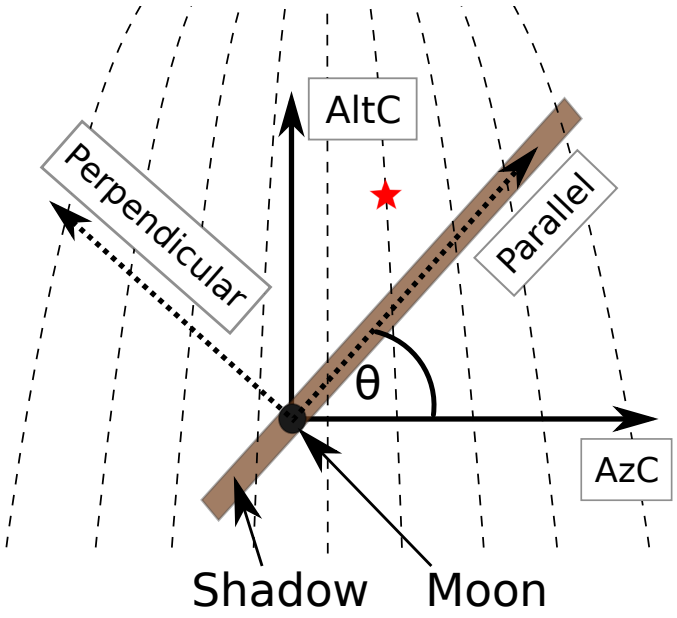

Figure 5.9.: A sketch showing the coordinate systems used, the "camera coordinates" ($AzC/AltC$) are a local systems for measuring absolute distances parallel and perpendicular to the az/alt system (lines of constant azimuth are shown by the long dash lines with the reduction in angular separation between the lines over emphasised for effect). The "shadow coordinates" (*Parallel*/*Perpendicular*) are the "camera coordinates" rotated around the Moon so that the *Parallel* direction is aligned with the shadow. The angle between the "shadow" and "camera" coordinates is θ and the star symbolises a shower position.

The stretch is calculated using the ± 1 TV shadow by taking the average (or some nominal) distance (D_{nom}) between the Moon and the shadow position for all events to be analysed. For each event, the separation between the ± 1 TV position and the Moon (D_{meas}) is calculated and the *Parallel* distance is multiplied by D_{nom}/D_{meas}. This ensures that a given rigidity particle is deflected by the same amount no matter where in the sky the observation was taken.

5.7. Feasibility Study

Before committing to conducting these observations a feasibility study was conducted to see whether an attempt at the standard observations would be successful. To make

that decision the following questions were addressed:

- What is the energy threshold of these observations?
- What is the sensitivity of these observations?
- How long will a detection take?
- How long is available for observations each year?

5.7.1. The Impact of RHV and UVF on Observations

Observations with UVF and RHV allow these observations to occur, but they significantly increase the energy threshold and reduce the effective area (Figure 5.10). Though this is not ideal it is would not be possible to conduct these observations without this. Thus this feasibility study requires careful consideration of these impacts and analysis must be conducted with this reduced performance.

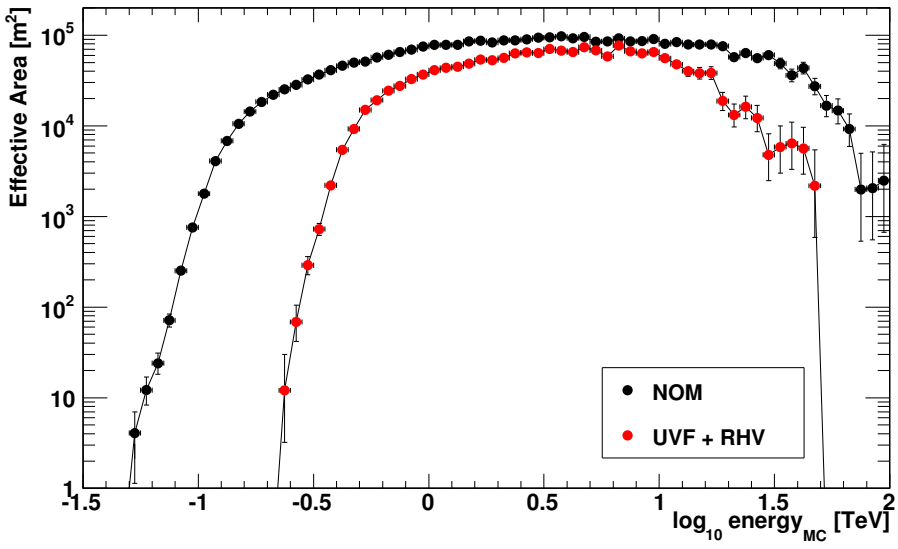

Figure 5.10.: The impact of UVF + RHV (red) is clearly visible on this high elevation, high noise effective area. There is a dramatic increase in the energy threshold and reduction in the area below $\sim$1 TeV. This effective area was produced with a reduced number of simulated showers and it is not known whether this or the impact on the reduced sensitivity is the cause of the reduction in the area at the highest energies.

5.7.2. Test Observations of the Crab Nebula

On 2013-03-19 the Crab Nebula was ∼3° from a 45% illuminated Moon. This provided an opportunity to test the sensitivity of the array in the proposed observing conditions (UVF + RHV). Four runs (Table 5.1) were taken using all of the available and "useful" time (it was decided not to follow the Crab down below 40° elevation since no observations will be taken below this elevation for the Moon shadow observations) with a total live time of 107.5 minutes. Observations were conducted using wobble mode (Section 3.4.7), 0.5° in the cardinal directions N and E. Due to the proximity to the Moon, S and W resulted in currents that were elevated above safe levels.

Table 5.1.: Crab nebula UVF + RHV observations ∼3° from a 45% illuminated Moon.

Run Number	Wobble Direction	Av. El. [deg]	Av. Az. [deg]	Dist. to Moon [deg]	Current [µAmp]	Live Time [min]
67481	E	63.3	255.5	3.53	6.56	27.5
67482	N	56.0	263.5	3.22	6.53	26.8
67483	E	49.1	267.8	3.21	6.05	26.4
67484	N	41.9	273.0	3.28	5.79	26.8

This data shows that (Table 5.2, Figure 5.11) even with both RHV and UVF the VERITAS array is still sensitive to γ-rays, though at a much-reduced level with a significance/sqrt(hour) of 6.82 rather than the normal which is around 35 at these elevations. This loss in significance is due to both a reduction in the effective area and a raise in the energy threshold (Figure 5.15). Unfortunately, with the statistics available (especially as the simulations only cover the zenith angle range of the first two runs), it is not possible to reconstruct a spectrum. These Crab Nebula observations have allowed for verification of the feasibility of observing with UVF + RHV close to the Moon. They have shown that, albeit at significantly reduced sensitivity, VERITAS is still able to detect the Crab Nebula.

5.7.3. Pointing Close to the Moon

Pointing close to the Moon results in a steep gradient in the NSB across the camera as is clearly visible in Figure 5.12. This observation was at ∼3° from the Moon and the currents are approximately inversely proportional to distance, pointing closer to the

Table 5.2.: Results of RHV + UVF observations of the Crab Nebula ∼3° from a 45% illuminated Moon.

Run Number Number	NOn	NOff	Alpha	Gamma Rate [min^{-1}]	Significance
67481	12	3	0.167	0.42	5.71
67482	12	7	0.143	0.41	5.17
67483	11	5	0.167	0.39	4.95
67484	7	7	0.143	0.22	3.40
Comb. (RR)	42	22	Mixed	0.36	9.64
Comb. (RBM)	37	26	0.080	0.32	10.56

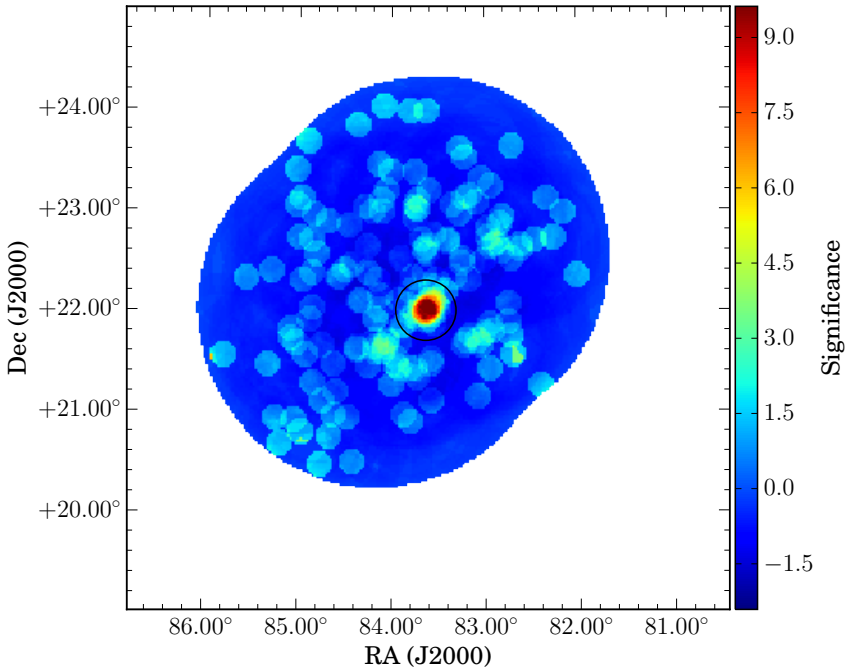

Figure 5.11.: The skymap for the UVF + RHV observations of the Crab Nebula, the Crab is clearly visible and the paucity of background counts is also apparent. The angular resolution and localisation is unaffected by the addition of the filters (within errors).

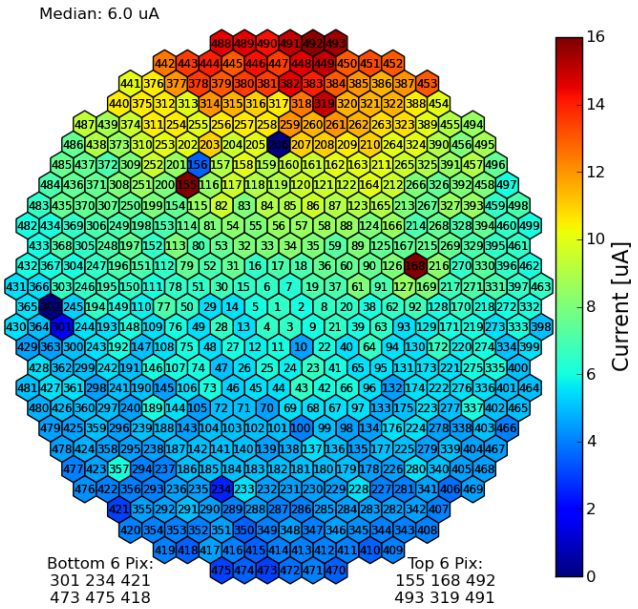

Figure 5.12.: The median currents for telescope 4 during the highest elevation test observation on the Crab Nebula, $\sim 3^{\circ}$ from the Moon and with UVF and RHV. The strong gradient across the camera is highly apparent.

Moon will thus present a significant challenge.

In order to determine how close it is possible to point to the Moon the currents in the individual pixels from the Crab Nebula test observations can be examined (Figure 5.13). The safety threshold of an average current $<15\,\mu$Amp and an individual pixel current of $<30\,\mu$Amp places a firm limit on the proximity that can be achieved. An upper limit of 15 pixels shut down due to currents $>30\,\mu$Amp was placed since pointing closer (and thus increasing the number of shut down pixels) has no net benefit since the part of the camera that is closer to the Moon is not active and reduces the collection area of the camera.

The challenge becomes particularly apparent when you consider the variation in the visual magnitude with illumination (Figure 5.14a) and the day of the cycle (Figure 5.14b).

Using the results in Figure 5.13 and scaling the currents using the data in Figure 5.14a it is possible to generate "fake" observations at different illuminations and to determine how close to the Moon it is possible to point without exceeding the safety limits. Performing these calculations (Table 5.3) (with a maximum of 15 pixels above the individual

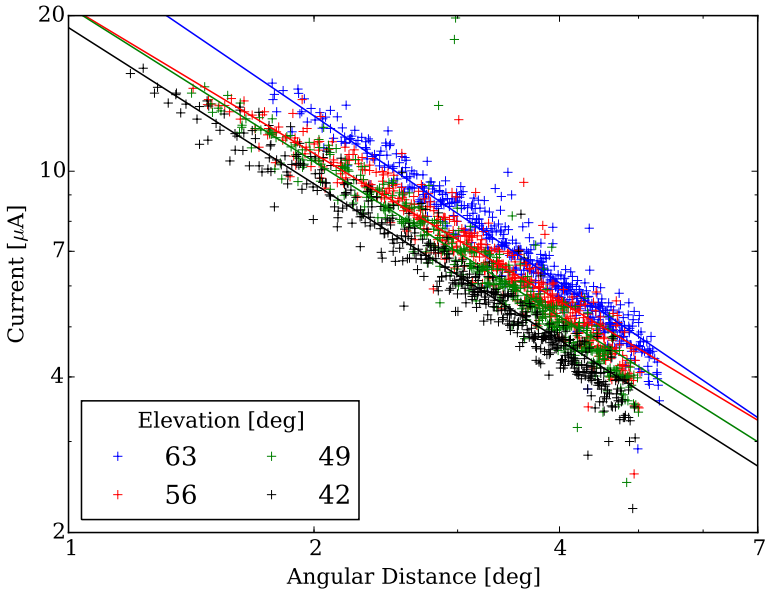

Figure 5.13.: The median currents for telescope 4 during the highest elevation test observation on the Crab Nebula, $\sim 3°$ from the 45% illuminated Moon and with UVF and RHV. The strong gradient across the camera is highly apparent.

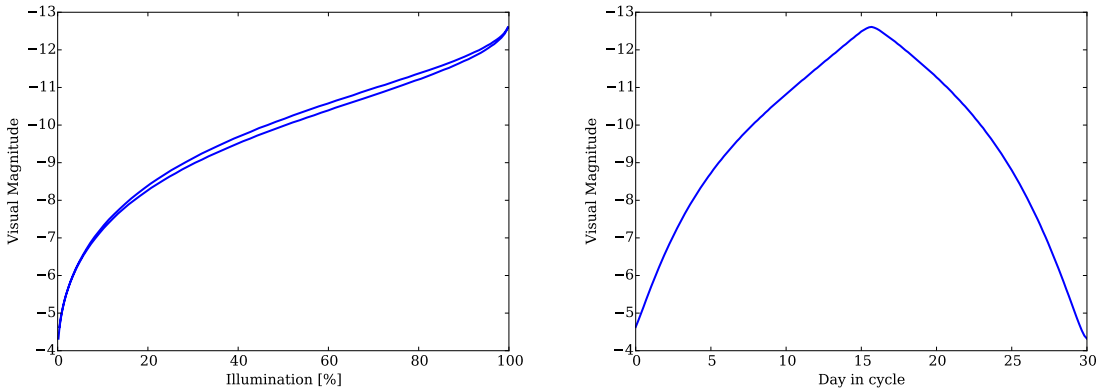

(a) Variation with Moon illumination.

(b) Variation with day of cycle.

Figure 5.14.: The variation in the visual magnitude of the Moon against the fraction of the illuminated surface visible from Earth and the day of the lunar cycle (from http://ssd.jpl.nasa.gov/horizons.cgi). Note the differences between a waxing and waning Moon due to the maria.

safety limit) provides useful data for determining the maximum amount of observing time in a given year. These results are approximations since the local observing conditions will have an impact. The most important is the presence of a haze which creates a halo around the Moon, dramatically increasing the currents in the camera when pointing close.

Table 5.3.: The closest pointing achievable given the Moon's illumination given the safety conditions imposed (mean and median $<15\,\mu$Amp and no more than 15 pixels with a current $>30\,\mu$Amp).

Moon Illum. [%]	Min. Distance From Moon [deg]	Mean [µAmp]	Median [µAmp]	Max [µAmp]	N [>30 µAmp]
52	3	14.3	12.7	34.7	11
49	2.75	13.2	11.4	36.8	11
45	2.5	11.9	9.9	40.4	11
37	2.25	10.2	8.1	48.1	11
30	2	8.1	5.9	74.1	11

5.7.4. Calculation of Observable Time

With knowledge of how close the telescopes can point to the Moon it is possible to determine how much observing time is available per year. This also depends upon the elevation requirements, since the energy threshold dramatically increases with angle from zenith above 30° as can be seen in Figure 5.15 where the threshold energy for normal (i.e. without UVF or RHV) observations is plotted against zenith angle. With UVF + RHV having an energy threshold that is significantly higher, this brings the energy threshold at a zenith angle of 45° to over 1.5 TeV, since I am targeting observations below 1 TeV and, even allowing for the fact that showers are detected at energies below the energy threshold, this presents an upper limit on the zenith angle of the observations.

As well as the zenith angle the amount of stretch in the shadow and the Moon's illumination dictate whether it is possible to point at the region of interest within the shadow, since a brighter Moon requires pointing further away from it without damaging the camera. These two factors can be combined and used, with the model of the shadow and by calculating the position of the Moon as seen by VERITAS throughout the 2014-

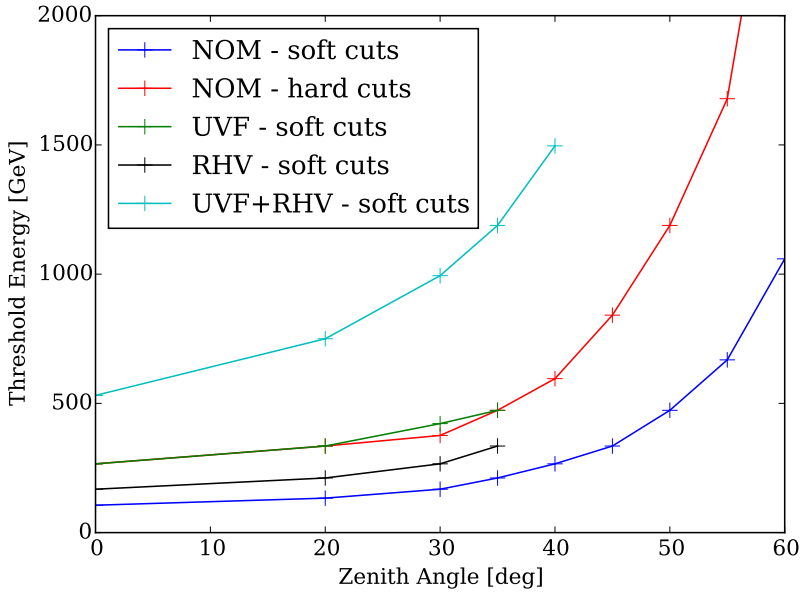

Figure 5.15.: For normal observations (i.e. without RHV or UVF, blue & red) the energy threshold clearly depends upon the zenith angle, increasing dramatically above 30°. With a higher energy threshold for UVF (green) and RHV (black) observations, this limits observations to below 45°.

2015 observing season, to determine an estimate of the total observing time available at 17.5 hours. The breakdown between shadows is dictated by which shadow is higher in the sky (lower energy threshold and further from the illuminated side of the Moon), of the 17.5 hours 8 are targeted at the positive shadow.

5.7.5. Observation Simulation

In order to test the observations feasibility and to develop the analysis technique a toy Monte-Carlo model was generated. For simplicity this model works in the "shadow" coordinate system described in Section 5.6.3. Cosmic rays are drawn from the following distributions:

Protons

The total cosmic-ray spectrum $I(E) = 1.8 \times 10^4 (E/1\,\mathrm{GeV})^{-2.7}$(Olive, 2014) is used, 79% are treated as protons (this is an approximation since the cosmic ray

composition varies with energy but, over the energy range I am interested in, this was deemed suitable). The assumed energy threshold was taken as 1.5 TeV.

Helium

Drawn from the same spectrum as the protons, the remaining 21% are treated as helium nuclei. The energy threshold was assumed to be the same as for protons.

Electrons

Drawn from an $I(E) \propto E^{-3}$ distribution with the normalisation fixed so that at 10 GeV the fraction of leptons is 0.01 (Olive, 2014). Of this leptonic flux, a constant fraction of 0.8 is taken to be electrons, this is at the lower end of what would be expected (Figure 5.3). The energy threshold is 500 GeV.

Positrons

Drawn from the same lepton distribution as the electrons, the remaining 20% are positrons (at the upper range of the expected positron flux).

Antiprotons

With an upper limit of 2×10^{-4} of the proton flux this was considered to be so low as to be ignored for this simulation.

For the first simulation I examined the hadronic cosmic ray shadows (proton and helium), this is referred to as the *hadronic simulation*. All of the thrown showers remain (this assumes that the fraction that trigger the array is constant for particle types and energies above the energy threshold, a reasonable assumption to first order). In the second simulation I examine the leptonic cosmic ray shadows (electron and positron), this is the *leptonic simulation*. To do this the effect of analysis cuts is included with 30% of the leptons passing the cuts and 0.5% of the hadrons (based upon standard numbers from analysing simulated showers).

The showers are thrown over a large area around the Moon, if the thrown position was within the Moon's radius then the shower was flagged and ignored from the analysis. The showers were then "propagated" through the magnetic field using the approximation of a constant deflection of 1° per TV (comparable to that which occurs when the Moon is due south and at an elevation of 50° at the VERITAS site) to determine where they would be observed at from the VERITAS site. After propagation, to simulate the angular resolution of VERITAS the locations are smeared by a Gaussian function of $\sigma = 0.3°$ for

the hadrons and 0.1° for the leptons. The lepton angular resolution is assumed to be the same as for γ-rays since the showers are almost indistinguishable, the hadron angular resolution is taken from Schönwald (2008).

Diffuse γ-ray showers which will provide an irreducible background are not included in this simulation. Since the ratio will be taken using observations at a variety of times with different background regions it is expected that they will largely cancel out though they will reduce the sensitivity. Also, though the γ-ray background at TeV energies is unknown but expected to be very low from observations by the *Fermi*-LAT at lower energies (Ackermann et al., 2015).

Simulation of the Hadronic Shadow

The *hadronic simulation* shown is smoothed with a Gaussian kernel of 0.1°* in Figure 5.16, the total number of counts falling within the FoV is $\sim 2.2 \times 10^6$. This corresponds to ∼30 hours of observations (assuming a post cleaning and reconstruction rate of 10 Hz). The proton shadow is the main deficit and ends before 1 TV with the helium shadow extending through the 1 TV shadow position due to its lower rigidity. With a power-law source spectrum the cut-off is rapid, producing more of a "blob" rather than a band, especially when combined with the impact of the angular resolution (0.3°), which stretches the shadow in all directions, and increases the difficulty in the analysis.

Due to the gradient across the camera in the NSB from the Moon's illumination it was decided not to use a standard reflected-region analysis since there will be a large variation in the NSB the different background regions are exposed too. Instead a modified version is to be used (and tested in this simulation) using single regions reflected above or below the midline of the camera, parallel to the shower axis (Figure 5.17). In the normal analysis the pointings will be alternated above and below the midline, whereas for this simulation and the background is taken from the single region above the midline of the FoV position as in Figure 5.16.

The significance is calculated using 0.39° radius *on* and *off* regions ($0.39 = \sqrt{0.25^2 + 0.3^2}$, the combination of the angular size of the Moon and the angular resolution) using Equation 3.2 (Subsection 3.4.8). The regions are positioned as in Figure 5.17. Calculating the significance along the shadow midline (Figure 5.18) shows the expected deficit, and is fit with a Gaussian function (chosen due to the similarity in the shape) with a peak

*This is done to make it easier to visualise, all the fitting is conducted to the raw data.

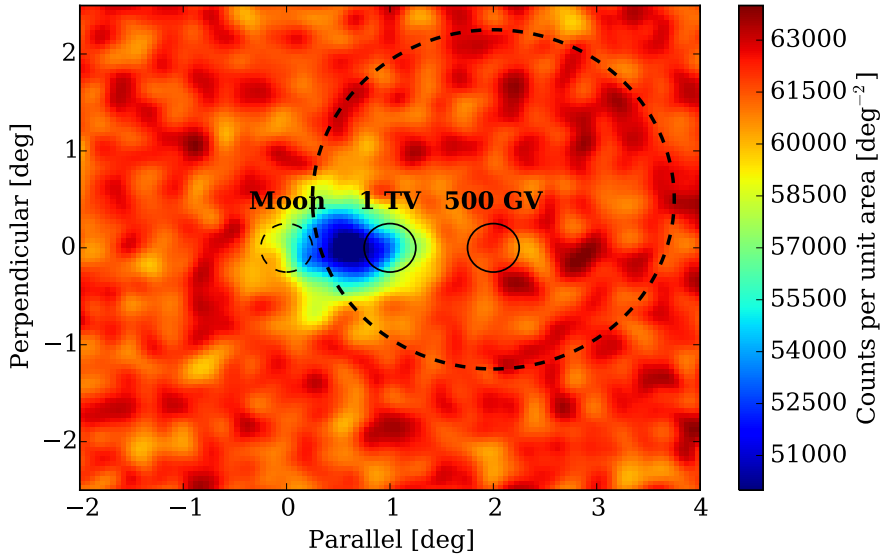

Figure 5.16.: A simulation of the expected skymap from observations of the hadron shadow (smoothed with a $0.1°$ Gaussian function for clarity). The number of events within the FoV (large dashed circle) is equivalent to ∼30 hours of observations. The Moon (small dashed circle) is also shown along with the $1\,\mathrm{TV}$ and $500\,\mathrm{GV}$ shadow positions. The proton shadow is visible close to the Moon, along with the helium shadow deflected further from the Moon due to its lower rigidity.

deficit at $(0.60 \pm 0.03)°$, of width = $(0.46 \pm 0.02)°$ and amplitude = -$(19.60 \pm 0.82)\,\sigma$*. The width of this fit shows the smearing effect of the different rigidities, especially with the helium component.

Simulation of the Leptonic Shadow

Using exactly the same setup as the *hadronic simulation* but with the number of showers of each type scaled by the effect of analysis cuts (30% for leptons, 0.5% for hadrons) produces the *leptonic simulation* show in the smoothed skymap (0.1°) in Figure 5.19.

*The errors are calculated using bootstrap method where an initial fit is conducted, the errors are calculated relative to this fit and the error distribution calculated. This is then used to generate 100 fake datasets using the starting data points and adding a random error drawn from the error distribution. These datasets are then fit and the standard deviation on the parameters is then taken as the error on the parameter as determined in the original fits

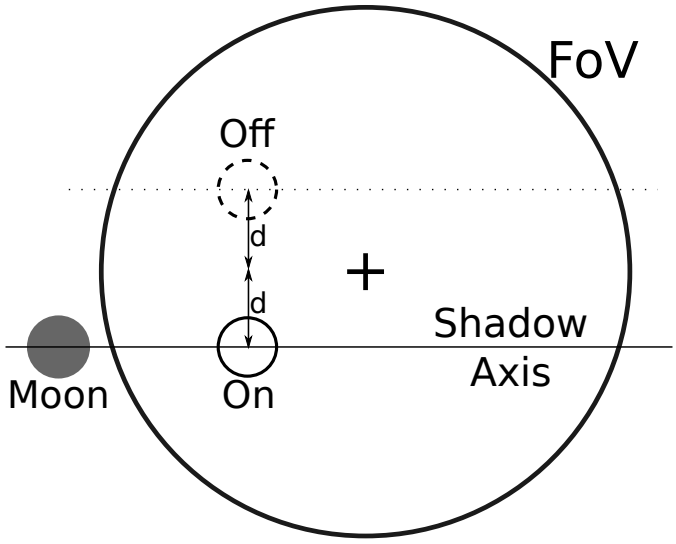

Figure 5.17.: The geometry used in the analysis of the Moon shadow data (not to scale). The shadow axis (solid line) is deflected by a distance, d, from the center of the camera (+), for a given test region, On. The background is estimated using a background region, Off, the same distance (dashed line) from the camera midline parallel to the shadow axis.

With reduced statistics the skymap is a lot noisier but a deficit is visible on the shower midline at a rigidity of $\sim -700\,\mathrm{GV}$. There are other deficits of similar magnitude but running multiple simulations shows this deficit remains, whereas the other deficits are effects of noise. Running a longer simulation improves the statistics, reducing these background fluctuations. Also, more advanced analysis methods which use the energy of the showers improve the sensitivity but, since the energy reconstruction is not available yet for the analysis of the hadronic showers they have not been included here.

