EXPLORING THE LIMITS OF LORENTZ INVARIANCE WITH VERITAS GAMMA-RAY OBSERVATIONS OF MARKARIAN 421

by

Scott Tyler Griffiths

A thesis submitted in partial fulfillment
of the requirements for the Doctor of
Philosophy degree in Physics
in the Graduate College of
The University of Iowa

August 2015

Thesis Supervisor: Professor Philip Kaaret

Graduate College
The University of Iowa
Iowa City, Iowa

CERTIFICATE OF APPROVAL

PH.D. THESIS

This is to certify that the Ph.D. thesis of

Scott Tyler Griffiths

has been approved by the Examining Committee for the thesis requirement for the Doctor of Philosophy degree in Physics at the August 2015 graduation.

Thesis Committee: _____
Philip Kaaret, Thesis Supervisor

_____
Steven Spangler

_____
Vincent Rodgers

_____
Randall McEntaffer

_____
Nepomuk Otte

To my wonderful wife Jacquelynn and my parents for their boundless love and support, and to John Andrews, Jr., who has always encouraged my love of science.

Oh, the depth of the riches and wisdom and knowledge of God!
How unsearchable are his judgments and how inscrutable his ways!
For who has known the mind of the Lord, or who has been his counselor?
Or who has given a gift to him that he might be repaid?
For from him and through him and to him are all things.
To him be glory forever. Amen.

— Romans 11:33-36

ACKNOWLEDGEMENTS

I am very grateful for the guidance and support provided to me by my advisor, Philip Kaaret. In addition to teaching me physics and astronomy and related academic knowledge, he taught me many valuable lessons about how to be a research scientist, such as teaching me how to collaborate with other scientists and engineers, encouraging me to be more assertive, and providing wisdom about the unavoidable politics of science. I am also exceedingly grateful for the funding he provided through research assistantships during the entire duration of my research career with him, and the many and varied research projects he allowed me to work on. I could not have finished this dissertation without his help and encouragement. I am very thankful for the four and a half years I spent learning from him.

I would like to acknowledge my friends and colleagues in the VERITAS collaboration for their help and support in my research. I would especially like to thank Nepomuk Otte, Martin Pohl, and David Hanna for their assistance with my Lorentz invariance violation research which constitutes the bulk of this thesis, as well as Jamie Holder and Gernot Maier for believing in me and writing letters of recommendation which opened the door to a post-doctoral position in Barcelona at the Institut de Física d'Altes Energies (IFAE). I would also like to thank Czar Benjamin Zitzer for always making startup observing shifts fun, and Jeff Grube for taking the time to give me a personal EventDisplay tutorial at the Adler Planetarium. I will miss the many kind people and congenial atmosphere within the VERITAS collaboration.

I cannot fail to also thank Joe Hill and Kevin Black for allowing me to work with them on the Gravity and Extreme Magnetism Small Explorer (GEMS) project. They afforded me

a very rare opportunity to be a part of the development of a NASA spacecraft. Some of the most memorable experiences of my graduate school career are the times I spent at Goddard Space Flight Center and Brookhaven National Laboratory working on GEMS. I am very grateful to Joe and Kevin for sharing their expertise with me and for their encouragement. I am also grateful to Kevin for going out of his way to get me my first lobster roll.

I would also like to thank my fellow graduate students and undergrads whom I had the pleasure of working with at Iowa. I would particularly like to thank Ryan Allured, Hannah Aytaç, and Ryan Daly for working with me on GEMS and its student-led subproject, the Bragg Reflection Polarimeter. I am very grateful to them for being excellent colleagues and for teaching me many valuable things. I would also like to thank Eva Smith for spending a summer helping me take data with the autocollimator.

Finally, I would like to thank my wonderful wife Jackie for supporting me during the preparation of this dissertation and for being an unending source of encouragement and joy. I would also like to thank my parents, Walt and Kim, and my sisters, Kristen and Jennifer, for their encouragement to pursue graduate studies and to finish them. Last but not least, I am grateful to Jesus for being my savior and friend. I am hopeful that He will be honored by this work.

ABSTRACT

The search for a theory of quantum gravity has persisted through the last century. Although many beautiful theories such as string theory and loop quantum gravity have been proposed, experimental evidence to support or refute these theories has been difficult to obtain. Searching for Lorentz invariance violation (LIV) is one of a limited number of experimental tests which can be used to search for evidence of quantum gravity since new physics may only be observable at energies well beyond those present in the most energetic astrophysical objects, which are far greater than the energies accessible in a terrestrial laboratory.

One method of searching for LIV is to look for energy-dependent time delays in the arrival of high-energy photons from distant astrophysical sources. We search for Lorentz invariance violation (LIV) using VERITAS, an imaging atmospheric Cherenkov telescope (IACT) located in southern Arizona. Significant TeV gamma ray flaring activity was detected from the blazar Markarian 421 on the night of February 17, 2010 (MJD 55244), which presented a good opportunity to search for delays in the energetic emission. We demonstrate the performance of two different dispersion estimation algorithms and apply these algorithms to our data to search for LIV. We find that while the emission from Markarian 421 contains significant variability, a necessary condition for an LIV detection, the presence of a constant background flux severely limits our sensitivity. We expect our findings to be useful for guiding future LIV studies, especially those using IACT data.

In the latter part of this work we discuss the alignment of ground-based gamma-ray telescopes and present a digital autocollimator which will be used in the alignment system

of a next-generation IACT. The configuration of our autocollimator enables measurement of the angle formed between the planar surface of a distant reflector and the line of sight over a range of $\pm 0.126^{\circ}$ with a precision better than 5 arcsec. We present a detailed description of the instrument and its data acquisition software that was used during laboratory testing.

PUBLIC ABSTRACT

The speed at which light travels through vacuum is thought to be an immutable constant of nature, unchanging regardless of who observes it and where. Evidence accumulated over the past century indicates that the maximum speed of light is always the same regardless of its energy (wavelength or frequency). However, it is possible that at very high energies, the speed of light in vacuum may change.

I use observations of the most energetic light, gamma rays, made with a telescope called VERITAS to look for arrival time differences between high and low energy light emitted by a jet of material streaming out of a supermassive black hole. If a difference were observed, it could help us find a theory of quantum gravity and would provide insight into the physics of the early universe. The way in which the black hole jet was emitting light at the time we observed it does not allow me to place any useful constraints on the constancy of the speed of light. However, the reasons why I was not able to generate a constraint are themselves a useful result.

I also discuss an instrument I designed and constructed to measure the tilt of a mirror located 9 meters away. I show that the instrument is capable of measuring tilts of approximately 1/1000 of a degree. This instrument will be used to help align the mirrors of a next-generation gamma-ray telescope.

TABLE OF CONTENTS

LIST OF TABLES

Table

LIST OF FIGURES

Figure

INTRODUCTION

The search for a theory of quantum gravity has persisted through the last century. Although many beautiful theories such as string theory and loop quantum gravity have been proposed, experimental evidence to support or refute these theories has been difficult to obtain. Lorentz invariance violation (LIV) is one of a limited number of experimental tests which can be used to search for evidence of quantum gravity. This is because new physics may only be observable at energies well beyond those present in the most energetic astrophysical objects, which are themselves far greater than the energies accessible in a terrestrial laboratory.

One method of searching for LIV is to look for energy-dependent time delays in the arrival of high-energy photons from distant astrophysical sources. In Part I of this work I search for Lorentz invariance violation (LIV) using VERITAS, an imaging atmospheric Cherenkov telescope (IACT) located in southern Arizona. Significant TeV gamma ray flaring activity was detected from the blazar Markarian 421 on the night of February 17, 2010 (MJD 55244), which presented a good opportunity to search for delays in the energetic emission.

I describe LIV and how it can arise in the context of quantum gravity in Chapter 1. A more condensed discussion of LIV which discusses how it is detected in practice can be found in the introduction to Chapter 4. Chapter 2 reviews blazars, which are a type of active galactic nucleus (AGN), and discusses the blazar Markarian 421. Chapter 3 discusses the Very Energetic Radiation Imaging Telescope Array System (VERITAS), which is the ground-based gamma-ray telescope that was used to take the observations of Markarian 421 used in this work.

Chapter 4 discusses the results of my search for LIV using VERITAS observations of Markarian 421. I demonstrate the performance of two different dispersion estimation algorithms using an artificial data set consisting of a single Gaussian pulse before applying the algorithms to the VERITAS data. I find that while the emission from Markarian 421 contains significant variability, a necessary condition for an LIV detection, the presence of a constant background flux severely limits our sensitivity. I expect my findings to be useful for guiding future LIV studies, especially those using IACT data.

In Part II, I discuss the alignment of ground-based gamma-ray telescopes. I first review the alignment system used on VERITAS in Chapter 5, which I spent several years operating and maintaining. I then present a digital autocollimator system in Chapter 6 that I designed and constructed. The autocollimator will be used in the alignment system of a prototype next-generation IACT with Schwarzschild-Couder optics.

Unlike a traditional autocollimator, my autocollimator is configured as a telescope focused at infinity. This configuration enables measurement of the angle formed between the planar surface of a distant reflector and the line of sight over a range of $\pm 0.126^\circ$ with a precision better than 5 arcsec. I present a detailed description of the instrument and its construction and discuss the data acquisition software that I wrote for use during laboratory testing.

Part I

VERITAS CONSTRAINTS OF LORENTZ INVARIANCE VIOLATION

CHAPTER 1

LORENTZ INVARIANCE VIOLATION

1.1 What is Lorentz Invariance Violation?

The introduction of special relativity by Einstein (1905) changed our fundamental understanding of space and time. It is built on two simple postulates: the principle of relativity, and the constancy of the propagation velocity of light. Relativity requires that two, non-accelerating observers should agree that the same laws of physics govern the dynamics of any mutually observed system irrespective of the observers' relative rotation, translation, and uniform motion. The latter postulate implies that those same non-accelerating (inertial) observers would also measure the same speed of light, regardless of how fast they themselves move. Thus a new way of relating the reference frames of inertial observers was required, one that would automatically preserve the constancy of the speed of light for all inertial observers, everywhere: the Lorentz transformation.

Note the tacit assumptions in the previous statement of the isotropy and homogeneity of space. Special relativity applies everywhere only in the case of an empty, flat universe with infinitesimally-massive or massless observers. In this case, Lorentz symmetry (the equivalence of two inertial reference frames under Lorentz transformation) is a global symmetry, valid everywhere and at all times. The implication of a global Lorentz symmetry is that spacetime is simply a static backdrop for interactions between matter, fixed and unchanging for all time. The geometry of this flat spacetime can be described with the Minkowski metric, which has a line element

$$ds^2 = -c^2dt^2 + dx^2 + dy^2 + dz^2. \tag{1.1}$$

Note that Lorentz symmetry applies globally in the SU(3) x SU(2) x U(1) standard model of particle physics (Tasson, 2014).

Our understanding of spacetime was radically changed again by the introduction of general relativity (Einstein, 1915, 1916) only a decade later. In the paradigm of general relativity, spacetime is no longer a static backdrop with no influence on the dynamics of matter. The presence of large quantities of matter deforms spacetime, giving it curvature. The motion of matter in response to the curvature of spacetime results in the phenomenology that we call gravity. Gravity is thus not interpreted as a force, but rather an interaction between matter and spacetime.

General relativity builds on the equivalence principle (Einstein, 1908), whereby the physical consequences of acceleration by an applied force and acceleration by a gravitational field are supposed to be equivalent (as are inertial mass and gravitational mass). In general relativity, the equivalence principle holds only locally in regions which are small enough so that curved spacetime can be approximated as flat (this can always be done if spacetime is a pseudo-Riemannian manifold, $ds^2 = g_{\mu\nu}dx^\mu dx^\nu$). In these small regions, the spacetime metric is Minkowskian and special relativity applies. We therefore conclude that Lorentz symmetry is local in general relativity, in the sense that any sufficiently small region of spacetime is approximately flat.

We have thus seen that Lorentz symmetry is a fundamental part of the universe we live in, at least to our present level of understanding. But what exactly is Lorentz symmetry, and what does it mean for Lorentz invariance to be violated?

As we have discussed, relativity requires that the laws of physics not be altered if

the observer is translated, rotated, or boosted to a constant velocity, assuming the absence of other factors (such as a gravitational field). That is, we should be able to perform a coordinate transform from $x \to x'$ of the form

$$x'^{\mu} = \Lambda^{\mu}_{\nu} x^{\nu} + C^{\mu}, \tag{1.2}$$

where Λ^{μ}_{ν} is the Lorentz transform tensor and C^{μ} is a constant vector representing a translation. If $C^{\mu} \neq 0$ (i.e., the transformation is affine), such a transformation is more properly called a Poincaré transformation (Misner et al., 1973). The Lorentz transformation is linear and includes rotations and boosts only. Lorentz invariance means that a quantity has the same value in every inertial reference frame (this should not be confused with Lorentz *covariance*, whereby tensor quantities in different inertial reference frames are related by Lorentz transforms, but are not in general expected to maintain the same value unless they are identically zero).

In this work, we take Lorentz invariance violation (LIV) to be a synonym for violations of the constancy of the vacuum speed of light, c, which is postulated by special relativity to be Lorentz invariant. It is important to note that we are discussing the speed of light in vacuum, since light is often slowed when passing through materials (the reduction in speed being determined by the refractive index of the material). Even in vacuum, light can be slowed by the presence of a gravitational potential according to the formula (Cheng, 2010)

$$c(r) = \left(1 + 2\frac{\Phi(r)}{c^2}\right) c, \tag{1.3}$$

where $c(r)$ is the speed of light at some distance r from a gravitational potential $\Phi(r) = -GM/r$ (note the minus sign). Thus we must confine our discussion to measurements in an

inertial frame, or otherwise account for known general-relativistic effects.

The constancy of the speed of light has been rigorously tested in both terrestrial and astrophysical settings (Aleksandrov, 1965; Alväger et al., 1964; Babcock & Bergman, 1964; Brecher, 1977; Filippas & Fox, 1964; Lovell et al., 1964; Michelson et al., 1935; Pérez, 2011; Zahejsky & Kolesnikov, 1966), and has been found in all cases to be constant within experimental errors. We will describe the unusual circumstances under which Lorentz invariance may be violated in the next section. Exploring LIV through tests of the constancy of the speed of light is especially popular among astrophysicists, partially because of the central role the speed of light plays in general/special relativity and because light is the oldest and most common tool for astronomical observations, but largely because of the work of Amelino-Camelia et al. (1998), which will be discussed later. We will also be working in an astronomical context, so it makes sense for us to adopt this viewpoint.

We note a discussion of LIV need not be restricted to violations of the constancy of the speed of light. In fact, restricting our attention to light only constitutes a rather narrow view since there are many quantities besides c that are Lorentz invariant, and many ways in which Lorentz invariance can be violated. Alan Kostelecký and his colleagues have championed this point of view. They have developed an effective field theory called the Standard Model Extension (SME; Colladay & Kostelecký 1997), which is a Lorentz-violating extension of the standard model. The SME is a framework for collecting limits on LIV from experiments using matter, photons, neutrinos, and gravity; the limits are updated annually in a comprehensive set of tables (Kostelecký & Russell, 2011). It is interesting to note that Lorentz invariance can be violated either in terms of rotations or boosts; most theorists

prefer the latter and hold rotational symmetry as somewhat sacrosanct since violations of rotational symmetry would drastically alter the known laws of physics. The SME allows for both. The interested reader is referred to articles by Kostelecký and collaborators; see also Tasson (2014) and references therein.

1.2 How might LIV arise in Quantum Gravity?

It is generally expected that new physics may occur at energies on the order of the Planck energy, $E_P = \sqrt{\hbar c^5/G} \approx 1.2 \times 10^{19}$ GeV, or, equivalently, the Planck length, $\ell_P = \sqrt{\hbar G/c^3} \approx 1.6 \times 10^{-35}$ m (see Gorelik (1992) for an interesting historical review). Planck units (Misner & Wheeler, 1957; Planck, 1900) combine the speed of light, the gravitational constant, and Planck's constant (which generally arises out of quantum phenomena), and are a natural choice for a length/energy scale where new physics would occur.

One of the few features of quantum gravity that most physicists agree on is the presence of quantum foam (Hawking, 1978; Wheeler, 1957). At scales larger than the Planck length, spacetime can be considered a smooth manifold, and the field equations of general relativity apply. However, at scales below the Planck length, spacetime is thought to contain all manner of exotic topologies and may be permeated throughout with tiny wormholes. At this scale, virtual pairs of electrons and positrons are thought to densely populate the vacuum, blinking into existence and then annihilating just as quickly.

At the Planck scale, the quantum foam may have some interaction with particles propagating through spacetime (Hawking et al., 1979). The quantum foam is expected to interact most strongly with scalar particles, the only known scalar particle being the Higgs

boson (Aad et al., 2012; Chatrchyan et al., 2012). However, there may be some small interactions between the quantum foam and fermions or vector bosons (e.g., photons) at energies below the Planck energy (Hawking et al., 1979, 1980). The nature of these interactions is unclear. Interactions of non-scalar particles with the quantum foam are suppressed by the Planck scale so that they are not normally observable. It may be the case that higher energy particles interact more strongly with the quantum foam than lower energy particles. If this is true, these small interactions may alter the propagation of particles in the vacuum in a manner that gives the vacuum a non-negligible, energy-dependent refractive index.

1.3 Astrophysical Tests of LIV

Many phenomena occur in the universe with energies that are orders of magnitude greater than those accessible by terrestrial laboratories. As a result, astrophysical observations are the best, and often only, means of probing physics at the highest energy scales. Since Lorentz invariance violation is expected to occur near the Planck energy ($\sim 10^{19}$ GeV), astrophysical observations offer the best means of searching for LIV. However, given that the most energetic photons yet observed have an energy of about 50 TeV (Tanimori et al., 1998), it initially seemed unlikely that astrophysical observations could provide any insight into physics occurring at the Planck scale.

A key observation of Amelino-Camelia et al. (1998) is that it is possible to probe the Planck scale using particles with energies well below the Planck energy. The Planck energy is a composite unit with dimensions of [Length]^2 [Mass] [Time]$^{-2}$ (mass can be exchanged for energy according to $E = mc^2$). Thus with a combination of short variability timescales, large

distances, and high energies (even if those energies are much less than the Planck energy) it is possible to probe the Planck scale with present technology. Amelino-Camelia et al. realized that even tiny modifications to the speed of light due to LIV may be observable as a time-offset between the arrivals of high and low energy photons if the time dispersion at the source was negligible and the photons traveled astronomical distances.

This reasoning can be expressed mathematically as follows, using notation that has been adopted by many authors (Abdo et al., 2009; Amelino-Camelia et al., 1998; Wagner, 2009). Consider the standard dispersion relation for light:

$$c^2 p^2 = E^2 \tag{1.4}$$

Assuming that quantum gravity effects are suppressed for energies well below the Planck energy, we can expand the energy in our dispersion relation perturbatively in the small parameter E/E_P:

$$c^2 p^2 = E^2 \left(1 + \xi \frac{E}{E_P} + \zeta \frac{E^2}{E_P^2} + \cdots\right) \tag{1.5}$$

where $\xi = \pm E_P/E_{QG}$ and $\zeta = \pm E_P^2/E_{QG}^2$ are constant, dimensionless coefficients that are usually assumed to be of order unity. Dividing the equation by the term in parentheses and taking the partial derivative with respect to the momentum p,

$$2E \frac{\partial E}{\partial p} = 2c^2 p \left(1 + \xi \frac{E}{E_P} + \zeta \frac{E^2}{E_P^2} + \cdots\right)^{-1} \tag{1.6}$$

Keeping terms only to second order, we have

$$v = \frac{\partial E}{\partial p} \approx c \left(1 - \xi \frac{E}{E_P} - \zeta \frac{E^2}{E_P^2}\right) \tag{1.7}$$

Note that we have tacitly assumed that the group velocity remains unaltered in the quantum gravity regime, which may or may not be the case. We thus find that the vacuum behaves

like a "medium" with an energy-dependent dispersion relation $n(E) = c/v(E)$. Moreover, when two photons of differing energies E_1 and E_2 are emitted from a source at the same time, the non-trivial vacuum dispersion will cause the photons arrival times to differ by

$$\Delta t = \frac{L}{v_2} - \frac{L}{v_1} = L\left(-\frac{v_2 - v_1}{v_1\, v_2}\right) \approx -\left(v_2 - v_1\right)\frac{L}{c^2} \tag{1.8}$$

Substituting Equation (1.7) for v_1 and v_2 on the right hand side of Equation (1.8), we get a simple expression for the expected propagation delay due to LIV:

$$\Delta t = \xi\frac{E_2 - E_1}{E_P}\frac{L}{c} + \zeta\frac{E_2^2 - E_1^2}{E_P^2}\frac{L}{c} + \mathcal{O}\left(\frac{E^3}{E_P^3}\right) \tag{1.9}$$

However, when the source of the high-energy photons is at cosmological distances, we must take into account the expansion of the universe, as has been pointed out by Jacob & Piran (2008). Cosmological expansion will not only change the way the distance to the source, L, is measured, it will also redshift the photons by a factor of a^{-1}, where a is the scale factor. For a standard ΛCDM cosmology, the time delay for a leading LIV term of order n is

$$\Delta t_n = \frac{1+n}{2H_0}\left(\frac{\Delta E}{\xi E_P}\right)^n \int_0^z \frac{(1+z')^n\, dz'}{\sqrt{\Omega_M(1+z')^3 + \Omega_\Lambda}} \tag{1.10}$$

where H_0 is the Hubble constant, z is the redshift, Ω_M is the current density parameter for matter (both baryonic matter and dark matter), and Ω_Λ is the cosmological constant. (Equation (1.10) appears in Jacob & Piran (2008), but for a more instructive derivation, see Martínez & Piran (2006).)

From Equation (1.9) it is clear that we should prefer sources that emit energetic photons and are far away, but another key feature that we require is sharp features in the light curve. Sharper features allow for less uncertainty in determining the time difference

between the same feature in two different energy bands. The quality of the source as it relates to its ability to probe LIV is sometimes quantified in terms of a "sensitivity factor," which is defined by Amelino-Camelia et al. (1998) as

$$\eta = |\Delta t^*|/\delta t \tag{1.11}$$

where δt is a measure of the time structure of the signal and Δt^* is the expected delay if $E_{QG} \sim E_P$. Larger sensitivity factors are better. For example, if we consider only the linear LIV term for a galactic source, then from Equation (1.9),

$$\Delta t^* \sim \frac{E}{E_P}\frac{L}{c} \tag{1.12}$$

and thus

$$\eta \sim \frac{E}{E_P}\frac{L}{c\,\delta t} \tag{1.13}$$

Many astronomical sources exist that would be expected to have high sensitivity factors. Perhaps the most obvious choice is prompt gamma-ray bursts (GRBs), which are the sources that were originally proposed by Amelino-Camelia et al. (1998). GRBs are extragalactic (van Paradijs et al., 1997), with redshifts of order unity ($\sim 10^{10}$ lightyears). They are very energetic, with energies up to 10's of GeV, and by definition of the fact that they are bursts, they generally have sharp light curve features with millisecond variability (note we are considering only prompt GRBs (Kouveliotou et al., 1993)). Amelino-Camelia et al. (1998) originally estimated a sensitivity factor of 0.63 (optimistically rounded up to unity) for a GRB at a distance of 10^{10} lightyears with a typical energy of 20 MeV and variability timescale of 1 ms. *Fermi*-LAT observations of GRB090510 (Abdo et al., 2009) at a distance of $z = 0.9$ contained a 31 GeV photon and a variability timescale of about 0.1

sec, yielding a sensitivity $\eta \approx 7$ to linear energy dispersion. GRB090510 has been used to set the best limits on both linear and quadratic LIV to date (Vasileiou et al., 2013).

Shortly afterwards, flares from active galactic nuclei (AGN) were proposed for LIV investigations by Biller et al. (1999) using a previously observed flare of Markarian 421 (Punch et al., 1992). AGN are also extragalactic sources, and their flares can occur on extremely rapid timescales (Gaidos et al., 1996). The observations of Markarian 421 used by Biller et al. had an average energy of ∼2 TeV and a variability timescale of 280 seconds. Assuming a distance to Markarian 421 of 1.1×10^{16} lightseconds, the observation yielded a sensitivity to linear LIV of $\eta \approx 0.01$. More recent observations have observed much faster variability timescales from more distant AGN, improving the sensitivity by an order of magnitude. These sources will be the primary topic of this dissertation and will be discussed in more detail in Chapter 2.

The final class of sources for use in probing LIV is pulsars. First proposed shortly after the other two classes of sources by Kaaret (1999), the Crab pulsar has historically been the most frequently used pulsar for LIV tests because of its brightness and its short 33 ms period, despite the fact that its relative proximity as a Galactic source (about 2 kpc or 6500 lightyears) is not ideal when searching for LIV. Kaaret (1999) used observations of the Crab pulsar taken by EGRET that had a median energy of 2.93 GeV in the highest energy band. More recent gamma-ray observations have revealed that the Crab pulsar emits pulsed radiation above 100 GeV (Aliu et al., 2011). Current sensitivity to linear LIV using the Crab pulsar is on the order of $\eta \sim 10^{-3}$. However, unlike bursting sources, pulsars can be observed continuously. Additional observations reduce the uncertainty on the locations

of the pulsar peaks, improving LIV sensitivity as the source is observed longer (Otte, 2011).

One obvious problem with the energy dependent dispersion/time-of-flight method is that it relies heavily on the assumption that the emission in two different energy bands occurs simultaneously. Without this assumption, the measurement of a time delay would more likely imply a source-intrinsic delay rather than the more exotic explanation of Lorentz symmetry breaking. Moreover, should a null result of zero delay be measured, the possibility remains of a "cosmic conspiracy" whereby the delay induced by LIV could be exactly canceled by an offset in emission times at the source (Biller et al., 1999). For AGN and GRB sources, the remedy to this issue is to survey a variety of sources at a variety of redshifts. A "cosmic conspiracy" would be highly unlikely given null results from a number of sources, and a time delay that increased as a function of increasing distance would provide strong evidence for LIV. For pulsars, a simpler method of disentangling source-intrinsic and LIV effects has been suggested by Otte (2011).

Table 1.1. List of Lorentz invariance violation papers and their lower limits on the quantum-gravity energy scale, E_{QG}. Papers before 2008 generally did not take into account cosmological expansion as pointed out by Jacob & Piran (2008), yielding an incorrect LIV threshold. The published limits are presented here; see Table 1 of Abramowski et al. (2011) for a similar table containing corrections to many of the early results.

Year	Target	Experiment	Linear E_{QG} (GeV)	Quad. E_{QG} (GeV)	Publication
1999	Mrk 421	Whipple	6.0×10^{16}		Biller et al. (1999)
1999	Crab Pulsar	*CGRO*-EGRET	1.8×10^{15}		Kaaret (1999)
1999	GRB 930131	*CGRO*-BATSE	8.3×10^{16}		Schaefer (1999)
2000	5 GRBs	*CGRO*-OSSE	$\sim 10^{15}$		Ellis et al. (2000)
2003	9 GRBs	*CGRO*-BATSE	6.9×10^{15}	2.9×10^{6}	Ellis et al. (2003)
2003	Crab Pulsar, Mrk 421 & 501	Various	3.66×10^{18}		Stecker (2003)
2004	GRB 021206	*RHESSI*	1.8×10^{17}	5.5×10^{7}	Boggs et al. (2004)
2006	15 GRBs	*Swift*-BAT, Konus-Wind	8.1×10^{16}	6.2×10^{6}	Martínez & Piran (2006)
2008	11 GRBs	*INTEGRAL*	4×10^{11}		Lamon et al. (2008)
2008	15 GRBs	*HETE-2*	2×10^{15}		Bolmont et al. (2008)
2008	35 GRBs	Various	1.4×10^{16}		Ellis et al. (2006); Ellis et al. (2008)
2008	Mrk 501	MAGIC	2.1×10^{17}	2.6×10^{10}	Albert et al. (2008)
2008	PKS 2155-304	H.E.S.S.	7.2×10^{17}	2.6×10^{10}	Aharonian et al. (2008)
2008	Mrk 501	MAGIC	3.0×10^{17}	5.7×10^{10}	Martínez & Errando (2009)
2009	GRB 080916C	*Fermi*	1.3×10^{18}		Abdo et al. (2009)
2009	GRB 090510	*Fermi*	1.49×10^{19}		Abdo et al. (2009)
2011	PKS 2155-304	H.E.S.S.	2.1×10^{18}	6.4×10^{10}	Abramowski et al. (2011)
2011	Crab Pulsar	VERITAS	3×10^{17}	7×10^{9}	Otte (2011)
2013	GRB 090510	*Fermi*-LAT	9.3×10^{19}	1.3×10^{11}	Vasileiou et al. (2013)
2015	PG 1553+113	H.E.S.S.	4.1×10^{17}	2.1×10^{10}	Abramowski et al. (2015)

CHAPTER 2

BLAZARS

2.1 Discovery and Identification

Blazars are some of the most energetic and rapidly variable objects in the universe. The prototype object, BL Lacertae (BL Lac hereafter), was originally identified via optical observations in the 1920's as a Galactic variable star in the constellation of Lacerta (Hoffmeister, 1929). BL Lac was known to be highly variable in the visual band, with an apparent magnitude varying from 13 to 16 (corresponding to factor of 16 variation in flux) within a period of a few days (Semakin, 1954). It was later identified as the optical counterpart to the radio source VRO 42.22.01 (Schmitt, 1968). VRO 42.22.01 was already known to be an unusual radio source: its linearly polarized emission suggested non-thermal radio emission, and its Faraday rotation hinted that it might be extragalactic (MacLeod & Andrew, 1968).

The properties of BL Lac indicated that it was likely a quasi-stellar object (QSO or quasar), similar to the Third Cambridge Catalog radio sources 3C 48 (Matthews & Sandage, 1963) and 3C 273 (Oke, 1963; Schmidt, 1963) which had only recently been discovered. The construction of large radio telescopes and the use of radio interferometry in the 1950-1960's provided the necessary angular resolution to associate radio sources with optical counterparts. It was then quickly realized that many of the brightest radio sources had optical counterparts that were unresolved, like stars, but with polarized (non-thermal) emission, variable fluxes, and spectra that were substantially different from those of stars. Moreover, redshift measurements of 3C 48 (Greenstein & Matthews, 1963) and 3C 273 (Schmidt, 1963) indicated that these objects were either extremely dense stars with large gravitational

redshifts, or, more likely, extremely luminous extragalactic objects.

However, further optical observations of BL Lac by (DuPuy et al., 1969) revealed that the (U-B) and (B-V) optical colors of BL Lac differentiated it from other types of QSOs. Furthermore, DuPuy et al. (1969) found that BL Lac possessed a continuous optical spectrum devoid of emission or absorption lines, which again differentiated it from other types of quasars whose spectra contained emission lines (Oke, 1967). These results were extended by Visvanathan (1969), who showed that the optical continuum was also strongly polarized and very likely a continuation of the polarized radio continuum. This was strong evidence that non-thermal synchrotron emission was responsible for both the radio and optical emission. The compiled evidence from radio, infrared, and optical measurements by MacLeod et al. (1971) cemented BL Lac as the archetype of a new class of astronomical object, and many BL Lac objects were discovered soon thereafter (Strittmatter et al., 1972). Nonetheless, an understanding of the nature of BL Lac objects remained elusive.

One of the major problems in determining the nature of BL Lac objects was the lack of spectral lines, which are generally necessary for obtaining a redshift. For example, Shapiro & Elliot (1974) speculated that BL Lac objects were not QSOs, but were instead nearby stellar-mass black holes in the Galactic disk. A breakthrough came when (Oke & Gunn, 1974) reported a redshift for BL Lac of $z = 0.07$ ($\sim$350 Mpc), which is still seen as accurate today. Oke & Gunn realized that most of the emission from BL Lac came from the nucleus of a typical galaxy. By sheer luck, Oke & Gunn were able to observe BL Lac's host galaxy during a period when the central emitter was particularly dim, and thus obtained the redshift of the host galaxy. This result confirmed earlier suspicions regarding

the extragalactic nature of BL Lac and indicated that the unusual emission from BL Lac was emitted by a singular object at the center of a typical galaxy (i.e., a galactic nucleus).

2.2 Blazar Jets

A revolution in our understanding of BL Lac objects occurred with the introduction of the concept of relativistic jets. Jet-like structures had been known to exist in association with M87 for quite some time (Baade & Minkowski, 1954; Curtis, 1918; Rees, 1978), but it wasn't until the seminal papers by (Blandford & Rees, 1974, 1978) and the more thorough treatment by Blandford & Königl (1979) that these jets were associated with relativistic outflows from active galactic nuclei (AGN). A major motivation for the introduction of relativistic jets was that the synchrotron brightness temperatures of some luminous objects were found to be on the order of 10^{15} K. It had been previously predicted that when the synchrotron brightness temperature exceeded 10^{12} K, rapid cooling would occur via inverse-Compton scattering—the so-called "inverse-Compton catastrophe" (Kellermann & Pauliny-Toth, 1969).

This embarrassing state of affairs was avoided by means of special relativistic Doppler beaming, which can make a source appear several orders of magnitude brighter when the relativistic jet is aligned with the observer's line of sight. The relativistic Doppler factor is

$$\delta = \frac{1}{\Gamma\left(1 - \beta\cos\theta\right)}, \tag{2.1}$$

where $\beta = v/c$, $\Gamma = \left(1-\beta^2\right)^{-1/2}$, and θ is the angle formed between the jet's velocity vector and the line of sight to the observer. Relativistic beaming can increase the observed gamma-ray luminosity from the true gamma-ray luminosity by a factor of δ^4 (Urry & Padovani, 1995), with δ typically in the range of $3 \leq \delta \leq 6$ for gamma-ray quasars (Dondi & Ghisellini,

1995) and $7 \leq \delta \leq 38$ for TeV blazars (Tavecchio et al., 1998). This massive enhancement of the actual source flux provides a convenient explanation for the large flux variability generally observed in blazars, since a small change in the source luminosity is amplified by a large amount. Relativistic beaming also provides a convenient solution to cases of apparent superluminal motion (Rees, 1966), and explains why counter-jets are often not observed in bright quasars (Bridle & Perley, 1984).

The structure and contents of AGN jets are still the subject of active research, but one model that has emerged is shown schematically in Figure 2.2. High resolution images of jets typically show a broad outflow region narrowing into an approximately cylindrical jet extending out from the AGN core (Marscher et al., 2008), sometimes out to several megaparsecs (Bagchi et al., 2014). This confinement may be due to a helical magnetic field (Vlahakis & Königl, 2004), which can accelerate and confine particles within the jet. This helical magnetic confinement field can be generated by a magnetic field threading the rotating

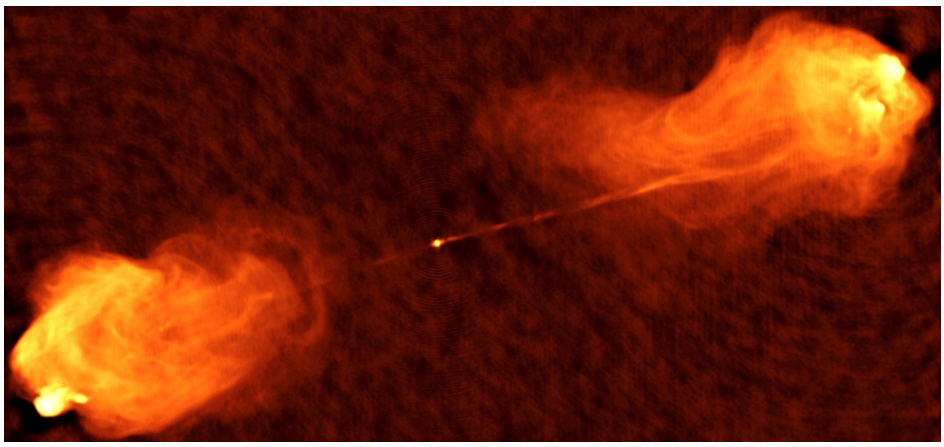

Figure 2.1. A 6-cm radio image of the radio galaxy Cygnus A. The galaxy is visible as a bright spot at the center of the image. Two well-collimated relativistic jets can be seen emerging from either side of Cygnus A. The interaction of the jets with the surrounding intergalactic medium generates large plumes of hot plasma called radio lobes. This image originally appeared in Perley et al. (1984) and is courtesy of NRAO/AUI.

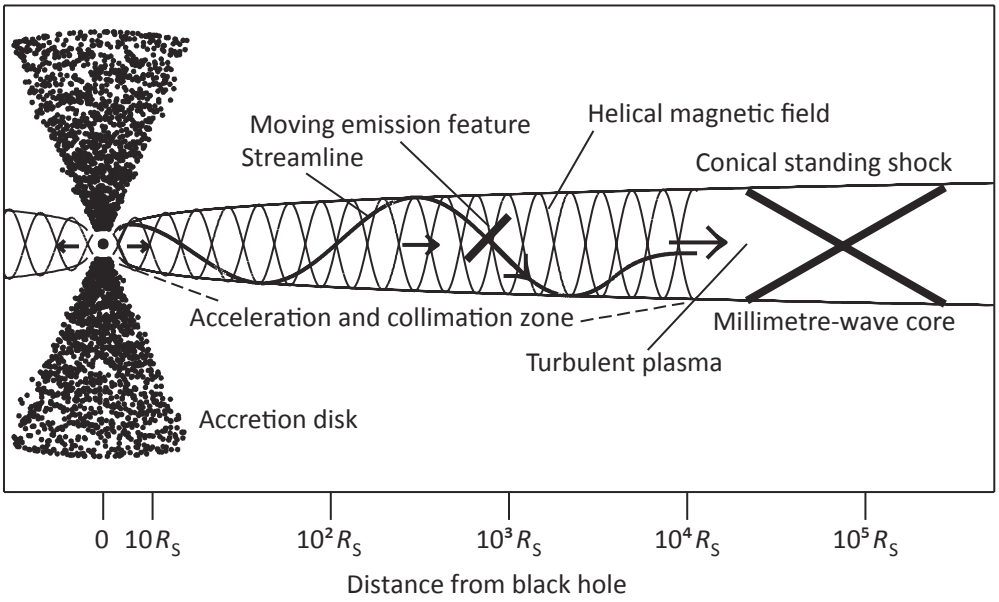

Figure 2.2. A proposed model of the inner structure of a blazar jet. The jets originate close to the black hole (shown at the left) normal to the plane of the accretion disk. A shock propagates down the jet via a helical streamline, while a helical magnetic field confines and collimates the jet. Blobs or knots of material move along the helical streamline at relativistic speeds and generate beamed emission. A standing shock wave and turbulent region may exist at the end of the jet. Image taken from Marscher et al. (2008).

accretion disk (Blandford, 1976; Blandford & Payne, 1982; Heinz & Begelman, 2000) or the black hole itself, if it is rotating (Blandford & Znajek, 1977). It is thought that the jets originate $\lesssim$ 10 Schwarzschild radii (R_s) from the black hole, and possibly as close as $5R_s$ (Doeleman et al., 2012).

One of the most important features for gamma-ray emission is the presence of relativistic blobs or knots of material moving down the jet axis (Ghisellini & Madau, 1996; Marscher et al., 2010). These blobs of material are considered compact on the scale of AGN, but are thought to be on the order of 5×10^{14} cm $\approx$ 30 AU (Tavecchio et al., 2011). Many AGN models consider the volume inside of the broad line region (BLR; see Figure 2.3) to be optically thick to gamma-rays due to the dense IR-UV photon field emitted by the accretion

disk (Jelley, 1966). The cross section for photon-photon interaction, $\gamma + \gamma \rightarrow e^- + e^+$, is

$$\sigma_{\gamma\gamma}(E, \epsilon, \phi) = \frac{1-\beta^2}{8}\left[2\beta(\beta^2-2) + (3-\beta^4)\ln\left(\frac{1+\beta}{1-\beta}\right)\right] \text{ barns}, \tag{2.2}$$

with $\beta(E, \epsilon, \phi) \equiv [1 - 2m_e^2/E\epsilon(1-\cos\phi)]^{1/2}$ (not to be confused with the relativistic $\beta = v/c$ used earlier), where E is the energy of the gamma-ray, ϵ is the energy of the low-energy photon, and ϕ is the angle formed between the two photon trajectories (Heitler, 1954). Taking the derivative of Equation (2.2) and equating it with zero, we find that a head on collision between two photons has the largest cross section when

$$\epsilon \sim \frac{2(mc^2)^2}{E} = 0.5\left(\frac{1\text{TeV}}{E}\right) \text{eV}, \tag{2.3}$$

i.e., a 1 TeV gamma-ray will most likely interact with 0.5 eV near infrared (NIR) photons. Thus gamma rays with energies greater than about 10 GeV are likely to be absorbed inside the BLR. Therefore the observed gamma rays likely originate outside of the BLR.

2.3 Blazars as Active Galactic Nuclei

It was later discovered that an optically-thick dusty torus generally surrounds the black hole and accretion disk of an AGN, and that some of the features of apparently different classes of AGN (beginning with Type 1 and 2 Seyfert galaxies) could be explained as the result of the viewing angle (Krolik & Begelman, 1988). The picture of AGN that would emerge as a result of further unification efforts (Antonucci, 1993; Urry & Padovani, 1995) is shown in Figure 2.3 and is the most widely accepted conception of AGN structure today. We thus find that blazars are not a unique type of AGN, but are instead AGN with active jets that are viewed along the jet axis (or nearly so). Jets are normally present in actively

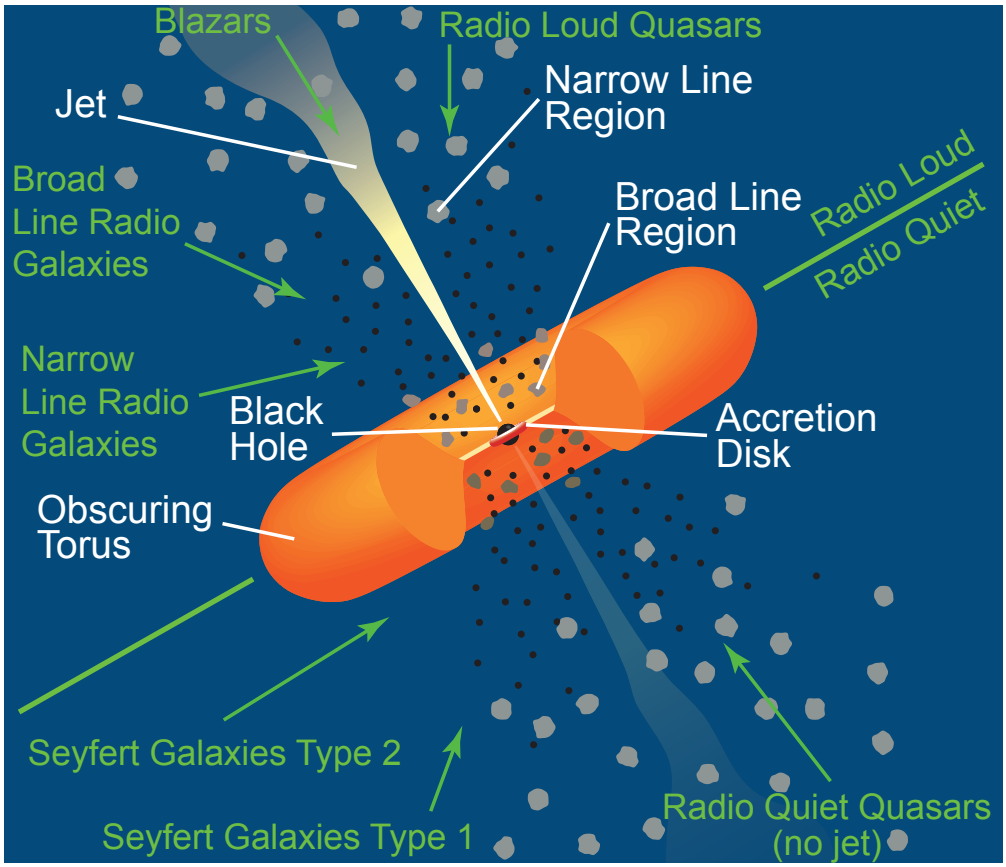

Figure 2.3. A unified picture of active galactic nuclei (AGN) as presented by Urry & Padovani (1995). Features of the AGN, including the central black hole, are labeled in white. The green labels indicate types of AGN that were initially believed to be distinct types of astronomical objects, but were later realized to be the same object viewed at a different angle (indicated by the green arrows) or in a different state (radio quiet and Seyfert galaxies are AGN without jets).

accreting AGN since the accretion is thought to power the jet (Allen et al., 2006; Blandford & Payne, 1982), although the spin of the black hole is also likely to play a role (Garofalo et al., 2010).

Although only a small fraction of AGN are oriented so that their jet points in our direction, their enhanced brightness as a result of Doppler beaming means that they are detected more often than other types of AGN in a flux-limited survey. As a result, several thousand blazars have been observed (Massaro et al., 2009).

2.4 Classification of Blazars by Broadband Emission Characteristics

As with the rest of astronomy, our understanding of blazars has evolved dramatically as new electromagnetic windows have been opened for exploration by the creation of new technologies and observatories. Blazars are not easily identifiable using optical measurements alone, which is why, for instance, BL Lacertae was originally mistaken for a variable star. As with BL Lacertae itself, radio observations are often essential in locating and classifying blazars. X-ray astronomy later proved to be another powerful tool to study blazars, and would again revolutionize our understanding of blazars.

BL Lacs appear as bright and highly-variable sources in X-rays (Mushotzky et al., 1978) so that BL Lacs can be selected using X-ray measurements only, independent of radio measurements. This led to a surface-level distinction between radio-selected and X-ray selected BL Lacs, differing only in their initial method of discovery. Surprisingly, it was later found that the spectral energy distributions (SEDs) of X-ray selected BL Lacs were intrinsically different than those of radio-selected BL Lacs, although with limited spectral coverage the reasons for this were initially unclear. Later multi-wavelength campaigns would measure the broadband SEDs of BL Lacs (Abdo et al., 2011) and would reveal a characteristic double-peaked spectrum (as in Figure 2.4) with either very weak or absent emission features.