Calculating the significances and fitting as in Section 5.7.5, but using a 0.27° ($= \sqrt{0.25^2 + 0.1^2}$) test region, shows the electron shadow has a peak deficit at -(1.46 ± 0.05)° with a width of (0.59 ± 0.07)° and an amplitude of -$(2.83 \pm 0.24)\,\sigma$ (Figure 5.20). If this method is used to detect the shadow then further observing time will be required.

With this number of simulated showers, no evidence of the positron shadow is visible. Increasing the number of showers by a factor of 10 shows a hint of the positron shadow at a significance of -1 to -2σ (the electron shadow is clearly visible with a significance of $\sim$-6σ), a factor of 100 and the positron shadow is detectable significance is about -5σ.

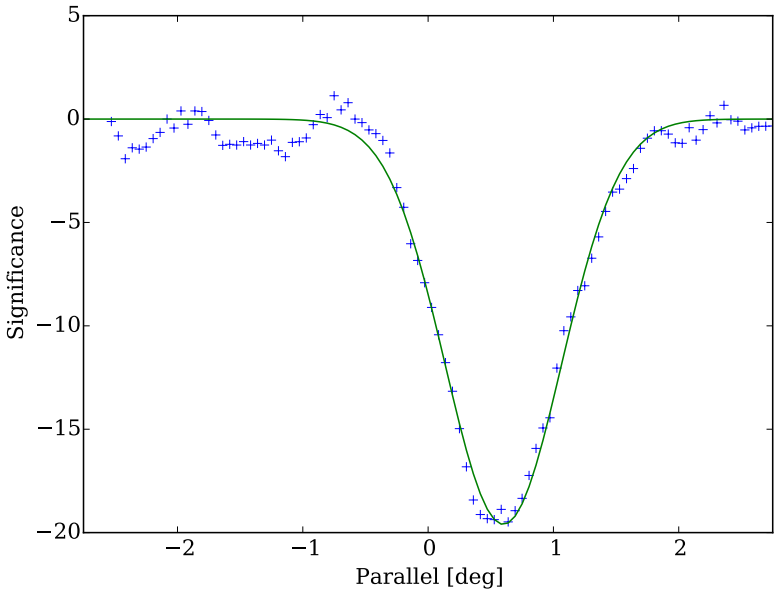

Figure 5.18.: The significance of the excess for circular test regions of radius $0.39°$ along the shadow midline. The background region is described in the text. The result is fit with a Gaussian distribution with a peak deficit at $(0.60 \pm 0.03)°$, width $= (0.46 \pm 0.02)°$ and amplitude $=$ -$(19.60 \pm 0.82)\,\sigma$.

Disentangling the positron shadow from the residual proton and helium shadows (after gamma/hadron cuts) will be difficult, since the shadows slightly overlap. However, since hadronic showers have noticeably different reconstructed energies from leptonic showers, the positions of the showers within a given energy band should be noticeably different. It will also be possible to use the different shapes of the shadows, with the larger PSF for hadronic showers the shadow will be comparatively smeared out. In short, more advanced techniques will be needed to detect the positron shadow, but there are several ways of improving the analysis even using existing tools, though it will also require a significant amount of time. However, if an upper limit can be put on the positron shadow strength and a firm measurement can be made of the electron shadow then it will be possible to place a useful upper limit on the positron fraction with significantly fewer hours than are required for a detection of the positron shadow.

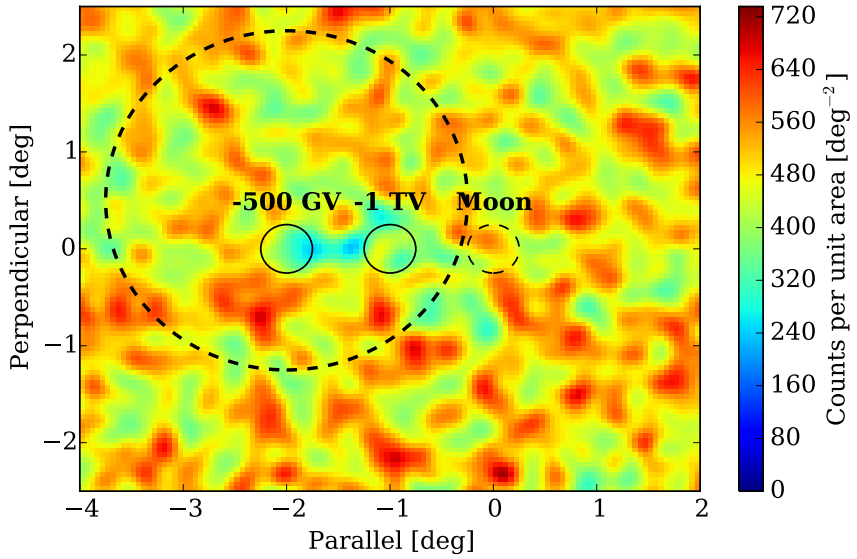

Figure 5.19.: A simulation of the expected skymap from observations of the lepton shadow (smoothed with a 0.1° Gaussian function for clarity). The number of events within the FoV (large dashed circle) is equivalent to $\sim$30 hours of observations. The Moon (small dashed circle) is also shown along with the -1 TV and -500 GV shadow positions. The electron shadow is visible between the -1 TV and -500 GV shadow position.

5.7.6. Feasibility Study Conclusions

This study shows that it is probably feasible to detect the electron and positron shadows with VERITAS but it will not be easy. Taking 10s of hours to detect either shadow (potentially hundreds for the positron shadow) it will require a multiyear investment as a maximum of 15-20 hours can be expected in any one year (less with losses due to bad weather). However, since an upper limit on the positron fraction in the 500 GeV to 1 TeV range could be potentially very interesting a detection may not be required, reducing the amount of time needed. Whilst the data is being taken work on the analysis techniques will continue, improving the sensitivity (for example work is under way to improve the gamma/hadron separation and the angular resolution). The addition of knowledge about the particles' energies will significantly improve the discrimination (since hadrons have a higher energy threshold) and this will reduce the time required, further work here is

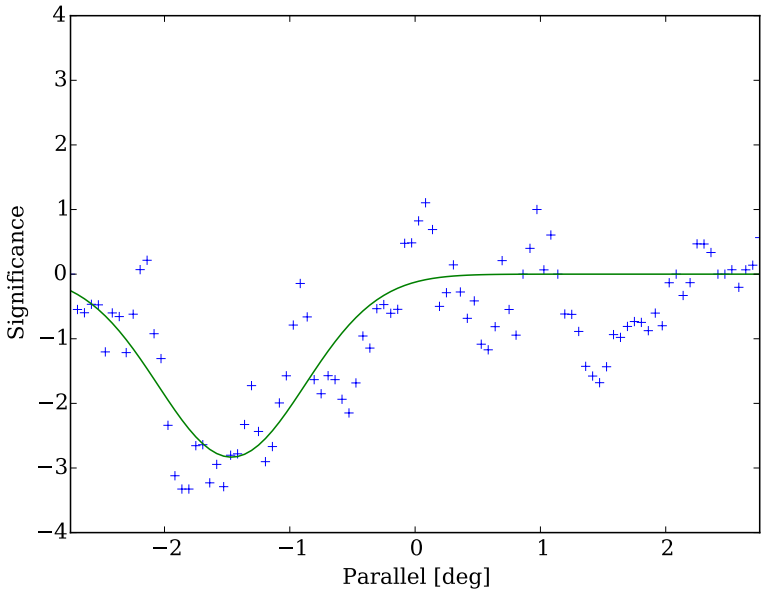

Figure 5.20.: The significance of the excess for circular test regions of radius $0.27°$ along the shadow midline. The background region is described in the text. The result is fit with a Gaussian function with a peak deficit at -$(1.46 \pm 0.05)°$ with a width of $(0.59 \pm 0.07)°$ and an amplitude of -$(2.83 \pm 0.24)\,\sigma$.

ongoing with γ-ray simulations now available UVF + RHV observations but not yet for cosmic ray simulations. The hadronic cosmic ray shadow is easier in terms of time required but, due to the higher energy threshold is closer to the Moon and thus the amount of time it will be in the FoV will be reduced. Analysis will be conducted as the data is taken to see if any evidence becomes apparent.

5.8. Observations

Observations were taken during the 2013-14 observing season totalled 326 minutes (out of a maximum of ∼900 minutes, Table 5.4), the lost time was due to the requirements of perfect weather and the challenges associated with setting up the observations. No data was taken during 2014-15 due to safety concerns, for the operators associated with installation of the filters, and for the telescopes pointing so close to the Moon, work is under way to rectify both of these.

Table 5.4.: Moon shadow observations during the 2013-14 observing season.

Date	Shadow	Moon [%]	Distance [deg]	Time [minutes]	Elevation [deg]	Notes
20131126	Pos	45	3.0	60	55	3 Telescope Run
20140107	Neg	41	2.5	48	47	–
20140205	Neg	35	2.0	27	48	–
20140206	Neg	45	2.5	63	53	–
20140407	Neg	50	3.0	100	54	–
20140506	Neg	42	2.5	28	48	–
Total	Pos	45	3.0	60	55	–
	Neg	44	2.6	266	51	–

Though this time is significantly less than that required to detect either shadow, the data has been analysed for this thesis to show that the tracking is working and to test the coordinate transformations. The data is analysed as described above using the quality cuts in Table 5.5, these are the same quality cuts as used for a hard-cuts analysis (Section 4.2) since I am trying to limit the impact of the optical brightness. Since I do not have UVF + RHV cosmic ray simulations the reconstructed energies and the MSW and MSL are not applicable, and are thus not used as a part of this analysis. When simulations are available they will be used to improve the discrimination between leptons and hadrons and the reconstructed energy will be used to improve the analysis. The positions of the showers are then transformed into shadow coordinates based upon the ± 1 TV shadow position (Section 5.6.3) with the *Parallel* distance normalised so that a ± 1 TV shower has the same *Parallel* distance for all runs, this ensures that the same rigidity shadows are at the same distance from the Moon.

Table 5.5.: Moon shadow quality cuts.

Cut	Value
Size	>1200 dc
nTubes	>5
Distance	$<1.43°$
nTels	≥ 2

The majority of the data was taken on the negative shadow. Looking at the skymap of all events (Figure 5.21) there is no hint of an electron shadow (as would be expected, the majority of the showers will be protons). There are so few events (15608) that (using the numbers from Subsection 5.7.5) after cuts around 50 electrons will remain, this is too few to detect the electron shadow.

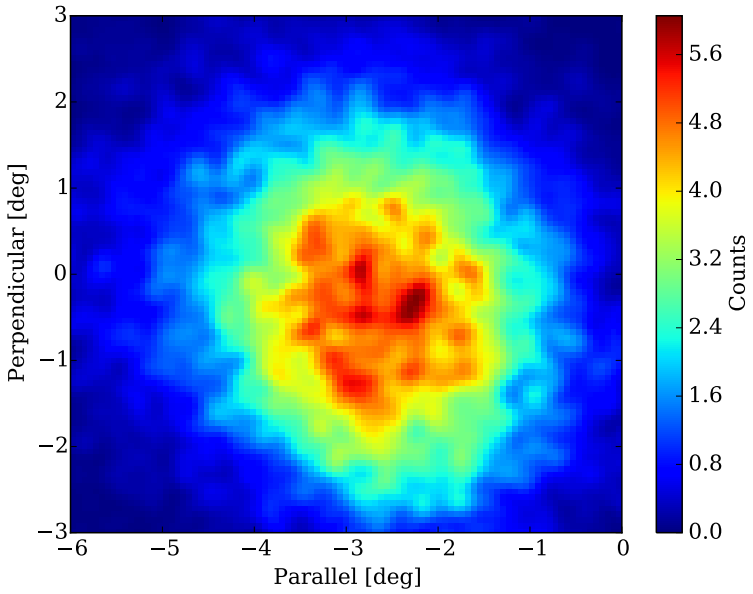

Figure 5.21.: The negative shadow observations, normalised to a $-1\,\mathrm{TV}$ shadow at $-0.9°$.

There is significantly less data on the positive shadow (Figure 5.22) so even though there is a greater chance of detecting the proton shadow this is not possible with the data available due to the limited statistics and the location of the shadow closer to the Moon (due to the higher energy threshold for hadronic showers). Only a handful of showers are expected to pass the gamma/hadron selection cuts so no positron shadow can be detected.

These results do not detect the cosmic ray shadows on their own (significantly more time is required to do that) or allow the calculation of an upper limit.

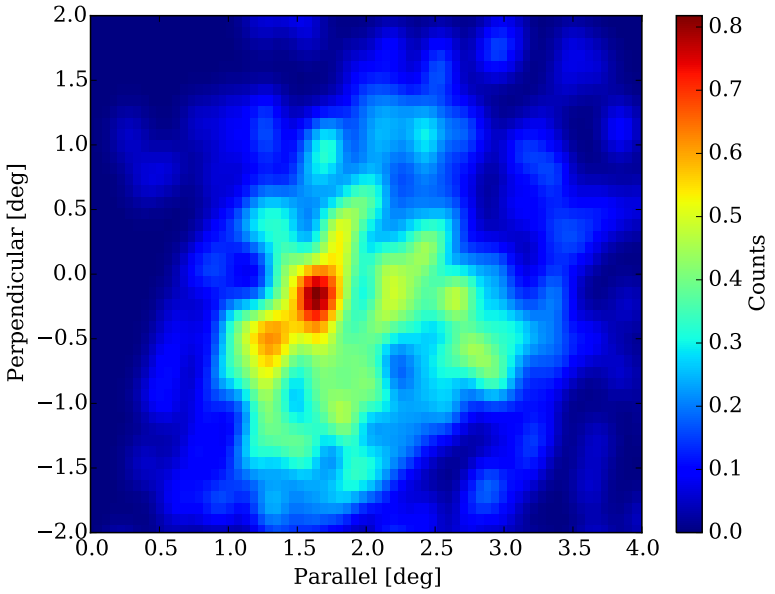

Figure 5.22.: The positive shadow observations, normalised to a 1 TV shadow at 1.46°.

5.9. Lunar Eclipse Observations

On 2014-04-15 VERITAS experienced a lunar eclipse, enabling additional observations to be taken during totality. Due to the nature of a total eclipse it is difficult to predict the brightness (it depends on the atmospheric properties in addition to the position of the Moon within the Earth's umbra). However, since the light that is reflected by the Moon in a lunar eclipse is light that has been refracted by the Earth's atmosphere onto the Moon it is strongly peaked towards the red end of the spectrum. This makes the UVF particularly effective since they are designed to only pass UV light and thus particularly effective at reducing the red light reflected from the Moon. Thus, it was decided to conduct the observations with UVF (though with the capability to switch to UVF and RHV if required) and observe the complete time of totality.

Three observing runs were taken, two of 30 minutes and one of 14 minutes, at quite low elevation (45 to 50°). During the majority of the first two observations (the 30 minute runs) the currents were low (<6 µA) and only raised in isolated pixels. The shorter 14 minute had higher currents and was not included in this analysis.

5.9.1. Assessing the Impact of the Optical Brightness

It is important to show that any deficit is not the result of the optical impact of the Moon. This can be in two forms:

- The RFB can raise the CFD thresholds of part of the camera. With the trigger rates recorded in this observation this will not be a factor as they are so low.
- As discussed in detail in Section 4.2 showers can be lost/turned away from the location of the source of the bright light due to the increased pedvar. This impact is quantified in this section.

In order to put a conservative upper limit on the impact of the optical brightness this analysis is done with quality cuts that are **worse** (standard medium, Table 5.6) at removing NSB impacts than those used in the analysis of the Lunar eclipse data. In reality the optical impact will be less that this.

A dark field of sky, with pedvars less than the darkest regions of the lunar eclipse observations (an observation of the Draco dwarf spheroidal galaxy), was padded* to the noise level of the Moon shadow observations. The two 30 minute lunar eclipse runs were broken into four, 15 minute time slices. The same (15 minute) Draco run is then padded to the pedvars in each of these time slices and the cosmic ray flux is measured across the camera. Any deficit is measured as a percentage of the cosmic ray flux without any padding ((flux with padding - flux without padding) / flux without padding) (Figure 5.23). The peak deficit (fraction of missing showers) in the cosmic ray flux was found in the first and last time slices (Figure 5.23a and Figure 5.23d, those most affected by the optical brightness) at <30%. The middle two time slices (Figure 5.23b and Figure 5.23c) show no deficit. When combining all four slices the peak deficit is 13.3%. With the quality cuts in the Lunar eclipse analysis (Table 5.6) this deficit drops to 12.1% and shrinks considerably in the area it covers.

5.9.2. Skymap Generation

The quality cuts chosen for the analysis of the lunar eclipse were chosen a priori based upon their known efficiency at reducing the optical impact whilst retaining sufficient

*Artificial noise is added to the existing data so that, on a channel by channel basis, the pedvars are raised to a higher level.

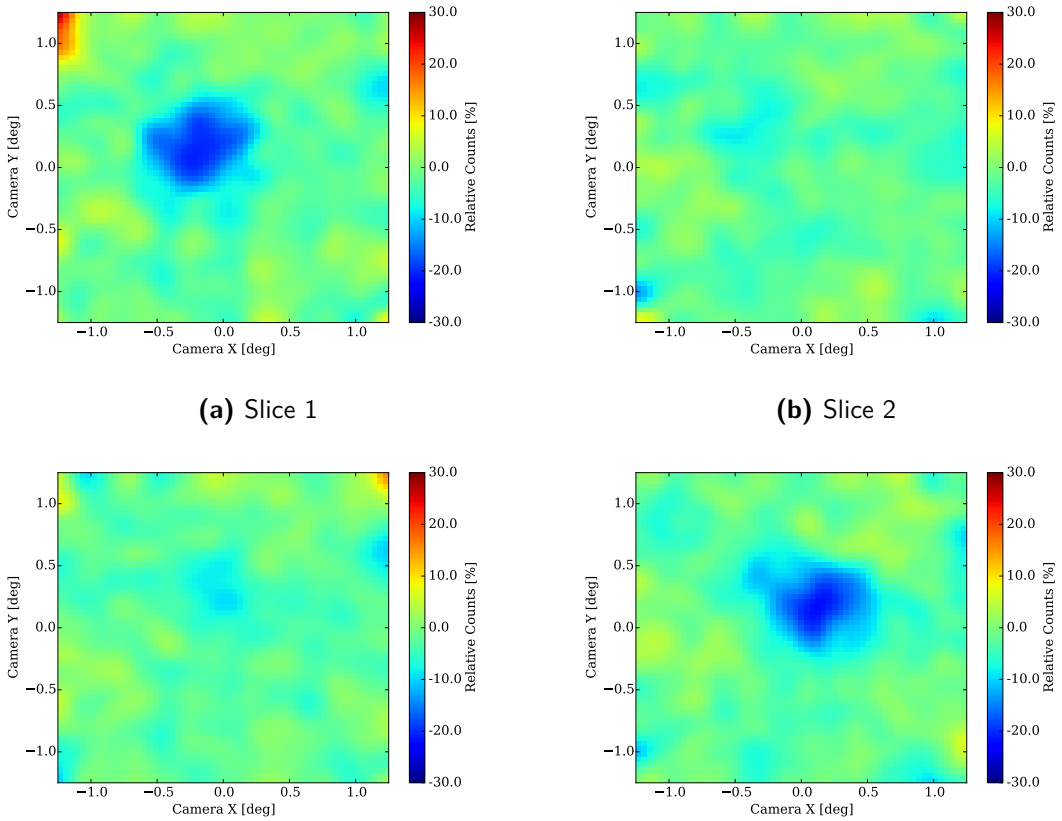

(a) Slice 1

(b) Slice 2

(c) Slice 3

(d) Slice 4

Figure 5.23.: Relative impact ((flux with padding - flux without padding) / flux without padding) of the optical brightness on the four (15 minute) time slices for the Draco run padded to the levels of the NSB in the Lunar eclipse observations. For the first and last slice the peak loss is less than 30% with no noticeable impact in the middle two slices.

showers to detect any deficit. The cost of these quality cuts is that they raise the energy threshold, reducing the statistics and the deflection in the shadow. Both the standard analysis and Lunar eclipse quality cuts are given in Table 5.6.

Table 5.6.: Moon shadow quality cuts.

Cut	Standard Medium	Lunar Eclipse Hard
Size	>700 dc	>1200 dc
nTubes	>5	>5
Distance	<1.43°	<1.43°
nTels	≥2	≥3

The initial analysis of this data was exactly the same (apart from the quality cuts) as for the standard observations (Subsection 5.7.5) with the positions of the showers transformed into shadow coordinates based upon a 1 TV shower (Section 5.6.3), the resulting skymap (without correction for the radial acceptance*) of the 9515 showers that passed the quality cuts is shown in Figure 5.24. This clearly shows a deficit of about 50% due to the shadow of the Moon centred on the Moon†.

5.9.3. Deficit Significance

To calculate a significance of the deficit as was done for the simulations requires selection of a background region. It is not possible to use the same approach as for the simulation as the shadow passes through the camera centre, instead a background region is defined (Figure 5.25) that avoids the shadow and the camera center where the shadow is expected. The test regions are circles of radius 0.30°‡ passing along the shower axis

*Similar plots have been done where the radial acceptance has been corrected for using data from other runs taken at identical elevations and azimuths, analysed using exactly the same method and with the same quality cuts. Overall its impact is negligible since, for hadronic cosmic rays with these cuts, the acceptance is almost flat across the FoV. Outside of the FoV with its steep dropoff in acceptance and after corrections the results outside the FoV are dominated by statistical fluctuations which confuse the eye hence they are not shown.

†Remember that the optical impact peaked at about 13% in the Draco analysis for quality cuts that were less effective at removing the optical impact, also, if only the middle two time slices are analysed (where no deficit is visible in the Draco analysis), though the statistics are much poorer, the deficit is still visible.

‡This smaller than that used in the hadronic simulation where the aim was to maximise the significance of the deficit, here it is to best constrain the shape of the deficit, thus a smaller radius is preferable.

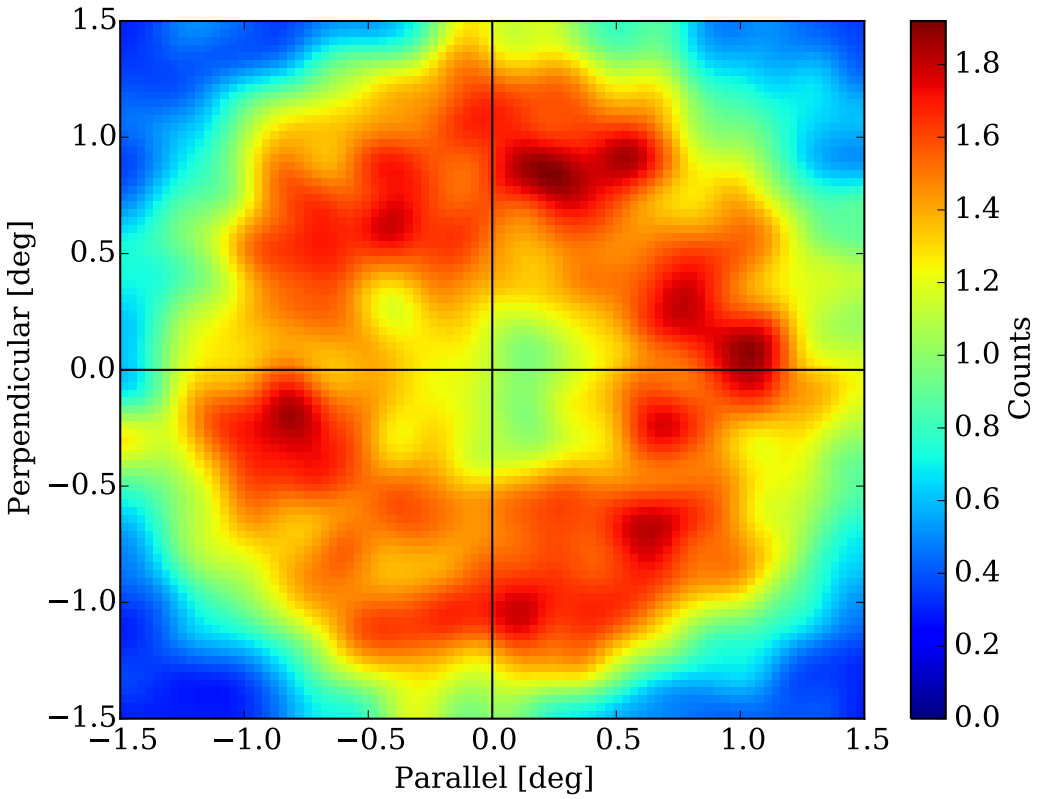

Figure 5.24.: The lunar eclipse skymap showing the deficit due the the shadow of the Moon.

as before. Fitting the central 1° region of the slice with a Gaussian distribution the peak deficit location is $(0.074 \pm 0.010)°$, with width = $(0.39 \pm 0.02)°$ and amplitude = -$(6.18 \pm 0.13)\,\sigma$. (Figure 5.26).

If gamma/hadron cuts are applied in an attempt to examine the electron shadow then only 32 counts remain and no conclusions can be drawn.

5.9.4. Antiproton Ratio Estimation

Though the statistics available are limited, an attempt can be made to measure the antiproton ratio. To do this set of simulated observations are conducted for different energy

The radius was chosen based upon the desire to have sufficient statistics whilst not increasing the width too much. This was done using a simulated dataset.

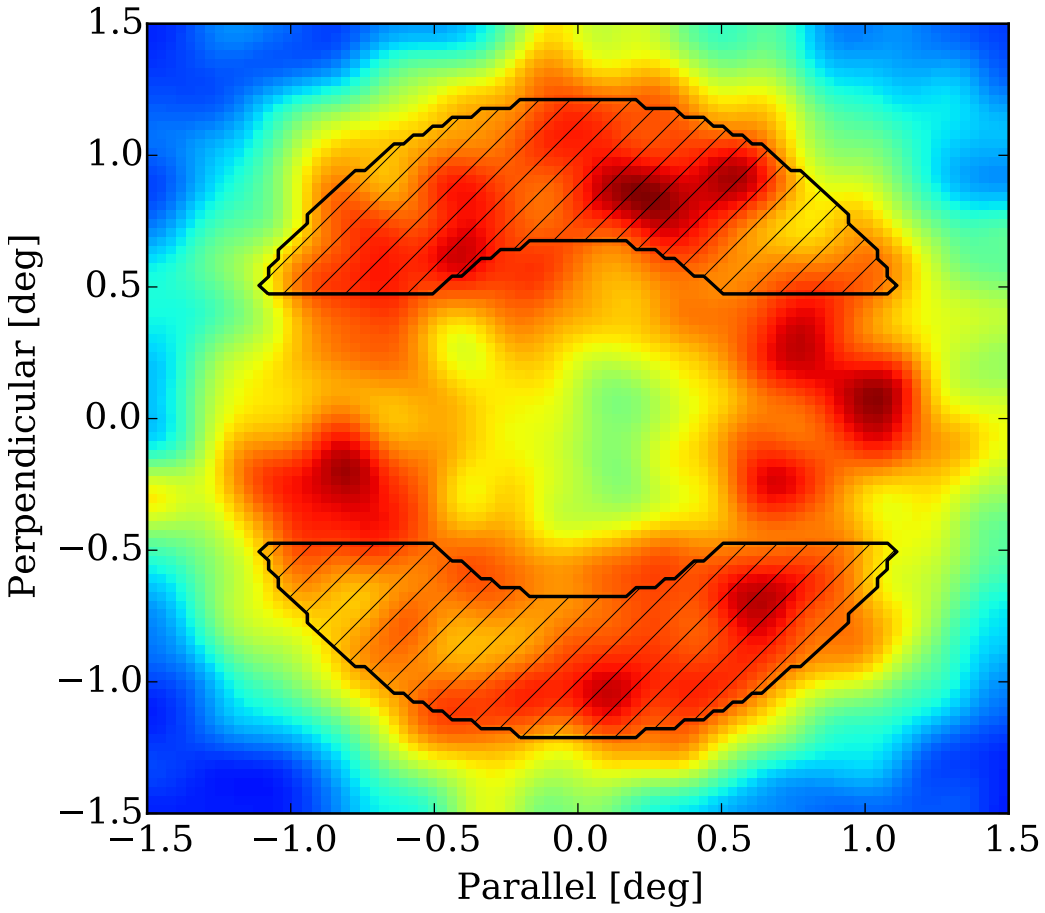

Figure 5.25.: The background region (hatched area) for the Lunar eclipse analysis, it was chosen to avoid the central region and band where a deficit is expected. The outer radius is limited by the drop off in the radial acceptance.