Early radio and optical measurements (particularly of polarization) indicated that the low energy SED peak results from synchrotron radiation emitted by energetic electrons (Oke et al., 1969; Visvanathan, 1969). The synchrotron emission can extend through UV and X-rays (Ledden et al., 1981; Urry et al., 1982; Worrall et al., 1981) up to soft gamma-ray energies. The high energy peak is thought to mainly result from inverse Compton scattering

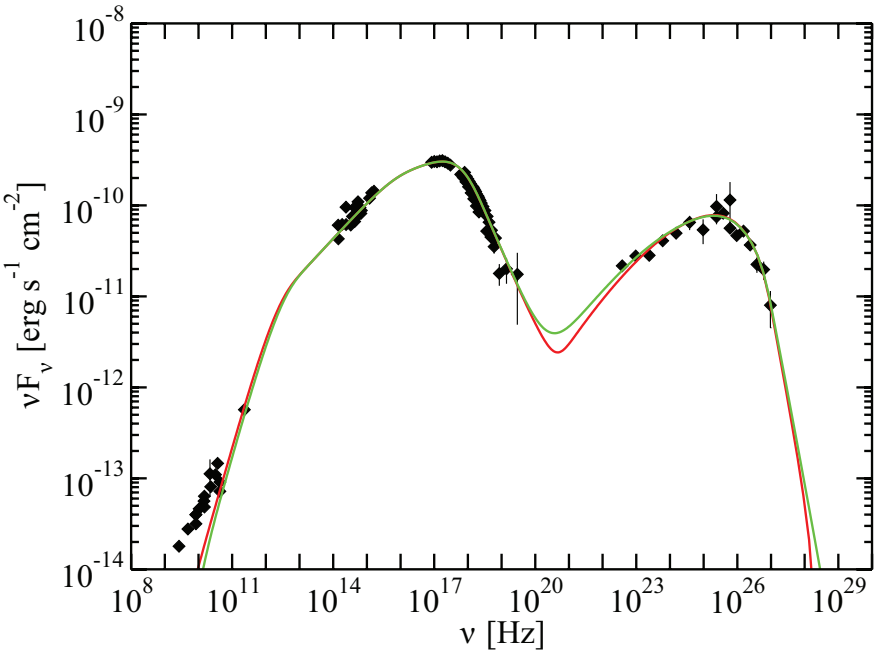

Figure 2.4. Multi-wavelength spectral energy distribution (SED) of Markarian 421. The SED extends from radio up to very high energies ($E > 100$ GeV), exhibiting the characteristic double-bump spectrum of blazars. The two continuous lines are synchrotron self-Compton (SSC) fits for time variability scales of 1 day (red curve) and 1 hour (green curve). Figure taken from Abdo et al. (2011).

of low-energy photons off of the population of high-energy electrons which produces the synchrotron radiation (Bloom & Marscher, 1993; Maraschi et al., 1992). Since the same population of electrons produces both peaks, the locations of the two peaks are related and move together in frequency space. Two scenarios are typically discussed which involve different populations of seed photons: synchrotron self-Compton (SSC), in which synchrotron photons are inverse Compton up-scattered by the same electrons that emitted them (Bloom & Marscher, 1993; Dermer & Schlickeiser, 1993; Jones et al., 1974; Maraschi et al., 1992); and external Compton (EC), which up-scatters photons produced by the accretion disk and possibly scattered or otherwise reprocessed by the broad line region (Ghisellini & Madau, 1996; Sikora et al., 1994). We note that hadronic emission (Mannheim, 1993; Mannheim &

Biermann, 1992) is also possible, but its consideration is outside the scope of our discussion.

The apparent difference between radio-selected and X-ray selected BL Lacs is now known to result from the location of the lower-energy synchrotron peak in frequency space: radio-selected BL Lacs typically exhibit a synchrotron peak in the IR-Optical bands, whereas X-ray selected BL Lacs typically peak in the UV or X-ray bands. This led to a reclassification by Padovani & Giommi (1995) of radio-selected BL Lacs as low-frequency-peaked BL Lacs (LBLs) and X-ray selected BL Lacs as high-frequency-peaked BL Lacs (HBLs) based on the location of the synchrotron peak in the SED. An additional class of intermediate-peaked BL Lacs (IBLs), which, as the name suggests, fall somewhere between LBLs and HBLs, was added later (Bondi et al., 2001).

In addition to the different types of BL Lac objecs, there are two general classes of blazars (Fossati et al., 1998): BL Lacs (LBLs, IBLs, HBLs), and flat-spectrum radio quasars (FSRQs). FSRQs, also known as optically violent variables (Cannon et al., 1971) and highly polarized quasars (Giommi et al., 1990), were identified by Angel & Stockman (1980) as blazars based on their polarization, optical variability, and similar spectral shapes. However, with the increasing availability of multi-wavelength data, it was realized that two classes of blazars have spectra that are shifted in frequency (Gear et al., 1994). Moreover, FSRQs typically have higher radio to optical polarization than BL Lacs, higher redshifts, more prominent gamma-ray peaks, and are more likely to exhibit emission lines (Massaro et al., 2011). The SEDs of FSRQs and LBLs are largely the same (Sambruna et al. (1996); Urry (1998), as in Figure 2.5).

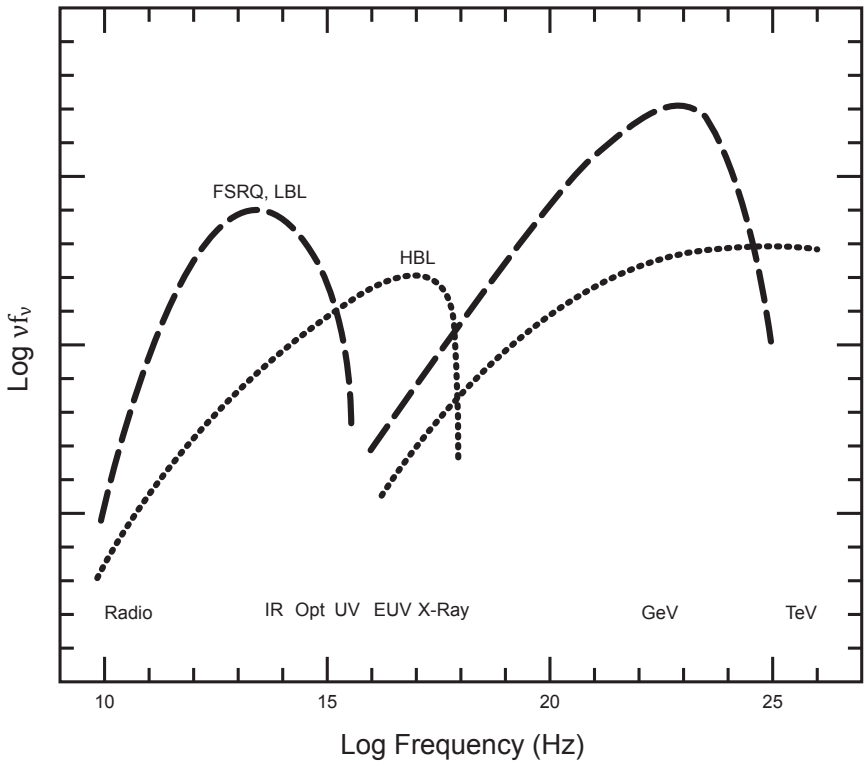

Figure 2.5. Example of typical multi-wavelength blazar spectral energy distributions (SEDs), extending from radio up to TeV gamma ray energies. The dashed line shows a characteristic SED for FSRQs and LBLs, while the dotted line shows a characteristic SED for HBLs. The low energy peak is mostly due to synchrotron radiation, while the high energy peak is thought to be the result of inverse Compton scattering. Figure taken from Urry (1998).

2.5 VHE Properties of Blazars

Although gamma-ray astronomy has been in development since the 1960's – see Weekes (1988) for a good review – it wasn't until the launch of the European gamma-ray satellite COS-B that a blazar (FSRQ) was detected in gamma rays (Swanenburg et al., 1978). The Challenger disaster in 1986 and the subsequent moratorium on U.S. space missions and the failure of the Gamma-1 telescope on the Soviet Gamma satellite resulted in little progress during the 1980's (Weekes, 1992). The 1990's saw the launch of the *Compton Gamma Ray Observatory* (CGRO), which operated in the 20 keV – 30 GeV band, and the

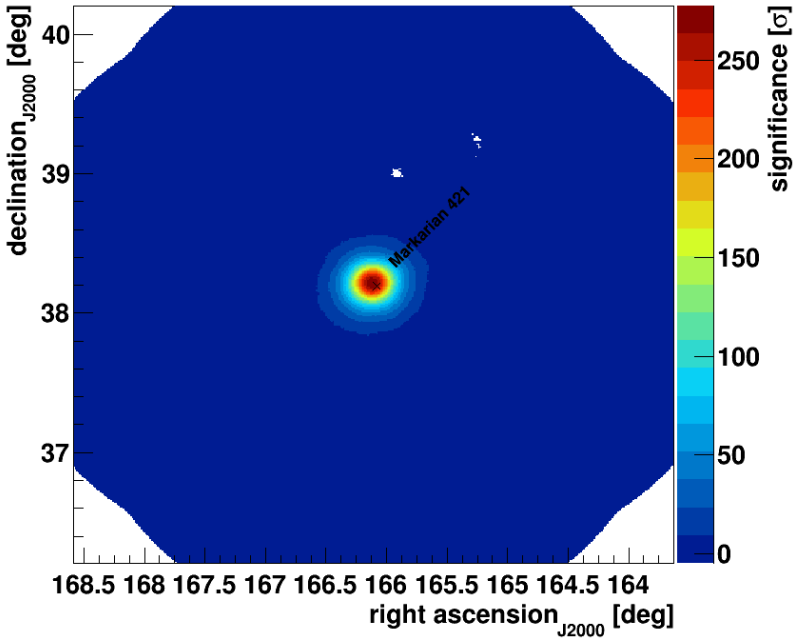

Figure 2.6. VERITAS skymap of Markarian 421. The source is located at the center of the skymap. The color scale shows the Li & Ma significance (Li & Ma, 1983), which is well above the nominal VERITAS 5σ detection threshold. The two irregularly-shaped white holes in the skymap above and to the right of the source are caused by bright stars.

maturation of ground-based imaging atmospheric Cherenkov telescopes (IACTs), which are typically sensitive in the $\sim$100 GeV – 10 TeV regime.

The blazar Markarian 421 was first observed by the Whipple and HEGRA collaborations (Petry et al., 1996; Punch et al., 1992), being only the second TeV source detected (after the Crab) at $> 5\sigma$. The detection of Markarian 421 by CGRO was published shortly thereafter (Lin et al., 1992). An all sky survey taken with CGRO's EGRET instrument later found that most of the extragalactic gamma-ray emitting sources in the sky are blazars. Moreover, it was discovered that most of the power output of blazars is emitted in the form of gamma rays (Fichtel, 1994; Maraschi et al., 1992). As noted in a survey of the development of blazar gamma-ray astronomy by Weekes (2003), had gamma-ray observatories existed earlier, it is likely that blazars would be considered primarily as gamma-ray emitters.

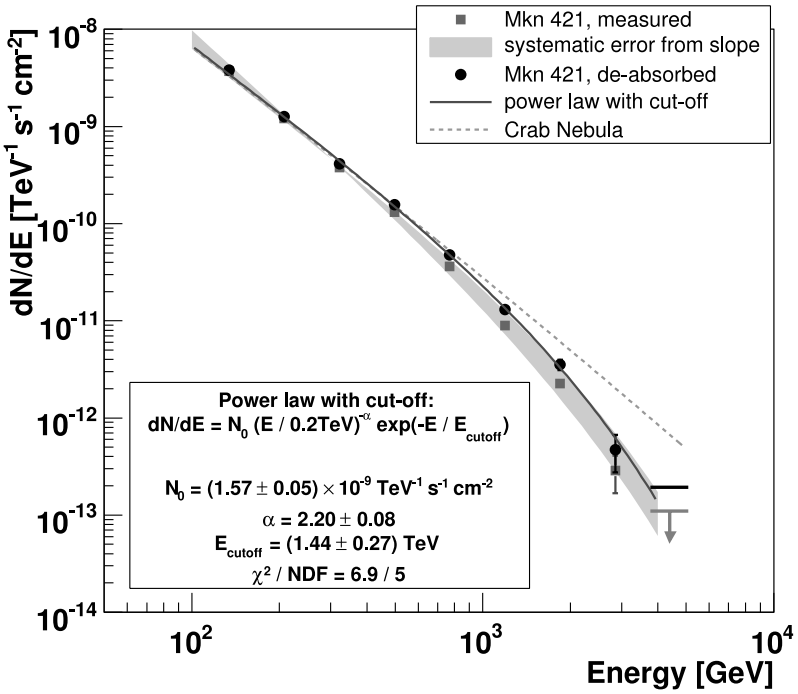

Figure 2.7. Spectrum of Markarian 421 fitted as a power law with an exponential cutoff (thin black line). The gray squares show the measured spectral points, while the black circles have been corrected for EBL absorption. The Crab Nebula's power law spectrum (dashed gray line) is shown for comparison. Image taken from Albert et al. (2007a).

The TeVCat VHE source catalog (`http://tevcat.uchicago.edu`) currently lists 59 VHE-detected blazars, 45 of which are high-frequency peaked BL Lacs (HBLs). TeVCat also contains seven intermediate-peaked BL Lacs (IBLs), one low-frequency peaked BL Lac (LBL), and five flat-spectrum radio quasars (FSRQs), as well as one blazar that has not yet been classified. The preponderance of HBLs and IBLs is a selection effect resulting from the VHE energy requirement ($E > 100$ GeV), since the inverse Compton peaks in the SEDs of these types of blazars are shifted in energy high enough to have a detectable flux in the VHE regime. LBLs and FSRQs have very soft VHE spectra (typically with a spectral index $\alpha \sim 4$), but the enhanced size of their high-frequency SED peaks relative to HBLs means that they can sometimes be detected at VHE energies, particularly when flaring.

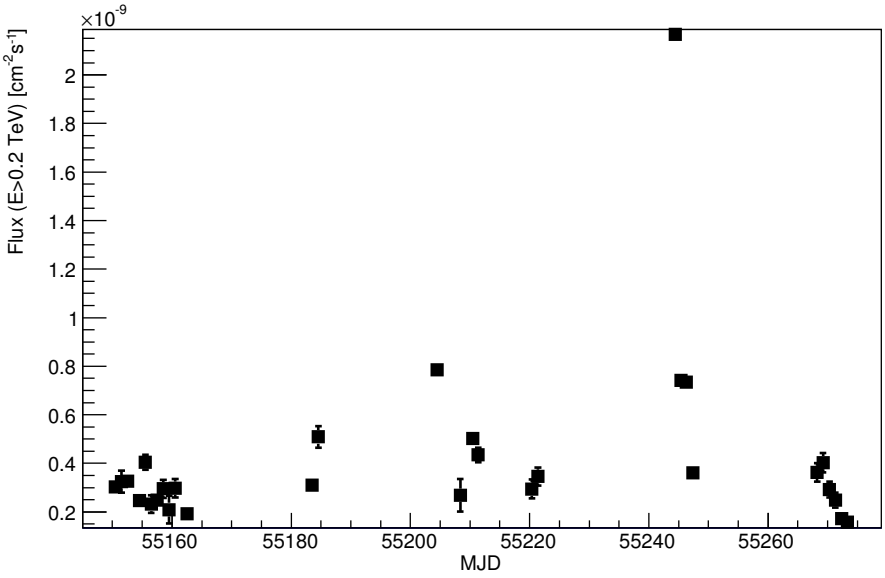

Figure 2.8. VERITAS light curve showing the flux from Markarian 421 above 0.2 TeV during 15 November 2009 – 19 March 2010 (MJD 55150 - 55274). The flaring activity on 17 February 2010 (MJD 55244) can be clearly seen; the flux on that day was approximately seven times the quiescent level.

Blazars appear as bright point sources to VHE instruments (see Figure 2.6), and generally exhibit power law spectra. Since blazars are extragalactic sources, it is generally necessary to correct for absorption by the extragalactic background light (EBL; Gould & Schréder (1967)) by de-absorbing the VHE spectrum using a model of the EBL (Franceschini et al., 2008). Nearby blazars, such as Markarian 421 and Markarian 501, are less affected by EBL absorption. Nonetheless, Markarian 421 has a somewhat unusual exponentially cutoff spectrum (see Figure 2.7), given by

$$\frac{dN}{dE} = \Phi_0 \left(\frac{E}{E_0}\right)^{-\Gamma} \exp\left(-\frac{E}{E_C}\right). \tag{2.4}$$

The exponential cutoff has several possible origins, which are described by Krennrich et al. (2001) and references therein.

Blazars are known for their rapid broadband variability (Urry, 1996) and large flaring episodes (Gaidos et al. (1996); see also Figure 2.8). Variability has been measured down to minute time scales (Aharonian et al., 2007; Albert et al., 2007b), which is smaller than the light crossing time for a supermassive black hole ($t_{lc} \sim 10^3 \left(\frac{M_{BH}}{10^8 M_\odot}\right)$ seconds, about 15 minutes for a $10^8 M_\odot$ black hole). This supports the idea that TeV emission, particularly when it has rapid variability, emanates from within the jet far away from the black hole (Agudo et al., 2011; Vovk & Babić, 2015).

CHAPTER 3

VERITAS AND THE IMAGING ATMOSPHERIC CHERENKOV TECHNIQUE

3.1 Overview of VERITAS

The VERITAS (Very Energetic Radiation Imaging Telescope Array System) observatory is located in southern Arizona at the Fred Lawrence Whipple Observatory. The observatory consists of an array of four imaging atmospheric Cherenkov telescopes (IACTs) which operate in tandem to image extensive air showers in the atmosphere. VERITAS is one of three such observatories currently operating in the world, the others being H.E.S.S. (located in Namibia, in southern Africa) and MAGIC (located on the Canary island of La Palma). A successor to VERITAS, MAGIC, and H.E.S.S. called the Cherenkov Telescope Array (CTA) is currently in development.

VERITAS is designed to image the bursts of Cherenkov light that are emitted from extensive air showers (Sec. 3.2). These showers are very brief (on the order of 10 ns) and dim, so although they occur frequently all across the sky, they cannot be seen with the naked eye. Each VERITAS telescope uses a camera composed of 499 photomultiplier tubes (PMTs) that are sampled by liquid-cooled analog-to-digital converters (Sec. 3.4). VERITAS requires that at least two of its telescopes image a shower (stereoscopic imaging) so that the shower's direction can be reconstructed (Sec. 3.5). The imaging is also necessary to separate gamma-ray showers from hadronic (protons and atomic nuclei) and muon showers, which occur thousands of times more frequently than gamma-ray showers.

The use of the atmosphere as a detecting medium means that VERITAS has a very

Figure 3.1. Photograph of VERITAS's four 12 meter diameter imaging atmospheric Cherenkov telescopes. VERITAS is located at the Fred Lawrence Whipple Observatory at the base of Mount Hopkins in southern Arizona. Image credit: Larry Ciupik.

large effective area of approximately 10^5 square meters. VERITAS can detect cosmic gamma-rays with energies of between $\sim$100 GeV – 50 TeV with an energy resolution of 15 - 20% and an angular resolution of approximately 0.1 degrees. VERITAS's cameras were upgraded in mid-2012 with higher quantum efficiency (35 - 40%) PMTs, which decrease VERITAS's energy threshold and increase its sensitivity. A downside of ground-based gamma-ray observatories is that they can only observe in good weather without clouds, rain, high winds, or lightning. Nonetheless, its location in southern Arizona means that VERITAS can typically observe a total of 700-1000 hours each year, even though a summer shutdown is required during Arizona's monsoon season.

3.2 Extensive Air Showers

Energetic astrophysical gamma-rays that enter the Earth's atmosphere will pair produce in the presence of atomic nuclei (which are needed to absorb momentum). The electron and positron will also interact with atmospheric nuclei, and the acceleration due to the Coulomb force will cause them to emit bremsstrahlung radiation. Since little energy will be

lost to nuclear recoil, the bremsstrahlung photons may be energetic enough to pair produce again. The positrons may also annihilate atmospheric electrons, although the cross section for annihilation is small until the positrons lose most of their energy. The original astrophysical gamma-ray thus loses its energy in an electromagnetic shower of electrons, positrons, and bremsstrahlung photons.

The generation of this cascade of particles, called an extensive air shower (EAS), requires the initial gamma-ray to have an energy exceeding about 80 MeV. A simple model that represents the features of extensive air showers surprisingly well was proposed by Heitler (1954) and is shown schematically in Figure 3.2. Each level is separated by one radiation length $X_0 \approx 37$ g/cm^2, which is both the mean distance over which a high-energy electron loses all but $1/e$ of its energy by bremsstrahlung, and 7/9 of the mean free path for pair production by a high-energy photon (Beringer et al., 2012). Extensive air showers typically occur at a height of $\sim$10 km, and the largest showers can contain in excess of 10^{10} charged particles (Matthews, 2005).

Muons, protons, and atomic nuclei impinging on the atmosphere can also generate

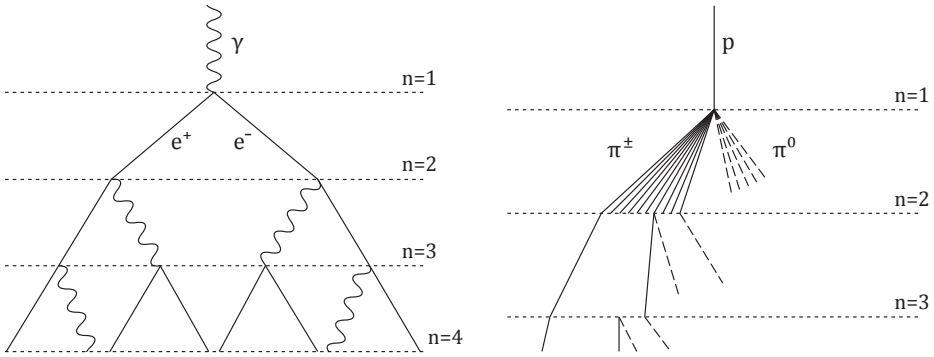

Figure 3.2. Schematic views of a gamma-ray initiated cascade (left) and a hadronic proton shower (right). Not all pion tracks are shown after the n=2 level; neutral pions do not re-interact. Taken from Matthews (2005).

extensive air showers, although the structure and constituents of the showers can be somewhat different. Astronomers are typically interested in photons because they have few (if any) electromagnetic interactions as they propagate from source to observer. However, it is important to understand the nature of showers from other particles, which are typically 10^3 times more numerous than the gamma rays of interest, so that this background can be removed.

Muons decay into electrons or positrons and neutrinos; the electrons and positrons then generate showers in the same manner as in the case of a gamma ray. Unlike gamma rays, however, energetic muons penetrate much deeper into the atmosphere. These lower altitude showers have a characteristic ring-like structure when imaged on the ground, rendering them distinguishable from gamma-ray showers.

Protons interacting with the atmospheric molecules will produce an initial shower of mostly pions, although baryons and kaons can also be produced. The pions will decay into gamma rays, muons, and neutrinos, resulting in a shower with a central hadronic core ringed by electromagnetic showers from the secondary gamma rays and muons. Atomic nuclei (mostly helium) produce large, extended showers of many types of particles. Showers from atomic nuclei are less structured than proton showers and appear highly irregular when imaged from the ground. Examples of these different types of air showers generated by Corsika, a Monte Carlo extensive air shower simulator, can be seen in Figure 3.3.

If the relativistic particles in the extensive air showers have velocities that exceed the speed of light in the atmosphere ($v > c/n$), they will emit Cherenkov radiation (Čerenkov, 1934, 1937). Note that the refractive index of "standard air" (ambient temperature of 15°C,

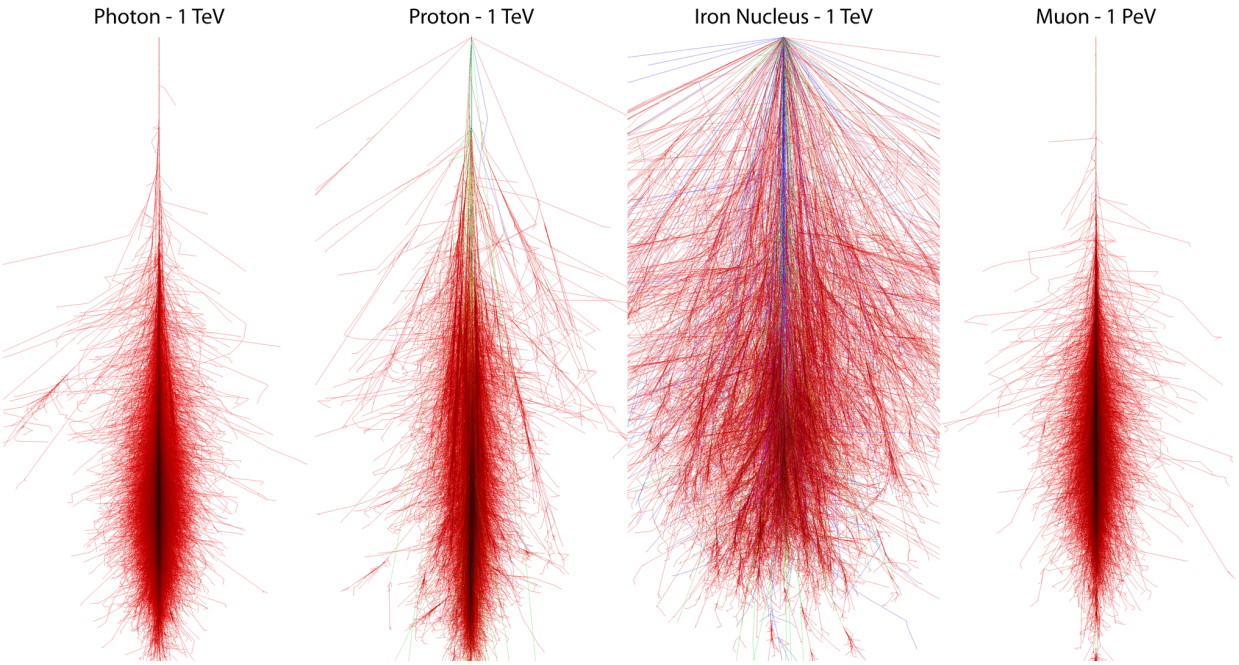

Figure 3.3. Images showing the vertical extent of extensive air showers simulated using CORSIKA. From left to right: photon, proton, iron nucleus, and muon showers. All particles have 1 TeV of kinetic energy except the muon, which has 1000 TeV. Red tracks are electrons, positrons, and gamma rays; green tracks are muons; blue tracks are hadrons. Images courtesy of F. Schmidt, "CORSIKA Shower Images," `http://www.ast.leeds.ac.uk/~fs/showerimages.html`.

pressure of 101.325 kPa, and carbon dioxide content of 0.045% by volume) is $n_{air} \approx 1.000283$ at a vacuum wavelength of 400 nm (Lide, 2005), whereas the refractive index of vacuum is defined to be $n_{vac} \equiv 1$. The radiation will be emitted in a forward cone with opening angle

$$\cos\theta_c = \frac{1}{\beta n}, \tag{3.1}$$

where $\beta = v/c$. For particles energetic enough to produce an air shower, many of the particles in the shower will have $\beta \approx 1$ so that $\theta_c \approx 1.4°$. The opening angle will increase somewhat towards the end of the shower as the energy per particle decreases. From the ground, the air shower can be imaged as a cone of Cherenkov light with an opening angle of about 1.5 degrees, the base of which is approximately 130 meters in diameter for photon air showers.

The wavelength of the Cherenkov light is given by the Frank-Tamm equation (Tamm & Frank, 1937) and has peak emission at UV-blue frequencies. The number density of Cherenkov light emitted by a particle of charge ze per unit path length per unit wavelength is given by

$$\frac{d^2N}{dx\,d\lambda} = \frac{2\pi\alpha z^2}{\lambda^2}\left(1 - \frac{1}{\beta^2\,n(\lambda)^2}\right), \tag{3.2}$$

where $\alpha \approx 1/137$ is the fine structure constant, $\beta \equiv v/c$, and $n(\lambda)$ is the wavelength-dependent index of refraction (see Section 27.7 of Nakamura & Group (2010)). The unabsorbed spectrum is thus proportional to λ^{-2}, cutting off in X-rays where the index of refraction drops below unity (Jelley, 1958). Including atmospheric absorption and scattering, which sharply attenuates short-wavelength radiation, the Cherenkov spectrum peaks near 330 nm.

Finally, we note that while extensive air showers occur fairly frequently in the atmosphere, they are very dim and last only a short time. A gamma-ray with a few TeV of kinetic energy will generate a shower containing on the order of 10^6 secondary particles. However, the Cherenkov light pool (which is the projection of the cone of Cherenkov light onto the ground) for a 1 TeV gamma ray will contain a flux of only 50 photons/m^2 (Weekes, 1988). This flux is integrated over a time of $\sim$10 ns, which is the time interval in which most of the Cherenkov light is emitted. This can be seen clearly in simulations of extensive air showers; a plot of relative photon arrival times vs. distance from the center of the light pool shown in Figure 3.4 shows that most of the photons are emitted within a 150 meter radius in a time span of less than 10 ns.

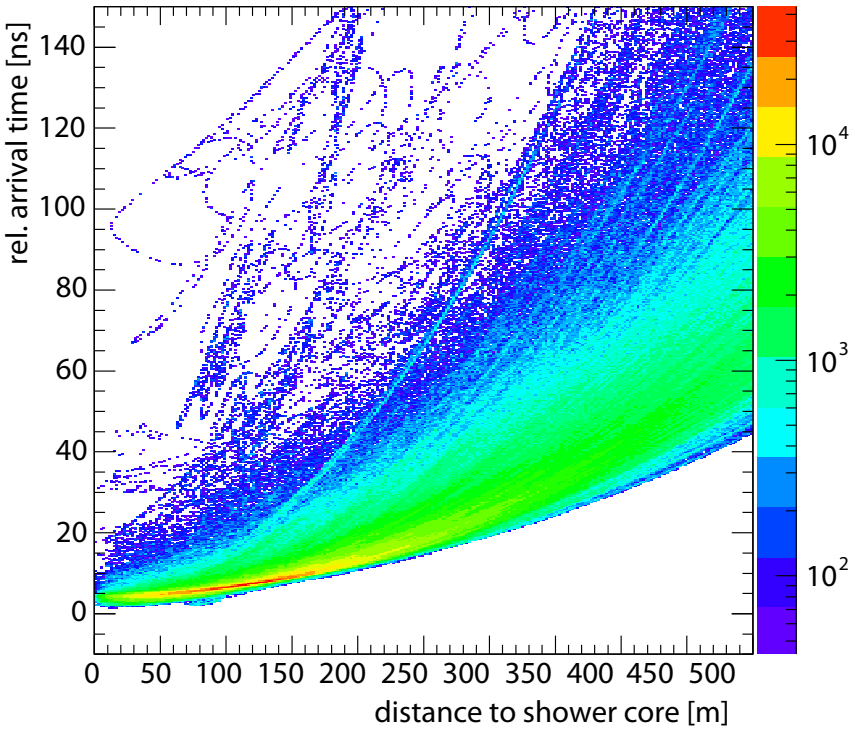

Figure 3.4. Plot of simulated photon arrival times vs. radial distance from the location of the shower core on the ground. Note that most of the Cherenkov light is collected within 10 ns. Image courtesy of Gernot Maier.

3.3 Davies-Cotton Optics

The dearth of Cherenkov photons from extensive air showers requires a large collecting area, which means a large reflecting surface. VERITAS uses 350 identical, hexagonal mirror segments (shown in Figure 3.5), each with an area of 0.322 m^2 (Holder et al., 2006), in a configuration known as a Davies-Cotton reflector. The 12 m diameter reflector has a total area of about 110 m^2. The anodized, aluminized mirrors are attached to a steel optical support structure using three adjustment screws, and can be removed for re-coating in an on-site laboratory so that they maintain their reflectivity.

The Davies-Cotton optical design was originally developed as a solar furnace for the U.S. Army (Davies & Cotton, 1957). Instead of a single, large reflecting surface, Davies-

Figure 3.5. Photo of VERITAS telescope 1 showing the Davies-Cotton reflector.

Cotton optics use an array of identical, concave, spherical mirror segments. These mirror segments are typically mass-produced by slumping glass over a smooth spherical mold and are subsequently aluminized on the front surface. A thin oxide layer (usually produced by anodizing) is typically added to the aluminized mirror surface to protect it from corrosion and wear. In the case of VERITAS, this coating is designed to maximize reflectivity in the UV-blue regime of the electromagnetic spectrum (350-450 nm) to optimally reflect Cherenkov light (see Figure 3.6). The mirror segments are attached to a rigid superstructure which is approximately spherical in shape called the optical support structure (OSS).

The reflective segments are polygonal in shape so that they can be tessellated to cover the entire surface of the OSS. When standing at the focal point, the edges of the mirror

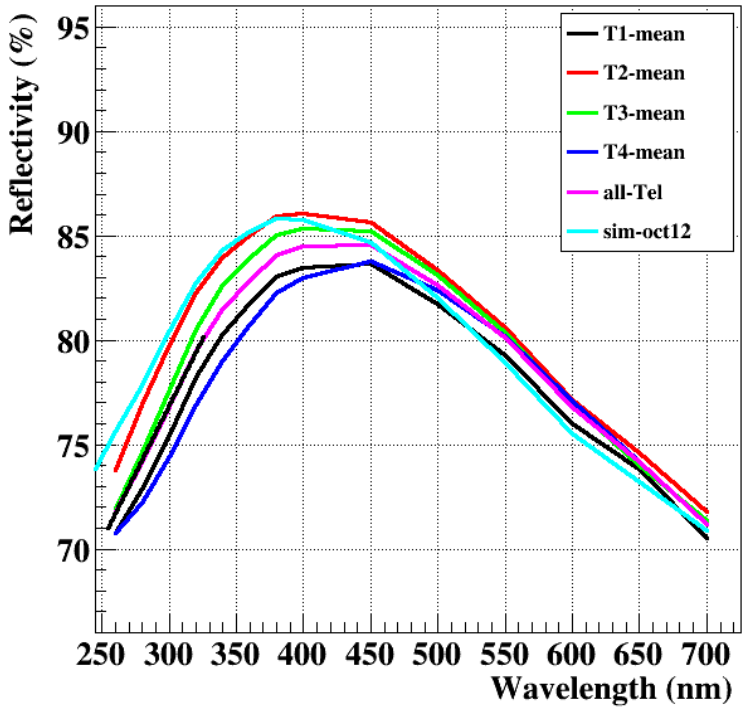

Figure 3.6. Plot of VERITAS mirror reflectivity for each of the four telescopes. Mirrors are removed from each VERITAS telescopes and tested in the on-site mirror lab. The plot shows mean reflectivity values for each of the four telescopes, the mean reflectivity for all of the telescopes, and the telescope reflectivity that is used in internal VERITAS simulations. Image credit: Emmet Roache.

segments are difficult to distinguish and the Davies-Cotton reflector appears as a single large mirror. For VERITAS, ray tracing simulations have shown that only 1.65% of the light striking the reflector is lost in the gaps between the reflectors (Fegan & Vassiliev, 2005). The original design of Davies and Cotton used square mirror segments for simplicity and cost savings, and the overall shape of the reflector was rectangular. In astronomical applications, it is more typical to use hexagonal reflectors that can be tiled into an approximately circular reflector.

The ideal Davies-Cotton reflector geometry is shown in Figure 3.7. An ideal Davies-Cotton reflector has infinitesimally small reflectors. Finite-sized facets introduce "facet aber-

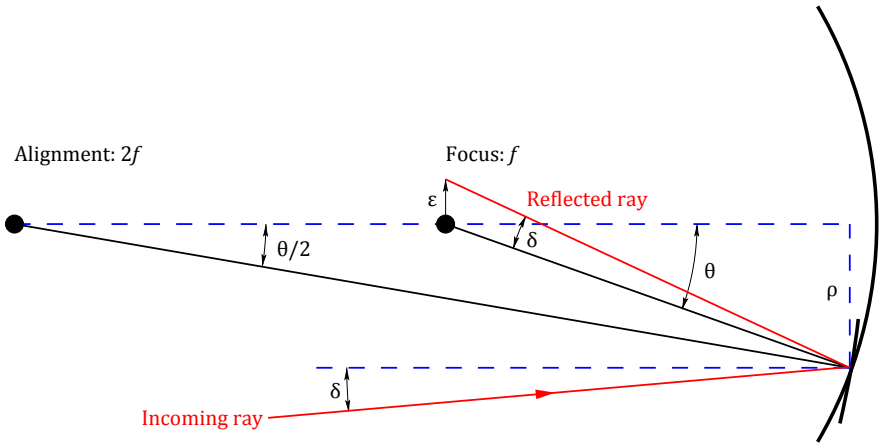

Figure 3.7. Ray diagram of a Davies-Cotton reflector in the ideal limit of an infinitesimally sized reflector. The reflector (shown as a straight line segment) is located at an angle θ from the optical axis (OA). The mirror segments are aligned with their surface normals pointed towards an alignment point located a distance $2f$ along the OA from the center of the reflector. An incident ray parallel to the OA with impact parameter ρ will form an image at the center of the focal plane at f, while a ray that is offset by an angle δ will form an image in the focal plane that is displaced by a distance ϵ. From an internal VERITAS memo (Fegan & Vassiliev, 2005); figure originally adapted from Lewis (1990).

rations" that increase the size of the point-spread function (PSF) both on- and off-axis. However, using larger facets makes the telescope less expensive to produce and assemble and makes the telescope easier to align.

Davies-Cotton reflectors have several properties that make them attractive for use in atmospheric-Cherenkov telescopes. Because they were designed for use as solar furnaces, Davies-Cotton reflectors are inexpensive to produce and are excellent light collectors. When used for atmospheric Cherenkov detectors, the reflector is typically designed as a $f/0.7 - f/1.2$ system (where the f-number is defined as $N = f/D$, with f the focal length and D the diameter of the reflector or lens, and written by convention as f/N). Optical systems with large f-numbers will produce dimmer images than an equal focal length system with a smaller f-number, exposure times being equal, since they cannot collect as much light.

Davies-Cotton reflectors also have better image quality than either spherical or parabolic reflectors (Fazio et al., 1968).

There are two downsides to the Davies-Cotton design. Historically, the most important deficit of the Davies-Cotton design is the fact that it is not isochronous: reflected light from an object does not arrive at the image plane at the same time. For VERITAS, the spread in the arrival times for photons falling on different parts of the reflector is about 5 ns (Krennrich et al., 2004). This time spread is approximately the same as the light collection time interval for an extensive air shower, and in practice it has not had a significant detrimental impact on the instrument. However, as the telescope size increases, the time spread becomes a significant problem, which is why larger telescopes (i.e. MAGIC, HESS-II, CTA-LST) prefer parabolic optics.

The other downside is the appearance of coma in off-axis images. The coma becomes rather severe towards the outer edge of the reflector, which limits the useful field of view to a few degrees (3.5° in the case of VERITAS). This is normally not a problem when observing known sources since even extended Galactic sources are generally less than a degree in extent. A small field of view does, however, cause difficulty when performing survey work, attempting to observe GRBs (which often have large errors in their initial positions, especially when the location comes from *Fermi*-GBM), or when observing very large nebulae. Future telescopes of a commensurate size to VERITAS may instead utilize Schwarzschild-Couder optics (see Figure 6.1 and surrounding discussion), which offers a wider field of view, better angular resolution, and better isochronicity.

3.4 VERITAS Cameras and Data Acquisition

Due to the short time duration ($\sim$10 ns) during which extensive air showers emit Cherenkov light, the VERITAS cameras must be extremely fast. Moreover, they should be optimized to detect blue-UV Cherenkov light. Photomultiplier tubes (PMTs) are extremely fast (rise time of $\sim$ 1 ns) and have reasonably good quantum efficiency (as high as $\sim$ 30% in the blue-UV range). Moreover, their size (about an inch in diameter) is well coupled with the plate scale of the Davies-Cotton optics, such that each PMT has an angular size of about 0.15°. In 2012, VERITAS was upgraded with high quantum efficiency PMTs. The original Photonis XP2970 PMTs, which had a quantum efficiency of 18-22%, were replaced with Hamamatsu R10560-100-20 MOD PMTs with a quantum efficiency of 32-34% (Kieda, 2011).

Each of the four VERITAS telescopes contains a camera consisting of 499 closely-spaced PMTs (see Figure 3.8). Semi-flexible hexagonal light cones (Winston, 1970) are placed in front of the PMTs to increase their light collection area. A custom wide-band preamplifier is physically mated to each PMT, and the whole assembly is encased in a metal enclosure to minimize noise and cross-talk. The preamplifier signals are conducted via 75 ohm coaxial cables to an instrumentation shed located next to each telescope. An approximately 1 kV bias voltage is supplied to each PMT via CAEN A1932AN high-voltage power supplies housed in CAEN SY1527 crates, also located in the instrumentation sheds. The high-voltage bias currents are carefully monitored by a custom FPGA-based current monitor system. In the event of a current spike, which often occurs as a result of bright star light, headlights, border patrol helicopter search lights, lightning, etc., the high-voltage to

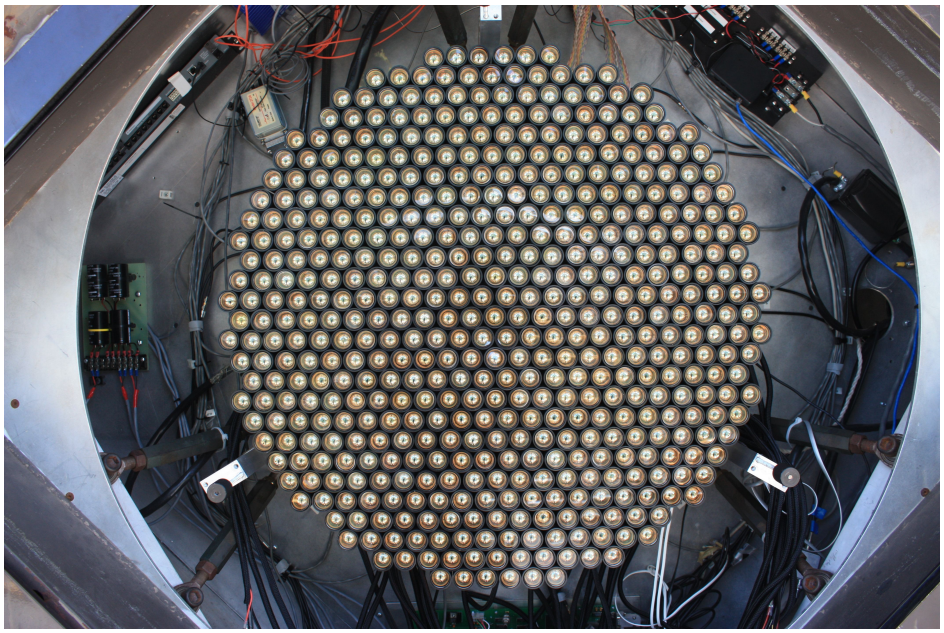

Figure 3.8. Photograph of a VERITAS photomultiplier tube (PMT) camera, which contains 499 closely-packed PMTs. Winston cones – thin hexagonal reflectors that are used to increase light collection area – are not shown in this image, but are usually located in front of the PMTs. The Winston cones create a near-seamless light collection area so that very little light that falls onto the camera is lost. This image was taken by Jamie Holder during the VERITAS camera upgrade in 2012.

the PMTs is immediately shut down to prevent permanent damage to the PMTs.

Bundles of ten coax cables carrying the pre-amplified signal from the PMTs are routed into custom flash analog-to-digital converter (FADC) modules (Buckley et al., 2003). Before readout by the flash ADCs, the signal may be amplified by an additional factor of six by the FADC analog front-end (see Figure 3.9), depending on the amplitude of the pre-amplified signal. The 8-bit FADCs sample the analog signal at a rate of 500 MS/s, buffering in a 32 μs ring buffer until an L3 trigger signal (discussed below) initiates readout. The result is a digital trace like the one shown in Figure 3.10.

Note that a constant, approximately 16 digital count (DC) pedestal signal is injected along with the pre-amplified signal. The night sky background (NSB) light can exhibit

both positive and negative fluctuations about an essentially constant level. Due to the AC coupling of the preamplifier signal, any negative fluctuations would normally be lost. Adding a small positive offset to the signal allows negative fluctuations in the NSB to be measured with a minimal impact on dynamic range. This pedestal level is measured by artificially triggering the L3 trigger at a rate of about 1 Hz in the absence of a real event and flagging the events as pedestals. These events have FADC traces that are essentially flat, and the mean value is taken as the pedestal level. The pedestal level is later removed from event FADC traces by the analysis software.

VERITAS uses three levels of triggering to reduce the amount of stored data to a manageable level (Weinstein, 2008). The first level, L1, triggers at the pixel level. It uses a constant fraction discriminator (CFD) to determine if the pixel has collected a sufficient amount of light (charge) to trigger. CFD triggering is independent of pulse height, and unlike traditional threshold triggering, it does not introduce a pulse height-dependent time delay. The second level of triggering, L2, is a pattern trigger that determines if a cluster of three or more adjacent pixels have triggered (Zitzer, 2013). The third and final level of triggering, L3, is an array-level trigger. The L3 is triggered when two or more telescopes concurrently issue L2 triggers, which indicates that two or more telescopes have imaged an extensive air shower. With this three-tiered triggering system, VERITAS typically triggers at a rate of a few hundred hertz.