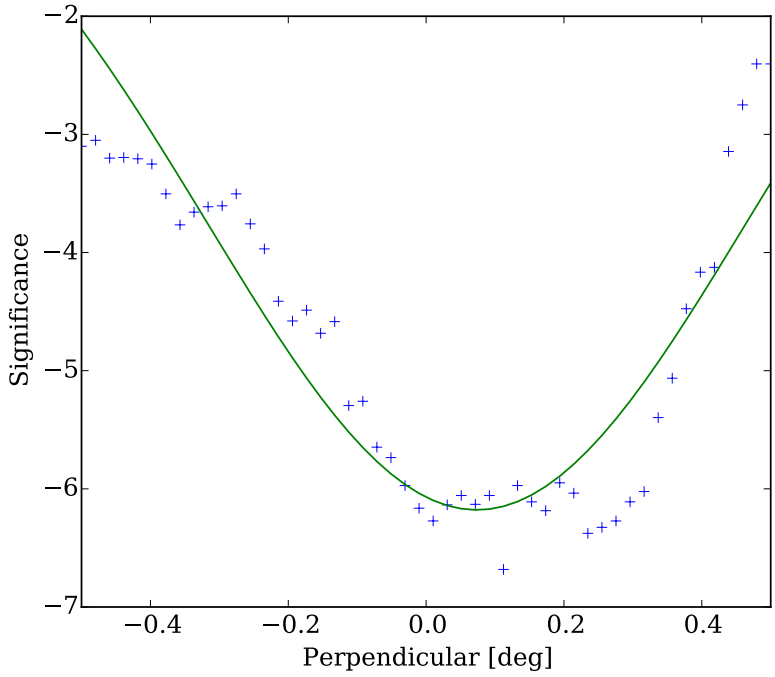

Figure 5.26.: The significance of the excess at each test position along the shadow midline. The Gaussian distribution fit has a peak deficit location of $(0.074 \pm 0.010)°$, with a width of $(0.39 \pm 0.02)°$ and an amplitude of -6.18 $\pm$ 0.13σ.

thresholds and antiproton ratios. These simulations are generated using the species and spectra as described in Subsection 5.7.5 (with the proton fraction split between protons an antiprotons depending upon the faction simulated). They are thrown assuming that the telescopes are pointing at a fixed Moon using a radial distribution function generated from the Draco observations with the NSB padded to the noise level in the lunar eclipse observations*. In contrast to the earlier simulations, where the showers are thrown and traced forwards from the Moon to the Earth, these are traced back from the Earth to the Moon to determine if their path would have passed through the Moon. A deflection of 1° per TV is used and the PSF is also included as before. Those showers that would have passed through the Moon are then excluded from the analysis. The total number of showers was chosen to match the number detected in the Lunar eclipse observations

*This is done to mimic the impact of the optical brightness.

(9515).

These simulations are then fit in exactly the same way as the data, with the significances calculated along the central band and the band is fit with a Gaussian function. The peak deficit and the width are then recorded. For each energy threshold and antiproton ratio, 400 simulations are thrown. From these simulations the mean and standard deviations are determined for the fit peak deficit locations and widths for each energy threshold and antiproton ratio. These are then plotted in Figure 5.27 which shows that, as would be expected the peak location and width get closer zero for increasing energy threshold whereas for increasing antiproton ratio the peak location gets closer to zero but the width increases.

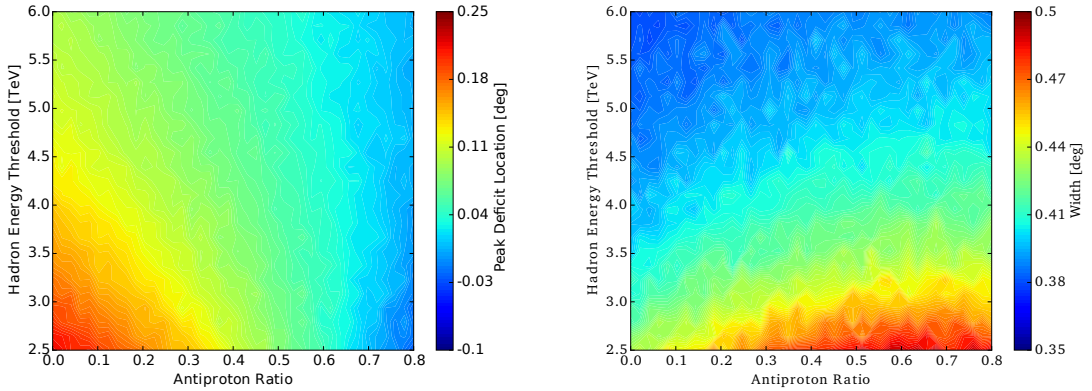

(a) Peak Location **(b)** Standard Deviation

Figure 5.27.: The simulated peak location and width of the fit to the significances for simulated Lunar eclipse observations.

These values can then be used to exclude regions of this phase space. A region is deemed to be excluded if it does not lie within one standard deviation of the measured value, that is, if the difference between the measured value and the simulated value is greater than the errors of that value added in quadrature. The result is the exclusion regions in Figure 5.28 where the red regions are the regions that are excluded at the 1σ level.

These results show that, at the 1σ level, the antiproton ratio has to be less than 0.8 across the entire energy range tested. This is not particularly constraining, especially in comparison to other limits which (in the few TeV range) are at antiproton ratios of

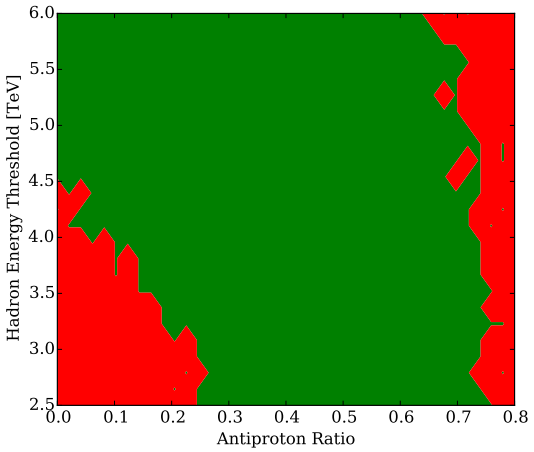

(a) Peak Deficit Location

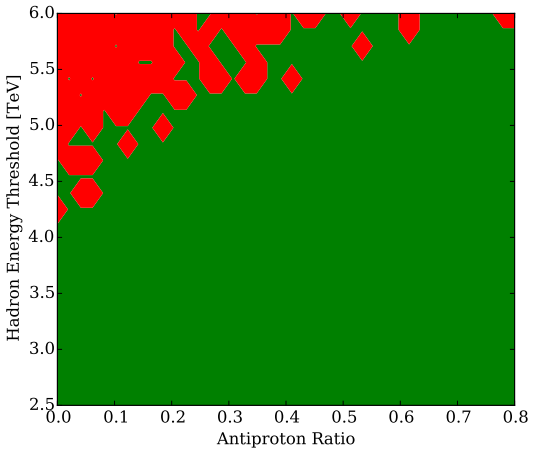

(b) Width

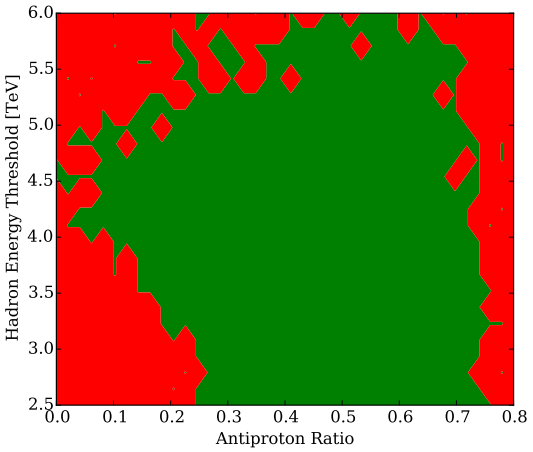

(c) Combined

Figure 5.28.: The combinations of antiproton ratio and energy threshold that are excluded (red) due to the peak deficit location and width of the fits to their significances.

about 0.1 (Figure 5.29). At this antiproton ratio a significant proportion of the tested energy thresholds are excluded, the only one that is not (4.5 TeV) is the about the same as the expected energy threshold assuming that cosmic rays have an energy that is 3× higher than γ-rays. It should be noted that these limits are only at the 1σ level, if a 2σ requirement is placed on them then all the phase space tested is allowed. Given that this is with only a single hour of observations, it does show that, if it is possible to conduct similar observations with CTA (which has an order of magnitude improvement in sensitivity) it should be possible to place upper limits that are better than the current ones in this energy range.

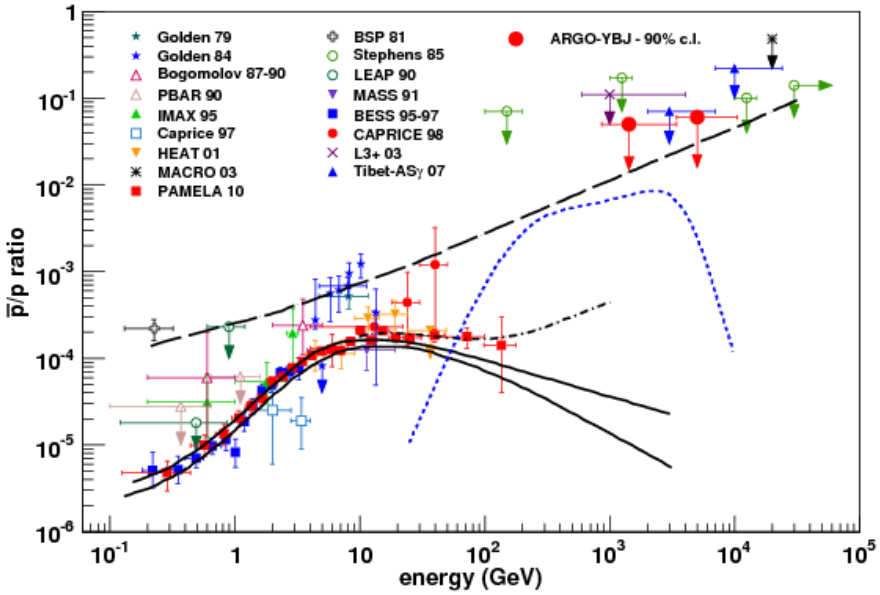

Figure 5.29.: Measurements and upper limits of the antiproton ratio (Bartoli et al., 2012) along with predictions (see Subsection 2.4.1 for explanations of the lines plotted). At low energies the turnover is an effect of the threshold for the formation of antiprotons from secondary particles and from the effects of solar modulation.

5.10. Conclusion

The EMIS technique presents an opportunity for IACTs to measure the positron fraction in the energy range 0.5 TeV to 1 TeV, with the potential to greatly increase the knowledge

about this quantity which has been of such interest in the last few years. Presented here was a feasibility study into making this measurement with VERITAS. To make a measurement with VERITAS will take many tens of hours, (if not hundreds for the positron shadow) taking many years but an upper limit could prove very informative. With the data from the 2013-14 season no evidence of a deficit is apparent, though the total observation time was low ($\sim$5.5 hours). Observing during a lunar eclipse shows a clear deficit in the cosmic ray flux due to the shadowing by the Moon. The significance of the deficit is calculated for a band along the shadow axis and fit with a Gaussian distribution. The peak deficit location is $(0.074 \pm 0.010)^\circ$ with a width of $(0.39 \pm 0.02)^\circ$ and amplitude of -$(6.18 \pm 0.13)\,\sigma$. At the 1σ level antiproton fractions >0.8 are excluded based upon the peak location and width and assuming that energy threshold is in the range of 2.5-6 TeV.

6. Diffuse γ-ray Emission from M 31

M 31 (the Andromeda Galaxy) is the closest spiral galaxy to the Milky Way (MW), at a distance of 780 kpc (Stanek & Garnavich, 1998), and is one of the prime targets for the study of the structure and behaviour of an entire galaxy. With a radius of 20 kpc, and an inclination of 78° (Draine et al., 2014), it has apparent dimensions of 3.2° by 1°* allowing its internal structure to be resolved at most wavelengths. As a result there are an extensive suite of observations of M 31 at many different wavelengths. VHE γ-ray emission from M 31 is expected due to the interaction of cosmic rays with the interstellar medium (using the mechanisms outlined in Section 2.5) to produce diffuse γ-rays, and from unresolved point sources within the galaxy. However, as yet, M 31 has not been detected in VHE γ-rays, though an upper limit on the VHE flux was presented by the HEGRA collaboration in Aharonian et al. (2003) at 3.3% of the Crab Nebula flux for individual point sources close to the center of the galaxy (this rises to ∼30% near the edges). HAWC has presented a preliminary upper limit for M 31 of less than 14% of the Crab Nebula flux using their 111 tank configuration (Brewer & Smith, 2015). A detection in HE γ-rays was published by the *Fermi*-LAT collaboration using 2 years of data (Abdo et al., 2010).

In this chapter I explore the expected emission from M 31, making a prediction of the expected HE and VHE flux based upon the comparative properties for M 31 and the MW. I then compare this prediction to an updated analysis of the *Fermi*-LAT data using a significantly larger data set and the improved data reduction techniques since the initial detection. Using this knowledge I then compare the results with an analysis of observations of M 31 by VERITAS and present a summary of the current VHE picture.

*http://ned.ipac.caltech.edu/

6.1. Predictions of M 31 γ-ray Emission

There are expected to be two main contributions to the VHE emission from M 31: diffuse emission from the interaction of cosmic rays with the interstellar medium (ISM) and the interstellar radiation field (ISRF), and unresolved point-source emission (I do not expect to see any individual point sources, placing a source with the same luminosity as the Crab Nebula in M 31 gives a flux of $(^{2\,\mathrm{kpc}}/_{780\,\mathrm{kpc}})^2 = 0.0007\%$ of the Crab Nebula). Due to its apparent size in relation to the VERITAS angular resolution, a detection by VERITAS will allow it to map the VHE γ-ray emission and thus I could determine where the emission is coming from. The emission could originate from regions where there are a number of potentially unresolved point sources, suggesting that it is the sources themselves, or cosmic rays that are only diffusing a short distance from the source population, that are producing the signal. Otherwise if the emission is more diffuse and coming from regions away from potential sources, then it will suggest that the emission is caused by cosmic-rays diffusing large distances through the galaxy from their sources. Since hadronic cosmic rays have a significantly longer lifetime than leptonic cosmic rays this would provide good evidence for a hadronic origin for the γ-ray emission. The γ-ray emission spectrum will also provide information about the underlying physics, in particular at around 0.1 GeV where detection of a low-energy cut-off (a "pion-bump") would provide significant evidence that the diffuse γ-ray emission is of hadronic origin.

6.1.1. Basis for M31 γ-ray Emission Model

It is widely believed that the γ-ray luminosity of normal galaxies (that is galaxies without an active galactic nucleus) should scale based upon a few key parameters which dictate the number of cosmic ray accelerators, the escape time of the cosmic rays and the amount of target material (see e.g. Pavlidou & Fields (2001)). The number of cosmic ray accelerators is predicted to scale with the star formation rate (SFR, Condon (1992)), since a higher star-formation rate leads to a higher supernova rate (SNRate) and the majority of cosmic ray acceleration is expected to come from SNRs or other, related, objects such as pulsars, PWN and (compact) binary systems. The escape time of cosmic rays will depend upon how they diffuse through the ISM of a galaxy. Since they are charged this will largely depend upon the strength of the (absolute and random components of the) magnetic field in the galaxy (Ptuskin, 2006). Target material is pri-

marily the amount of ISM (for pion decay) and the strength of the interstellar radiation field (ISRF) (for inverse Compton (IC) emission). Of these the pion decay is expected to dominate, though the IC component could be particularly important at high energies and close to sources (Ptuskin, 2006).

Working from these assumptions a number of predictions have been made on the HE γ-ray flux for a number of galaxies, most notably in Pavlidou & Fields (2001) where they use the following scaling relationship.

$$F_\gamma(>100\,\mathrm{MeV}) = 2.34\times10^{-8}\frac{SNRate_{galaxy}}{SNRate_{MW}}\frac{M_{gas}}{10^8 M_\odot}\left(\frac{d}{100\,\mathrm{kpc}}\right)^{-2}\ \text{photons cm}^{-2}\ \text{s}^{-1} \tag{6.1}$$

That is, they predict that the flux should scale linearly with both the supernova rate (SNRate) and the mass of gas (principally hydrogen). Other (similar) calculations are performed in Persic & Rephaeli (2012) and Rephaeli & Persic (2013), where they use their results to determine the cosmic ray density in the galaxy, and many other publications, all of which draw largely similar conclusions for the relationship between the diffuse γ-ray flux, the SFR and the mass of hydrogen (M_H).

Using the parameters in Table 6.1 and existing (published) measurements of the flux from normal and starburst galaxies (taken from the 3FGL, Acero et al. (2015)), I have derived a scaling relationship that compares the galaxy's luminosity with its SFR. In addition to the standard galaxies that were used in Abdo et al. (2010), the two Seyfert galaxies NGC 4945 and NGC 1068 are also included, although there is a debate in the literature whether the emission is dominated by the central black hole or the star-forming regions (see e.g. Lenain & Walter (2011)). They are plotted here on the assumption that the emission is dominated by the star formation region but it should be remembered that there will be a contribution from the central region.

It can be seen that there is quite a strong correlation with the SFR (Figure 6.1a). The best fit though is not linear with SFR, rather it scales as $SFR^{1.28}$. This is to be expected since the initial assumption is that $F_\gamma \propto SFR \times M_H$ and since the SFR will depend upon M_H I expect a scaling similar to that which is observed (Persic & Rephaeli, 2011). Though the correlation with the number of hydrogen atoms (Figure 6.1b) is much lower, there is a general trend with the luminosity scaling as $M_H{}^{1.66}$.

Using the SFR scaling on the starburst galaxies M 82 and NGC 253 it is possible to predict the expected VHE γ-ray flux from M 31 using

Table 6.1.: Properties used to compare the galaxies.

Galaxy	Distance [kpc]	MHI [$\times 10^8 M_\odot$]	MH2 [$\times 10^8 M_\odot$]	SFR [$M_\odot\,yr^{-1}$]	HE Luminosity [erg s^{-1}]
M 31	780 ± 33^a	73 ± 22^a	3.6 ± 1.8^a	$0.35 - 1^a$	$(4.28 \pm 0.68) \times 10^{38b}$
MW	–	35 ± 4^a	14 ± 2^a	$1 - 3^a$	8.0×10^{38d}
LMC	50 ± 2^a	4.8 ± 0.2^a	0.5 ± 0.1^a	$0.20 - 0.25^a$	$(5.75 \pm 0.35) \times 10^{37b}$
SMC	61 ± 3^a	4.2 ± 0.4^a	0.25 ± 0.15^a	$0.04 - 0.08^a$	$(1.8 \pm 0.15) \times 10^{37b}$
M 82	3630 ± 340^a	8.8 ± 2.9^a	5 ± 4^a	$13 - 33^a$	$(1.69 \pm 0.26) \times 10^{40b}$
NGC 253	3940 ± 370^a	64 ± 14^a	40 ± 8^a	$3.5 - 10.4^a$	$(1.61 \pm 0.27) \times 10^{40b}$
NGC 4945	3700^c	$9.64^{+0.10}_{-0.40}{}^c$		4.7^c	$(2.14 \pm 0.36) \times 10^{40b}$
NGC 1068	16700^c	$9.71^{+0.11}_{-0.19}{}^c$		50^c	$(1.63 \pm 0.31) \times 10^{41b}$

[a] Abdo et al. (2010)

[b] Acero et al. (2015)

[c] Rephaeli & Persic (2013)

[d] Strong et al. (2010)

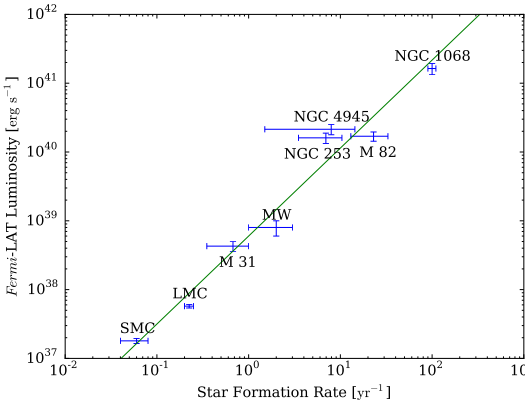

(a) γ-ray Luminosity vs. SFR overlaid with Luminosity $\propto (\mathrm{SFR})^{1.28}$

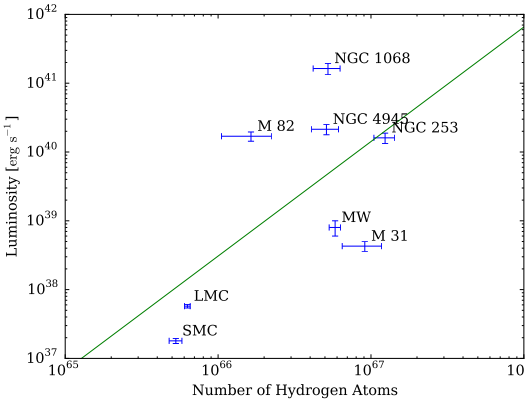

(b) γ-ray Luminosity vs. Number of Hydrogen Atoms overlaid with Luminosity $\propto (\mathrm{N_H})^{1.66}$.

Figure 6.1.: HE (*Fermi*-LAT) γ-ray luminosity compared with SFR and the number of hydrogen atoms (fits constrained to pass through origin).

$$Flux_{M31(pred)} = Scaling\,Factor \times Flux_{galaxy(measured)} \tag{6.2}$$

$$= \left(\frac{SFR_{M31}}{SFR_{galaxy}}\right)^{1.28} \times \left(\frac{distance_{galaxy}}{distance_{M31}}\right)^{2} \times Flux_{galaxy(measured)} \tag{6.3}$$

This scaling has been calculated using measurements of HE γ-ray using the *Fermi*-LAT. However, since it is the same processes for VHE emission, the same scaling law should apply. This means that a prediction of the VHE M 31 flux can be made based upon the existing VHE detections of the starburst galaxies M 82 and NGC 253 (Table 6.2).

Table 6.2.: Predictions of the M 31 VHE γ-ray Flux - the measured VHE γ-ray flux from starburst galaxies scaled using Equation 6.3.

Galaxy	Measured Flux @ 1 TeV	M 31 Predicted Flux @ 1 TeV	
	[TeV^{-1} cm^{-2} s^{-1}]	[TeV^{-1} cm^{-2} s^{-1}]	[% Crab]
M 82[a]	$(3.3 \pm 0.7) \times 10^{-13}$	$(0.78 \pm 0.17) \times 10^{-13}$	0.23 ± 0.05
NGC 253[b]	$(0.96 \pm 0.15) \times 10^{-13}$	$(1.24 \pm 0.27) \times 10^{-13}$	0.37 ± 0.08

[a] Acciari et al. (2009a)

[b] Abramowski et al. (2012)

The VHE flux is predicted to be a few tenths of one percent of the Crab Nebula flux, spread over the entire galaxy.

6.1.2. Generating an Emission Model

Existing Models

No existing models of the VHE emission from M 31 have been found in the literature. Over the years various dark matter predictions have been made (e.g. (Falvard et al., 2004) and (Fornasa et al., 2007)), these are old models that make predictions for specific experiments - Celeste and *Fermi*-LAT respectively. As well as the emission from the central regions the galaxy where the dark matter density is particularly high, this emission is predicted to extend beyond the visible part of the galaxy out into the halo where, though the density is low, the overall mass of dark matter is quite high. The prediction by Falvard et al. (2004) is most applicable to the VERITAS energy range and predicts

integrated fluxes (above 30 GeV) less than $1 \times 10^{-12}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$, extrapolating this up an order of magnitude with an optimistic powerlaw spectrum of index -2 brings these predicted fluxes down by a factor of 100 to less than $1 \times 10^{-14}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$, significantly lower than the expected astrophysical emission from Table 6.2 and the expected sensitivity in a 50 hour observation.

A more recent prediction has been done for the HAWC experiment (Abeysekara et al., 2014). This prediction shows that M 31 is a prime dark matter candidate for HAWC with a J factor (a measure of the astrophysical dark matter source strength) which is comparable with the best of the dwarf spheroidal galaxies. With a comparable energy range to the analysis used in this work, and a comparable J-factor to Segue which VERITAS has observed for over 100 hours, it is unlikely that VERITAS will detect the dark matter emission from M 31 in 50 hours. However, it is planned to determine an upper limit on the dark matter flux from M 31 and add M 31 to the existing VERITAS combined dark matter upper limits programme (Zitzer, 2015a).

Model Usage in this Work

M 31 is often considered a sister galaxy to the MW and thus, to a reasonable approximation, a model derived for one galaxy can be used to model the other. The approach chosen is to build up a model based upon the detailed modelling that is available for the MW and then vary the parameters accordingly to generate a model for M 31. The "standard" model for cosmic and γ-ray emission from the MW is GALPROP.

> *"GALPROP is a numerical code for calculating the propagation of relativistic charged particles and the diffuse emissions produced during their propagation. The GALPROP code incorporates as much realistic astrophysical input as possible together with latest theoretical developments. The code calculates the propagation of cosmic-ray nuclei, antiprotons, electrons and positrons, and computes diffuse γ-rays and synchrotron emission in the same framework."*[*]

Thus, it can be used to generate a predicted γ-ray spectrum for emission from the entire MW using all the available data. It is possible to generate many different models within GALPROP by tuning the input parameters. Doing this could be a thesis within itself, thus the model chosen (54_z04LMS, Figure 6.2) was taken from Strong et al.

[*] http://galprop.stanford.edu/

(2010) since it reproduces the *Fermi*-LAT diffuse spectrum well, in particular at the highest energies. At this point it is important to note that, as measured by the *Fermi*-LAT, at the highest energies the total HE γ-ray emission from all the detected point sources (including extragalactic sources) is approximately equal to the total galactic diffuse emission (from the MW) (Ackermann et al., 2012a).

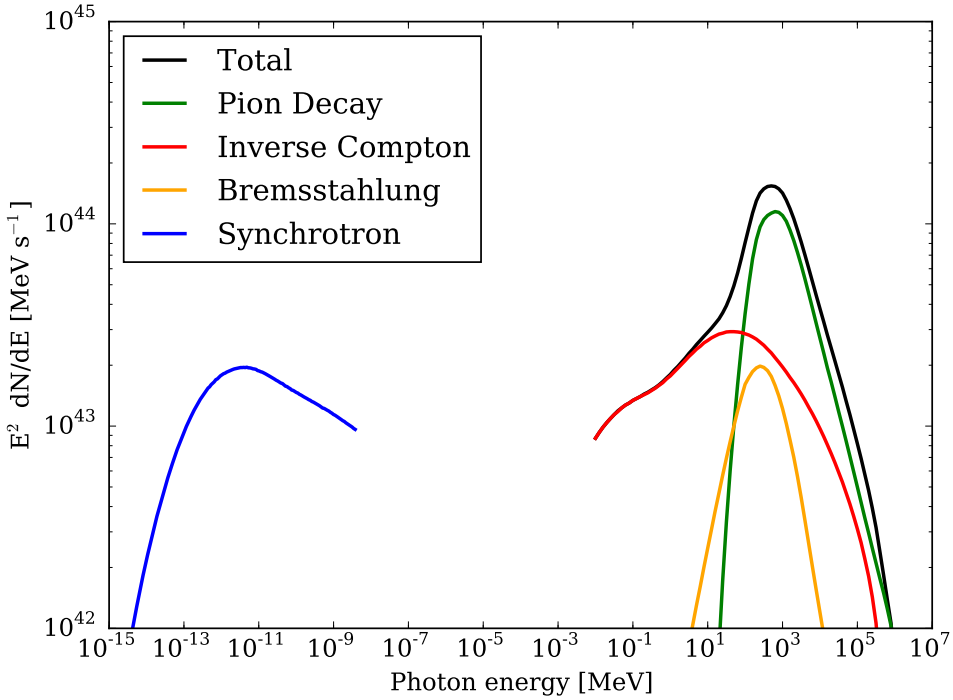

Figure 6.2.: A model of the luminosity of the Milky Way by GALPROP, 54_z04LMS from Strong et al. (2010).

Varying GALPROP to generate a model for M 31 is not a viable option (GALPROP is a large code that includes many measurements of the MW, many of which would not be available for M 31, even if time did permit such work), instead the GALPROP model is used as the basis to generate a fit using the radiative models from the code *naima* (Zabalza, 2015). In contrast to GALPROP, *naima* is a much simpler code that calculates the γ-ray emission given input cosmic-ray spectra and other properties of the galaxy (such as the number of hydrogen atoms, magnetic field strength, and the emission volume). The code is built up from existing models of the radiative processes and solves individual equations for each process. For example it uses the Kafexhiu et al.

(2014) parametrisation of the energy spectra and production rates of γ-rays from p-p interactions to describe the γ-ray emission from pion decay, given an input proton cosmic-ray spectrum and the number of hydrogen atoms within the galaxy. Thus, though the code itself is a new undertaking, and does not yet appear in the literature (work using it was presented by the H.E.S.S collaboration at the 34th ICRC, The Hague, 2015), it is built using existing and well-respected models of the processes involved. For the purposes of this work, where it is being used to generate a model that matches GALPROP model of the MW and then altered to reflect the differences in M 31 it was considered a suitable option, especially considering the paucity of alternative models available.

Milky Way Modelling

A *naima* model was generated to match the GALPROP model (Figure 6.2). The cosmic ray spectra were initially matched to the spectra measured in the solar neighbourhood, then the low-energy end of the electron spectrum was adjusted to match the GALPROP bremsstrahlung (since this is affected by Solar modulation). There is likely to be significant variation from these spectra (in particular the lepton (= electron + positron) spectrum) within the MW, however, since an average spectrum is required this is was deemed suitable, especially as the match to the GALPROP γ-ray spectrum was reasonable. The input parameters are listed in Table 6.3. They were chosen based upon measured properties of the MW and then adjusted to match the GALPROP model (Figure 6.3). There are inevitably assumptions and approximations in doing this, but, since I am using such a simple model to recreate the emission from an entire galaxy, avoiding approximations such as these is impossible. However, the model agrees well with the GALPROP results (and it should be remembered that there is noticeable variation between different GALPROP models of the MW) and thus it will be used as the basis for the rest of this work.

M 31 Modelling

The simplest model simply the cosmic ray spectra using Equation 6.3 and taking the ratio of the SFRs as 0.5 (MW SFR $= 1 - 3$, M 31 SFR $= 0.3 - 1$ (Abdo et al., 2010)). Overplotting the spectral points from Abdo et al. (2010) in Figure 6.4 shows that this simple approximation provides a very good agreement. From here on this will be called the NAIMA MODEL. Abdo et al. (2010) discuss the possible contamination from the hard

Table 6.3.: *naima* Input Parameters

Parameter	MW	M31
Galaxy Properties		
Radius	20 kpc	
Thickness	3 kpc	
Mass hydrogen	$63 \times 10^8 M_\odot$	
Mass HI	$22 \times 10^8 M_\odot$	
B field strength	7 µGauss	
Cosmic Ray Spectra		
Proton Particle Distribution	Broken Power Law	
Amplitude	$1.25 \times 10^{42}\,\mathrm{eV}^{-1}$	$5.17 \times 10^{40}\,\mathrm{eV}^{-1}$
E_0	1 TeV	
α_1	2.7	
α_2	3.0	
Power law break	3.0 PeV	
Electron Particle Distribution	Exponential Cutoff Broken Power Law	
Amplitude	$1.25 \times 10^{44}\,\mathrm{eV}^{-1}$	$5.17 \times 10^{40}\,\mathrm{eV}^{-1}$
E_0	100 GeV	
α_1	1.7	
α_2	3.17	
Power law break	1.15 GeV	
Exponential cutoff	1 TeV	
Radiation Fields		
CMB Temperature	2.72 K	
CMB Energy Density	$0.3\,\mathrm{eV\ cm}^{-3}$	
FIR Temperature	5000 K	
FIR Energy Density	$1.2\,\mathrm{eV\ cm}^{-3}$	
NIR Temperature	70 K	
NIR Energy Density	$0.5\,\mathrm{eV\ cm}^{-3}$	

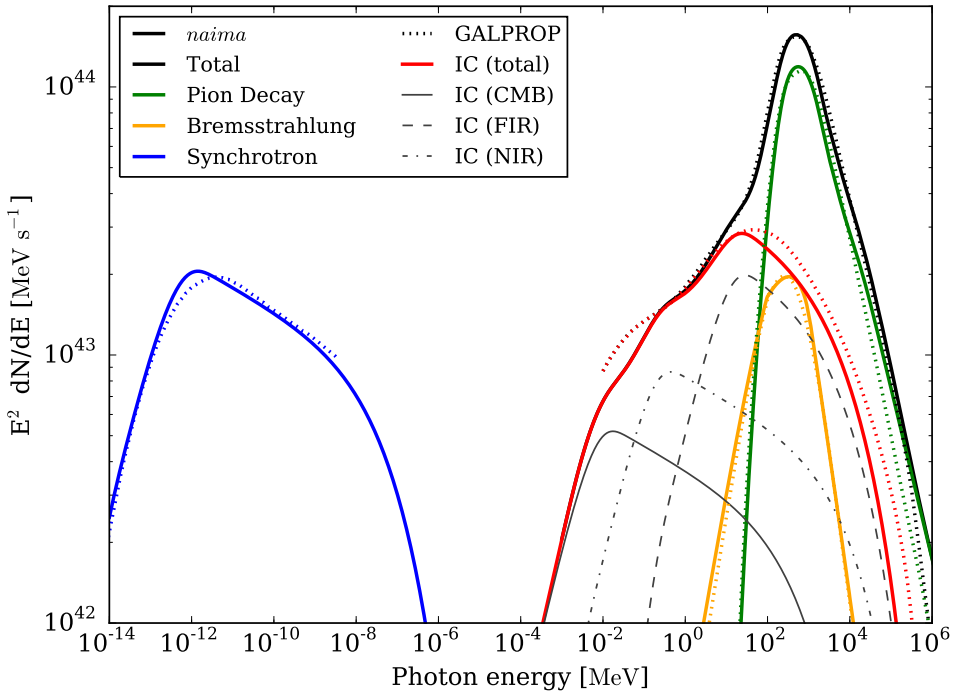

Figure 6.3.: *naima* MW Model (solid lines) compared with the GALPROP model (dotted lines, Figure 6.2). N.b. the IC components are only shown for the *naima* model.

spectrum blazar *1ES 0037+405* in the highest energy point, at that time it was not detected as a source in its own right but there was a cluster of high-energy photons associated with its position. This blazar has since been detected and thus it is likely that this point is an over estimate of the M 31 flux in that energy band. The work in Section 6.2 will provide more information.

Since the majority of the emission is expected to be from pion decay (with the expected long confinement time of protons producing a roughly uniform proton cosmic ray density) then it is expected that the emission will broadly trace the density of the target material. HI maps (Figure 6.5, Braun et al. (2009)) show that the bulk of this material is contained within the star-forming ring $\sim$10 kpc from the galactic center. Overplotting the locations of SNRs identified in M31DEEPXMM (Stiele et al., 2011) shows that the majority are also located within this ring, with the result that both the cosmic-ray sources and the target material are located here. Also, the location of the SNRs is likely to trace the location of other point γ-ray sources, thus they too are likely to lie in the star-forming

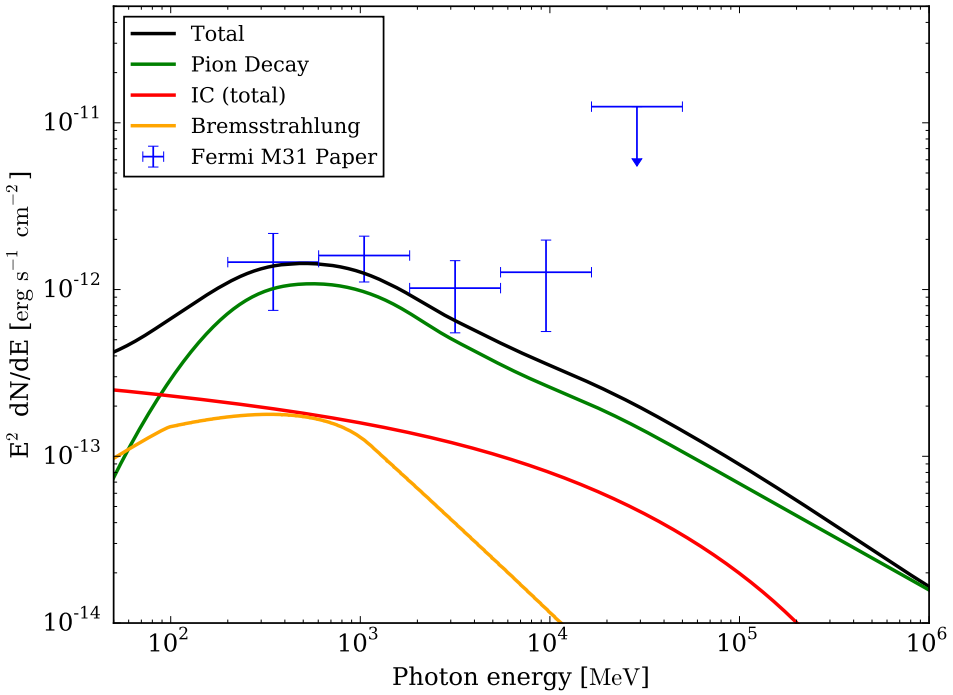

Figure 6.4.: The NAIMA MODEL compared to the data points from Abdo et al. (2010) showing good agreement at low energies, at higher energies there is expected contamination from the hard spectrum blazar *1ES 0037+405*.

ring. Thus, I would expect the bulk of the emission to come from this region.

6.2. Updated and Improved *Fermi*-LAT Analysis

Following its initial detection (Abdo et al., 2010), M 31 has appeared in both the 2FGL (J0042.5+4114, Nolan et al. (2012)) and the 3FGL (J0042.5+4117, Acero et al. (2015)), in both cases as a point source. Detections have also been published in Li et al. (2013) where they examine dark matter limits from M 31 and in Pshirkov et al. (2015), where they claim emission from a spatially extended halo around M 31. Since the initial detection, significant improvements have been made in the data analysis and the amount of observing time has also increased. As well as improving the analysis of M 31 directly this has also resulted in an improvement in the background modelling, for example, in the initial paper there was a suggestion that the hard spectrum blazar *1ES 0037+405*

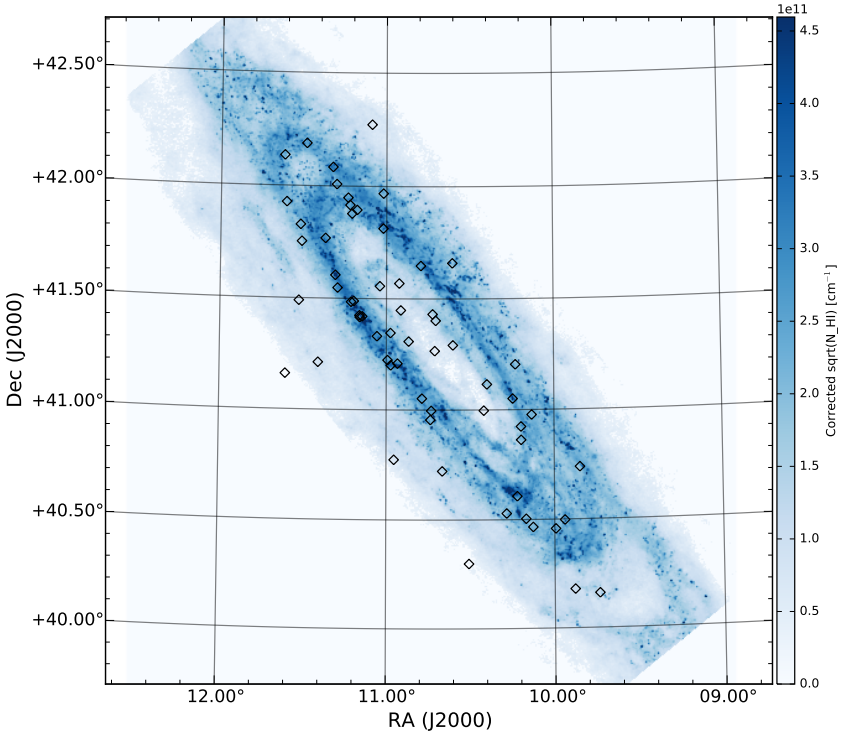

Figure 6.5.: A skymap of the extinction-corrected HI column density in M 31 (Braun et al., 2009). Overplotted are the locations of SNRs identified in M31DEEPXMM (Stiele et al., 2011)

was contributing at the highest energies to the M 31 flux. Since then this source has been identified as a source in its own right and included in the background model (it first appeared in the 2FGL (Nolan et al., 2012)). This improvement in the background model should improve the accuracy of the M 31 measurement.

6.2.1. Data Analysis

As a part of this work I have conducted an updated analysis of the *Fermi*-LAT data, using the Fermi Science Tools version v9r33p0* and the data in Table 6.4 to conduct a

*http://fermi.gsfc.nasa.gov/ssc/data/analysis/software/

Table 6.4.: Data selection conditions for the *Fermi*-LAT analysis

Start Time	2008-08-04 15:43:37
End Time	2015-02-11 06:44:58
Energy Range [GeV]	0.1 to 300
Region Center (RA, Dec (J2000))	10.6847°, 41.2687°
Data Radius	30°
ROI	15°
Data Type	Pass 7
IRF	P7REP_SOURCE_V15
Event Class	2
Zenith Angle	<100°
Quality Cuts	DATA_QUAL==1, LAT_CONFIG==1, ABS(ROCK_ANGLE)<52°
Bin size	0.1°

binned, *Pass 7* analysis. The analysis was conducted using the *LATAnalysisScripts*[*].

At the simplest level *Fermi*-LAT data can be analysed by looking at the raw counts map (Figure 6.6). Looking at this data, there is a suggestion of an excess associated with M 31, but it is impossible to draw any firm conclusions. Rather, a full analysis needs to be conducted.

Due to the number of sources detected by *Fermi*-LAT and its poor PSF, the analysis is rather different from that employed by VERITAS, where sources are sparse and clustering of events is sufficient to identify them. Rather, a model is generated which contains all the information (spectrum, position, spatial extent) for all of the proposed sources and a likelihood that this fits the observed data is then calculated. The optimum model is the one that produces the highest likelihood (hence the name for this type of fitting, maximum likelihood analysis). In order to compare models the test statistic (TS $= -2ln(^{L_{max,0}}/_{L_{max,1}})$, where $L_{max,0}$ is the maximum likelihood value for the null hypothesis and $L_{max,1}$ is for the new hypothesis) is calculated, the higher the TS the more favoured the new model is over the original model. As a basic rule of thumb, the square root of the TS is approximately equal to the detection significance for a given

[*] http://fermi.gsfc.nasa.gov/ssc/data/analysis/user/

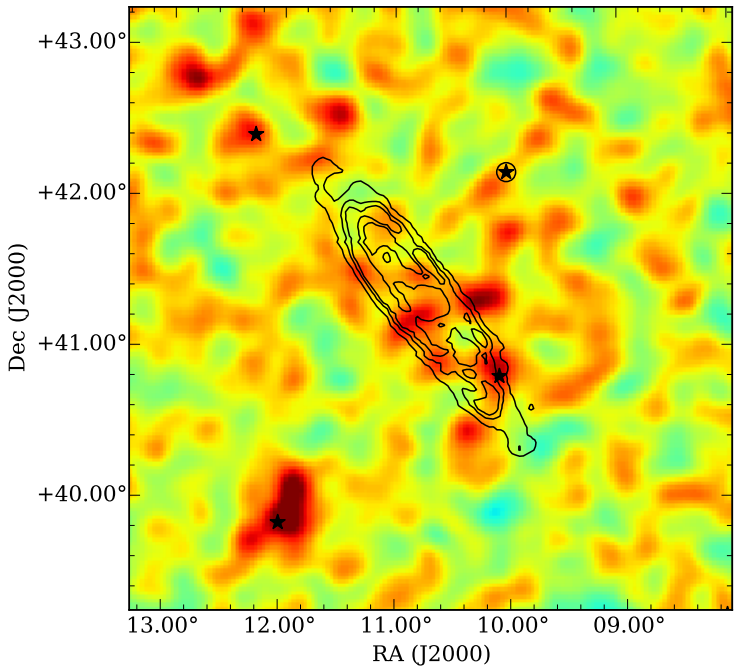

Figure 6.6.: A smoothed (0.1° Gaussian) map of the raw number of counts used in the *Fermi*-LAT analysis. Overlaid are contours of the IRAS 100 µm emission as in Figure 6.7a and the stars indicate the position of point sources from the 3FGL (Acero et al., 2015) and an additional source that was added (star in circle) as described later in this section.

source. For full details of the analysis method see the Cicerone*

Four different models of M 31 were tested in this work

Point A point source as published in the 3FGL.

Raw A template to test for extended emission generated using the same Improved Reprocessing of the IRAS Survey (IRIS) 100 micrometer far infrared map (Figure 6.7a) that was used in the original *Fermi*-LAT detection paper (Abdo et al., 2010).

Ring The *Raw* template but with the central "blob" removed (Figure 6.7b) to test

*http://fermi.gsfc.nasa.gov/ssc/data/analysis/documentation/Cicerone/

whether emission is better fit with a ring structure as seen in the hydrogen maps.

Both Both the *Ring* and the *Point* templates in the same fit, this is to test the relative importances of the two models.

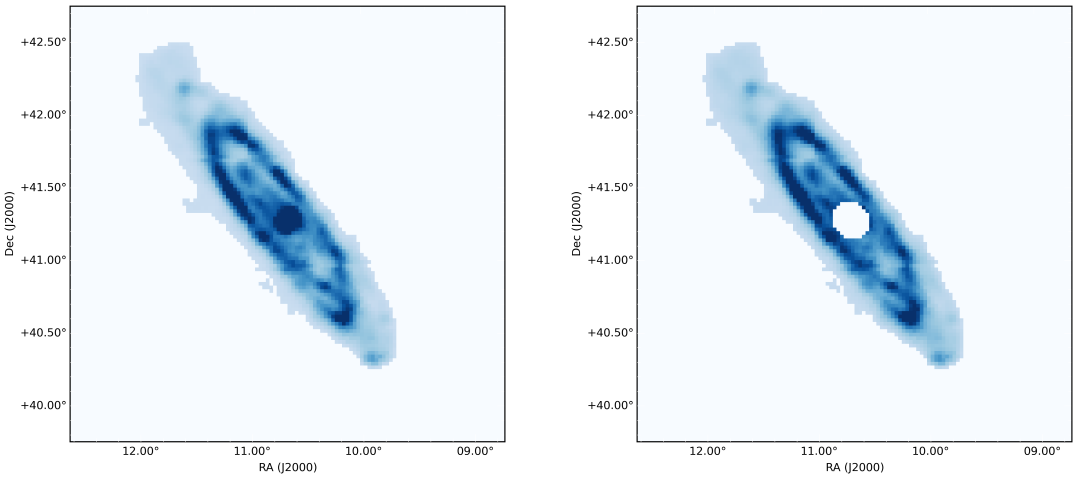

(a) IRIS *Raw* Template **(b)** IRIS *Ring* Template

Figure 6.7.: The extended templates used in the *Fermi*-LAT analysis derived from the IRAS Survey (IRIS) 100 µm far infrared map of M 31 (Miville-Deschenes & Lagache, 2005) using the method outlined in http://fermi.gsfc.nasa.gov/ssc/data/analysis/scitools/extended/extended.html

The background model was generated using make3FGLxml* using the galactic diffuse model *gll_iem_v05_rev1*, the isotropic diffuse model *iso_source_v05* and point sources from the 3FGL catalogue (Acero et al. (2015), all sources outside the ROI (15°) were fixed to the 3FGL values).

After an initial fit using the *Raw* template, a TS map was generated over the central 7.5°×7.5° to test for additional sources (a TS map places a test source at each grid point and conducts a fit with all other parameters in the model fixed, it then calculates the TS between the model with the test source and the base model). This identified an additional point source with a TS 27.83 (approximately 5.3σ) located at the position of

*http://fermi.gsfc.nasa.gov/ssc/data/analysis/user/

the peak in the TS map ((RA, Dec (J2000)) = (10.01°, 42.17°), associated with the excess counts visible in Figure 6.6) which was added to the background model as a power-law point source with the index and amplitude left free to vary.

This background was used to conduct the fit for all four models with a power-law spectrum of the form

$$\text{Flux} = \text{Prefactor} \times \left(\frac{\text{Energy}}{572.78\,\text{MeV}} \right)^{\text{Index}} \tag{6.4}$$

where the normalisation energy (572.78 MeV) is the pivot energy for M 31 in the 3FGL (Table 6.5).

Table 6.5.: Results of the *Fermi*-LAT analysis of the M 31 region.

Model		TS	Prefactor [$\text{cm}^{-2}\ \text{s}^{-1}\ \text{MeV}^{-1}$]	Index
Point		59.28	$(1.90 \pm 0.06) \times 10^{-12}$	-2.50 ± 0.02
Raw		64.34	$(2.51 \pm 0.41) \times 10^{-12}$	-2.23 ± 0.10
Ring		62.98	$(2.51 \pm 0.43) \times 10^{-12}$	-2.21 ± 0.10
Both	(Ring)	19.94	$(1.36 \pm 0.65) \times 10^{-12}$	-2.14 ± 0.17
	(Point)	17.94	$(1.01 \pm 0.18) \times 10^{-12}$	-2.42 ± 0.18

6.2.2. Source Extension

Table 6.5 shows that there is marginal evidence that M 31 is best fit by an extended source, either the *Raw* or the *Ring* model (at the 2.24 and 1.92σ level respectively). To explore the shape of M 31 further a TS map was generated using the result of the *Raw* model fit. Two maps were generated, one with M 31 in the model (Figure 6.8a) and one without (Figure 6.8b) (but with all the other model parameters kept the same). Figure 6.8b clearly shows the need for an extended source at M 31, though it does not agree particularly well with the IRAS 100 µm contours, in particular, showing a displacement away from the center of the galaxy and not showing such an elongated shape as expected. There is no evidence of any further sources within the test region.

One challenge when analysing extended sources is removing the contribution of diffuse emission from within M 31 from the galactic diffuse model, the model of the diffuse emission from the MW that is used. The galactic diffuse model was generated using

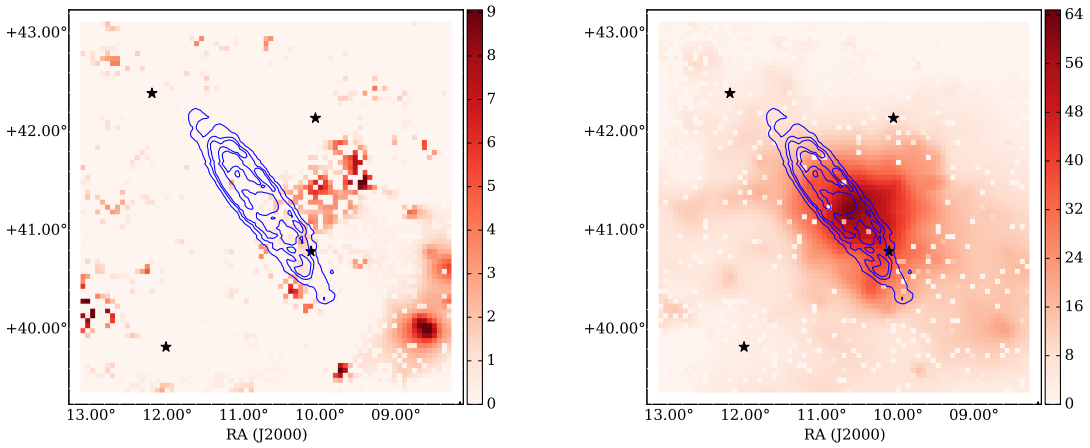

(a) With M 31 *Raw* in the model

(b) Without any M 31 model

Figure 6.8.: TS maps of the M 31 region with and without M 31 included in the model. Overlaid are contours produced from the *Raw* model of M 31 (derived from the IRAS $100\,\mu\mathrm{m}$ observations). M 31 is clearly visible as an extended source though it does not align perfectly with the IRAS contours. The white points are where the TS map failed to converge for that point.

spectral line surveys of HI and CO (as a tracer of H2) to derive the distribution of interstellar gas in galactocentric rings. Also, infrared tracers of dust column density were used to correct column densities in directions where the optical depth of HI was either over- or under-estimated. Due to this method it is possible that, since M 31's emission is expected to be due to interactions between cosmic rays and the ISM and the gas maps used in the generation of the diffuse model have not necessarily removed the contribution of M 31, some of the emission from M 31 is treated as diffuse emission from within the MW and is included in the galactic diffuse model. The inverse-Compton emission model was generated for the central galaxy using *Fermi*-LAT data itself, due to this method the *Fermi*-LAT Collaboration publish a guide to show where they do not recommend conducting extended source analyses (Figure 6.9), M 31 lies outside this region*.

*for more detail on how the galactic diffuse model was generated see http://fermi.gsfc.nasa.gov/ssc/data/access/lat/Model_details/FSSC_model_diffus_reprocessed_v12.pdf

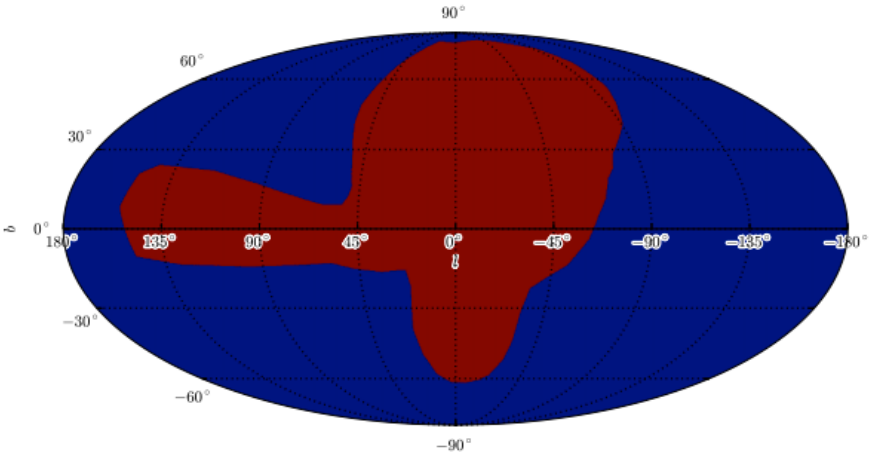

Figure 6.9.: *Fermi*-LAT Extended Analysis Exclusion Region. The red region shows the central galaxy where *Fermi*-LAT data has been used for calculating the inverse Compton component of the background and thus it is not recommended to conduct extended ($>2°$) analyses, M 31 lies outside this region at $(l, b) = (121.17°, -21.57°)$

In the original *Fermi*-LAT paper they undertook significant work to remove contamination of the diffuse model by M 31, without access to all of the constituents for the building up of the diffuse model it is not possible to do this. To determine if this has had an impact a few checks are conducted, the first was to see whether there is any significant disagreement between the spectrum measured in this analysis with the published analysis. With improved analysis tools and longer observing time I would expect to see a reduction in the errors on the flux rather than a change in the flux, a change in the flux would suggest that there is either an issue in the background modelling or elsewhere in the analysis. This will be discussed in Subsection 6.2.3. Comparing both the hydrogen maps and the IRIS maps with the diffuse emission maps at different energies shows no correlation, this lack of correlation suggests that any effect is small. Since I cannot do any more than this with the tools available, and from the comparison between the maps the effect looks small, I will have to leave this as a caveat on this work.

6.2.3. Spectrum

The *Raw* model was chosen as the model for detailed spectral analysis. The binning chosen was the same as that used in the original detection paper with an additional bin

covering 100 MeV to 200 MeV and with the highest-energy bin extended to cover the energy range used in this analysis but not in the original (from 50 GeV to 300 GeV). The spectral points for M 31 were calculated by breaking the analysis into different energy bins and running a complete analysis in that energy range. Both the original and the updated spectral points are plotted in Figure 6.10, the good agreement between the points at all but the highest energy (where a contribution from *1ES 0037+405* is expected in the points from the original paper) strongly suggests that the concerns over contamination from the galactic diffuse model into the extended source are not important.