Upon triggering, data from the FADC modules is processed by a VME data acquisition (VDAQ) system consisting of five VME crates, which aggressively buffers the data before dispatching it to a dedicated server called the Event Builder. The Event Builder assembles

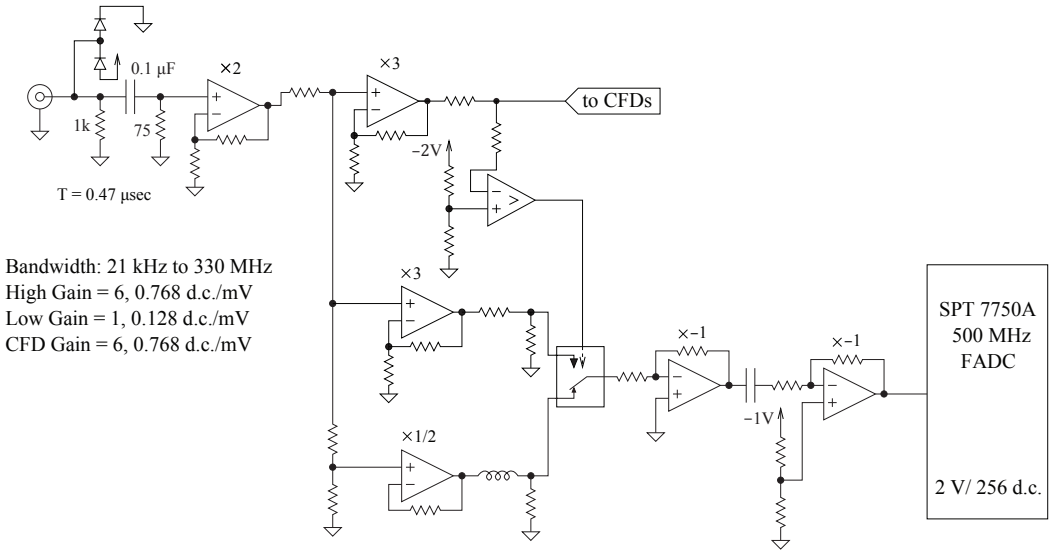

Figure 3.9. Schematic of the FADC analog front-end. Starting from the coax cable on the left, the signal passes through a double-diode protection circuit and a high-pass CR filter and is then amplified by a factor of two. The signal is then amplified by a factor of three and split off into a constant fraction discriminator (CFD; the L1 trigger) and a comparator, which controls an analog switch. The analog switch determines whether a high-gain (x3) or low-gain (1/2) non-inverting amplifier is used. The signal is then buffered before being digitized by a flash ADC. Image credit: Paul Rebillot and Jim Buckley.

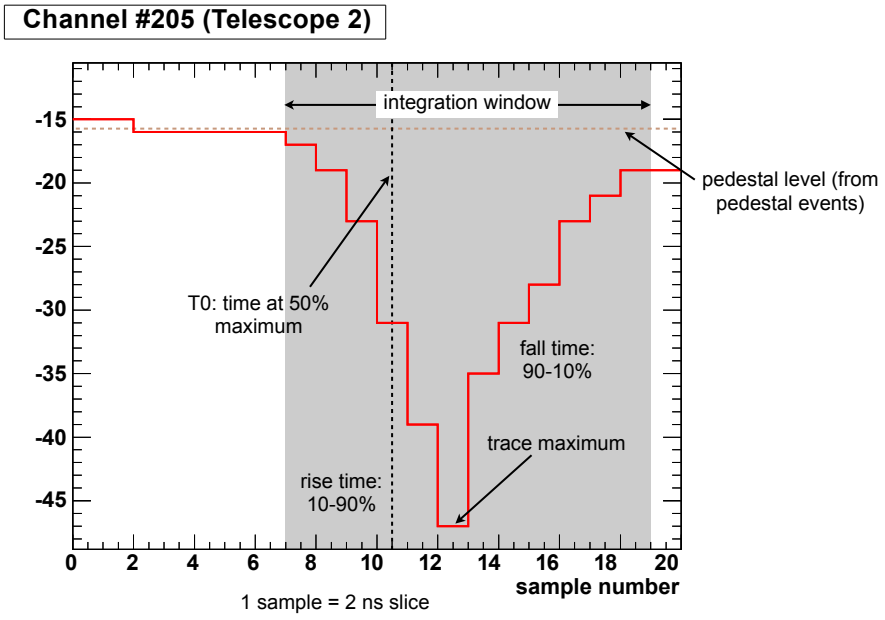

Figure 3.10. Example of an FADC trace showing the location of T_0, the fiducial alignment time of the trace, and the integration window. The vertical axis is in units of digital counts (DC). Since the FADCs sample at a rate of 500 MS/s, each sample represents a 2 ns time slice. Image credit: Gernot Maier.

the data from the VME crates, creating a set of events for that telescope. The data from all four telescopes' Event Builders is collected by a central server called the Harvester, which creates array events and generates the final data products. For more information on the VERITAS data acquisition system, see Hays (2007).

3.5 Shower Imaging and Reconstruction

Images of extensive air showers appear as ellipses when projected onto the focal plane of the telescope. Figure 3.11 shows in detail how an extensive air shower is reflected by the telescope mirrors into the focal plane, where the PMT camera is located. The major axis of the ellipse is aligned with the direction of the air shower. The shower axis, which runs through the center of the shower and is approximately equal to the direction of the progenitor

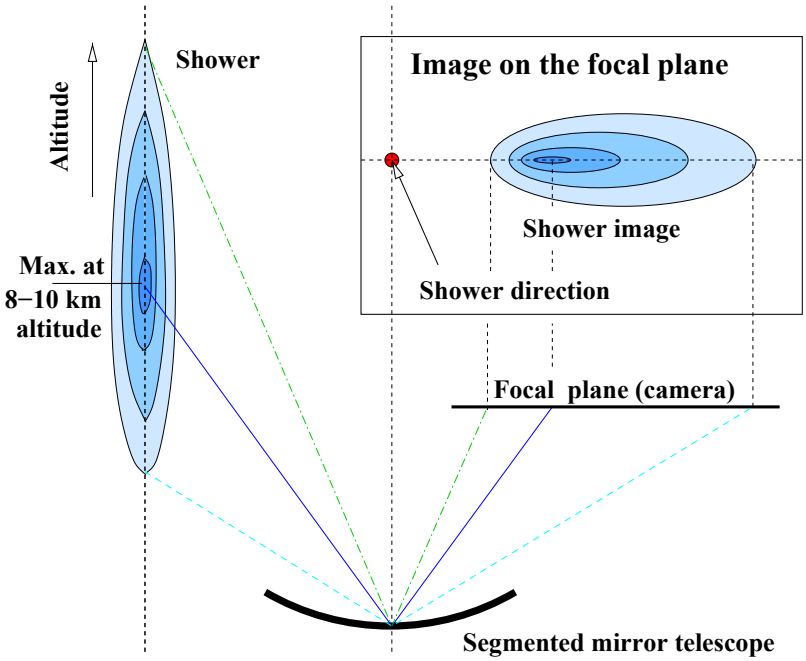

Figure 3.11. Illustration of how an extensive air shower (blue teardrop shape at left) is imaged by ground-based telescope. The image is reflected by the telescope's segmented reflector into the telescope's focal plane (inset box); the projection of the shower onto the focal plane appears as an ellipse. Image taken from Völk & Bernlöhr (2009).

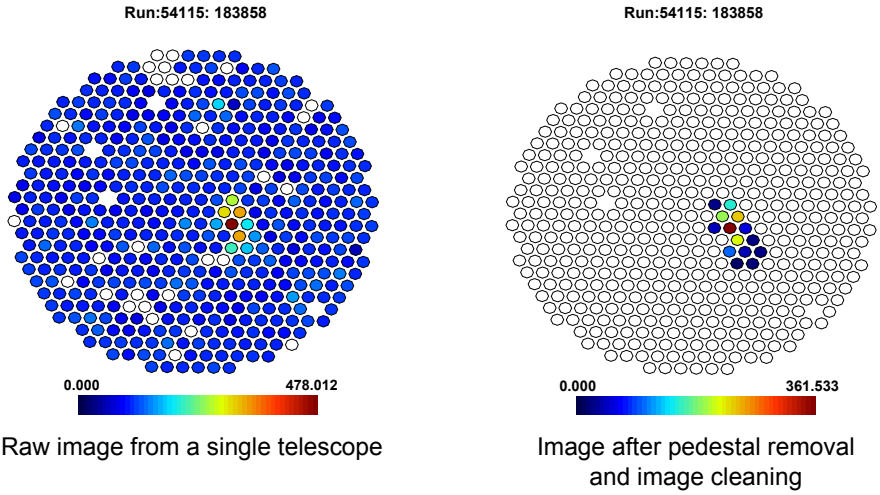

Figure 3.12. Illustration of the image cleaning process. The color scale shows the charge in each PMT (in units of digital counts), which is obtained by integrating the FADC trace (see Figure 3.10). *Left*: Raw image obtained from the VERITAS data acquisition system. *Right*: Cleaned image containing only picture pixels and boundary pixels. Image credit: Brett McArthur.

gamma ray, intersects the line extending from the major axis of the shower (shown as a red dot in Figure 3.11).

The raw image from the data acquisition system, however, bears little resemblance to an ellipse. Before the shower reconstruction can proceed, the pedestals need to be subtracted (recall that the pedestals are artificially added offsets of about 16 digital counts that allow for positive and negative fluctuations in the AC coupled signal to be observed). In addition, pixels which are unrelated to the image of the shower should be removed since they can interfere with the parameterization of the image. This process is called cleaning the image (see Figure 3.12). Several image cleaning methods exist (Benbow & Collaboration, 2005; Bretz, 2005; Daniel, 2008; Shayduk et al., 2005). The method typically used by VERITAS is to first find "picture pixels" with a signal five times the pedestal RMS value and keep those. Pixels adjacent to the picture pixels with a signal 2.5 times the pedestal RMS, called

"boundary pixels," are also kept. All other pixels, including picture pixels with no adjacent picture or boundary pixels, are removed since these are most likely noise.

Once the image is cleaned, the image is parameterized in terms of its centroid and the angle of its major axis. These are used to reconstruct the shower direction and core location as discussed below. In addition, the shape and brightness of the image are parameterized using Hillas parameters (Hillas, 1985). The Hillas parameters are a robust set of analytically computed properties of the image. The original Hillas parameters are shown in Figure 3.13. Most of the original Hillas parameters are still used, especially width, length, distance, and size (not shown), and a few more have been added based on the results of more recent Monte Carlo simulations. Table 3.1 contains a list of the Hillas parameters currently in use by the EventDisplay VERITAS data analysis package.

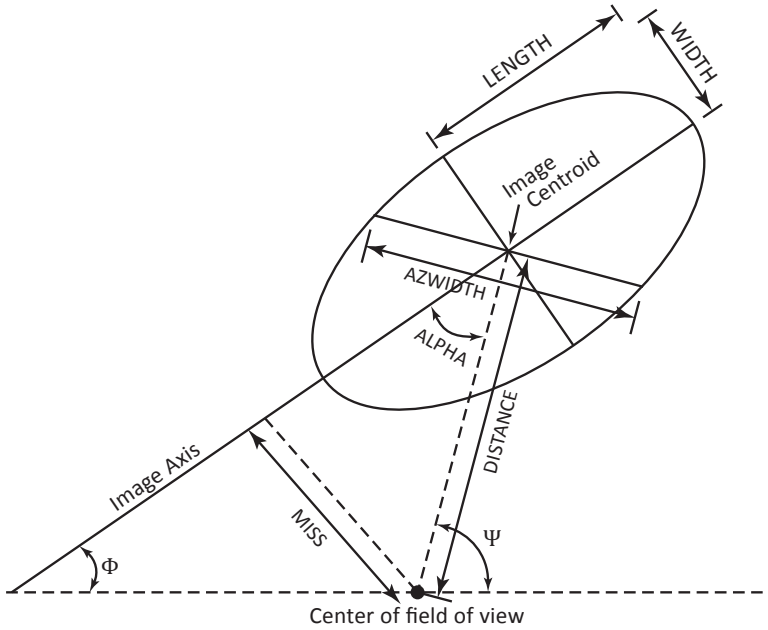

Figure 3.13. Diagram showing the original Hillas parameters (Hillas, 1985). Although image parameterization techniques have since evolved, most of the original Hillas parameters are still used, the idea behind the parameterization of shower images has remained the same. Image taken from Fegan (1997).

Table 3.1. Description of the Hillas parameters currently used in VERITAS analysis.

Parameter	Description
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 center of the field of view of the camera
size	The total integrated light content of the shower; a measure of the energy of the progenitor particle
fui	The fraction of the image size under the image ellipse; can be used to remove less compact images
loss	The fraction of the image size contained in edge pixels; can be used to remove strongly truncated images at the camera edge

The image parameterization is particularly important for separating gamma rays from the cosmic ray background, which we will not discuss it in detail here (see, e.g., Krawczynski et al. (2006) for details). It is also used to determine the energy of the gamma ray (or other particle) that initiated the extensive air shower, which will be discussed in Section 3.6. Finally, we note in passing that it is possible to avoid parameterizing the image using template methods (de Naurois & Rolland, 2009; LeBohec et al., 1998; Parsons & Hinton, 2014), which have been shown to improve sensitivity in many cases, but are generally reserved for specialized analyses.

With the image parameterization in hand, the most common method of reconstructing the shower direction and core location is shown in Figures 3.14 and 3.15. There are essentially four unknown quantities, which are the x, y locations of the points P and S in Figure 3.14. If four telescopes image the shower, the shower direction can be reconstructed completely by using only the image centroids and the locations of the telescopes on the ground. It is also

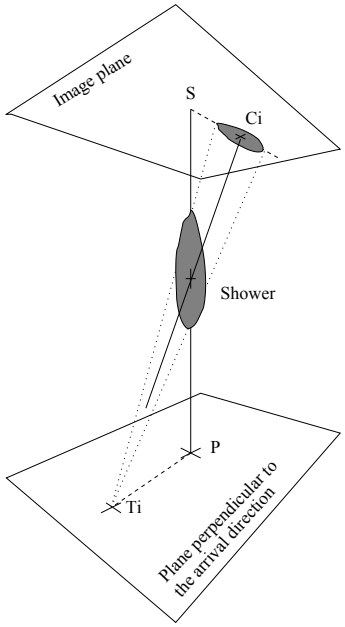

Figure 3.14. Diagram illustrating the geometry of extensive air shower reconstruction. The bottom plane is normal to the shower axis. The point P is the projection of the shower core onto the ground, while T_i is the projected location of telescope i. The image plane contains the image centroid C_i and the (reconstructed) source location in the image (see Figure 3.15). This cartoon shows only a single telescope (T_i); by combining images from two or more telescopes, the location and direction of the shower can be reconstructed via trigonometry. Image taken from LeBohec et al. (2005).

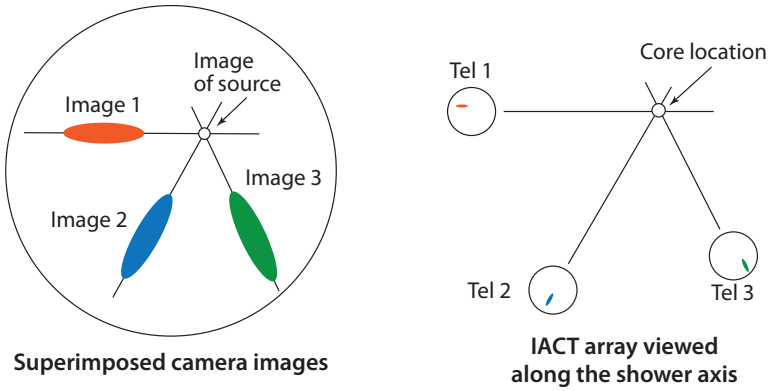

Figure 3.15. Diagram showing the two planes from Figure 3.14 in more detail with multiple telescopes. *Left*: Images from each telescope that captured the shower are superimposed to reconstruct the location of the source (point S in Figure 3.14). *Right*: The locations of the telescopes are projected into a plane perpendicular to the shower axis. The core location (point P in Figure 3.14) is determined by extending the image axes from the projected locations of the telescopes and determining their point of intersection. Image taken from Daum et al. (1997).

possible to reconstruct the showers using two or three telescopes, but the reconstruction is somewhat less robust (LeBohec et al., 2005).

Much of the variation in the shower reconstruction methods lies in fact that the image axes don't perfectly intersect as they do in Figure 3.15. Usually, the points of intersection are calculated as the mean of the pairwise intersection locations. However, many alternatives exist, several of which are outlined in Hofmann et al. (1999). The algorithm we have just discussed is Hofmann et al.'s Algorithm 1, which is the default algorithm used by the two VERITAS analysis packages, VEGAS and EventDisplay (Cogan, 2008; Daniel, 2008).

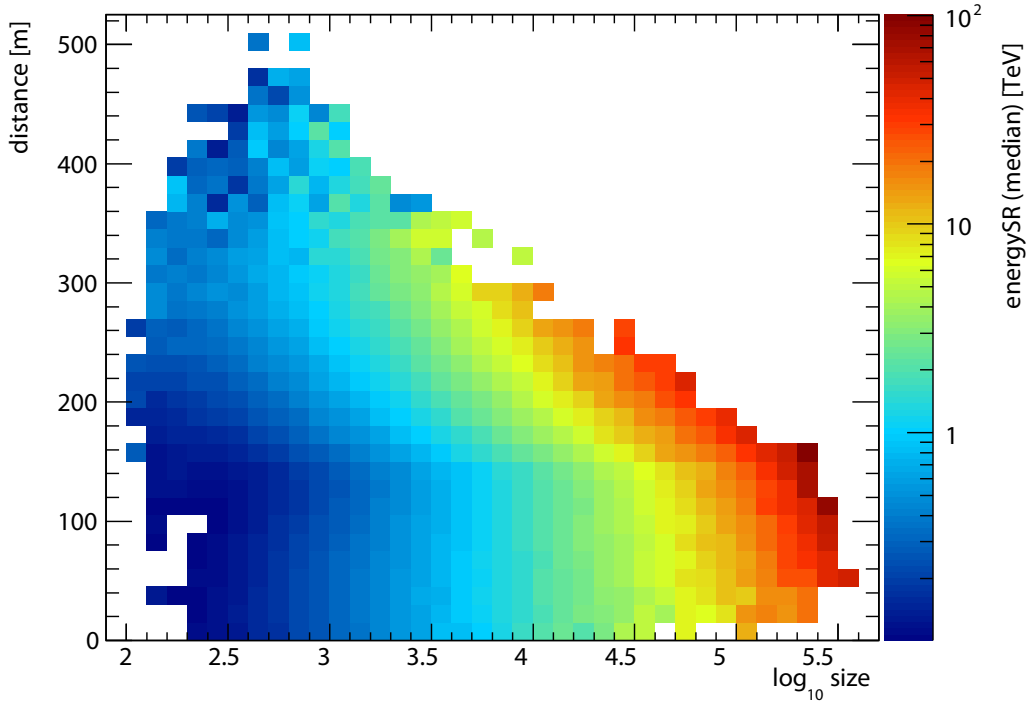

Figure 3.16. Example of an energy reconstruction lookup table for a fixed zenith angle, azimuth, wobble offset, night sky background noise level, and telescope number. The color scale represents the reconstructed energy, the x-axis is the shower's size parameter, and the y-axis is the distance to the shower core on the ground (i.e., the impact parameter). Image credit: Gernot Maier.

3.6 Energy Reconstruction

The energy of a particle that initiates an extensive air shower is approximately proportional to the total amount of Cherenkov light emitted by the shower. Therefore, the energy of the progenitor particle can be estimated from the Hillas parameterization of the size of the image. However, there is also some dependence on the observing conditions: the particular telescope (indicated by telescope number), the direction of the telescope with respect to the shower (zenith angle, azimuth, and wobble offset), the distance to the shower core, and the night sky background level. A set of seven-dimensional lookup tables covering all of these parameters is generated using Monte Carlo simulations (see Figure 3.16 for a 2D projection of this table). The zenith angle, azimuth angle, and wobble offset are stored in the VERITAS database during observations; the size, impact distance, and noise level can be estimated from the data. It is then a simple matter to use the lookup tables to estimate the energy of the primary particle.

CHAPTER 4

TESTING LORENTZ INVARIANCE USING TEV OBSERVATIONS OF THE LARGE FEBRUARY 2010 FLARE OF MARKARIAN 421

4.1 Introduction

Quantum field theory and general relativity have each enjoyed unparalleled success individually, but a grand unification of these two disparate theories has not materialized despite nearly a century of concerted effort. Several elegant quantum gravity (QG) theories have been put forward, but these have been guided primarily by symmetry and elegant mathematics as they are difficult to test experimentally. This is predominantly because QG effects are expected to have observable consequences only at energy scales approaching the Planck energy, $E_P = \sqrt{\hbar c^5/G} \approx 10^{19}$ GeV. Lorentz invariance violation (LIV) near the Planck energy can arise in a wide variety of QG theories (Amelino-Camelia et al., 1997; Bojowald et al., 2005; Gambini & Pullin, 1999; Kostelecký & Samuel, 1989) and is one of the few testable outcomes currently available.

A key insight made by Amelino-Camelia et al. (1998) is that QG effects, while suppressed by the Planck scale, may alter the propagation velocity of high energy photons enough to produce a measurable time delay when integrated over astronomical distances (see Section 1.3). Despite the existence of more thorough theoretical frameworks, such as the Standard Model Extension (Colladay & Kostelecký, 1997; Kostelecký & Mewes, 2009), it is standard to follow the original approach of Amelino-Camelia et al. (1998) to express the vacuum dispersion relation for photons as a series expansion:

$$E^2 \approx p^2c^2 \left[1 - \sum_{n=1}^{\infty} s_{\pm} \left(\frac{E}{E_{QG}}\right)^n\right], \tag{4.1}$$

where c corresponds to the speed of light in vacuum for low energy photons and $E_{QG} \equiv \xi E_P$ is the QG energy scale. The value of $s_\pm$ on the right hand side depends on whether LIV effects are expected to decrease $(+1)$ or increase (-1) photon velocity. Considering only the leading term of order n, the dispersion-induced time delay (Jacob & Piran, 2008) for a particle of red-shifted energy E traveling through an expanding universe is

$$\Delta t = s_\pm \frac{1+n}{2H_0} \left(\frac{E}{E_{QG}}\right)^n \int_0^z \frac{(1+z')^n \, dz'}{\sqrt{\Omega_m (1+z')^3 + \Omega_\Lambda}} \tag{4.2}$$

where z is the redshift of the photon emitting source and Ω_m, Ω_Λ, and H_0 are the present cosmological parameters. The QG energy scale and $s_\pm$ parameter (which are unknown) as well as the distance (redshift) to the source and cosmological parameters (which are known) are customarily absorbed into a single LIV parameter

$$\tau_n \equiv s_\pm \frac{1+n}{2H_0} \left(\frac{1}{E_{QG}}\right)^n \int_0^z \frac{(1+z')^n \, dz'}{\sqrt{\Omega_m (1+z')^3 + \Omega_\Lambda}}. \tag{4.3}$$

Equation (4.2) then takes on the particularly simple form

$$\Delta t = \tau_n \, E^n. \tag{4.4}$$

Typically only the linear $(n = 1)$ and quadratic $(n = 2)$ terms—hereafter linear and quadratic LIV—are discussed because currently available measurements do not constrain higher order terms. Moreover, we will neglect any consideration of photon helicity. *Fermi*-LAT observations of GRB 090510 have produced a linear LIV limit (denoted by a subscript 1) of $E_{QG,1} > 7.6 \, E_P$ (Vasileiou et al., 2013), strongly disfavoring the possibility of linear LIV below the Planck scale. However, the *Fermi*-LAT limit on quadratic LIV of $E_{QG,2} > 1.3 \times 10^{11}$ GeV, is still several orders of magnitude away from the Planck energy.