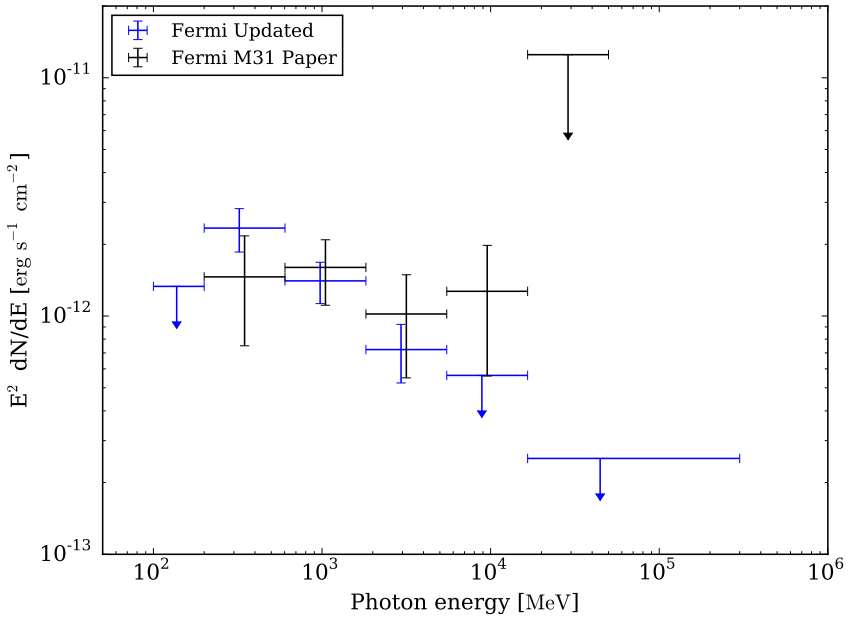

Figure 6.10.: *Fermi*-LAT M 31 spectra from the original paper and the updated analysis.

The resulting differential spectral fit is

$$\text{Flux} = (2.51 \pm 0.33_{\text{stat}}) \times 10^{-12} \times \left(\frac{\text{Energy}}{572.78\,\text{MeV}}\right)^{-(2.22\pm0.10_{\text{stat}})} \text{cm}^{-2}\ \text{s}^{-1}\ \text{MeV}^{-1} \quad (6.5)$$

with a 0.1 GeV to 300 GeV flux of $(1.0 \pm 0.13_{\text{stat}}) \times 10^{-8}\text{cm}^{-2}\ \text{s}^{-1}$. This is comparable to the spectrum published by Abdo et al. (2010):

$$\text{Flux} = (3.10 \pm 0.56_{\text{stat}}) \times 10^{-12} \times \left(\frac{\text{Energy}}{572.78\,\text{MeV}}\right)^{-(2.1\pm0.2_{\text{stat}})} \text{cm}^{-2}\ \text{s}^{-1}\ \text{MeV}^{-1} \quad (6.6)$$

with a 0.1 GeV to 100 GeV flux of $(1.10 \pm 0.47_{\text{stat}} \pm 0.20_{\text{sys}}) \times 10^{-8}$ cm^{-2} s^{-1}.

The differences in these results (for example the softening of the spectral index) can be attributed to either the increased statistics or the improvement of the background modelling. The principal change in the background model is the inclusion of the hard-spectrum blazar *1ES 0037+405* which has an effect at the highest energies, its inclusion explains the fact that the penultimate energy bin turned into an upper limit since in the earlier work the contribution from *1ES 0037+405* would have been included in M 31. Since the NAIMA MODEL clearly shows curvature at low energies a power law is not a suitable fit for the entire energy range, rather, a fit is conducted over the range above this curvature, 700 MeV. A complete reanalysis of the data above 700 MeV gives the spectrum

$$\text{Flux} = 2.11 \pm 0.40_{\text{stat}} \times 10^{-14} \times \left(\frac{\text{Energy}}{4384\,\text{MeV}}\right)^{-(2.44\pm0.23_{\text{stat}})} \text{cm}^{-2}\ \text{s}^{-1}\ \text{MeV}^{-1} \quad (6.7)$$

where 4384 MeV is the decorrelation energy of the fit. The 0.7 GeV to 300 GeV flux is $(9.0 \pm 1.7) \times 10^{-10}$cm^{-2} s^{-1}. This spectrum shows, as expected from looking at the NAIMA MODEL, a softening of the spectral index.

Overlaying these data points on the NAIMA MODEL generated earlier (Figure 6.11a) shows a slight tension with the second-lowest energy bin. The data points show a sharper drop-off than the NAIMA MODEL, this could suggest that the inverse Compton component is too high, this is discussed further in Subsection 6.2.5. Overlaying the butterflies (Figure 6.11b) for the two fits (>100 MeV and >700 MeV) shows that the constraint of a power-law fit probably artificially reduces the amplitude of the >100 MeV fit, whereas the >700 MeV fit, which is fitting to a region of the spectrum that more closely resembles a power-law, is a much better fit and has a better agreement between its index and that of the model. With the higher-energy cut, the softer index agrees with the NAIMA MODEL much more closely, suggesting that the model is at least a reasonable representation of the spectral shape (though it should be noted that the errors at the highest energies are very large).

6.2.4. VERITAS Prediction

These results can be extrapolated to VHE energies, taking a reference energy of 1 TeV, Table 6.6.

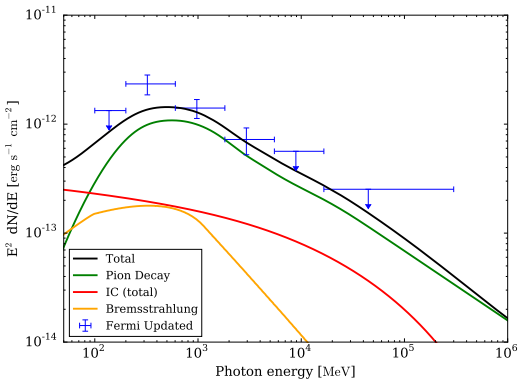

(a) The spectral points from the *Fermi*-LAT observations of M 31 compared with the NAIMA MODEL.

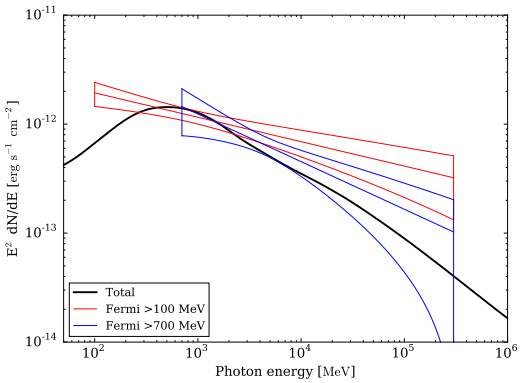

(b) The butterflies from the *Fermi*-LAT observations of M 31 compared with the NAIMA MODEL.

Figure 6.11.: Update *Fermi*-LAT analysis of M 31 using the *Raw* template.

Table 6.6.: Prediction of the M 31 differential flux at 1 TeV and integral flux from 300 GeV to 30 TeV by extrapolating the *Fermi*-LAT spectra. The Crab spectrum is the power-law from Aharonian et al. (2006).

Fermi-LAT	Differential Flux		Integral Flux	
Low Energy Cutoff [MeV]	[cm^{-2} s^{-1} TeV^{-1}]	[% Crab]	[cm^{-2} s^{-1}]	[% Crab]
100	1.59×10^{-13}	0.461	5.65×10^{-13}	0.376
700	3.69×10^{-14}	0.107	1.46×10^{-13}	0.0969

6.2.5. Updated *naima* M 31 Model to Match *Fermi*-LAT Observations

With the improved *Fermi*-LAT observations there is no major reason to edit the NAIMA MODEL, though there is a tension at lower energies where the *Fermi*-LAT data points show a higher peak in the 100-200 MeV range. To fit this would require an increase in the pion decay component of the γ-ray emission and a reduction in the IC component. The reduction in the IC component is justified since M 31 has a lower infrared density than the MW (Walterbos & Schwering, 1987) which will reduce this component. The addition of a larger bremsstrahlung component would also improve the low energy shape, as would changes to the spectral indices and relative normalisations of the the cosmic ray spectra. In short there are many possible changes that could be made to tweak the

model to match the observed spectrum, many of which have scientific justification. With this large degeneracy and limited information to discriminate between them, for now it it chosen to keep the existing model by with a note that, the discrepancy at the lower energy could have a number of different sources and with further data and improved analysis is required to constrain this.

6.2.6. *naima* VERITAS Predictions

Using the NAIMA MODEL it is possible to make predictions of the VHE γ-ray flux. The predicted differential flux at 302 GeV is $2.73 \times 10^{-13}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}\mathrm{TeV}^{-1}$. The integral (0.3-30 TeV) flux is $4.76 \times 10^{-14}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$ or 0.03% of the Crab Nebula flux.

6.2.7. *Fermi*-LAT Conclusions

The *Fermi*-LAT analysis has shown that there is a source consistent with the position of M 31, there is evidence at the 2.2σ level that the source is extended with no preference between the two extended source models tested. Extrapolating the *Fermi*-LAT spectra to VHE energies gives a predicted VHE flux of $3.69 \times 10^{-14}\,\mathrm{cm}^{-2}\ \mathrm{s}^{-1}\ \mathrm{TeV}^{-1}$ using the spectrum produced from the dataset with a higher energy threshold, this is the preferred spectrum as I expect that there will be a change in the spectral parameters at lower-energies due to the dominant emission mechanism, pion decay. The lower energy upper limit is highly suggestive of pion emission, further work at lower energies to explore this is desirable, in particular following the release of the upgraded data processing *Pass 8* by the *Fermi*-LAT Collaboration which will have significant improvements in the PSF at lower energies as well as an increase in the effective area at higher energies.

6.3. VERITAS Analysis

6.3.1. Analysis Methodology

M 31 is a spatially-extended, optically-bright source, thus it presents significant challenges for analysis since the standard techniques are optimised to analyse optically dark point sources.

6.3.2. Spatial Extension Techniques

The simplest way to cope with a spatially-extended source is to increase the *on* region radius. By increasing the size of this region you increase the amount of signal that is included in the region but you also increase the amount of background. This background is estimated using *off* regions (using the techniques outlined in Subsection 3.4.7), by increasing the size of the *on* region, you place a penalty on α, the ratio of the acceptance of the *on* region to the acceptance of the *off* region. This reduction in α reduces the statistical significance of any excess detected (since the error on the background counts increases). Increasing the value of *thetaSq* increases the sensitivity for extended sources but it also increases the correlation between bins in the skymap and increases radius of the exclusion regions around bright stars and (potential) sources, reducing the background. If the value of *thetaSq* is too high then the significance distribution (a histogram of the significances in the skymap bins outside of the exclusion regions) is not as well fit by a Gaussian function of mean = 0, $\sigma = 1$ as it should be for random background. For this work I used a value of $thetaSq = 0.04$ based on tests on independent data as this was the largest value that could be used before the significance distribution diverged from a Gaussian function of mean = 0, $\sigma = 1$.

In addition, the extended nature of M 31 requires that a large fraction of the FoV has to be excluded from use in the calculation of the background. This reduces α and thus further reduces the sensitivity of the observations. Techniques exist which are better at characterising the background in the presence of significantly extended sources such as M 31:

On-Off

This is a useful method of extended sources that cover a significant fraction of the FoV, but would require double the observing time to provide the *off* dataset. In addition, due to the optical brightness of M 31 there would be significant differences between the *on* and *off* runs, making this technique much less reliable.

Template Background

This method uses the events that were rejected by the mean scaled width (MSW) cut to characterise the background. Though this method would have many benefits it would have required significant work to implement within VEGAS and thus was not employed.

3D Maximum Likelihood Method

This technique is currently in development and (similar to the technique used in the analysis of *Fermi*-LAT data) does a maximum likelihood calculation to determine whether there is a source present or not (Cardenzana & Collaboration, 2015). This technique has the potential to significantly increase the sensitivity of IACTs to extended sources, however, at the time of this analysis it was still undergoing testing and thus was not employed on this data.

In lieu of these limitations it was decided to employ two different *thetaSq* cuts, *Small* (*thetaSq* = 0.01) and *Large* (*thetaSq* = 0.04). The *Small* analysis will provide some sensitivity for extended sources, the *Large* analysis will provide increased sensitivity but at the cost of a greater correlation in the skymap (see Appendix B for an example of a comparison).

6.3.3. Test Region Determination

For each of the *thetaSq* cuts, sets of test regions were predefined using the *Raw* model as this was found to be the best fit in the *Fermi*-LAT analysis. These positions were chosen to reduce the number of trials given the constraint of having to exclude the optically bright central region (discussed in Section 6.5) and not wanting any overlap between the test regions. These positions are shown in Figure 6.12a and Figure 6.12b.

To determine the background counts and acceptance an elliptical ring surrounding M 31 was used (Figure 6.13).

6.3.4. Source Flux Calculation Routine

To determine the total flux emitted by M 31 the data from the test regions needs to be combined. The method to do so is outlined here, the same method was also used to calculate the scaling relationship for the Crab raster scan (Appendix B) where the scaling relationship is applied to the flux normalisation to determine the flux from the whole scan region. This method has been devised to use the VEGAS output files and thus uses much of the VEGAS terminology.

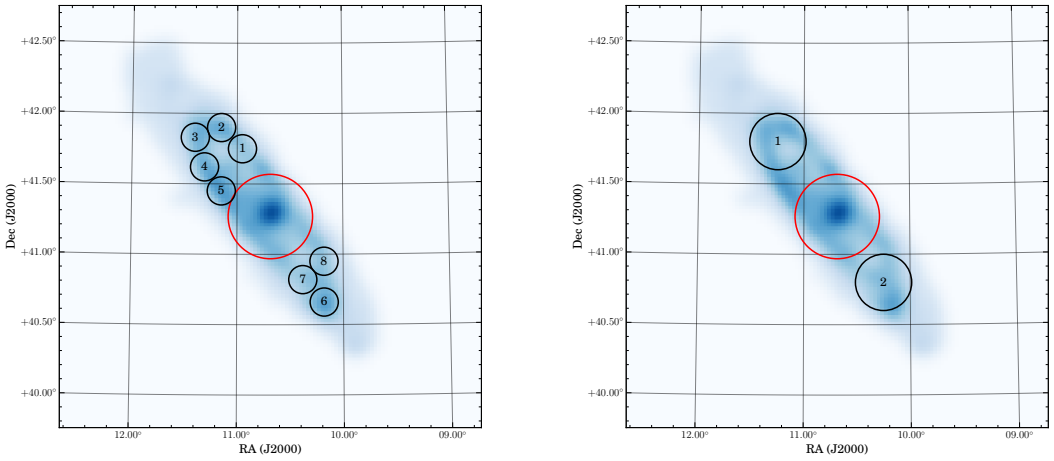

(a) *Small* Test Regions **(b)** *Large* Test Regions

Figure 6.12.: The test positions within M 31 shown against the IRAS contours smoothed with the VERITAS PSF (0.1°)

Excess Counts

To calculate the excess counts requires knowledge of the N_{on}, N_{off} and α, the ratio of the acceptance from the *on* to the *off* regions. The *on* and *off* regions depend upon the source in question, for M 31 they are explained in Subsection 6.3.3. N_{on}, N_{off} are then just the sum of the counts within those regions for the complete observation. The way VEGAS works, to correctly calculate α, each combination of array configuration (pre/post telescope 1 move and post upgrade) and season (winter/summer atmosphere) must be treated independently (these combinations are referred to as *groups* in VEGAS parlance). The weighting for each group is determined using the number of events *in total*, across the whole FoV that are detected during the data collected in that group as a fraction of the total number of events recorded across the whole FoV for the entire dataset. The weighted acceptance is then summed for each combination, for the *on* and *off* regions (including the correction for bin area since the skymaps are tangential projections). α is then the ratio of the total *on* acceptance to the total *off* acceptance. From this, the number of excess events ($N_{on} - \alpha \times N_{off}$), the significance of the detection (using Equation 3.2) or on upper limit on the number of counts can be calculated.

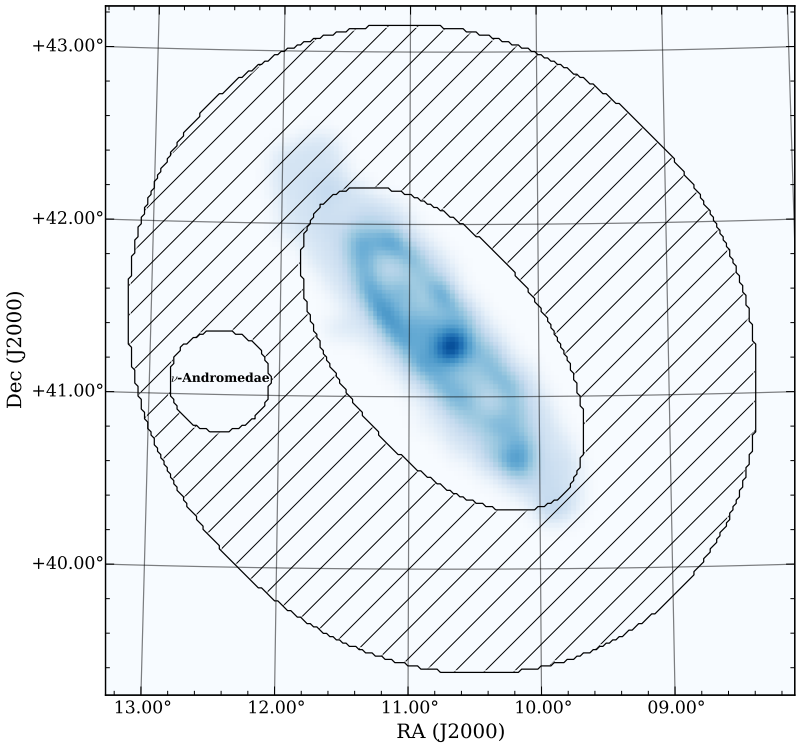

Figure 6.13.: The background region (hatched) used in the VERITAS analysis. ν-Andromadae is excluded from the background analysis.

6.4. Effective Areas

VEGAS (and VERITAS in general) only has effective areas generated from a point source. When the effective area is generated some of the γ-rays which originate from the source are reconstructed outside of the test region. These lost showers correct for the fact that, provided the simulations have the same angular resolution as the observations, the same fraction of the observed flux is reconstructed outside the test region and thus not counted. However, for a slightly-extended source that is fully contained within the point source test region, the fraction of the flux that is reconstructed outside of the test region is not the same as for a point source, since some of the emission is closer to the edge of the test region, the fraction lost varies. This is overcome by extending the test

region size (from a radius of e.g. 0.1° to 0.2°) which reduces this effect to a level within the systematic errors.

For a significantly-extended source that extends beyond the point-source test region other effects come into play, as well as γ-rays originating from inside the test region being reconstructed outside the test region, γ-rays from outside the test region are also reconstructed inside the test region. In this case increasing the test region radius does not work (unless the entire source can be contained within a single test region with sufficient boundary that no source leaks out, this is normally neither practical nor desirable as it will reduce the significance of the detection by reducing the signal-to-noise) since there will always be some flux leakage. To overcome this, an estimate of the flux that is "lost" for both the point source used in the effective area generation and the extended source that is being measured needs to be determined. This is done by convolving a point source and a *model* of the expected emission from the extended source with the angular resolution of the array*. For each of these I then calculate the fraction of counts within the test region before and after the convolution with the angular resolution - the *point/extended source measured fraction* ($MF_{P/E}$).

$$MF_{P/E} = \frac{\text{Flux measured with test region}}{\text{Flux originating within test region}} \tag{6.8}$$

The correction factor (CF) is then the ratio of these measured fractions

$$CF = \frac{MF_E}{MF_P} \tag{6.9}$$

The effective area is then "corrected" using this factor, this is done for each *on* region (to first order this does not need to be done for each *group*, though there have been PSF improvements over time, provided the average acceptance for each test region in each group is roughly the same this will not be a major issue).

*For this work a constant angular resolution (68% containment radius) of 0.1° with no variation with energy or shower size was chosen based on the results of Park (2015) and my own work looking at the Crab Nebula using the analyses used in this work, for analyses with a lower energy threshold (and thus an angular resolution that varies with energy/shower size) an energy dependent angular resolution will be required

6.5. Estimating the Total Flux

We are interested in calculating the differential flux from M 31. To do this we assume a spectrum for the emission from the whole of M 31 (assuming, for simplicity, a power law spectrum)

$$S_{M31} = N_{M31} E^{-\alpha} \tag{6.10}$$

where S_{M31} is the M 31 spectrum and N_{M31} is the normalisation constant of the spectrum. From this we are measuring the spectra from different test regions (r), which emit a fraction of the flux from the entire source,F_r. That is

$$\begin{aligned} S_r =& N_r E^{-\alpha} & (6.11)\\ =& F_r N_{M31} E^{-\alpha} & (6.12) \end{aligned}$$

Given a test input spectrum (T) we want to estimate the counts from all the source regions combined, $C_{R,T}$,

$$C_{R,T} = \sum_r C_{r,T} \tag{6.13}$$

where $C_{r,T}$ are the counts from a single test region.

The counts from an individual region depend upon $S_{r,T}$, the effective area from that region, EA_r, and the live time, τ.

$$\begin{aligned} C_{R,T} =& \sum_r \int \tau S_{r,T} E A_r \, \mathrm{d}E & (6.14)\\ =& \sum_r \int \tau F_r N_{M31,T} E^{-\alpha} E A_r \, \mathrm{d}E & (6.15)\\ =& N_{M31,T} \sum_r \int \tau F_r E^{-\alpha} E A_r \, \mathrm{d}E & (6.16) \end{aligned}$$

Now, for a given number of observed counts, $C_{R,Obs}$, we can determine the observed normalisation constant, $N_{M31,Obs}$, since

$$\begin{aligned} \frac{N_{M31,Obs}}{N_{M31,T}} =& \frac{C_{R,Obs}}{C_{R,T}} & (6.17)\\ N_{M31,Obs} =& \frac{C_{R,Obs} N_{M31,T}}{C_{R,T}} & (6.18) \end{aligned}$$

Thus, substituting Equation 6.16 into Equation 6.17, we can determine the normalisation constant, $N_{M31,Obs}$,

$$N_{M31,Obs} = \frac{C_{R,Obs}}{\sum_r \int \tau F_r E^{-\alpha} E A_r \, \mathrm{d}E} \tag{6.19}$$

and thus we have determined the differential flux from M 31.

Optical Brightness

The optical brightness of M 31 presents a challenge to the analysis of the data. It is one of the few VERITAS targets where the brightness is sufficient that the PMT anode currents can be used to conduct (very crude) optical astronomy (Figure 6.14). Section 4.2 addresses the exploration of methods to overcome regions of high and varying NSB in detail. Following on from the conclusions drawn in that work it was decided to use two cuts and cleaning combinations: medium cuts with *Threshold* cleaning (Thresh) and hard cuts with *EventDisplay* (ED) cleaning. For the *Threshold* cleaning, the choice of threshold was based on the desire to maintain as much of the FoV as possible unaffected by the optical brightness but without raising the threshold so high that the image cleaning is removing too many of the showers. Also, since the same thresholds will be used for all the runs with a particular season and atmosphere, the choice had to work for them all, even though there is significant variation between some of the runs (largely depending upon whether the observations were taken in true dark time or whether they were taken with a partially-illuminated Moon above the horizon, Figure 6.15). The standard pedvar for dark nights is 10 dc which, when the standard cleaning thresholds are applied (5 for picture, 2.5 for boundary) give cleaning thresholds of 50 dc for the picture and 25 dc for the boundary pixels. For the 2014 season, most of the data was taken with a partially illuminated Moon (and raised CFD thesholds Subsection 4.3.1) which raises the pedvar. To account for this, the pedvar was raised to 12 dc giving cleaning thresholds of 60, 30 dc.

Even with these two techniques in use there is still a significant deficit in the center of M 31 due to the optical brightness. This is visible by looking at plots of the events that are rejected by the gamma/hadron selection cut as shown in Figure 6.16 (note: these events are not nearly as well reconstructed as the events that pass the selections cuts so the PSF is much larger, estimated to be about 0.4°, thus the optical impact is spread over a much larger area). However, since the majority of the emission is expected to

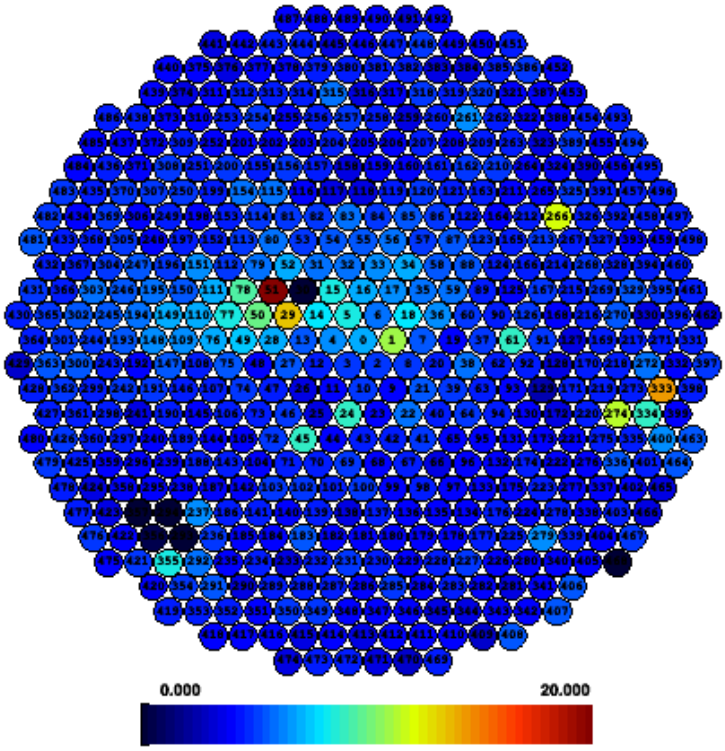

Figure 6.14.: A plot of the median anode currents from one of the 2010 observations of *1ES 0037+405*, M 31 is clearly visible in the currents, as is the bright star ν-Andromedea.

come from the star forming ring this is less of an issue and thus it is possible to exclude this region from the analysis.

6.5.1. Observations

The VERITAS data (Table 6.7) was combined from a number of different, nearby, observation targets, namely M 32, *1ES 0037+405* and M 31 itself. It has been taken over a number of years, with all three array configurations (prior to the move of telescope 1, after the move of telescope 1 but before the camera upgrade and after the upgrade to the cameras). As a result the data requires careful treatment when analysed, in particular when all the data is combined.

The data used in the M 31 analysis (Table 6.7) was subjected to a detailed data-quality assessment prior to use, in particular for periods of time affected by clouds

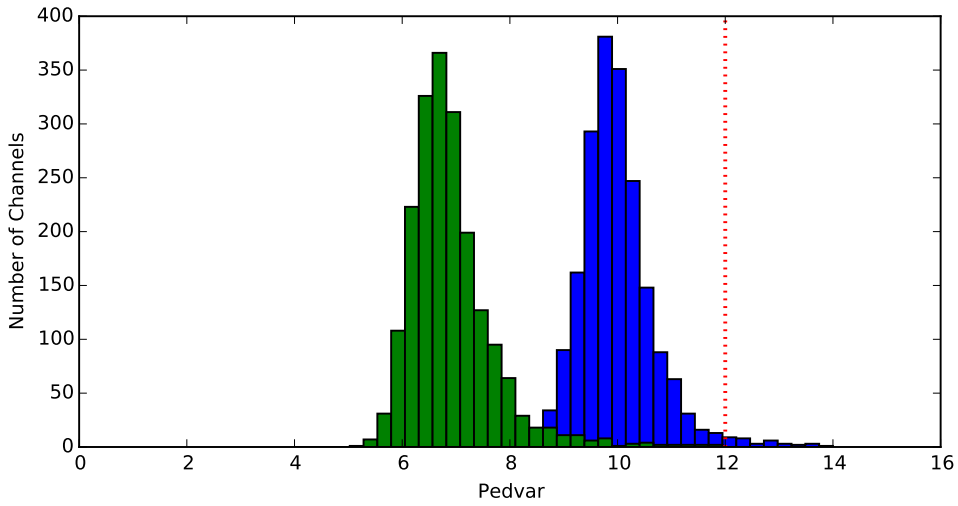

Figure 6.15.: Pedvars from two runs taken during the 2014 observations showing the variation between runs. This season had runs both with (blue) and without (green) a partially illuminated Moon. The threshold (red line, $12\,\mathrm{d\,c}$) was chosen so that, for the majority of the data, only a small number of channels had to be excluded (though this meant that for the subset of data without the affect of the Moon the thresholds were marginally too high, raising the energy threshold to a higher level than could be achieved with a lower threshold and a dedicated analysis for those runs).