Observations made with ground-based very high energy (VHE; E > 100 GeV) gamma-ray observatories are ideal for investigating quadratic LIV because they can observe photons with energies up to tens of TeV. Higher energy photons permit greater sensitivity to LIV, especially for quadratic LIV where time delays depend on the energy squared. Several types of astrophysical sources can be used for LIV time-of-flight measurements with ground-based VHE observatories, including strong flares from blazars (Abramowski et al., 2011; Albert et al., 2008; Biller et al., 1999) and pulsations from VHE-emitting pulsars (Kaaret, 1999; Otte, 2011). Gamma-ray bursts have not yet been observed at VHE energies, and only one VHE-emitting pulsar has been detected, the Crab pulsar (PSR B0531+21).

We utilized VERITAS observations of a strong VHE flare from the high-frequency peaked BL Lac object Markarian 421 ($z = 0.031$, $d \approx 130$ Mpc). Markarian 421 (Mrk 421 hereafter) was the first blazar – an active galactic nucleus (AGN) with its jet oriented along our line of sight – detected at TeV energies (Punch et al., 1992) and is well known for its strong, rapid VHE flares (Gaidos et al., 1996). VERITAS detected strong flaring activity from Mrk 421 on the night of February 17, 2010 (MJD 55244), with a peak flux approximately ten times greater than the Crab nebula (one of the brightest VHE sources and the canonical VHE standard candle) at energies > 300 GeV. Although Mrk 421 is one of the closest TeV blazars, this tremendous flux at very high energies provides a good opportunity to search for quadratic LIV.

In Section 4.2 we outline the methods used to collect and reduce the VERITAS data. In Section 4.3 we briefly describe how we estimate the amount of dispersion in our data using a dispersion cancellation method or the PairView algorithm. In Section 4.4, we present the

results of our analysis. Finally, in Section 4.5 we relate our results to other work that has been done in the field.

4.2 Observations and Data Reduction

VERITAS is an array of four imaging atmospheric Cherenkov telescopes (IACTs) that measures gamma rays with an energy range of ∼100 GeV - 30 TeV and an energy resolution of approximately 15% at 1 TeV (Holder et al., 2008). The array records stereoscopic images of extensive air showers initiated by gamma rays impinging on Earth's atmosphere. The images are used to reconstruct the gamma ray's energy and direction using the procedure detailed in Acciari et al. (2008), which we briefly describe here.

Images of the extensive air showers are parameterized using a moments analysis (Hillas, 1985) to determine the initiating particle's direction. Comparison of the parameters with Monte Carlo simulations is used to determine the energy of the gamma rays and to separate gamma rays from background cosmic rays, which are three orders of magnitude more numerous. During the process of separating gamma rays from the cosmic ray background, we prioritized including as many source photons as possible over background rejection, leading to a signal-to-background ratio of approximately 14.

Observations of the Mrk 421 February 17, 2010 flare consist of five hours of data collected during clear conditions, with a cumulative significance of 190 σ and an average gamma-ray emission rate of 93.9 gammas/min. The light curve is shown in Figure 4.1. The observations were made using wobble mode (Daum et al., 1997), whereby the telescope pointing is offset from the source location by 0.5° in one of four orthogonal directions. The

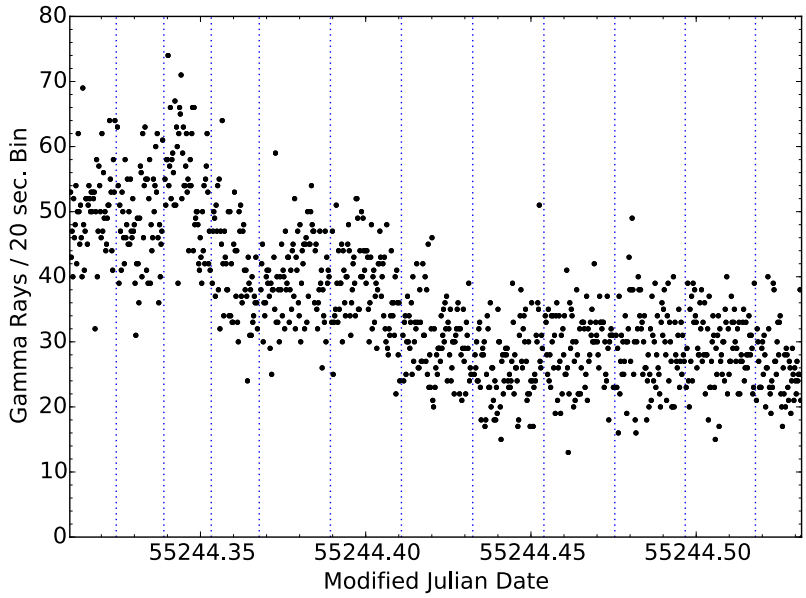

Figure 4.1. Light curve from February 17, 2010 (MJD 55244) plotting the number of gamma rays with $E > 200$ GeV. Vertical dotted lines mark the transitions between runs.

observations are, therefore, divided up into twelve 20-30 minute observing runs, proceeding through the four wobble directions cyclically. We number the runs based on the partitions shown in Figure 4.1, starting with 1 on the left and ending with 12 on the right. We cut the first and last 100 seconds from each run in order to avoid deleterious effects due to run boundaries in the LIV analysis.

Due to temporary hardware problems in one of the telescopes, the observations were made with only three of the four telescopes. This increases the size of the gamma-ray point spread function, but otherwise should have minimal impact on the data. The five-hour span of the observations required Mrk 421 to be observed at zenith angles ranging from 7° - 56°. At large zenith angles (>30°), the detection threshold increases due to a decrease in the apparent size of the shower, while the effective area for higher energy photons increases (Konopelko et al., 1999). In order to achieve the same energy threshold throughout the

observation, we removed events with energies below 200 GeV for all of the runs, increasing the energy threshold by a factor of two.

Detecting dispersion requires that structure due to rapid variability be present in the light curve; this is generally the case in bursting or flaring astrophysical events. Variability within each run is quantified by applying the χ^2 test to determine whether the binned light curve of the run is consistent with being constant. We applied the χ^2 test to each run using time binning of 15, 30, 60, and 120 seconds; the results are shown in Table 4.1. Based on the results of the χ^2 test we found run #10 to be the most variable. Therefore we chose to use only this run for the LIV analysis.

To illustrate the difference in variability between the runs, we show light curves from the most variable (#10) and least variable (#1) runs in Figure 4.2. Run #10 exhibits significant variability, especially in the interval of 41400 – 41900 seconds elapsed since midnight UTC. The χ^2 statistic indicates that this run is variable, having less than a 2.5% probability the data is compatible with constant emission, regardless of the binning used. Conversely, run #1 is highly compatible with constant emission.

4.3 Dispersion Estimation

If we assume that any dispersion present in the light curve can be attributed to LIV, the time delay can be expressed using the simple functional form given in Equation (4.4). Defining $\Delta t = t_{obs} - t_0$ to be the difference between the time our telescope observes a photon and the time we would have observed it in the absence of LIV, Equation (4.4) becomes

$$t_0 = t_{obs} - \tau_n E^n, \tag{4.5}$$

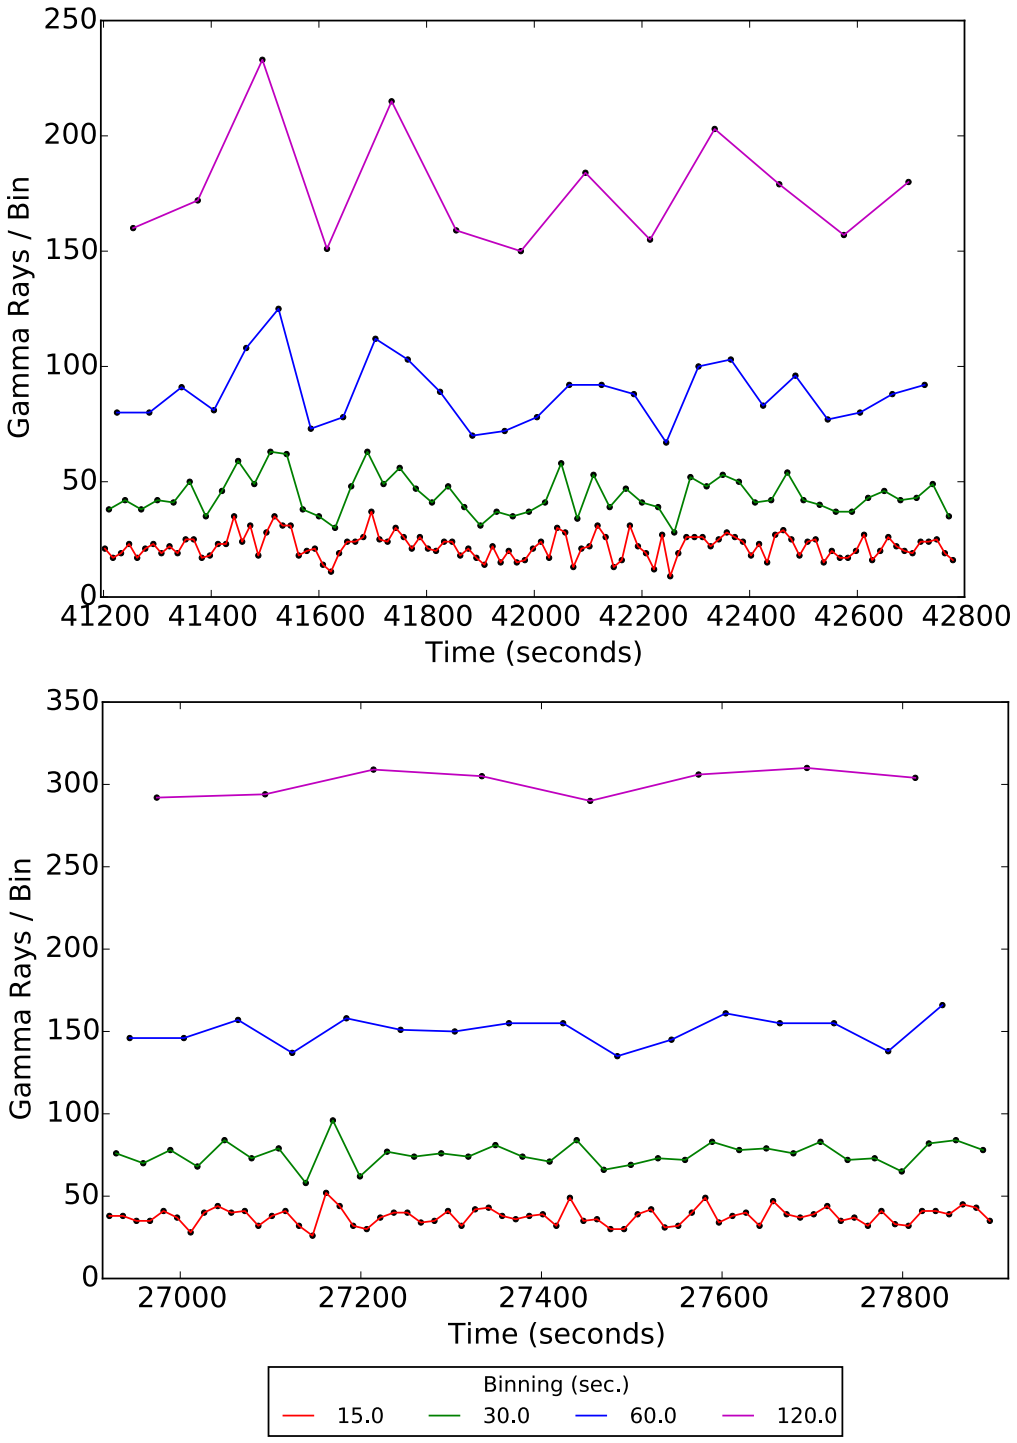

Figure 4.2. Light curves from the most variable run, #10 (top), and the least variable run, #1 (bottom), plotting the number of gamma rays with $E > 200$ GeV. Several choices of time binning are used: from bottom to top, 15, 30, 60, and 120 second binning. Data points are shown with black points; the colored lines connected the data points are intended as a visual aid.

Table 4.1. Results of the χ^2 test applied to each run's light curve using 15, 30, 60, and 120 time binning. The χ^2 statistic over the number of degrees of freedom (NDF) and the probability P that the χ^2 value occurred by chance are listed for each time binning.

	15 sec.		30 sec.		60 sec.		120 sec.	
Run	χ^2/NDF	P (%)	χ^2/NDF	P (%)	χ^2/NDF	P (%)	χ^2/NDF	P (%)
1	47.11 / 65	95.4	23.29 / 32	86.9	7.83 / 15	93.0	1.48 / 7	98.3
2	52.72 / 65	86.3	36.75 / 32	25.8	25.16 / 15	4.8	17.49 / 7	1.4
3	80.93 / 65	8.8	53.10 / 32	1.1	27.46 / 15	2.5	23.72 / 7	0.1
4	79.00 / 65	11.4	47.67 / 32	3.7	30.40 / 15	1.1	13.59 / 7	5.9
5	102.46 / 105	55.2	56.51 / 52	31.0	30.24 / 25	21.5	17.12 / 12	14.5
6	97.63 / 105	68.2	55.23 / 52	35.4	34.15 / 25	10.5	17.50 / 12	13.2
7	92.15 / 105	81.0	48.91 / 52	59.6	19.30 / 25	78.3	10.11 / 12	60.7
8	125.19 / 105	8.7	83.59 / 52	0.4	43.68 / 25	1.2	16.49 / 12	17.0
9	131.35 / 105	4.2	54.79 / 52	36.9	30.75 / 25	19.8	13.56 / 12	32.9
10	135.53 / 105	2.4	83.26 / 52	0.4	55.33 / 25	0.05	46.64 / 12	0.001
11	99.54 / 105	63.2	52.75 / 52	44.5	19.97 / 25	74.8	12.11 / 12	43.7
12	64.34 / 65	50.0	36.06 / 32	28.4	23.80 / 15	6.9	7.85 / 7	34.6

where $n = \{1, 2\}$ corresponds to linear or quadratic LIV. Unfortunately, we have no knowledge of t_0, so we require some method of estimating the LIV dispersion parameter τ_n solely from the observed photon arrival times and energies. We used two different methods to estimate the dispersion in our light curves. Neither method requires binning the data in either energy or time, which artificially limits the attainable resolution and reduces repeatability (reproducing identical results would require identical bin sizes and locations), and neither method used in this paper has any free parameters and thus they require no tuning.

The first method is a dispersion cancellation (also called sharpness maximization) algorithm based on the work of Scargle et al. (2008). The dispersion cancellation method applies Eq. 4.5 as a correction to time- and energy-tagged events and estimates τ_n via a trial and error process. Our algorithm differs from the method described in Scargle et al. (2008) in how the cell-based light curve is constructed and the probability distribution that

is used. As a result of these differences, and because this method is complicated, we describe this method in detail in Section 4.3.1. The second method, PairView, is much simpler to implement and is exactly the same as the algorithm presented in Vasileiou et al. (2013). We will give only a brief overview of the method in Section 4.3.2; interested readers are strongly encouraged to read the thorough description given in Vasileiou et al. (2013).

4.3.1 Dispersion Cancellation Algorithm

Detecting dispersion requires that structure due to rapid variability be present in the light curve; this is generally the case in bursting or flaring astrophysical events. The effect of dispersion in a light curve can be illustrated by considering a collection of concurrently emitted photons with a variety of energies, which produces a light curve consisting of a narrow peak. Dispersion broadens this peak, reducing its sharpness. If the amount of dispersion is known, the light curve can be dedispersed, restoring the original narrow peak. Real light curves are more complicated than a single narrow peak; we rely on the presence of structure in the light curve and attempt to correlate any observed features.

The dispersion cancellation algorithm seeks to maximize the sharpness in the observed light curve by iteratively applying Eq. 4.5 to the observed time series using a sequence of trial τ_n values. This is done for both linear and quadratic LIV ($n = 1, 2$). For each τ_n, we calculate the sharpness of the light curve using Shannon entropy as a metric for sharpness. We then select the value of τ_n which minimizes the Shannon entropy (minimizing Shannon entropy is equivalent to maximizing sharpness), and use this value as an estimator $\hat{\tau}_n$ for the dispersion present in the light curve. The uncertainty of the estimator is assessed using

bootstrapping, which will be described in Section 4.3.3. In the rest of this section, we describe our dispersion cancellation algorithm in detail.

We first select a value for n depending on whether we are searching for linear or quadratic LIV. Next, we apply Eq. 4.5 to our time- and energy-tagged events to produce a dedispersed light curve. Instead of representing the light curve using a binned histogram (which we wish to avoid for reasons described at the start of Section 4.3), we construct a light curve out of "cells" containing a single photon each. Each cell is assigned a width, δt, equal to the time difference between successive photons. Specifically, starting with the first photon in a time-ordered list, which we assign a time t_0, we compute the time difference $\delta t = t_{i+1} - t_i$, and assign to each δt-width cell the energy of the photon at time t_i. The energy of the last photon in the list is not used, but since there are many photons in the list this does not adversely affect the results.

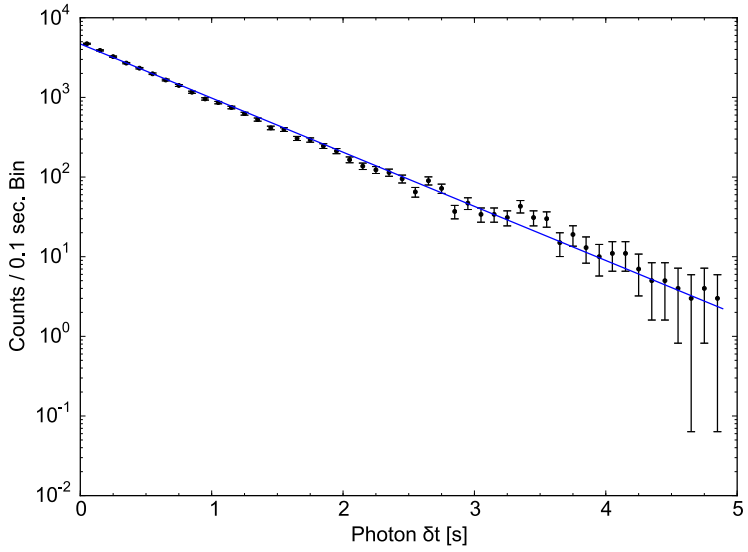

Figure 4.3. Histogram of photon arrival time differences with 0.1 second histogram bins. The error bars are computed as $\sigma' \approx 1 + \sqrt{N + 0.75}$ to maintain accuracy for small N and correspond to Gaussian 1σ errors (Gehrels, 1986). The solid line shows the normalized exponential distribution of Equation (4.6).

We expect photon emission from an astrophysical source to be a Poisson process. This implies that the time interval between successive events should be exponentially distributed. Comparison with the data in Figure 4.3 shows that the cell widths follow an exponential distribution that is normalized over the interval $0 \leq t \leq \infty$:

$$p(\delta t) = \lambda\, e^{-\lambda \delta t}, \tag{4.6}$$

where λ is the average count rate over the entire light curve. We therefore find that Equation (4.6) can be used to provide an estimate of the probability of a cell of width δt occurring in a light curve.

The sharpness of a light curve can be quantified in terms of its Shannon entropy (Shannon & Weaver, 1949):

$$S = -\sum_i p_i \log p_i. \tag{4.7}$$

Note that maximum Shannon entropy occurs when events are distributed uniformly across a time interval; conversely, minimum Shannon entropy requires all of the events in a finite-sized light curve to occur at one time. We therefore need to minimize the Shannon entropy in order to maximize light curve sharpness.

The minimization is performed using a parameter sweep over some reasonable range of τ_n values. We used a range of $-100 \leq \tau_1 \leq 100$ s/TeV with steps of 0.1 s/TeV for linear LIV, and a range of $-10 \leq \tau_2 \leq 10$ s/TeV^2 with steps of 0.01 s/TeV^2 for quadratic LIV. For each τ_n value in the range, we iterate through each time- and energy-tagged event (t_{obs}, E_{obs}) and construct a hypothetical dedispersed light curve containing times $\hat{t}$:

$$\hat{t} = t_{obs} - \tau_n E_{obs}^n. \tag{4.8}$$

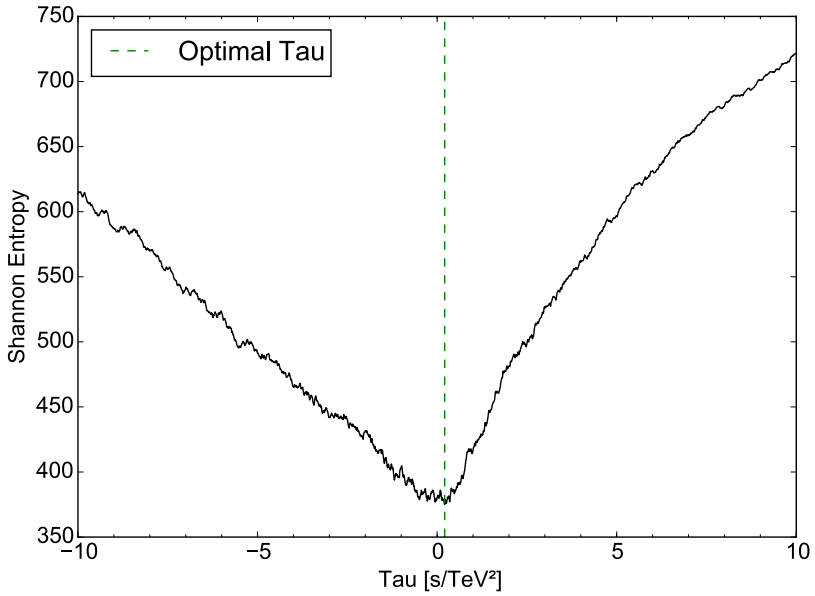

Figure 4.4. Example plot of the Shannon entropy as a function of the LIV dispersion parameter τ_2 for quadratic LIV. The value of the estimated LIV dispersion, $\hat{\tau}_2$, indicated by the vertical dashed line, minimizes the Shannon entropy and thus maximizes sharpness in the light curve.

A cell representation is constructed and the Shannon entropy is calculated for this hypothetical light curve according to the procedure described above. The hypothetical light curve with the smallest Shannon entropy (sharpest features) is assumed to be a reasonable estimate of the light curve one would measure in the absence of LIV – compare Equation (4.8) with Equation (4.5) – so the value of τ_n which produces the hypothetical light curve with the smallest Shannon entropy is used as our estimator for the LIV dispersion, which we call $\hat{\tau}_n$. An example plot of the Shannon entropy versus τ_2 is shown in Figure 4.4; a dashed vertical green line indicates the location of $\hat{\tau}_2$ where the Shannon entropy is minimized.

To summarize, the functioning of our dispersion cancellation algorithm can be outlined as follows:

1. Choose a model for the LIV time delay – linear ($n = 1$) or quadratic ($n = 2$).

2. Pick a trial value of τ_n for the chosen LIV model.
3. Dedisperse the list of observed times and sort into time order.
4. Generate an unbinned cell representation of the light curve.
5. Calculate a probability for each cell from a normalized probability distribution.
6. Compute the Shannon entropy of the light curve.
7. Iterate steps 2-6 over a set of trial values of τ_n.
8. Select the τ_n that minimizes the Shannon entropy and use it as an estimator $\hat{\tau}_n$ for the LIV dispersion.