.

(Subsection 3.4.1). Any time periods where the effects of clouds were visible in the L3 rate or in the FIR were removed, as were any periods when the hardware issues were affecting the performance of the telescopes. The total observing time is quoted in terms of live time, this also takes into account the dead time of the array during data taking and subtracts that time accordingly.

6.5.2. Results

No evidence for emission from M 31 was detected, either visibly in the skymaps (Figure 6.17) or by looking at the significance distributions (a histogram of the significance of each bin in the skymap where non-statistical fluctuations would result in an excess of bins with a statistical significance >0, an excess in bins with a significance >3 would

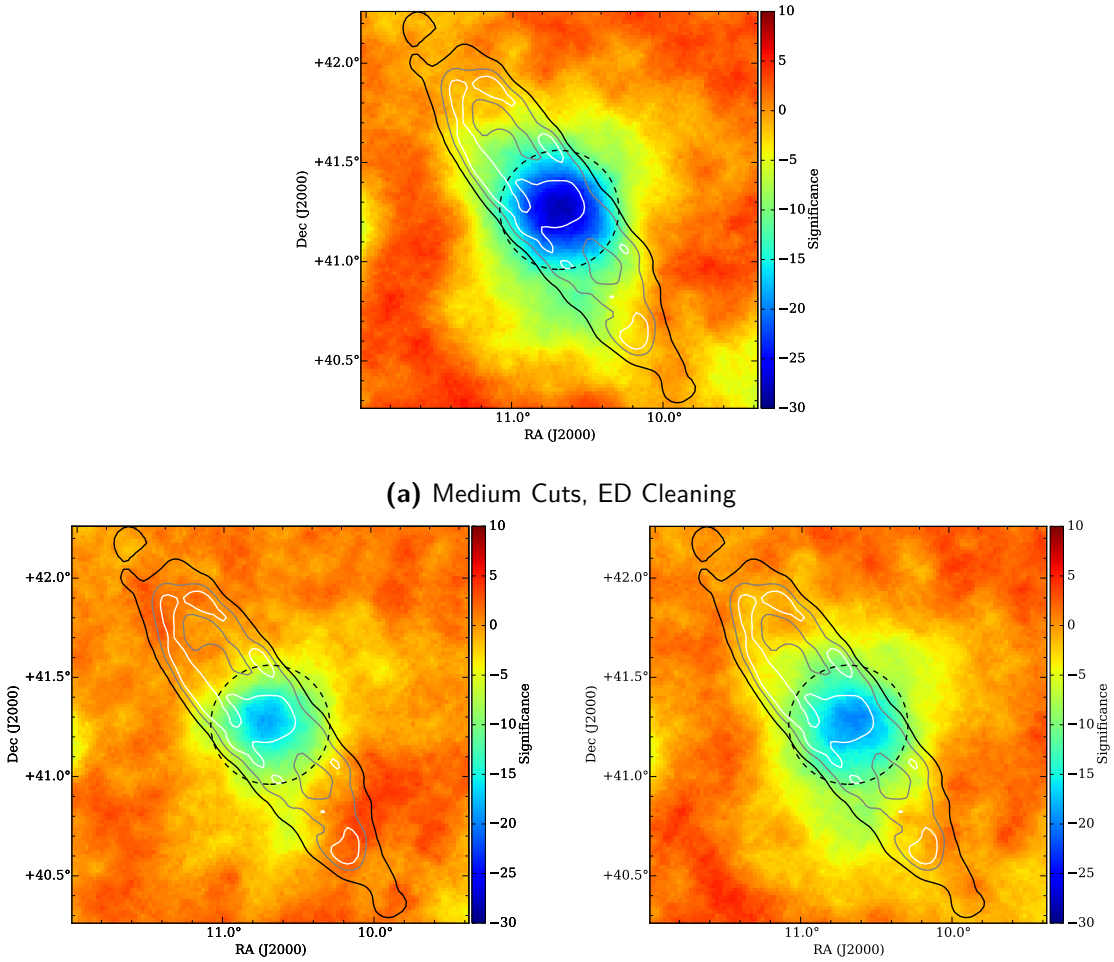

(a) Medium Cuts, ED Cleaning

(b) Medium Cuts, Threshold Cleaning

(c) Hard Cuts, ED Cleaning

Figure 6.16.: Skymaps generated using events that fail the gamma/hadron selections cuts. These were used to test the impact of the image cleaning and cuts for reducing the impact of the optical brightness without imposing an additional trials penalty on the analysis by tuning the cuts/cleaning using the analysed data. Threshold cleaning has clearly reduced the size of the hole with a smaller reduction by hard cuts and ED cleaning. With gamma/hadron cuts applied the angular resolution improves by a factor of $\sim$4, reducing the impact to within the central excluded region (dashed circle).

Table 6.7.: The VERITAS data used in the analysis of M 31

Target	Year(s)	Array Type	Total Live Time [hours]
M 31	2014	Post Camera Upgrade	18.43
M 31	2011	Post T1 Move	14.57
1ES 0037+405	2009-10	Post T1 Move	11.47
M 32	2008-9	Pre T1 Move	10.22
Total	—	—	54.69

be particularly suggestive, an excess >5 is a detection, Figure 6.18)). The *Small* significance distributions are compatible with a Gaussian distribution of mean 0 and standard deviation 1 as would be expected from a dataset where there is no signal. The medium cuts, Thresh cleaning, *Large* significance distribution has a standard deviation that is larger than 1 (1.192 ± 0.007), this is not uncommon with analyses that use extended *thetaSq* values due to the reduction in α and the greater correlation between bins, increasing the background ring size improved this but has not completely eradicated the problem. Looking at the sky map a couple of large patches are visible which are responsible for this fluctuation but they do not align with either M 31 or any of the *Fermi*-LAT detected sources (black stars).

From these results I can conclude that there are no individual sources or source regions that are detected in the FoV. Instead, I can combine the test regions to generate an estimate of the significance of the combined test regions, for the four analyses the combined excesses/significances are given in Table 6.8. Apart from the medium cuts, *Large* result (where one of the test regions lies squarely on an excess which is consistent with a statistical fluctuation) these results are all consistent with no signal being present.

Table 6.8.: VERITAS results from the combined test regions.

Test Region Size / Cuts & Cleaning	*Small*		*Large*	
	Excess	Significance	Excess	Significance
hard & ED	-35	-0.73	42	0.89
medium & Thresh	70	0.79	206	2.34

The statistical upper limits were calculated using the bounded method of Rolke (Subsection 3.4.13) at the 95% level and are an upper limit on the excess counts from the

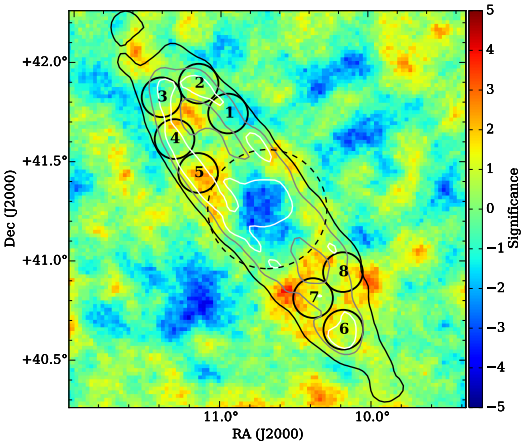

(a) Medium Cuts, Threshold Cleaning, *Small*

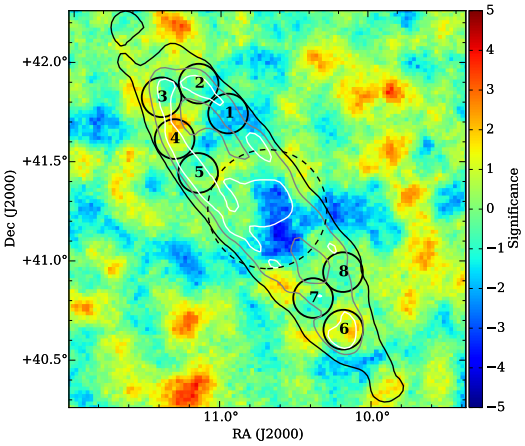

(b) Hard Cuts, ED Cleaning, *Small*

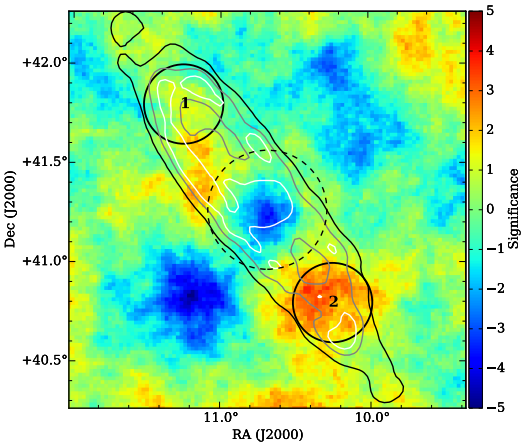

(c) Medium Cuts, Threshold Cleaning, *Large*

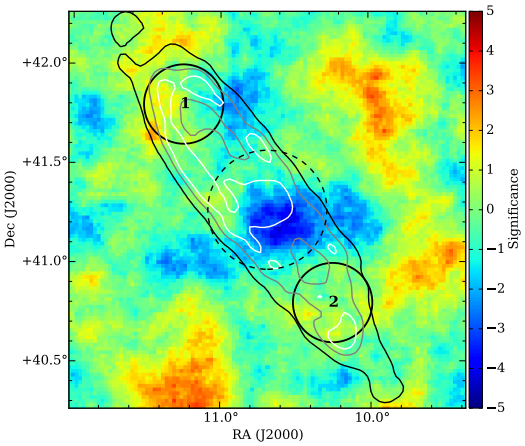

(d) Hard Cuts, ED Cleaning, *Large*

Figure 6.17.: Skymaps of the complete VERITAS data set on M 31

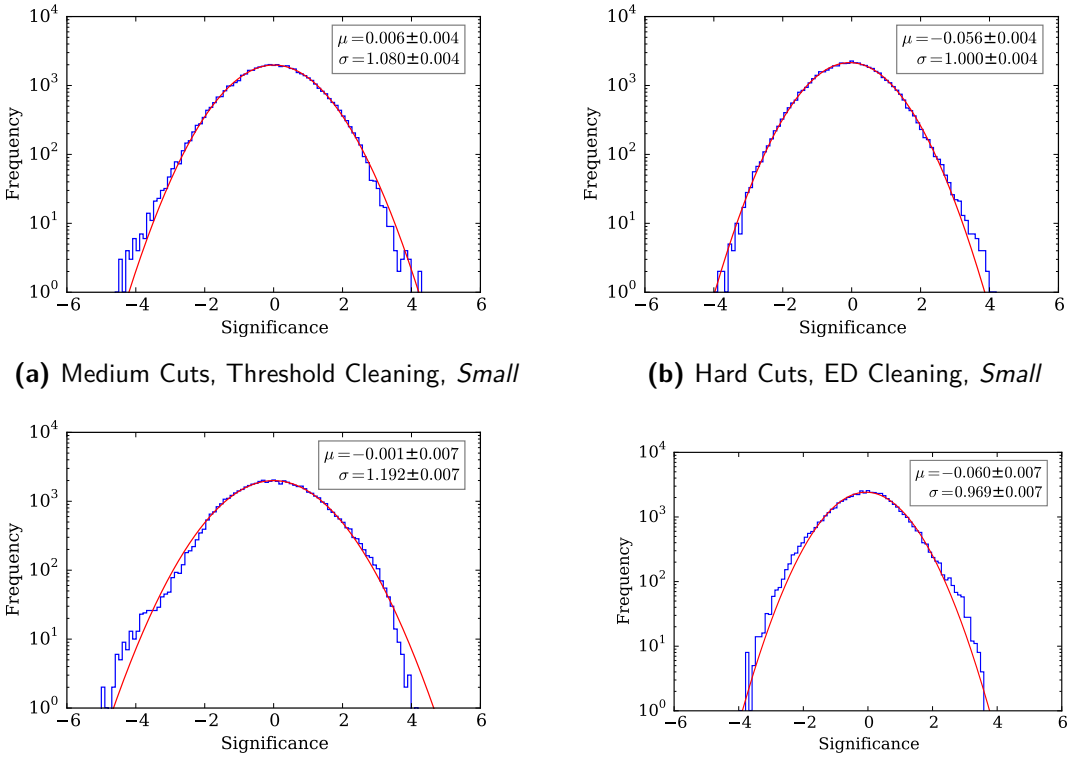

(a) Medium Cuts, Threshold Cleaning, *Small*

(b) Hard Cuts, ED Cleaning, *Small*

(c) Medium Cuts, Threshold Cleaning, *Large*

(d) Hard Cuts, ED Cleaning, *Large*

Figure 6.18.: Significance Distributions of the complete VERITAS data set on M 31 less exclusion regions around ζ-Tauri and the center of M 31. The red lines are the best fit Gaussian distributions which should have a mean of 0 and a σ of 1 if the background is correctly modelled.

test position. The differential flux upper limit was calculated using the assumption of a power-law spectrum with an index of -2.5. These upper limits have been scaled from those derived from the test regions to give an upper limit on the total flux from M 31. The comparison with the Crab Nebula was calculated using the H.E.S.S. power-law spectrum (Aharonian et al., 2006)

$$\frac{\mathrm{d}N}{\mathrm{d}E} = (3.45 \pm 0.05) \times 10^{-11} \left(\frac{E}{1\,\mathrm{TeV}}\right)^{-(2.63\pm0.02)} \mathrm{cm}^{-2}\ \mathrm{s}^{-1}\ \mathrm{TeV}^{-1} \tag{6.20}$$

The differential upper limits are calculated at the minimum safe energy (the lowest energy at which the bias in the reconstructed energy is less than 10%) and the integral upper limits are calculated between the minimum safe energy and 30 TeV. The upper

limits are given in Table 6.9 and compared with the *Fermi*-LAT results and the NAIMA MODEL in Figure 6.19.

Table 6.9.: VERITAS M 31 flux upper imits.

Cuts	Cleaning	Region Size	Min Safe Energy [GeV]	Differential UL [$\mathrm{GeV^{-1}\ cm^{-2}\ s^{-1}}$]	Integral UL [$\mathrm{cm^{-2}\ s^{-1}}$]	[% Crab]
hard	ED	*Small*	416.9	6.9×10^{-15}	1.9×10^{-12}	2.2
		Large	346.7	2.7×10^{-14}	6.2×10^{-12}	5.2
med	Thresh	*Small*	263.0	4.9×10^{-14}	8.5×10^{-12}	4.6
		Large	263.0	7.1×10^{-14}	1.2×10^{-11}	6.6

6.6. Prediction Comparisons

Table 6.10 shows the predictions I have made of the M 31 VHE integral flux between 300 GeV and 30 TeV as a percentage of the (Aharonian et al., 2006) Crab Nebula flux.

Table 6.10.: The predicted M 31 integral flux between 300 GeV and 30 TeV as a percentage of the Crab Nebula flux. The Crab spectrum is the power-law from Aharonian et al. (2006).

Prediction	300 GeV to 30 TeV Integral Flux [% Crab Nebula]
M 82 Scaled	0.23
NGC 253 Scaled	0.28
Fermi-LAT 700 MeV Extrapolated	0.097
NAIMA MODEL	0.03

This shows that, even working with the current upper limit which has an integral flux of 2.2% of the Crab Nebula flux between 416.9 GeV and 30 TeV there is still a reasonable way to go to constrain these predictions. The current upper limits are in the ∼3% of the Crab Nebula flux range, thus, to gain the factor of 20 improvement that is expected to be required (the more optimistic predictions are all at about the 0.15% Crab Nebula flux level) would require an estimate factor of 400 increase in the observing time, a total of about 2000 hours. However, at the moment only about 25% of the expected emission is

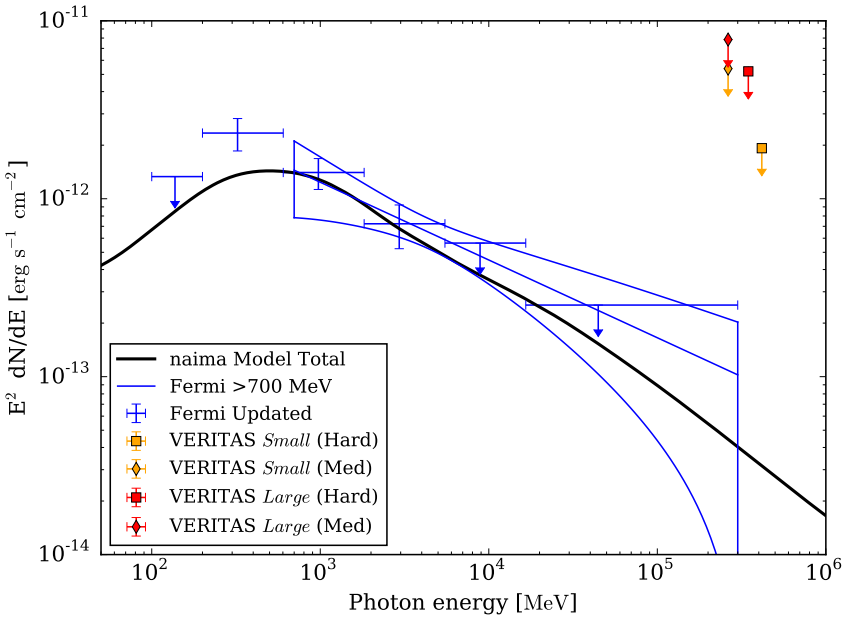

Figure 6.19.: The VERITAS upper limits compared with the *Fermi*-LAT spectra and the *naima* model.

being integrated used to generate the upper limit. With the improved analysis methods in the pipeline it (Subsection 6.3.2) this will be able to be increased to at least 50% (using the template background or maximum likelihood methods and diffuse simulations to generate the effective areas), reducing the time required to 500 hours. Additional improvements are being developed to increase the sensitivity of the analysis, they are unlikely to enable VERITAS to detect M 31 but they will reduce the upper limit.

These upper limits do, however, rule out any additional, hard spectrum, component and are the most constraining VHE upper limits to date for the extended emission from M 31.

6.7. Conclusions

54.7 hours of VERITAS observations of M 31 are presented, giving a 95% confidence level upper limit on the total VHE emission from M 31. By combining the flux from multiple regions an upper limit of total flux is calculated for different sets of cuts and test region sizes, Table 6.11.

Table 6.11.: VERITAS M 31 flux upper imits.

Cuts	Cleaning	Region Size	Min Safe Energy [GeV]	Differential UL [TeV^{-1} cm^{-2} s^{-1}]	Integral UL [cm^{-2} s^{-1}]	[% Crab]
hard	ED	*Small*	416.9	6.9×10^{-15}	1.9×10^{-12}	2.2
		Large	346.7	2.7×10^{-14}	6.2×10^{-12}	5.2
med	Thresh	*Small*	263.0	4.9×10^{-14}	8.5×10^{-12}	4.6
		Large	263.0	7.1×10^{-14}	1.2×10^{-11}	6.6

These limits are significantly above any model predictions of the total flux from M 31 but they, combined with the lack of any evidence of from the skymaps, show that there are no extremely bright sources within M 31 or any regions that show anomalous emission. It is predicted that, in order to detect M 31, ∼500 hours of observations will be required along with the implementation of a more advanced analysis method for extended sources.

In the *Fermi*-LAT, 7 years of data shows that emission is likely to be extended, though the impact of the background model requires careful consideration. At the lowest energies, there is evidence for a turnover in the spectrum, suggestive of pion emission being the source of the HE γ-ray emission.

A simple model of the different emission mechanisms based upon the MW but scaled to the parameters of M 31 fits the *Fermi*-LAT data well, at the energies of VERITAS it predicts at least 3 orders of magnitude less flux than the upper limit (though the contribution of point sources could be greater than that from the diffuse emission at these energies).

7. Discussion and Concluding Remarks

In this thesis I have looked at two ways to study cosmic rays with IACT telescopes, I have conducted a feasibility study into using the EMIS technique to measure the positron fraction and antiproton ratio, investigating and implementing the changes needed to conduct these observations with VERITAS and conducted and analysed the initial observations. I have observed the γ-ray emission from M 31 due to the interaction of cosmic rays with the ISM with *Fermi*-LAT and VERITAS, producing the most detailed study to date of this emission. In addition I have presented work I have conducted to support these observations, principally work on reducing the impact of the night-sky-background (NSB), both at low levels where variations (e.g. from stars) impact upon the image reconstruction or higher levels due to the Moon which prevent standard observations.

7.1. Technical Work

The work I have conducted on reducing the impact of the NSB has resulted in an improved image cleaning technique which uses the timing between the channels to improve the background rejection, this has increased the sensitivity of the analysis by a few percent. I also showed that, in order to reduce the optical impact of stars, a simpler image cleaning technique that uses a single pair of cleaning thresholds across the camera, throughout the run, reduces their optical impact, increasing the area within the FoV that is available for study.

My primary technical work has been with the bright moonlight programme where I have helped to increase the duty cycle of the instrument by allowing observations during periods when previously the PMT currents were above the safety threshold. In this work I have verified two new observing modes, reduced high voltage (RHV) and UV filter (UVF). RHV reduces the high voltage supplied to the PMTs to 81% of the nominal value, reducing the gain by a factor of 2.87, and thus reduced the current through the last dynode. This reduces its rate of damage. UVF introduces a UV bandpass filter (Schott

UG11) in front of the PMTs, reducing the amount of optical light that is incident upon the PMTs, protecting the photocathodes and improving the signal-to-noise. RHV mode provides a 13% boost in yearly exposure above 250 GeV and the UVF mode provides a 16% boost in yearly exposure above 500 GeV. These modes are used in standard operation and RHV has been used to detect the blazar *1ES 1727+502* in a flaring state, the first evidence of variability of the source (Archambault et al., 2015).

7.2. Moon Shadow Feasibility Study

Since the detection of the rise in the positron fraction at energies greater than 10 GeV it has become one of the most interesting measurements in cosmic ray physics. The origin of its unexpected rise has led to a lot of theoretical work and many different predictions. To constrain these theories higher energy measurements in the $\sim 0.5 - 1$ TeV range are required.

In this work I presented a feasibility study into whether it will be possible to observe the positron fraction with VERITAS in this energy range and the analysis of the data already taken. By pointing close ($2 - 3°$) to a partially illuminated Moon using UVF and RHV I attempted to use the EMIS technique (detecting the shadow of the Moon in the cosmic ray flux after deflection by the Earth's magnetic field) to detect the deficit for different rigidities of both leptonic and hadronic cosmic ray showers. No deficit was detected with the ∼5.5 hours of data taken so far, the feasibility study showed that at least 30 hours would be required to detect the electron shadow, and a few hundred for the positron.

During the 2014-04-15 lunar eclipse, observations were conducted with the array pointed directly at the Moon using the UVF observing mode. A deficit in the cosmic ray flux associated with the Moon's shadow is detected, the first such detection by an IACT using PMTs (the CLUE experiment has previously detected a deficit using a wire-proportional counter (Dokoutchaeva et al., 2001)). An upper limit is placed on the antiproton fraction of 0.8, this is over the energy range 2.5-6 TeV from 1 hour of observations.

7.2.1. Future Prospects

This work has shown that it could be feasible to use IACTs to detect the Moon's shadow but it will be very difficult. In reality it will likely require a telescope that is designed for such work, with good background rejection, fast timing (to aid background rejection), a low energy threshold and a camera designed to be UV sensitive and completely optically blind. With satellites reaching the high energy limit of their measurements due to limits on collecting area, the possibility of CTA being such an observatory is being actively investigated by that collaboration (Vandenbroucke et al., 2013) and the VERITAS results and operational knowledge are contributing to this work.

7.3. M 31 HE and VHE Observations

M 31 is one of the most well-studied objects in the sky and a known emitter of HE γ-rays (Abdo et al., 2010). This makes it an obvious target for VHE γ-rays observations as a detection would allow for detailed study due to the plethora of multiwavelength observations. It is especially an interesting object as, if it shows the expected extension, it will be significantly larger than VERITAS' angular resolution and thus the internal structure can be studied. 54.7 hours of observations with VERITAS did not detect any emission. By combining the flux from multiple regions an upper limit of total flux is calculated for different sets of cuts and test region sizes, Table 7.1. Based on predictions of the M 31 VHE flux from the *Fermi*-LAT observations, a model of the expected diffuse γ-ray emission and by scaling starburst galaxies flux, it is expected that ~500 hours of observations will be required to detect M 31 along with a more advanced extended source analysis.

Table 7.1.: VERITAS M 31 flux upper imits.

Cuts	Cleaning	Region Size	Min Safe Energy [GeV]	Differential UL [$\mathrm{GeV^{-1}\ cm^{-2}\ s^{-1}}$]	Integral UL [$\mathrm{cm^{-2}\ s^{-1}}$]	[% Crab]
hard	ED	*Small*	416.9	6.9×10^{-15}	1.9×10^{-12}	2.2
		Large	346.7	2.7×10^{-14}	6.2×10^{-12}	5.2
med	Thresh	*Small*	263.0	4.9×10^{-14}	8.5×10^{-12}	4.6
		Large	263.0	7.1×10^{-14}	1.2×10^{-11}	6.6

A reanalysis of the *Fermi*-LAT data using six and a half years of data detects M 31 and shows that there is marginal evidence of extended emission at the 2.24σ level. Looking at the spectral points, there is evidence of a turnover at low energies, suggestive of a pion bump. The points agree well with the GALPROP model of the MW luminosity scaled to the star formation rate of M 31 (as would be expected since the galaxies are very similar), with the GALPROP model predicting a VHE flux $\sim$1.5 orders of magnitude less than the VERITAS upper limits. Due to the curved nature of the low-energy end of the GALPROP spectrum, the *Fermi*-LAT spectrum was fit above 0.7 GeV using a power law. The fit was $(2.11 \pm 0.40) \times 10^{-14}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}\,\mathrm{MeV}^{-1}$ at a decorrelation energy of 4.384 GeV and with a power law index of -(2.44 ± 0.23). The $0.7 - 300$ GeV flux is $(9.0 \pm 1.7) \times 10^{-10}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$.

7.3.1. Future Prospects

This work has shown the potential for very interesting results if M 31 is detected in VHE γ-rays in the future. Both CTA and HAWC present opportunities to detect and study M 31. Due to M 31's declination, CTA will only be able to observe it with its northern site, which has approximately an order of magnitudes increase in sensitivity over VERITAS for point sources (Figure 1.1). With a larger FoV (improving sensitivity for extended sources), telescopes that are better designed for reducing the impact of the NSB (faster electronics, better timing properties) and further improvements in sensitivity from improvements in analysis methods, it might be possible to detect and resolve M 31 with a deep exposures. HAWC has presented a preliminary upper limit for M 31 of less than 14% of the Crab Nebula flux using their 111 tank configuration (Brewer & Smith, 2015). The full HAWC detector, with 3-years of data, will have an order of magnitude greater sensitivity than their presented result, an better sensitivity to the predicted extended emission from M 31 than the VERITAS analysis presented here.

7.4. Publication Plans

This thesis will produce three publications, all of which have already been presented at the 34th International Cosmic Ray Conference (ICRC) in the The Hague, The Netherlands 2015.

VERITAS Observations under Bright Moonlight (Griffin, 2015)

VERITAS Observations of M31 (the Andromeda Galaxy) (Bird, 2015c)

Observing the Cosmic Ray Moon Shadow with VERITAS (Bird, 2015a)

Papers for each of these three are in production at the moment and will be published in the near future.

In addition, the following was presented at the 34th ICRC,

Raster Scanning the Crab Nebula to Produce an Extended VHE Calibration Source (Bird, 2015b)

Though there are no plans to publish this on its own it is being used in the verification of other analysis methods for the measurement of extended sources.

7.5. Final Remarks

There are significant challenges for both of these observations but, with improving technology and longer observing times both could be accomplished in the next few years. They both present opportunities to study the behaviour of cosmic rays, in particular how they behave on a galactic scale. This will constrain our existing models and help to solve one of the longest running questions in physics *"How are galactic cosmic rays accelerated?"*

A. Hardware Zero Suppression

A major part of the 2012 upgrade to the VERITAS array was the replacing of the PMTs with those of a higher quantum efficiency (QE). This increases their sensitivity, increasing the trigger rate and thus increasing the deadtime, potentially to undesirable levels (desired level <10%, predicted >30%). As part of the preparation for the upgrade an investigation was conducted into two methods that could reduce the deadtime by changing the operating modes of the FADCs to speed up the readout of the data.

1. Zero Suppression

   In zero suppression mode, when the event is read out by the DAQ from the buffer (Subsection 3.1.4), only those channels which have a peak voltage above a certain threshold are read out, the rest are passed over and not read out (Figure A.1). The advantage is that fewer channels are read out, at the cost of information lost from the removed pixels.

2. QADC Mode

   In QADC mode, instead of reading out the charge sample by sample the integrated charge across the whole window is read out, dramatically speeding up the readout process. The disadvantage of this is that a fixed, wide readout window is used which dramatically reduces the sensitivity to smaller signals in the channels, removing the benefit of having FADCs.

Zero suppression is highly preferable to QADC mode since it maintains significantly more information and thus it was decided to investigate that method first and only investigate QADC mode if zero suppressed failed to provide sufficient gains. Two tests were conducted to determine the impact of zero suppression and the preferred level of suppression to use. In the first test, the array was artificially triggered at different rates to measure how the deadtime varied as a function of the trigger rate. The second test takes existing data and applies software zero suppression to it which allows the level

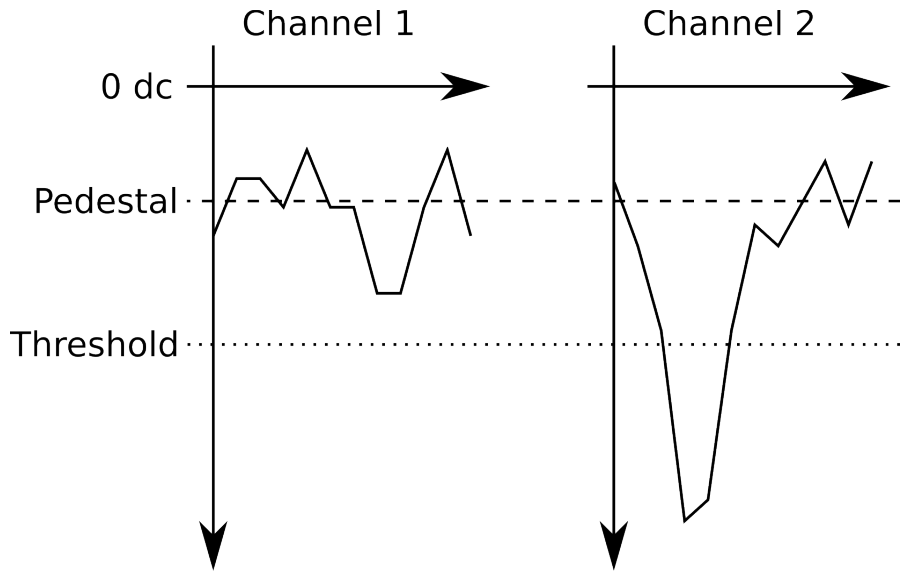

Figure A.1.: When a trigger is received the channel is only read out if it crosses the predefined threshold. In this figure channel 1 fails to cross the threshold and thus is not read out whereas channel 2 does and is read out.

of suppression for a given threshold to be determined. The artificial triggering of the array and the software zero suppression were both conducted prior to the upgrade to investigate the potential gains. Using this information, after the upgrade a threshold was chosen and the results verified using a single data run taken with zero suppression enabled. This allowed confirmation of the initial results and comparison of the data quality.

A.1. Artificial Triggering of the Array

Throughout these tests the telescopes were triggered using a quasi-random pulse generator. The telescopes were triggered at a variety of external rates from 300 Hz to 1000 Hz, an external rate of 500 Hz resulted in a readout rate of ∼450 Hz and a deadtime of 19%. The number of pixels that were suppressed was then varied from 0 to 100%, to do this, the pixels thresholds were set to one of two levels:

- 0 dc, the channel is always read out and it is *unsuppressed*.
- 255 dc (maximum level of suppression), the channel is never read out and it is

suppressed.

The deadtime was measured using online software (qlmonitor) after the telescopes had been allowed to settle at the new rate.

It was found that there is a deadtime overhead to running the telescopes with zero suppression turned on (FADC mode) in comparison to running with the capability turned off (FULL mode)(Figure A.2). This is caused by the additional overhead of reading out the board headers when operating FADC mode. The FADC channels are on boards in groups of 10, in FULL mode, when all channels are always read out, the channels are directly read out. However, FADC mode requires the boards header to be read to determine which channels require reading out, before reading out the channels.

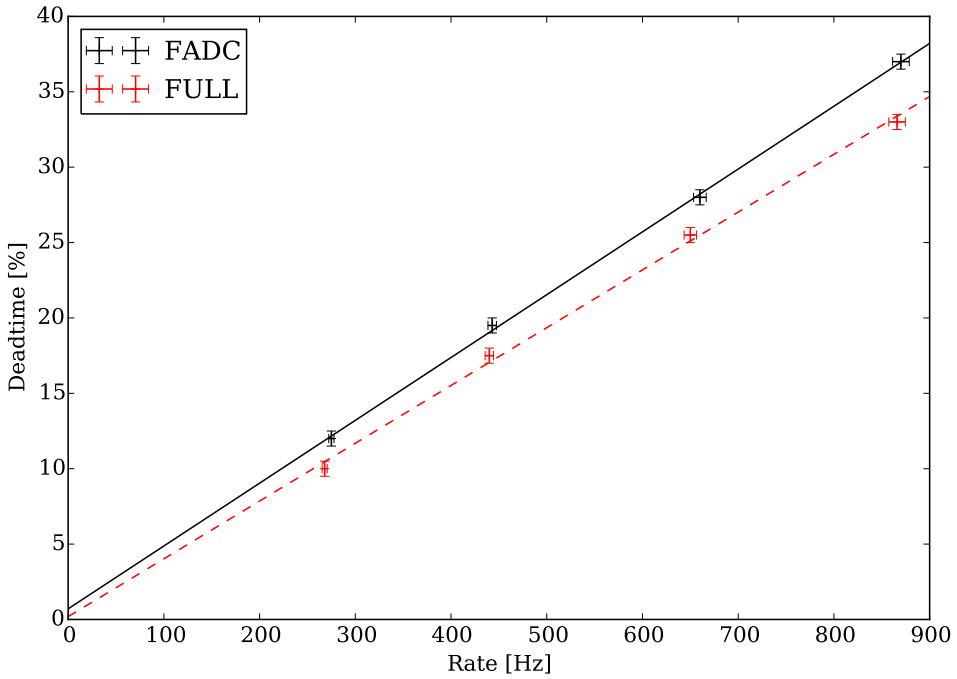

Figure A.2.: A plot of deadtime against rate for FULL (red, dashed) and FADC mode (black, solid) with all channels suppressed. The additional overhead of reading out the board headers in FADC mode is clearly visible. These measurements were taken with a 5 word readout window (a word is 4 samples long, each sample lasting 2 ns) which was the default prior to the PMT upgrade.

Apart from the results presented in Figure A.2, all the tests were taken with all the boards in the telescopes in FADC mode and the channels either suppressed (threshold = 255 dc) or unsuppressed (threshold = 0 dc). The suppressed channels were evenly distributed across the boards within an individual telescope and were the same across the three telescopes investigated. For a given level of suppression the channels suppressed were kept the same.

For a variety of trigger rates, the deadtime was measured with 0, 25, 50 75 and 100% of the channels suppressed (for comparison, in FULL mode, at 450 Hz the deadtime is ∼17.5%) (Figure A.3). This shows that, as expected, with an increase in the level of suppression, the deadtime decreased, with the deadtime being proportional to the rate.

The results presented in Figure A.4 show how the deadtime varies for a single rate (450 Hz), for comparison running in FULL mode results in a deadtime of ∼17.5%. Running with ∼60 % of the channels suppressed reduces the deadtime by ∼ 1/3.

A.2. Software Zero Suppression

To test the impact of zero suppression on the data analysis a raw data file is read in, zero suppression is applied to all events other than pedestal events (the result is the same as hardware zero suppression) and then the results are recorded in a new file. The impact of software zero suppression will vary with NSB, thus the tests the threshold level will have to be set for a compromise level, in this case a Galactic FoV was chosen since this will have a reasonably high NSB rate. This method enabled a number of different levels of suppression to be tested on different sources and also on simulations. To test the impact of the zero suppression on the analysis, the suppressed and unsuppressed files were both analysed up until the end of the image cleaning stage of the analysis (Subsection 3.4.3). The suppressed and unsuppressed analyses were then compared, the number of channels suppressed against threshold is plotted in Figure A.5 showing the level of suppression increases with threshold as expected. At a low threshold the channels suppressed are channels that are only contain noise and the percentage of levels suppressed is a linear function of the threshold. Above 6 dc the majority of channels have been suppressed and the gradient decreases as only those channels which are imaging the shower remain, many of which have a signal far greater than 6 dc.

This analysis also allowed for the comparison between the post cleaning (Subsection 3.4.3) *Size* and *nTubes* of the same image with and without zero suppression ap-

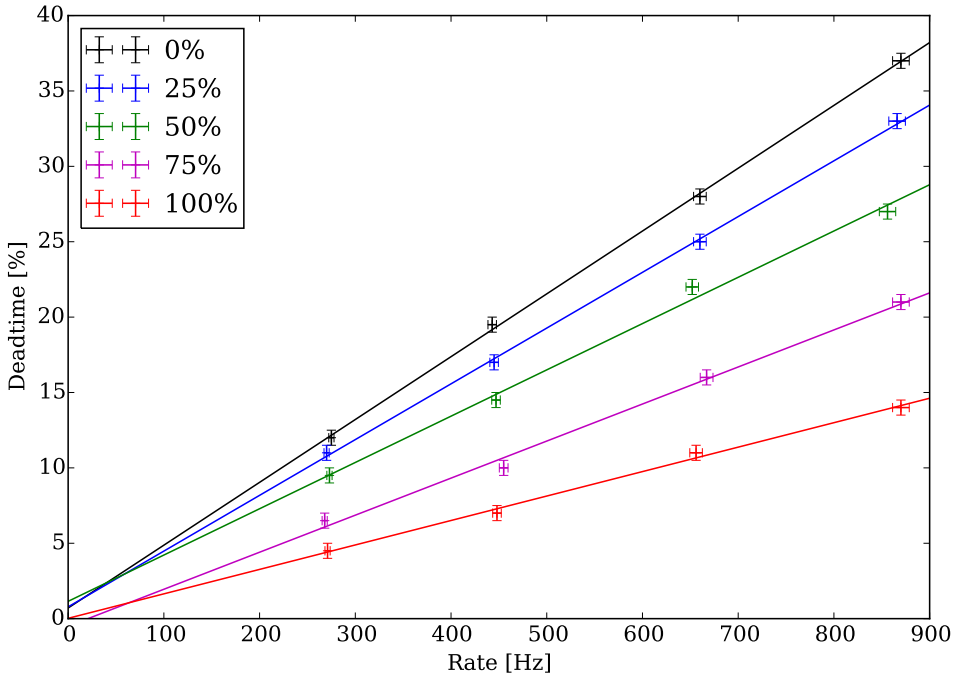

Figure A.3.: A plot of deadtime against rate for different levels of suppression (0%, black; 25%, blue; 50%, green; 75%, magenta and 100%, red) with a readout window of 5 words (the default before the 2012 upgrade). For all levels of suppression as the rate increases the deadtime increases linearly. The external trigger rates were kept the same for each of the levels of suppression, the array response rate did not change with increased deadtime due to the quasi-random trigger used where the smallest gap between triggers was larger that the deadtime per trigger.

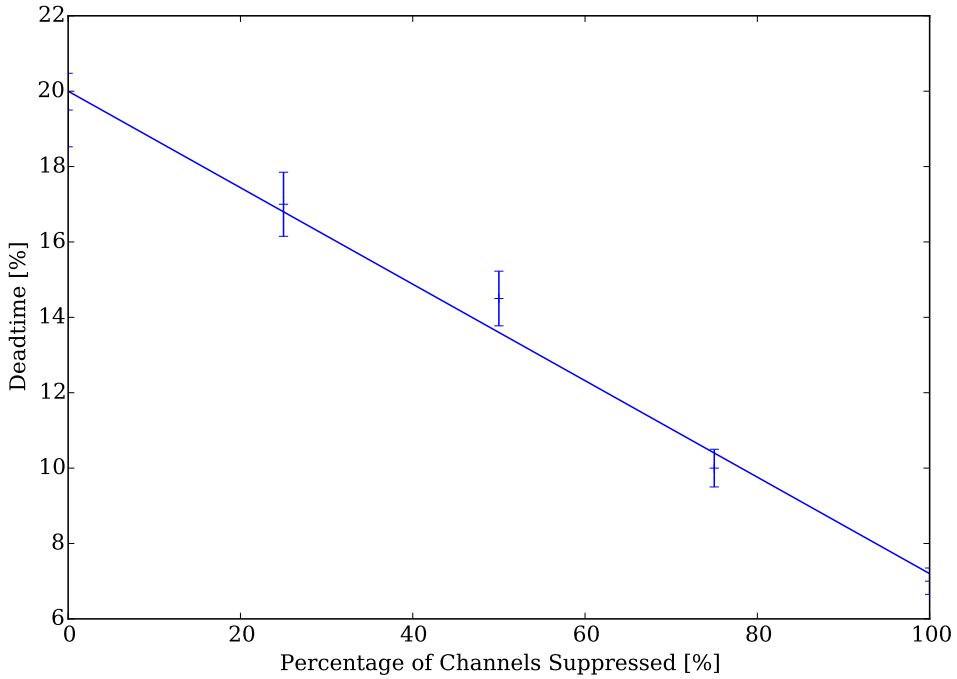

Figure A.4.: Deadtime vs. Number of Channels Suppressed for a a fixed rate of 450 Hz with an external trigger

plied. The percentage differences for all the events in a single run were averaged and the results plotted against the percentage of channels suppressed Figure A.6. This shows that at low thresholds, there is very little difference in the post cleaning analysis. However, as the percentage of channels suppressed increases the difference increases quite dramatically, with the difference in *nTubes* rising by a factor of nearly 7 between 75% and 90% suppression.

From this, using the pre-upgrade PMTs a threshold of 5 dc (∼60% suppression) was identified as the optimum. These initial tests also identified the major causes of deadtime within the array and allowed for work to be conducted to reduce the impact of these bottlenecks.

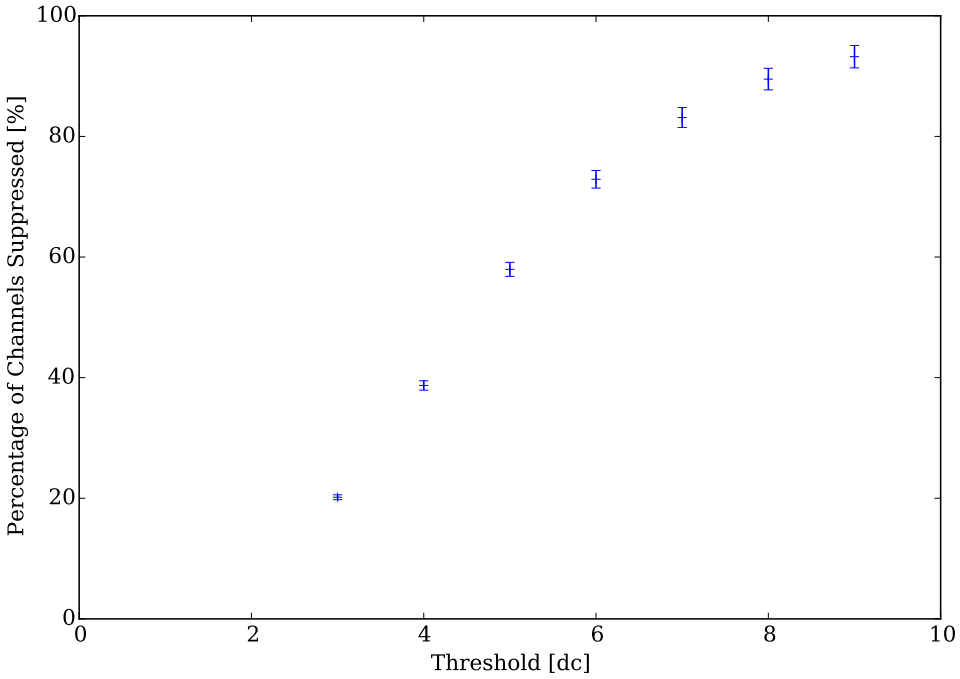

Figure A.5.: The percentage of channels suppressed vs. the suppression threshold for the software suppression of a single (20 min) pre-upgrade run.

A.3. Analysis Verification

Following the upgrade the trigger rate for the array was ∼450 Hz with a deadtime of ∼15%. In order to confirm the earlier work and to determine a suitable threshold tests were conducted on actual data. To do this, a pair of runs was taken on the Crab Nebula, the first without zero suppression applied, the second with a threshold of 7 dc, (scaled up from the pre-upgrade optimum 5 dc based upon the measured change in the single photo-electron pulse height). The first run was also software zero suppressed for comparison with the hardware zero suppression and verification of the earlier work. These runs were analysed using standard lookup tables (that is, without zero suppression applied) and thus the zero-suppressed results are liable to have larger errors than normal.

The results of this test (Table A.1) show the improvements in the deadtime that occurred due to the improvements in the readout software by the time these tests were

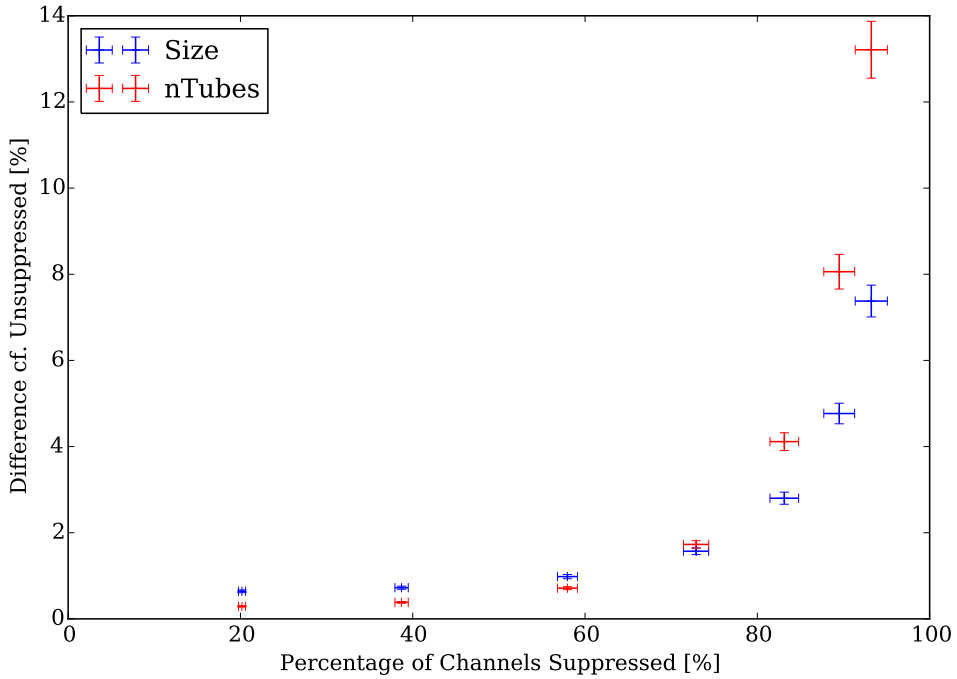

Figure A.6.: The percentage difference between the post cleaning size (blue) and number of tubes (nTubes, red) for unsuppressed, pre-upgrade showers and the same shower software zero suppressed at thresholds of 3 dc to 9 dc.

taken, based on the earlier work a trigger rate of ∼450 Hz had an expected deadtime of ∼17.5%. Hardware zero suppression shows a further reduction in the deadtime and thus an increase in the significance (S) for the run. Correcting for the resultant live time (LT) difference by calculating $S/\sqrt{LT}$ shows that there is an increase in the significance beyond that expected from the livetime difference, since no such gain is shown in the software suppression it suggests that this is a result of night to night observation differences and/or a statistical effect (the differences between the L3 rates is due to the reduction in the deadtime).

Table A.1.: Comparison between a standard (20 minute) Crab Nebula observation with the same run software zero suppressed at 7 dc and a hardware zero suppressed run (7 dc) at matched sky conditions (weather and zenith angle) to the standard run. The trigger rate in both runs was ∼450 Hz

Suppressed	L3 Rate	Deadtime [%]	Live Time [minutes]	Significance	Significance / [sqrt(Live Time (hrs))]
None	452	15.7	16.88	22.25	41.95
Software	452	15.7	16.88	21.06	39.71
Hardware	462	11.1	17.80	23.68	43.48

A.4. Conclusions and Outcomes

- These results showed that there is a reduction in the deadtime if zero suppression is used, the additional events increase the statistics collected and thus improve the sensitivity of the array. However, there is a cost to applying zero suppression, information is lost and there is no way to recover this data at a later date.
- During these tests issues in the data acquisition software were identified and corrected, after these changes the deadtime at normal operating rates reduced to less than 10%, a level that is comparable to the deadtime before the upgrade.
- At a deadtime of ∼10%, the benefit of zero suppression is only a few percent (absolute) reduction in the deadtime. Given that the other improvements reduced the benefits of running with zero suppression and with the risk of loosing information that may be useful for advanced analysis methods currently in development it was decided not to run the array in zero suppression mode but to remain running with the standard (FULL) mode.
- The fact that zero suppression barely affected the sensitivity of the array at low levels shows that the current analysis is less sensitive than the hardware to small signals and thus there is room for improvement.

B. Raster Scanning the Crab Nebula To Produce an Extended Calibration Source

Since its initial discovery as a VHE γ-ray source, the Crab Nebula has proved to be a very useful calibration source for the verification of systems and processes and the development of new techniques. Its main advantages are that it is bright, with a constant flux (at least to a level that variability has thus far not been detected by any imaging atmospheric Cherenkov Telescope (IACT) (O'Faolain de Bhroithe, 2015)) and visible by all the major observatories (allowing for verification between them). Since the majority of sources observed by IACTs are point sources, the Crab Nebula (which appears to IACTs as a point source) provides a suitable calibration source for those objects. However, there is no similar source that is bright, has a constant flux and with a known, extended morphology available to act as an extended calibration source.

I developed a novel method of generating a fake extended source by conducting a raster scan over the Crab Nebula, this generates an extended source of known flux and morphology allowing for the verification of extended source analyses. This method is complementary to existing techniques where a simulated source is added to a true background, since it is less reliant on a detailed understanding of the response of the systems away from the centre of the field of view.

B.1. Methodology

The raster scan is adapted from the software used for the VERITAS mirror alignment (McCann et al., 2010) and the star raster scan for relative throughput measurements (Griffin & Hanna, 2013). It works by adjusting the pointing corrections of telescopes, thus the telescope still believes that it is pointing in the original direction though it now has an offset. This allows for the standard analysis method Section 3.4 to be used with no changes, (note that the VERITAS pointing monitors - a set of optical star trackers

that are used to correct for any pointing errors - are not used as they will correct for the offset introduced by the scan). The pattern followed by the scan can be adjusted and the code has been set up so the four telescopes follow the same pattern at the same time.

The observations are conducted in the standard way, with the Crab Nebula offset from the camera center by 0.5° in one of the cardinal directions (N, S, E or W) (*Wobble* mode, Section 3.4.7) to enable background estimation from the rest of the camera. Data collection is then started before the raster scan commences and the run is longer than the length of the scan, this additional time at the start and the end of the run is removed using time cuts. The pointing of the array is recorded for each raster position and used to determine the actual pattern followed. From that the apparent location of the Crab Nebula is determined.

In this initial test, two patterns that are loosely rectangular in shape were used, with different orientations and offset directions. In the future, the raster program will be updated to generate a box in RA/Dec that will allow for multiple scans to be overlaid.

B.2. Scan Patterns

Two scan patterns were performed (one wobble N (Figure B.1a) thus the scan does not cross the camera center and one wobble S (Figure B.1b) where it does cross the camera center) with a dwell time of 10 s on each point. With 189 pointings in total this gives a scan time of 31.5 minutes. The pattern was designed with sufficient statistics and a noticeable asymmetry in the shape for verification that the scan has produced the desired result.

With a "standard" point source analysis, using the same set up as used in this analysis (but without the extended exclusion regions) the Crab Nebula is detected at a significance of $\sim 40\sigma$ per sqrt(hour), thus for a 30 minute run I would expect a detection at $\sim 30\sigma$. However, since the Crab Nebula will be "smeared" over an area, approx 1° by 0.5°, the fraction of the flux that falls within a test region will be significantly below this, thus though I expect to get a significant detection, it will be significantly lower than that from the point source.

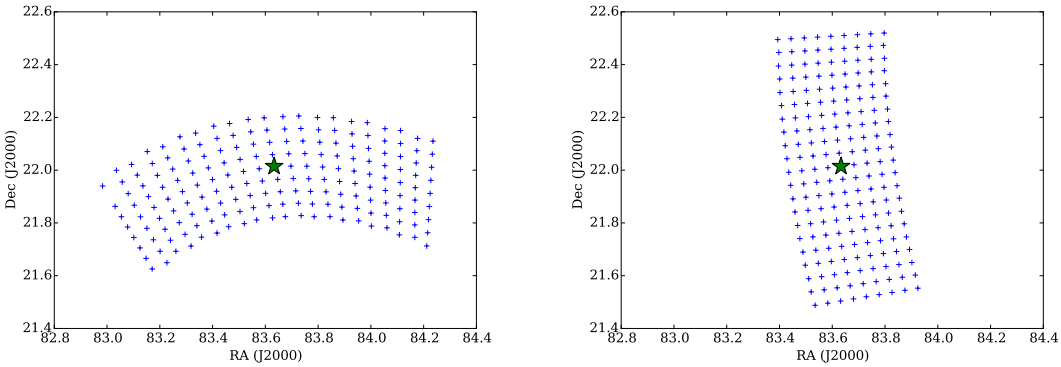

(a) *Pattern 1* **(b)** Pattern 2

Figure B.1.: The raster patterns that were employed. The large star indicates the Crab Nebula position, the blue crosses are the effective positions of the Crab Nebula for each point in the scan.

B.3. Skymaps

The standard VERITAS analysis is used with standard gamma/hadron selection cuts to produce skymaps using the ring background method (where the background region is taken from a ring surrounding the test position Section 3.4.7). In generating the skymap the whole of the area covered by the effective Crab Nebula positions and identical area centred on ζ-Tauri are excluded from the background model. For the weaker stars in the FoV only the nominal positions are excluded since a similar large exclusion region significantly reduces the background available and the impact of the star is diminished since its effect is smeared out over a larger area.

Two different test region radii are tested, 0.1° and 0.3°, the skymaps of these for each of the patterns are shown in Figure B.2 and Figure B.3. They clearly show an extended source that is coincident with the shape scanned out by the raster process, thus showing that the method has effectively produced an extended test source. The peak significance in the maps with a test region radius of 0.3° is 10-12σ, a reasonable agreement with the simple calculation above, though there are noticeable variations across the scan due to statistical fluctuations.

The edges of the scan, where the statistics are lower, show larger statistical fluctuations, this is "smeared out" by the larger test regions creating a smoother shape but

they have also had the effect of reducing the significance of the edges of the scan. This reduces the apparent size of the source with the high significance region significantly away from the ends of the scan. The significance of the detection also depends upon the size of the test region, with a larger test region containing a larger signal, increasing the significance of the detection. If the test region is made larger still then the significance decreases due to the reduction in the background area available (with a test region radius of 0.5° the peak significance for *Pattern 2* drops to about 8).