4.3.2 PairView Algorithm

The PairView algorithm (Vasileiou et al., 2013) computes the spectral lag (arrival time difference between a high and low energy photon) weighted by either the energy difference (for linear LIV) or the difference of the squared energy (for quadratic LIV). For a data set consisting of N photons with arrival times $t_1, t_2, ..., t_N$ and arrival energies $E_1, E_2, ..., E_N$, the procedure is to calculate the weighted spectral lags $l_{i,j}$ for the $N(N-1)/2$ pairs of photons with indicies $i > j$:

$$l_{i,j} = \frac{t_i - t_j}{E_i^n - E_j^n}. \tag{4.9}$$

In the event that two photons have identical energies (within instrumental resolution), the photon pair is not included in the $l_{i,j}$ in order to avoid division by zero.

The principle behind the PairView method can be motivated by considering a thought experiment in which an astrophysical source emits many photons of differing energies at the same time. In the presence of LIV dispersion described by τ_n (see Equation (4.3)), the

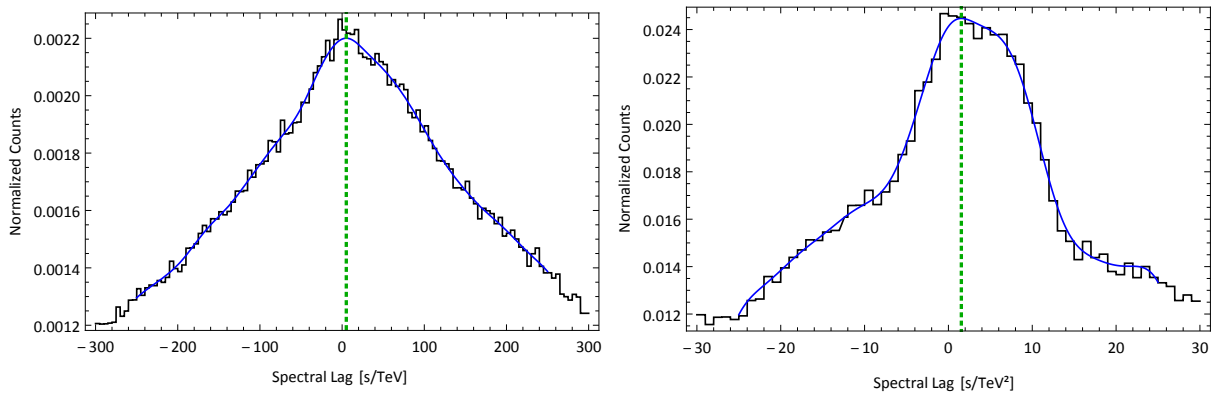

Figure 4.5. Plots of the spectral lag distributions $l_{i,j}$ for linear (left) and quadratic (right) LIV. The black line with steps is a standard histogram of the $l_{i,j}$ distribution, while the smooth blue curve is a kernel density estimation (KDE). The KDE curve is truncated at the edges for clarity. The green vertical dashed line shows the location of the peak of the KDE distribution, which is used as the estimate $\hat{\tau}_n$ of the LIV dispersion.

photons will not arrive at a distant detector all at once, but will instead be spread into a finite-width distribution whose width is controlled by the amount of dispersion. In this simple case the $l_{i,j}$ will be distributed as a delta function centered exactly on τ_n (note that the $l_{i,j}$ and τ_n have the same units). Alternatively, if the hypothetical astrophysical source were to instead emit a pulse of photons with some small spread in time, the $l_{i,j}$ distribution would also be broadened into a finite-width distribution, but would remain centered on $l_{i,j}$.

The application of the PairView method is thus very simple: choose a time interval and LIV order ($n = 1, 2$), calculate the $l_{i,j}$ according to Equation (4.9), then determine the location of the peak of the $l_{i,j}$ distribution. Kernel density estimation (KDE) with a Gaussian kernel is used to represent the distribution so that arbitrary binning can be avoided. Following Vasileiou et al. (2013), we used ROOT's TKDE class to perform the kernel density estimation. As a practical matter, the $l_{i,j}$ distribution may be extremely broad due to events with nearly identical energies. Since information about the dispersion is contained in the

peak of the distribution, which we found to be centered near zero, we restrict the $l_{i,j}$ to a range of $-300 \leq l_{i,j} \leq 300$ s/TeV for linear LIV and $-30 \leq l_{i,j} \leq 30$ s/TeV^2.

4.3.3 Bootstrap Uncertainty Analysis

Each algorithm produces an estimate of the LIV dispersion present in the data, which we call $\hat{\tau}_n$ (the hat indicates that it is an estimator for τ_n, where n is the LIV order). The uncertainty with which $\hat{\tau}_n$ estimates the true LIV dispersion is assessed using bootstrapping (Efron, 1979; Efron & Tibshirani, 1993). During the bootstrap analysis, the events are resampled by randomly shuffling the photon energies while their arrival times are kept fixed. For example, given a set of events $\{(t_1, E_1), (t_2, E_2), (t_3, E_3), (t_4, E_4)\}$, the resampled events might be $\{(t_1, E_2), (t_2, E_1), (t_3, E_4), (t_4, E_3)\}$. For a data set containing N events, this results in $N!$ possible distinct synthetic data sets with an identical time structure and energy

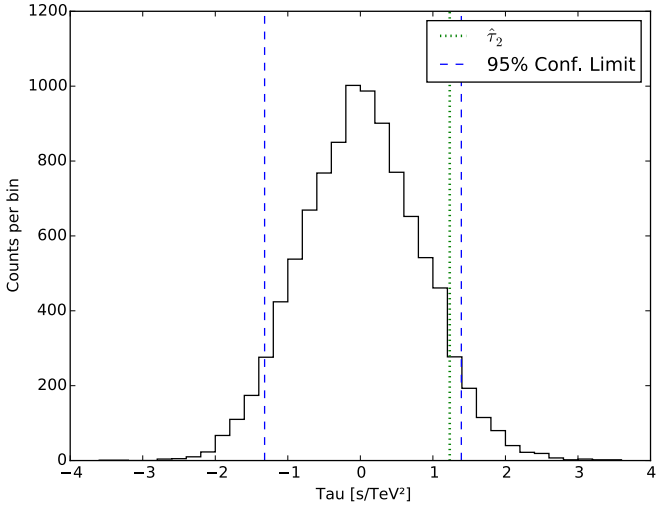

Figure 4.6. Example of a PairView bootstrap distribution for quadratic LIV. The veritical green dotted line marks the location of $\hat{\tau}_2$, which is determined from the $l_{i,j}$ distribution shown in Figure 4.5. The vertical dashed blue lines indicate the uncertainties of $\hat{\tau}_2$ at the 95% confidence level.

distribution.

The bootstrap distributions are obtained by repeating the procedures used to produce Figure 4.4 and Figure 4.5 ten thousand times, each time using a different pairing between the photon arrival times and energies. An example histogram of the ten thousand $\hat{\tau}_n$ measurements for quadratic LIV is shown in Figure 4.6. Note that the location of $\hat{\tau}_2$ from Figure 4.5 is carried over into Figure 4.6 and is demarcated by a vertical dashed green line. The estimate of the LIV dispersion comes from the real data (as opposed to the resampled data), while the uncertainties on $\hat{\tau}_n$ are derived solely from the bootstrap distributions and are evaluated as one-sided 95% confidence limits.

The presence of dispersion implies a functional relationship between the arrival times and energies; shuffling the energies breaks this relationship and therefore should erase any dispersion that may be present in the data. We therefore expect the mean of bootstrapped

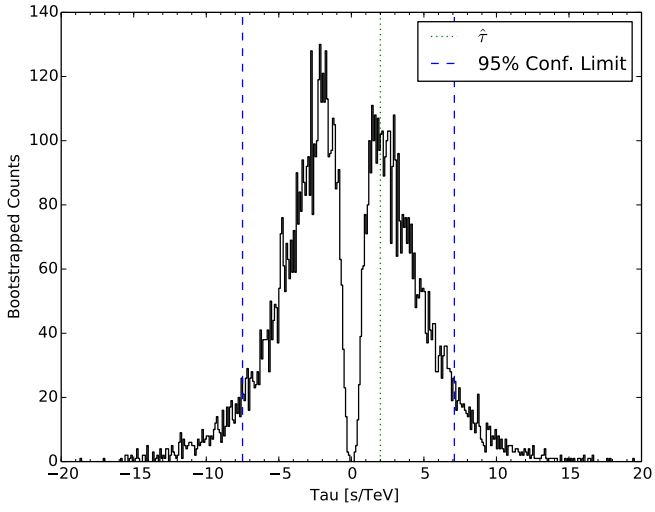

Figure 4.7. Example of a bimodal distribution resulting from including the edges of the run, which artificially add structure to the data. This distribution was generated from run #2 with the dispersion cancellation algorithm.

$\hat{\tau}_n$ distributions to be close to zero. Nonetheless, the time structure of the light curve may cause artificial dispersion to be present in every bootstrapped distribution, in which case the $\hat{\tau}_n$ distribution will be bimodal and symmetric about $\hat{\tau}_n = 0$ as in Figure 4.7.

We found this to be a common occurrence for data with a flaring structure on top of a constant background in which the rise from zero flux to the constant background was included in the light curve. This adds a sharp structure to the data which is in reality an observational artifact. Great care should be taken to avoid this; we will discuss the methods we employed for handling observation boundaries in Section 4.4.1. However, a bimodal distribution can also result from a sharp light curve feature which is actually present in the data (see, e.g., Figure 14 in Vasileiou et al. (2013)), so the existence of a bimodal bootstrap distribution does not automatically imply a problem with the data.

4.3.4 Algorithm Validation

In order to assess how well the dispersion estimation algorithms perform, we artificially inject dispersion into resampled data and run it through the algorithm. Ideally the output of the algorithm (which we call the reconstructed dispersion) should be exactly equal to the applied dispersion, without bias and with no dependence on the amplitude of the injected dispersion. In practice, the ability of the algorithm to reconstruct the dispersion is expected to depend strongly on the amount of structure (i.e., variability) present within the data. We therefore utilize synthetic data sets generated using the resampling procedure used with the bootstrap, which have the same time structure and energy spectrum as the real data.

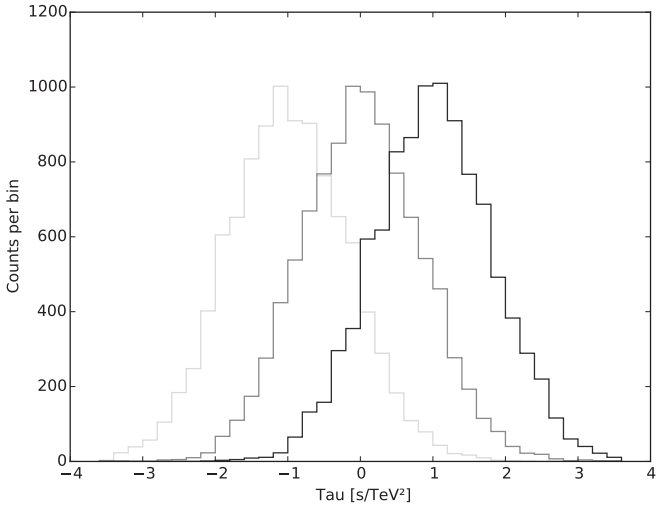

Figure 4.8. Example $\hat{\tau}_2$ distributions generated via bootstrapping with the PairView algorithm. Three distributions are shown for three values of $\tau_{2,applied}$: from left to right, or light to dark, $\tau_{applied} = -1.0$ s/TeV^2, 0 s/TeV^2, and 1.0 s/TeV^2.

The procedure for verifying the performance of the algorithms is as follows. The data (either simulated or real) are first resampled using the bootstrapping procedure described in Section 4.3.3, which generates a new synthetic data set while removing any dispersion that might have already been present. Next, we apply dispersion to the data using the inverse of Equation (4.5); that is, for each photon in the synthetic light curve with index i, we compute a shifted time

$$t'_i = t_i + \tau_{applied}\, E_i^n, \tag{4.10}$$

where $\{t_i, E_i\}$ are the time and energy of the i^{th} photon, n is the LIV order, and t'_i is the shifted time in the presence of the applied dispersion $\tau_{applied}$. The (t'_i, E_i) are then run through the algorithm to produce an estimate of $\tau_{applied}$. This procedure is typically repeated ten thousand times to yield a distribution of $\hat{\tau}_{applied}$. Figure 4.8 shows three examples of these distributions for quadratic LIV with $\tau_{applied} = -1.0$ s/TeV^2, 0 s/TeV^2, and 1.0 s/TeV^2.

For each value of $\tau_{applied}$, we compute the mean and standard deviation of the corresponding distribution. The mean of the $\hat{\tau}_{applied}$ distribution is an estimate of how well the algorithm can estimate the dispersion in a light curve; we call this the reconstructed τ_n or $\tau_{reconstructed}$. The standard deviation of $\tau_{reconstructed}$ gives an estimate of the uncertainty with which the algorithm can reconstruct the applied dispersion. Several plots of $\tau_{reconstructed}$ vs. $\tau_{applied}$ will be shown in Section 4.4.

Ideally, $\tau_{reconstructed} \approx \tau_{applied}$, and the uncertainties of the $\tau_{reconstructed}$ should be small. However, if there is insufficient structure in the light curve, the reconstructed and applied dispersions will not be equal (which is manifest as a slope differing from unity when plotting $\tau_{reconstructed}$ vs. $\tau_{applied}$) and the uncertainties will be large. Thus the $\tau_{reconstructed}$ vs. $\tau_{applied}$ plot can be used to provide a diagnostic of both the efficacy of the algorithm and the suitability of the data for LIV investigations.

Before discussing the results of our data analysis, we first show that the dispersion estimation algorithms described in Section 4.3 can provide an accurate estimation of the dispersion present in an artificial data set consisting of a single Gaussian pulse using the technique described in Section 4.3.4. The Gaussian pulse is centered in an interval 1500 seconds long (to match the length of run #10) and has a standard deviation of 100 seconds – approximately the width of the largest flares in run #10. We draw a sample of 500 times to produce the light curve shown in Figure 4.9, and then draw a random sample of 500 energies from the list of 2500 energies used in run #10.

The results of processing the Gaussian pulse through both dispersion estimation algorithms are shown in Figure 4.10 for quadratic LIV. The dispersion cancellation algorithm

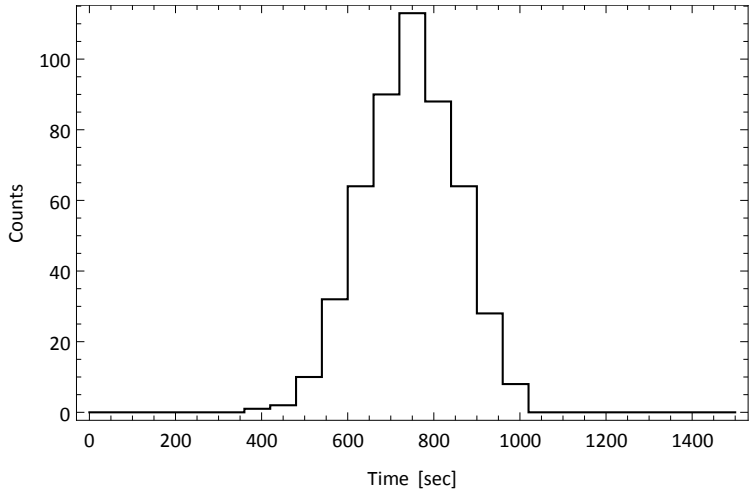

Figure 4.9. Histogram of the Gaussian distribution used to test the dispersion estimation algorithms. The bins are 60 seconds wide, and the pulse contains 500 events.

results are shown in the left column, while the PairView algorithm results are shown on the right. The top row of plots shows that both algorithms pick out a dispersion of approximately zero, which is expected since the energies were randomly sampled from the photon energies in run #10. The middle row shows histograms of bootstrap distributions with means that are nearly zero (also expected; see Section 4.3.3) and very narrow widths. The narrow widths indicate that the algorithms are capable of estimating the dispersion in the data with very little uncertainty – approximately 2×10^{-5} s/TeV^2 for the dispersion cancellation algorithm and about half of that for PairView.

The most important plots in Figure 4.10 are those shown in the bottom row, which were produced using the method described in Section 4.3.4. These plots show that if dispersion were present in the data, both algorithms can accurately estimate the amount of dispersion in the data with very little uncertainty ($< 3 \times 10^{-5}$ s/TeV^2). The dispersion cancellation algorithm appears to systematically overestimate the applied dispersion by 1×10^{-6} s/TeV^2, but this is simply the resolution of the scan used to determine the minimum Shannon en-

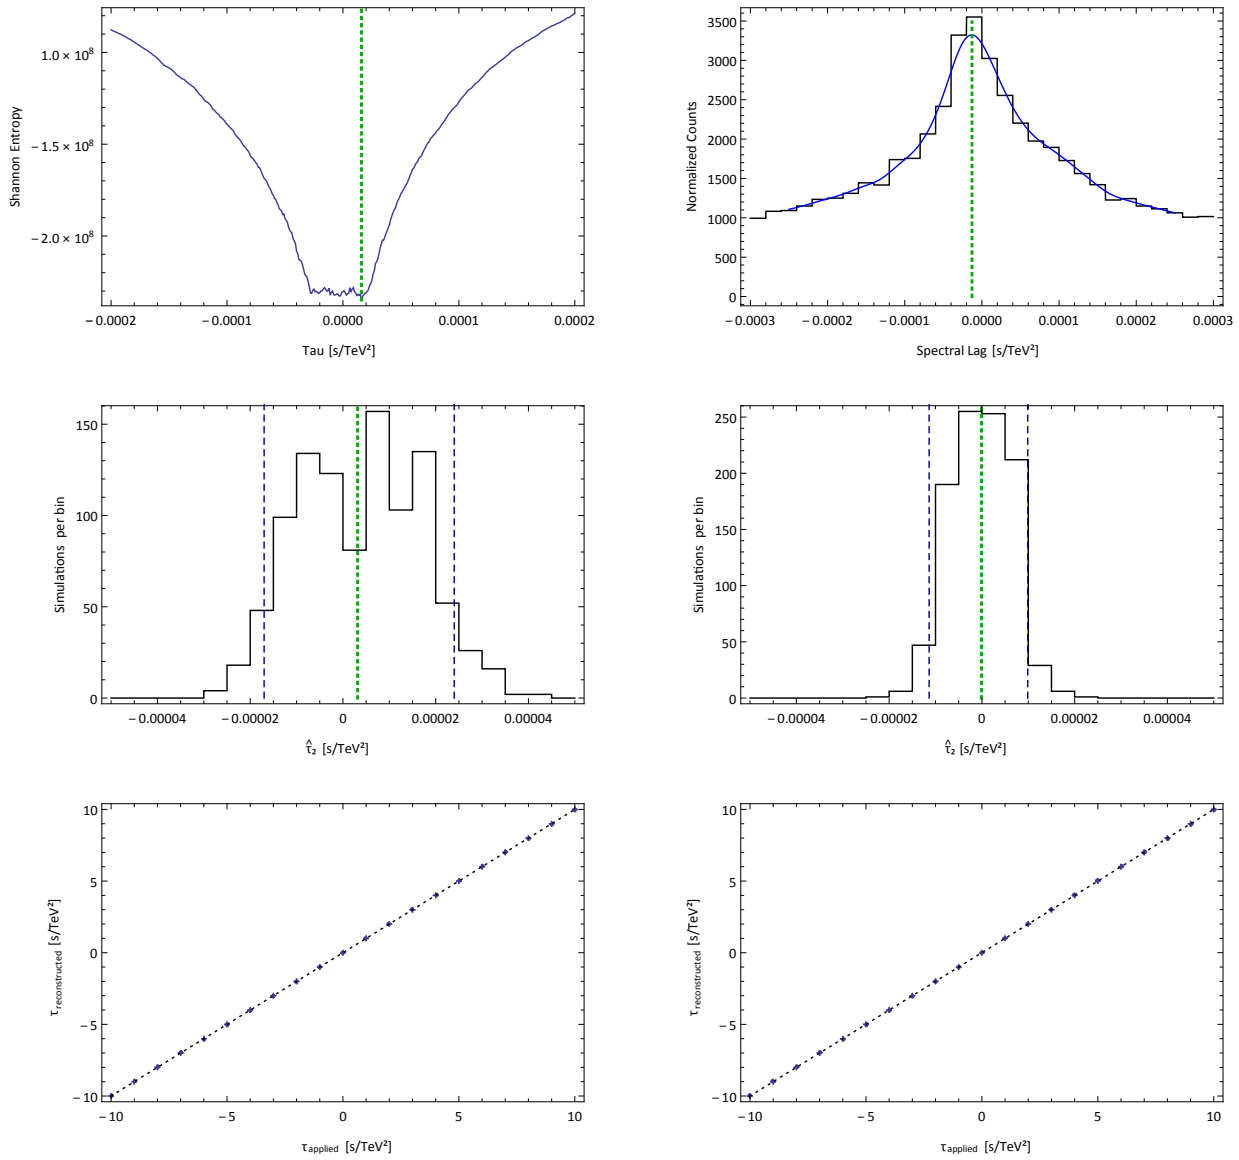

Figure 4.10. Results of testing the dispersion estimation algorithms using a single Gaussian pulse for quadratic LIV. Plots in the left column are from the dispersion cancellation algorithm; plots in the right column are from PairView. Top row: dispersion cancellation scan (left) and spectral lag distribution (right). Middle row: bootstrap distributions of $\hat{\tau}_2$, with a green dotted vertical line in the center to indicate the mean of the distribution, and two blue dashed vertical lines to the left and right indicate one-sided 95% confidence intervals. Bottom row: plots of $\tau_{reconstructed}$ vs. $\tau_{applied}$ with error bars showing the uncertainties at the 95% confidence level. The points fall within 10^{-6} s/TeV^2 of the line, and the error bars are too small to see.

tropy (top left plot in Figure 4.10). A finer resolution scan would remove this bias at the cost of computation time, but this is unnecessary since the means of the distributions fall well within the uncertainties.

4.4 Results

In this section, we first discuss considerations for handling observations containing flaring on top of a constant background. This type of data requires careful handling because the boundaries of the observation, where the flux varies from zero to a constant level as a result of starting or stopping the observation, can appear as sharp edges during the dispersion estimation analysis and introduce structure that is not really there. We then discuss our analysis of the most variable run in our data set and show how the estimated LIV dispersion can be used to constrain the quantum gravity energy scale, E_{QG}.

4.4.1 Techniques for Handling a Constant Background

Using a source with flares occurring on top of a constant background necessitates careful handling of the boundaries of finite duration observations. To illustrate why, assume that an astrophysical source emits a pair of gamma rays simultaneously, one with a lower energy and one with a higher energy. Also assume for the moment (without loss of generality) that $s_{\pm} = +1$ in Equation (4.3); that is, we are working with subluminal LIV. In this case, the LIV dispersion has the effect of slowing high-energy photons, delaying the observed arrival time of high energy photons with respect to the arrival time of a lower energy photon that was emitted concurrently at the source. If our hypothetical pair of gamma rays would normally arrive just before we start our observations in the absence of LIV, then in the

presence of LIV we may find that we observe the high energy gamma ray, since it is delayed, but not the low energy gamma ray. Likewise, if the pair of gamma rays would arrive just before we cease observing without LIV, then with LIV we may not observe the high energy photon. We will call these pairs of high and low energy photons counterparts.