B.4. Spectral Reconstruction

A spectral analysis of the region centred on the true Crab position (the central point of the scan) is shown in Figure B.4. Each spectrum was fit with a power law of the form $\frac{\mathrm{d}N}{\mathrm{d}E} = N_0 \left(\frac{E}{0.6\,\mathrm{TeV}}\right)^{-\alpha}$ TeV^{-1} cm^{-2} s^{-1} (Table B.1). The *Pointed* spectrum is taken from the VERITAS spectrum presented in (Meagher, 2015) but with the fit conducted over the range 0.2 to 1.5 TeV.

Run	Significance [σ]	N_0	α
Pointed	15.62	$(1.28 \pm 0.04) \times 10^{-10}$	2.30 ± 0.01
Pattern 1	11.68	$(5.3 \pm 0.6) \times 10^{-11}$	2.4 ± 0.2
Pattern 2	8.93	$(4.9 \pm 0.7) \times 10^{-11}$	2.6 ± 0.3

Table B.1.: Spectral parameters for the central region (test region radius = 0.3°) of the two raster scan patterns (*Pattern 1, 2* compared to the VERITAS Crab Nebula spectrum (*Pointed*) fit over the energy range 0.2 to 1.5 TeV.

As expected the fluxes from the raster observations are significantly lower than for the standard Crab Nebula observation since only a fraction of the flux now originates within the test region.

An estimate of the fraction of the flux that I would expect can be made using the method in Subsection 6.3.4 which gives an estimate of N_0 from the entire scan of $(1.1 \pm 0.1) \times 10^{-10}/(9.8 \pm 1.4) \times 10^{-11}$, these are close to the nominal value.

In the future, by taking more scans and stacking the results the statistical errors will reduce significantly, this will enable more detailed testing of this method of estimating the flux from the entire object and further verification of the extended source analysis.

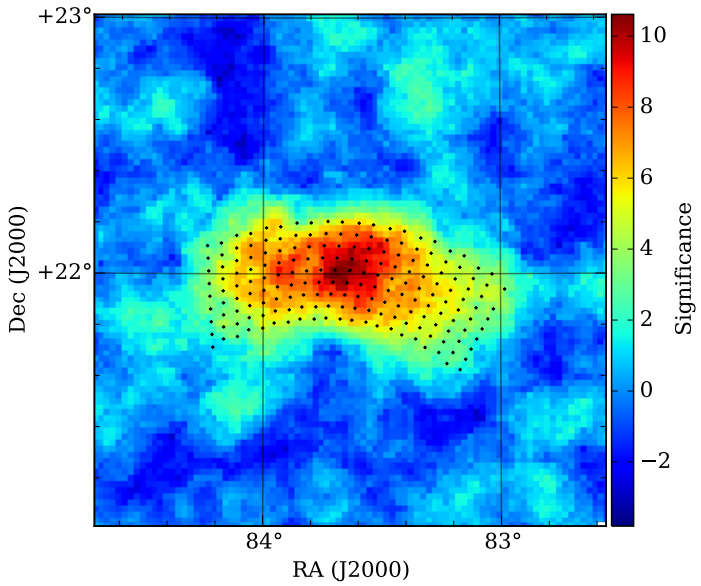

(a) Test region radius $= 0.1°$

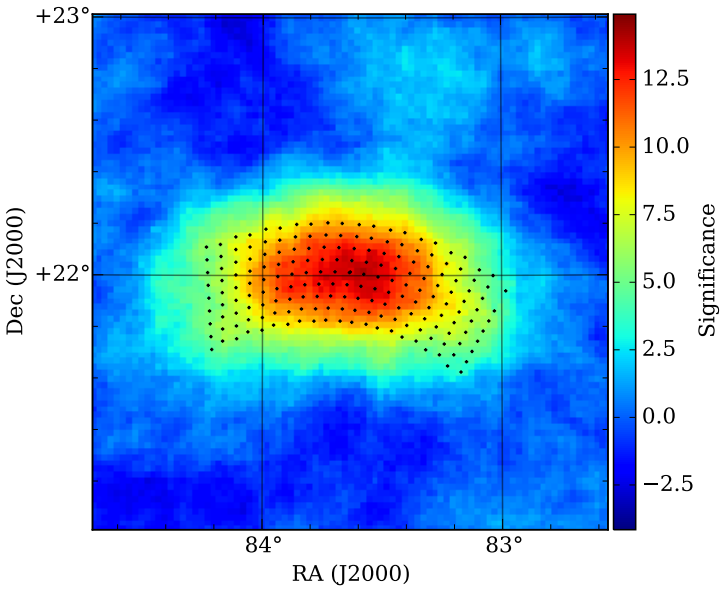

(b) Test region radius $= 0.3°$

Figure B.2.: Skymap of the Crab Raster Scan for *Pattern 1*, the nominal Crab positions are the black dots.

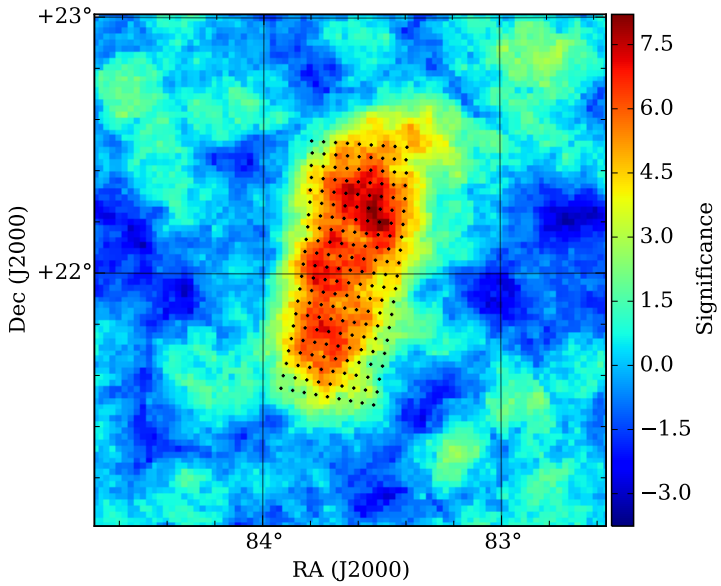

(a) Test region radius $= 0.1°$

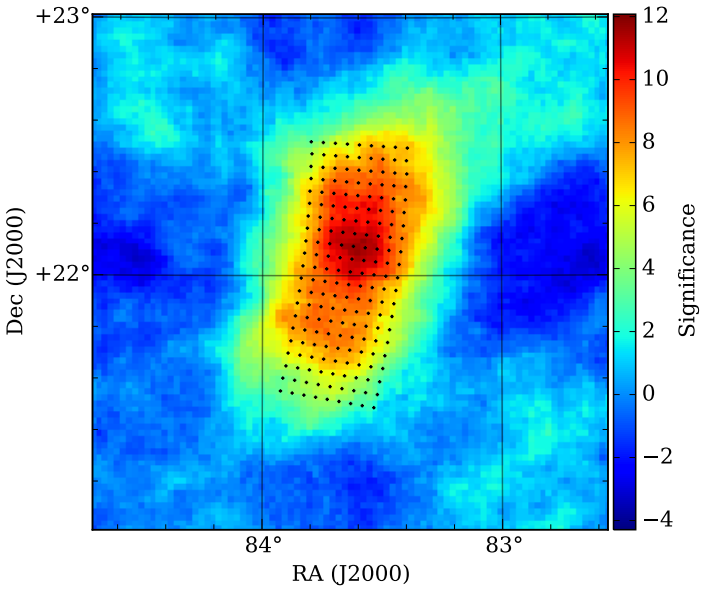

(b) Test region radius $= 0.3°$

Figure B.3.: Skymap of the Crab Raster Scan for *Pattern 2*, the nominal Crab positions are the black dots.

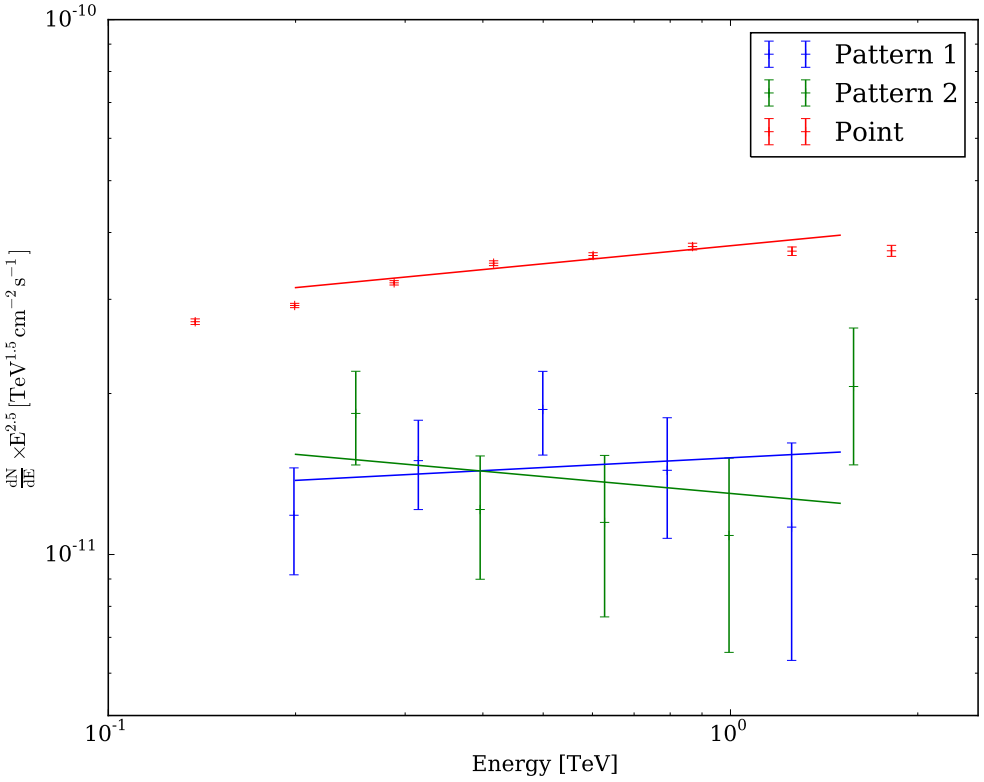

Figure B.4.: Spectrum from the center of the scan in *Pattern 1* (red) and *Pattern 2* (blue) compared with the VERITAS Crab Nebula spectrum fitted over the energy range $0.2 - 1.5\,\mathrm{TeV}$, *Pointed* (black). This analysis was conducted with a test region radius of $0.3°$ for the raster scan patterns.

B.5. Conclusions

By conducting a raster scan across the Crab Nebula it is possible to generate an extended source. Since the scan follows a known pattern the morphology of the source is known and thus it can be used for testing extended source analysis methods. The central, 0.3° radius region within each scan was fit with a power law of N_0 (*Pattern 1*/*Pattern 2*) = $(5.3 \pm 0.6) \times 10^{-11}/(4.9 \pm 0.7) \times 10^{-11}\,\mathrm{TeV}^{-1}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$, $\alpha = 2.4 \pm 0.2/2.6 \pm 0.2$, E_0 = 0.6 TeV. Comparison with a model of the scan showed that this was expected to be 46.8%/50.8% of the total flux from the scan, thus the N_0 for the whole region is estimated at $(1.1 \pm 0.1) \times 10^{-10}/(9.8 \pm 1.4) \times 10^{-11}\,\mathrm{TeV}^{-1}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. For comparison, the VERITAS Crab Nebula spectrum fitted over the same energy range has a normalisation of $(1.28 \pm 0.04) \times 10^{-10}\,\mathrm{TeV}^{-1}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ and an index of -2.30 ± 0.01. With improvements to the raster code, multiple scans will be combined to improve the statistics of the method and enable its use for improved testing of spectral reconstruction and morphology for extended objects.

Bibliography

Aartsen, M. G., Abbasi, R., Abdou, Y., et al. 2013, Science, 342, 1242856

Abbasi, R., Ackermann, M., Adams, J., et al. 2009, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 601, 294

Abdo, A. A., Allen, B., Aune, T., et al. 2008, The Astrophysical Journal, 688, 15

Abdo, A. A., Allen, B. T., Aune, T., et al. 2009, The Astrophysical Journal, 700, L127

Abdo, A. A., Ackermann, M., Ajello, M., et al. 2010, Astronomy & Astrophysics, 523, L2

Abdo, A. A., Allen, B. T., Atkins, R., et al. 2012, The Astrophysical Journal, 750, 63

Abe, K., Fuke, H., Haino, S., et al. 2008, Physics Letters B, 670, 103

Abeysekara, A., Alfaro, R., Alvarez, C., et al. 2014, Physical Review D, 90, 122002

Abraham, J., Aglietta, M., Aguirre, I., et al. 2004, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 523, 50

Abramowski, A., Acero, F., Aharonian, F., et al. 2012, The Astrophysical Journal, 757, 158

Abramowski, A., Aharonian, F., Benkhali, F. A., et al. 2014, Physical Review D, 90, 122007

Accardo, L., Aguilar, M., Aisa, D., et al. 2014, Physical Review Letters, 113, 121101

Acciari, V. A., Aliu, E., Arlen, T., et al. 2009a, Nature, 462, 770

Acciari, V. A., Aliu, E., Aune, T., et al. 2009b, The Astrophysical Journal, 707, 612

Acero, F., Ackermann, M., Ajello, M., et al. 2015, The Astrophysical Journal Supplement Series, 218, 23

Acharya, B., Actis, M., Aghajani, T., et al. 2013, Astroparticle Physics, 43, 3

Ackermann, M., Ajello, M., Atwood, W. B., et al. 2012a, The Astrophysical Journal, 750, 3

Ackermann, M., Ajello, M., Allafort, A., et al. 2012b, Physical Review Letters, 108, 011103

—. 2013, Science, 339, 807

Ackermann, M., Ajello, M., Albert, A., et al. 2015, The Astrophysical Journal, 799, 86

Adriani, O., Barbarino, G. C., Bazilevskaya, G. a., et al. 2009, Nature, 458, 607

Adriani, O., Barbarino, G. C., Bazilevskaya, G. A., et al. 2010, Physical Review Letters, 105, 11

Ageron, M., Aguilar, J., Al Samarai, I., et al. 2011, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 656, 11

Aguilar, M., Alcaraz, J., Allaby, J., et al. 2007, Physics Letters B, 646, 145

Aguilar, M., Alberti, G., Alpat, B., et al. 2013, Physical Review Letters, 110, 141102

Aharonian, F., Akhperjanian, A., Beilicke, M., et al. 2004, Astronomy & Astrophysics, 421, 529

Aharonian, F., Akhperjanian, A. G., Bazer-Bachi, A. R., et al. 2006, Astronomy & Astrophysics, 457, 899

Aharonian, F. A., Akhperjanian, A. G., Beilicke, M., et al. 2003, Astronomy & Astrophysics, 400, 153

Ahlen, S. P., & Tarlé, G. 2014, arXiv:1410.7239

Albert, J., Aliu, E., Anderhub, H., et al. 2007, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 583, 494

Aleksić, J., Alvarez, E., Antonelli, L., et al. 2012, Astroparticle Physics, 35, 435

Archambault, S., Archer, A., Beilicke, M., et al. 2015, The Astrophysical Journal, in Prep, arXiv:1506.06246

Atwood, W. B., Abdo, A. A., Ackermann, M., et al. 2009, The Astrophysical Journal, 697, 1071

Baade, W., & Zwicky, F. 1934, Physical Review, 46, 76

Baring, M. G., Ellison, D. C., Reynolds, S. P., Grenier, I. A., & Goret, P. 1999, The Astrophysical Journal, 513, 311

Bartoli, B., Bernardini, P., Bi, X., et al. 2012, Physical Review D, 85, 022002

Barwick, S. W., Beatty, J. J., Bhattacharyya, A., et al. 1997, The Astrophysical Journal, 482, L191

Bird, D. J., Corbato, S. C., Dai, H. Y., et al. 1995, The Astrophysical Journal, 441, 144

Bird, R. 2015a, in Proceedings of the 34th ICRC, PoS(ICRC2015)852

Bird, R. 2015b, in Proceedings of the 34th ICRC, PoS(ICRC2015)854

Bird, R. 2015c, in Proceedings of the 34th ICRC, PoS(ICRC2015)851

Blasi, P. 2013, The Astronomy and Astrophysics Review, 21, 70

Blasi, P., & Amato, E. 2012, Journal of Cosmology and Astroparticle Physics, 010

Blasi, P., & Serpico, P. D. 2009, Physical Review Letters, 103, 081103

Boezio, M., Pearce, M., Picozza, P., et al. 2009, New Journal of Physics, 11, 105023

Bottino, A., Donato, F., Fornengo, N., & Salati, P. 1998, Physical Review D, 58, 123503

Braun, R., Thilker, D. A., Walterbos, R. A. M., & Corbelli, E. 2009, The Astrophysical Journal, 695, 937

Brewer, I., & Smith, A. 2015, Bulletin of the American Physical Society, Volume 60,

Brun, R., & Rademakers, F. 1997, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 389, 81

Buffington, A., Orth, C. D., & Smoot, G. F. 1975, The Astrophysical Journal, 199, 669

Cardenzana, J. V., & Collaboration, f. t. V. 2015, in Proceedings of the 34th ICRC

Cawley, M. F., & Weekes, T. C. 1995, Experimental Astronomy, 6, 7

Cholis, I., & Hooper, D. 2013, Physical Review D, 88, 023013

Cirelli, M., Kadastik, M., Raidal, M., & Strumia, A. 2009, Nuclear Physics B, 813, 1

Colin, P., Tridon, D. B., Britzger, D., et al. 2009, in Proceedings of the 31st ICRC

Colin, P., Tridon, D. B., Ortega, A. D., et al. 2011, in Proceedings of the 32nd ICRC

Condon, J. J. 1992, Annual Review of Astronomy and Astrophysics, 30, 575

Cowsik, R., Burch, B., & Madziwa-Nussinov, T. 2014, The Astrophysical Journal, 786, 124

Daugherty, J. K., Hartman, R. C., & Schmidt, P. J. 1975, The Astrophysical Journal, 198, 493

Daum, A., Hermann, G., Heß, M., et al. 1997, Astroparticle Physics, 8, 1

Dokoutchaeva, V., Kartashov, D., Malakhov, N., et al. 2001, Il Nuovo Cimento C, 24CS2, 669

Donato, F., Maurin, D., Brun, P., Delahaye, T., & Salati, P. 2009, Physical Review Letters, 102, 071301

Donato, F., Maurin, D., Salati, P., et al. 2001, The Astrophysical Journal, 563, 172

Draine, B. T., Aniano, G., Krause, O., et al. 2014, The Astrophysical Journal, 780, 172

Drury, L. O. 2012, Astroparticle Physics, 39-40, 52

Drury, Â. 1991, Monthly Notices of the Royal Astronomical Society, 251, 340

Falvard, A., Giraud, E., Jacholkowska, A., et al. 2004, Astroparticle Physics, 20, 467

Fanselow, J. L., Hartman, R. C., Hildebrad, R. H., & Meyer, P. 1969, The Astrophysical Journal, 158, 771

Fegan, D. J. 1997, Journal of Physics G: Nuclear and Particle Physics, 23, 1013

Fermi, E. 1949, Physical Review, 75, 1169

Finlay, C. C., Maus, S., Beggan, C. D., et al. 2010, Geophysical Journal International, 183, 1216

Fomin, V. P., Lamb, R. C., Lewis, D. A., & Punch, M. 1991, in Proceedings of the 22nd ICRC

Fornasa, M., Taoso, M., & Bertone, G. 2007, Physical Review D, 76, 043517

Fruck, C. 2015, in Proceedings of the 34th ICRC, PoS(ICRC2015)859

Gaisser, T., & Stanev, T. 2008, Physics Letters B, 667, 254

Gaisser, T. K. 1990, Cosmic Rays and Particle Physics (Cambridge University Press)

Gaisser, T. K. 2010, in Proceedings of XVI International Symposium on Very High Energy Cosmic Ray Interactions

Giesen, G., Boudaud, M., Genolini, Y., et al. 2015, 12

Ginzburg, V. L., & Syrovatskii, S. I. 1964, The Origin of Cosmic Rays (New York: Macmillan)

Golden, R. L., Mauger, B. G., Horan, S., et al. 1987, Astronomy & Astrophysics, 188, 145

Greisen, K. 1966, Physical Review Letters, 16, 748

Griffin, S. 2015, in Proceedings of the 34th ICRC, PoS(ICRC2015)989

Griffin, S., & Hanna, D. 2013, in Proceedings of the 33rd ICRC

Guo, Y.-Q., Feng, Z.-Y., Yuan, Q., Liu, C., & Hu, H.-B. 2013, New Journal of Physics, 15, 013053

Hall, J., Vasilev, V., Kieda, D., et al. 2003, in Proceedings of the 28th ICRC

Hamaguchi, K., Moroi, T., & Nakayama, K. 2015, Physics Letters B, 747, 523

Hanna, D., McCann, A., McCutcheon, M., & Nikkinen, L. 2010, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 612, 278

Hays, E. 2007, in Proceedings of the 30th ICRC

Heo, J., & Kim, C. 2013, Physical Review D, 87, 013007

Hess, V. F. 1912, Physikalische Zeitschrift, 1084

Hillas, A. M. 1984, Annual Review of Astronomy and Astrophysics, 22, 425

Hillas, A. M. 1985, in Proceedings of the 19th ICRC, Vol. 3, 445–448

Hinshaw, G., Larson, D., Komatsu, E., et al. 2013, The Astrophysical Journal Supplement Series, 208, 19

Holder, J. 2005, in Proceedings of the 29th ICRC

Holder, J., Acciari, V. A., Aliu, E., et al. 2008, in AIP Conference Proceedings (AIP), 657–660

Hörandel, J. R. 2008, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 588, 181

Hubble, E. P. 1929, The Astrophysical Journal, 69, 103

Kafexhiu, E., Aharonian, F., Taylor, A. M., & Vila, G. S. 2014, Physical Review D, 90, 19

Kawai, H., Yoshida, S., Yoshii, H., et al. 2008, Nuclear Physics B - Proceedings Supplements, 175-176, 221

Kelner, S., Aharonian, F., & Bugayov, V. 2006, Physical Review D, 74, 034018

Khangulyan, D., Aharonian, F. A., & Kelner, S. R. 2014, The Astrophysical Journal, 783, 100

Kieda, D. B. 2011, in Proceedings of the 32nd ICRC

Kounine, A. 2012, International Journal of Modern Physics E, 21, 1230005

Lenain, J.-P., & Walter, R. 2011, Astronomy & Astrophysics, 535, 5

Li, T., & Ma, Y. 1983, The Astrophysical Journal, 272, 317

Li, Z., Yuan, Q., & Xu, Y. 2013, arXiv:1312.7609

Lin, R., & Hudson, H. 1976, Solar Physics, 50, doi:10.1007/BF00206199

Linden, T., & Profumo, S. 2013, The Astrophysical Journal, 772, 18

Longair, M. S. 2011, High Energy Astrophysics, 3rd edn. (Cambridge University Press)

Lorenz, E., & Wagner, R. 2012, The European Physical Journal H, 37, 459

M. Boudaud, S. Aupetit, S. Caroff, et al. 2015, Astronomy & Astrophysics, 20

Margiotta, A. 2014, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 766, 83

Maurin, D., Melot, F., & Taillet, R. 2014, Astronomy & Astrophysics, 569, A32

Mayr (Marius), S. 1614, Mundus Jovialis [The World of Jupiter]

McCann, A., Hanna, D., Kildea, J., & McCutcheon, M. 2010, Astroparticle Physics, 32, 325

Meagher, K. J. 2015, in Proceedings of the 34th ICRC, PoS(ICRC2015)792

Meegan, C., Lichti, G., Bhat, P. N., et al. 2009, The Astrophysical Journal, 702, 791

Meeus, J. 1991, Astronomical Algorithms, 2nd edn. (AbeBooks)

Messier, C. 1774, Mémoires de l'Académie Royale des Sciences, 435

Meyer, M., Horns, D., & Zechlin, H.-S. 2010, Astronomy & Astrophysics, 523, A2

Miville-Deschenes, M., & Lagache, G. 2005, The Astrophysical Journal Supplement Series, 157, 302

Mohanty, G., Biller, S., & Carter-Lewis, D. 1998, Astroparticle Physics, 9, 15

Mostafa, M. 2013, in Proceedings of the 33rd ICRC

Mueller, D., & Tang, K.-K. 1987, The Astrophysical Journal, 312, 183

Nolan, P. L., Abdo, A. A., Ackermann, M., et al. 2012, The Astrophysical Journal Supplement Series, 199, 31

O'Faolain de Bhroithe, A. 2015, in Proceedings of the 34th ICRC, PoS(ICRC2015)788

Olive, K. 2014, Chinese Physics C, 38, 090001

Otte, A. N. 2011, in Proceedings of the 32nd ICRC

Park, N. 2015, in Proceedings of the 34th ICRC, PoS(ICRC2015)771

Pavlidou, V., & Fields, B. D. 2001, The Astrophysical Journal, 558, 63

Perkins, J. S., & Maier, G. 2009, in Proceedings of the 2009 Fermi Symposium

Persic, M., & Rephaeli, Y. 2011, Nuovo Cimento della Societa Italiana di Fisica C, 34, 217

—. 2012, Journal of Physics: Conference Series, 355, 012038

Pomarède, D., Boyle, P. J., Urban, M., et al. 2001, Astroparticle Physics, 14, 287

Porter, T. A., Johnson, R. P., & Graham, P. W. 2011, Annual Review of Astronomy and Astrophysics, 49, 155

Pretz, J. 2015, in Proceedings of the 34th ICRC, PoS(ICRC2015)025

Pshirkov, M. S., Vasiliev, V. V., & Postnov, K. A. 2015, 5

Ptuskin, V. 2006, Journal of Physics: Conference Series, 47, 113

Rephaeli, Y., & Persic, M. 2013, in Astrophysics and Space Science Proceedings, 34th edn. (Springer), 193–209

Reynolds, P. T., Akerlof, C. W., Cawley, M. F., et al. 1993, The Astrophysical Journal, 404, 206

Rolke, W. A., López, A. M., & Conrad, J. 2005, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 551, 493

Rowell, G. P. 2003, Astronomy & Astrophysics, 410, 389

Schönwald, A. 2008, PhD thesis, Humboldt University of Berlin

Smith, A. 2015, in Proceedings of the 34th ICRC, PoS(ICRC2015)787

Stanek, K. Z., & Garnavich, P. M. 1998, The Astrophysical Journal, 503, L131

Stasak, D. 2015, in Proceedings of the 34th ICRC, PoS(ICRC2015)411

Stephens, A. A. 1985, Astronomy & Astrophysics, 149, 1

Stiele, H., Pietsch, W., Haberl, F., et al. 2011, Astronomy & Astrophysics, 534, 51

Strong, A. W., Moskalenko, I. V., & Ptuskin, V. S. 2007, Annual Review of Nuclear and Particle Science, 57, 285

Strong, A. W., Moskalenko, I. V., & Reimer, O. 2000, The Astrophysical Journal, 537, 763

Strong, A. W., Porter, T. A., Digel, S. W., et al. 2010, The Astrophysical Journal, 722, 7

Vandenbroucke, J., Humensky, B., Holder, J., & Ong, R. A. 2013, in White Paper for the 2013 Snowmass Community Summer Study

Völk, H. J., & Bernlöhr, K. 2009, Experimental Astronomy, 25, 173

Wakely, S., Cui, W., Pizlo, F., et al. 2003, in Proceedings of the 28th ICRC

Walterbos, R. A. M., & Schwering, P. B. W. 1987, Astronomy and Astrophysics, 180, 27

Weekes, T. C., Cawley, M. F., Fegan, D. J., et al. 1989, The Astrophysical Journal, 342, 379

Weinstein, A. 2008, in Proceedings of the 30th ICRC

Winston, R. 1970, Journal of the Optical Society of America, 6

Wissel, S. A. 2007, in Proceedings of the 30th ICRC

Zabalza, V. 2015, Naima: Derivation of non-thermal particle distributions through MCMC spectral fitting

Zatsepin, G. T., & Kuz'min, V. A. 1966, Journal of Experimental and Theoretical Physics Letters, 4

Zitzer, B. 2013, in Proceedings of the 33rd ICRC

Zitzer, B. 2015a, in Proceedings of the 34th ICRC, PoS(ICRC2015)1225

Zitzer, B. 2015b, in Proceeding at the 5th International Symposium on High Energy Gamma-Ray Astronomy