Both dispersion estimation methods are ideally applied in situations where all photons in a flaring episode are observed. This is not possible if the source has a constant background. Instead, we choose an uninterrupted observation window which is not near the start or end of an observation, then attempt to include all of the counterpart photons to those that occur inside our window. Due to the wobble mode of observation (see Section 4.2), the observations are separated into 20 - 30 minute runs with gaps (on the order of a minute long) between runs while the telescopes move to the next wobble position. This necessitates choosing a time interval within a run.

This interval is determined by "padding" the run start and stop times by some amount x; e.g., our interval is (run start time + x, run end time - x). The amount of padding x is chosen by considering the highest energy photon within the run (for run #10, this is 11.35 TeV) and the range of possible τ_n we expect to find. The LIV delays are largest for quadratic LIV (see Equation (4.4)), so making a rough estimate that $|\tau_2| \leq 1$ s/TeV^2 and rounding our maximum photon energy to 10 TeV, we choose a padding of $(1 \text{ s/TeV}^2)(10 \text{ TeV})^2 = 100$ seconds.

This padding is also useful for dealing with an additional concern, which we alluded to at the end of Section 4.3.3. The presence of a constant background can artificially introduce structure which is not related to the flaring structure on top of the constant background,

which is what we are interested in. If not handled properly, the dispersion estimation algorithms will pick up the edges of the observation where the flux increases rapidly from zero, just before the observation starts, to the constant background level. Essentially the entire observation will be reduced to one large square pulse. Since the dispersion estimation algorithms are looking for this type of sharp variability, they will produce artificially good results.

We note that these issues are not discussed in the literature because the flaring structure in the light curves used in nearly all LIV investigations is not sitting on top of a constant background. GRBs, which are the most commonly used sources for LIV studies (see Section 1.3), are ephemeral and can be captured in their entirety with a single observation. Pulsar light curves contain pulses on top of a constant background, but the epoch folding process eliminates issues resulting from finite-length observations. Finally, in three out of four published LIV studies using AGN flaring, the light curves of the flares appear to go to zero at the edges – see Figure 1 of Biller et al. (1999), Figure 2 of Abramowski et al. (2011), and Figure 6 of Abramowski et al. (2015).

Albert et al. (2008) is the only other study to use an AGN flare on top of a constant background. They use an energy cost function (ECF) method that is essentially a dispersion cancellation method (it uses an ECF in place of the Shannon entropy as a metric for sharpness), in addition to a likelihood method (Martínez & Errando, 2009). Their approach is to choose a single flare contained completely within a run, which avoids the problems discussed above. Our data does not contain a single, large flare like the data in Albert et al. (2008); instead, we have a series of small flares. We show the results of using a single flare from run

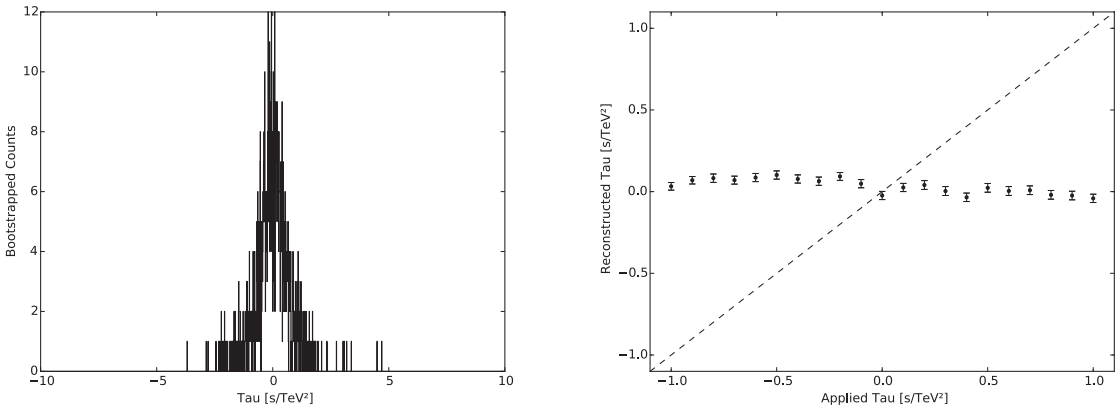

Figure 4.11. Plots of the bootstrap distribution (left) and reconstructed vs. applied dispersion (right) for a single flare from run #10, occurring at 41500 sec. elapsed since midnight UTC on MJD 55244. The dashed line in the plot on the right indicates perfect reconstruction (see Section 4.3.4). Both plots are generated using the dispersion cancellation method for quadratic LIV. Although the bootstrap distribution is very sharp, the flatness of the reconstructed vs. applied dispersion plot indicates it is not possible to reliably estimate the dispersion in these data.

#10 (occurring at approximately 41500 sec. in Figure 4.2) in Figure 4.11; the flatness of the plot on the right indicates that selecting a single sub-flare from our data is not feasible.

4.4.2 Results from Most Variable Run

Here we present the results searching for dispersion in the most variable run in our data set, run #10. We cross-check our results using both of the dispersion estimation methods. Because the flaring structure sits on top of a constant background, we are careful to handle the data using the boundary handling methods outlined in Section 4.4.1.

We first assess the ability of our algorithms to determine the amount of dispersion present in the data. This is done using the reconstructed vs. applied dispersion technique in Section 4.3.4. In Section 4.3.4, we showed that with a Gaussian pulse the algorithms were easily able to reconstruct the amount of dispersion present in the data. However, as we

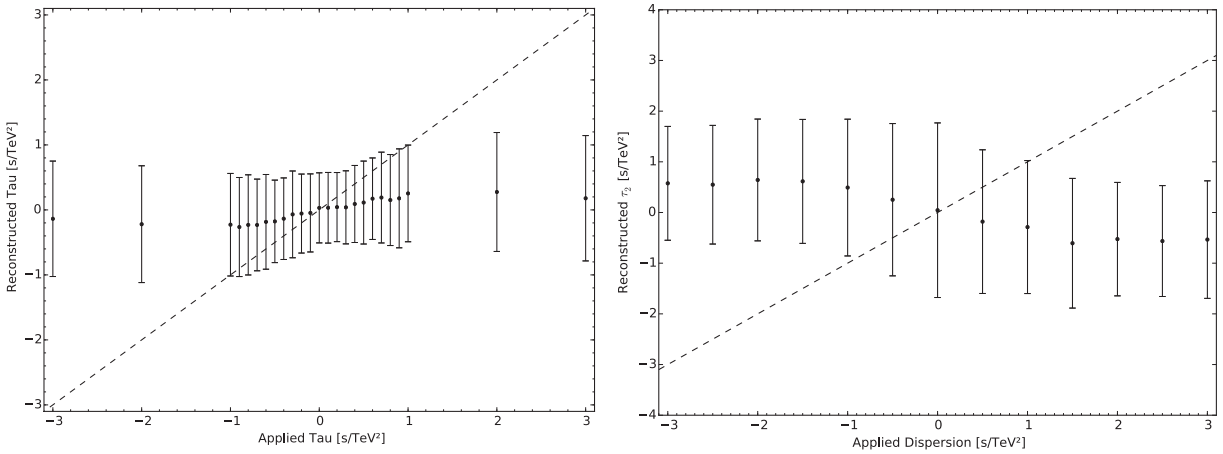

Figure 4.12. Reconstructed vs. applied dispersion plots for the most variable run in the data set, run #10. The dispersion cancellation method (left) produces marginal results in the window $-1 < \tau_2 < 1$, while the PairView method (right) fails completely.

noted in Section 4.3.4, the ability of the algorithms to determine the amount of dispersion in the data depends strongly on the amount of structure in the data. Therefore, the plots of reconstructed vs. applied dispersion can be used to determine the suitability of data for an LIV investigation. These plots are shown in Figure 4.12, where the dashed diagonal line represents perfect reconstruction. Both algorithms provide poor reconstruction of the applied dispersion; the PairView algorithm completely fails to reproduce the expected slope of $+1$.

The dispersion cancellation method manages to produce marginal results over the restricted range $-1 < \tau_2 < 1$. These results are unacceptable, but for the purpose of demonstrating the method, we fit a line to the dispersion cancellation reconstructed vs. applied dispersion plot and convolve it with the dispersion cancellation bootstrap. Fitting a line to the reconstructed τ_2 values produces a fit of $y = 0.258x - 0.0158$. Although this fit is only really valid in a range of about $-0.25 < \tau_2 < 0.25$, we apply the fit to each value

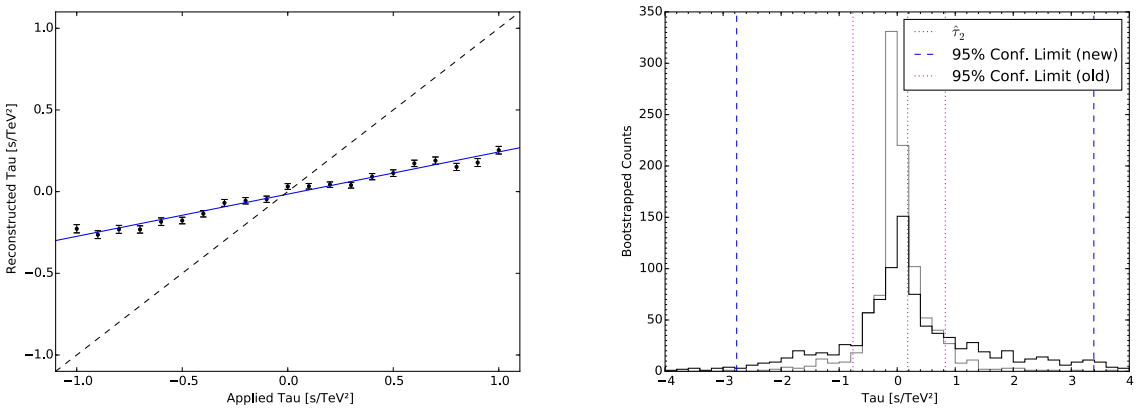

Figure 4.13. The marginal response of the dispersion cancellation algorithm over a narrow range of τ_2 values is convolved with the bootstrap distribution to produce ersatz uncertainties that will be used to demonstrate the calculation of quantum gravity energy scale limits. *Left*: Reconstructed vs. applied dispersion plot for run #10 fit with a straight line in the interval $-1 < \tau_2 < 1$. The fitted line is used to model the algorithm response. *Right*: Bootstrap distribution convolved with the algorithm response (shown with a solid black line), which has the effect of spreading the original bootstrap (solid gray line). The old uncertainties are shown with dotted vertical magenta lines, while the new uncertainties are shown with dashed vertical blue lines.

of $\hat{\tau}_{2,old}$ in the bootstrap distribution to produce new values $\hat{\tau}_{2,new} = (\hat{\tau}_{2,old} + 0.0158)/0.258$. The results of this process are shown in Figure 4.13. From the convolved bootstrap, we derive upper and lower limits on the estimated quadratic energy dependent dispersion of $\tau_{lower} = -2.77\text{s/TeV}^2$ and $\tau_{upper} = -3.39\text{s/TeV}^2$.

To obtain limits on the QG energy scale from the $\hat{\tau}_n$ values, we use the definition of τ from Equation (4.3) and solve for the QG energy scale E_{QG}:

$$E_{QG} = \left[\frac{s_\pm}{\tau_n} \frac{1+n}{2H_0} \int_0^z \frac{(1+z')^n \, dz'}{\sqrt{\Omega_m (1+z')^3 + \Omega_\Lambda}} \right]^{1/n} . \tag{4.11}$$

For the cosmological parameters, we use the recent values obtained from *Planck* measurements of the CMB: $H_0 = 67.3$ km/s/Mpc, $\Omega_m = 0.315$, and $\Omega_\Lambda = 0.685$ (Ade et al., 2014). The redshift of Mrk 421 is $z = 0.031$. The 95% confidence upper limit τ_{upper} is used to place

lower limits on the QG energy scale for subluminal LIV ($s_{\pm} = +1$), whereas τ_{lower} is used for superluminal LIV ($s_{\pm} = -1$). LIV limits from threshold effects strongly constrain superluminal LIV (Stecker, 2003), but we will calculate limits for subluminal LIV and superluminal LIV for completeness. For quadratic LIV, we find a superluminal limit of $E_{QG,2} > 8.88 \times 10^{10}$ GeV and a subluminal limit of $E_{QG,2} > 8.02 \times 10^{10}$ GeV. We reiterate again that these results are not reliable and are for demonstration purposes only.

4.5 Conclusions

We have searched for photon dispersion arising from LIV in VHE observations of a large flare of the blazar Markarian 421. We established that the most variable run in our data set, run #10, contains significant variability, but that the flaring structure sits on top of a constant background. We described two methods for estimating the dispersion in our data and showed that both methods are capable of accurately estimating the dispersion in an ideal data set consisting of a single Gaussian pulse. We also discussed techniques we developed specifically for analyzing data sets containing flaring structure on top of a constant background, a situation which is not discussed in detail elsewhere in published literature. Finally, we applied our methods to the most variable run in our data set.

Unfortunately, we found that the presence of a constant background severely limited our ability to accurately estimate the dispersion in our data. The reason that flares on top of a constant background are not generally discussed in the literature is likely because such data is generally not desirable for LIV investigations unless the flare is exceptionally strong. Nonetheless, we attempted to correct for the poor response of our dispersion cancellation

algorithm and generated results for the purpose of demonstrating how to use a dispersion estimate to calculate a lower limit on the quantum gravity energy scale. Our subluminal quadratic LIV result is only slightly worse than the current best limit from *Fermi*-LAT of $E_{QG,2} > 1.3 \times 10^{11}$ GeV (Vasileiou et al., 2013), slightly better than the H.E.S.S. limit from PKS 2155-304 of $E_{QG,2} > 6.4 \times 10^{10}$ GeV (Abramowski et al., 2011), and a factor of approximately 4 better than the quadratic limits set by MAGIC observations of Markarian 501 Albert et al. (2008). However, our result cannot be considered reliable.

Our main result is in the development and testing of methods to estimate the dispersion in AGN flaring observations containing a constant background for the purpose of searching for LIV. This work represents the first application of the dispersion cancellation method described by Scargle et al. (2008) and the PairView algorithm (Vasileiou et al., 2013) to blazar flare data from an imaging atmospheric Cherenkov telescope (IACT). We tested a modified method to represent the light curve using cells and an alternative probability distribution than those used in Scargle et al.'s original DisCan algorithm and found that our modifications produced good results when used with light curves without a constant background. Although modifications to the dispersion cancellation technique have been discussed by others (e.g., de Almeida & Daniel (2012)), the methods we presented differ from those discussed elsewhere in the literature. We have also discussed several problems we discovered in the course of our investigation related to the presence of the constant background in our data, and we have presented options for avoiding or ameliorating its effects.

An obvious extension to this study would be to further characterize flaring data on a constant background. We have shown that our dispersion estimation algorithms work

well in the absence of a constant background, and that they do not work well when a large background is present. It is clearly possible to estimate the dispersion of a large flare on top of a constant background, since this has been done in Albert et al. (2008) using a dispersion cancellation technique similar to the one we present here. So how large of a background can the dispersion estimation algorithms tolerate? The answer to the question can likely be found through the use of extensive Monte Carlo simulations, but this would require a significant amount of time since dispersion cancellation algorithms are computationally expensive.

In light of our results, we recommend that further LIV investigations should only be undertaken by IACT collaborations provided that certain requirements are met in advance. First, we note that observations of a very large flare do not guarantee that the data will be suitable for an LIV investigation. Ideally, a single observation would see the flux increasing from zero, varying significantly, and then dropping back down to zero again before the observation ends. Second, there should not be breaks in the observation while the telescope wobbles to a new location. One way to achieve this would be to make observations after a bright AGN flare from a single wobble position with very long runs to minimize gaps in the data. Data from other wobble positions could be taken on subsequent nights. A better alternative might be to employ orbit-mode observations (Finnegan, 2011) in which the telescope wobbles continuously. Finally, only segments of observations that contain variability at minute or faster time scales should be used for an LIV analysis; an enhanced level of flux does not necessarily imply that the data will be useful for an LIV analysis.

Part II

ALIGNMENT OF GAMMA-RAY TELESCOPES

CHAPTER 5

VERITAS POINTING MONITORS

5.1 Introduction

The VERITAS pointing monitor (VPM) system is designed to provide accurate measurements of the pointing of each of the four VERITAS telescopes. The system is designed to supplement the pointing information provided from the mechanical encoders in each telescope's positioner. The positioners are RPM-PSI's model PG-4000 (Gibbs et al., 2003) and are specified to have a backlash of $< 0.01^\circ$ (36 arcsec) and an encoder precision of 0.014° (50 arcsec) for the azimuth encoders and 0.011° (40 arcsec) for the elevation encoders. Adding these precisions in quadrature, the total precision with which the position of the telescopes can be determined using only the encoders is not expected to be better than 64 arcsec. As will be discussed in Section 5.4, the VPM system can determine the pointing of each telescope with an accuracy of 20 arcsec or better under typical observing conditions.

In this chapter, we will describe the hardware and software components of the VPM system. We will then discuss the performance of the VPM system in terms of how accurately it can measure the location on the celestial sphere that the telescopes are pointed towards.

5.2 Pointing Monitor Hardware

The VPM system uses two Prosilica EC1380 CCD cameras per telescope with a Firewire 400 interface for data and power. The cameras use a Sony ICX-285 CCD sensor with a 2/3 format and a resolution of 1360 x 1024 (1.4 megapixels). Each pixel is 6.45 x 6.45 microns in size and can be sampled with a depth of 8 or 12 bits. The camera uses a global

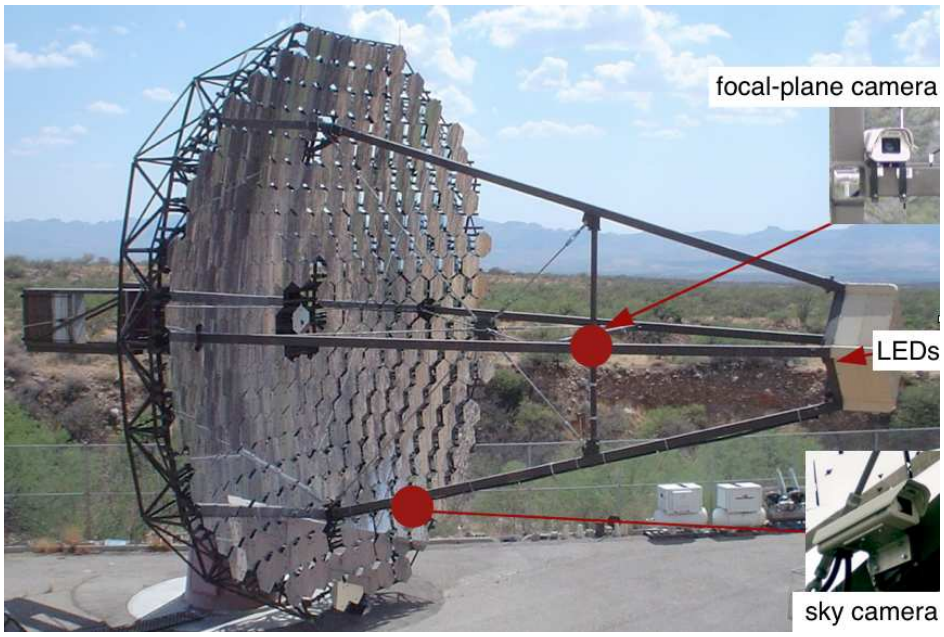

Figure 5.1. Photo showing a VERITAS telescope with the locations of the three hardware components of the VPM system. Figure taken from Grube et al. (2011).

shutter (as opposed to a rolling shutter), which provides a better image when the telescope is in motion and a more uniform exposure when the exposure time is long. The cameras attach (via C-mount) to Pentax F/1.4 lenses. A 50 mm lens (Pentax C25011TH) is used for the focal plane camera, which sits on the quadrupod (four-arm) crossbar at a distance of 4.9 meters away from the telescope focal plane. A long 75 mm lens (Pentax C27509KA) is used for the sky cameras, yielding a pixel scale of 17.74 arcsec/pixel and a field of view of 6.7 x 5.0 degrees. The cameras are contained inside Pelco weatherproof enclosures and mounted to the telescope via custom-machined brackets.

Each telescope has a single VPM computer mounted to the optical support structure (OSS) of the telescope on the non-reflecting side of the mirrors. The use of ruggedized outdoor computers is necessitated by the Firewire interface, which has a maximum cable length of 4.5 meters (specified by the IEEE 1394 standard).Cables running to and from the

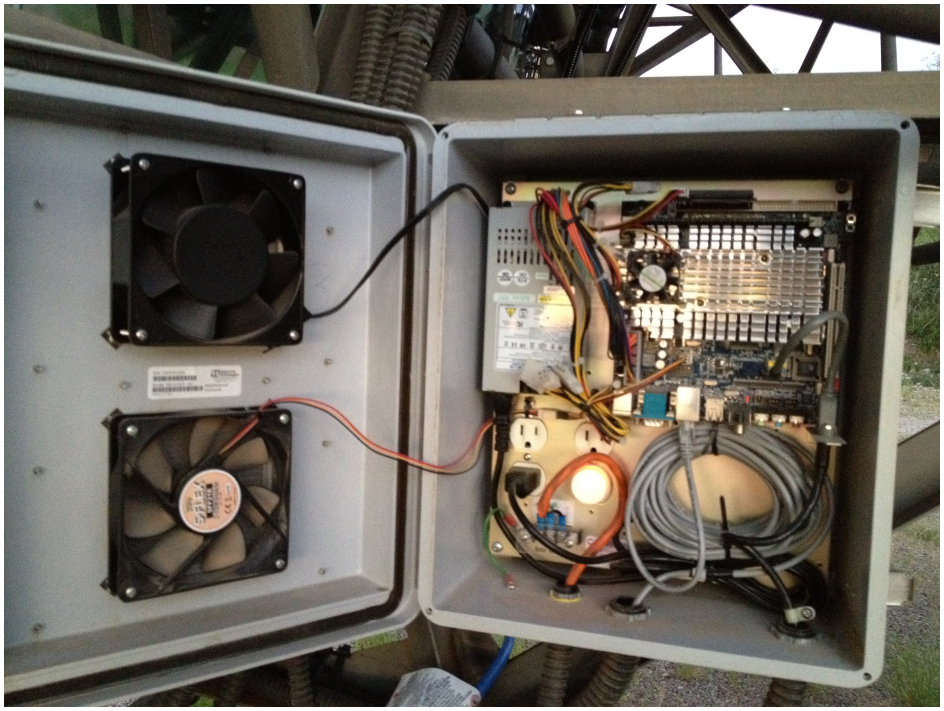

Figure 5.2. Photo of a VPM control computer *in situ*. The computer is attached to the OSS via aluminum angle brackets (not visible). Conduits for power, ethernet, and Firewire (from left to right) can be seen at the bottom of the enclosure. A small light can be seen in the bottom right 120V receptacle, which provides light at night if observers need to open the enclosure.

computer are housed in black plastic conduit to protect the cables from exposure and from wildlife (which sometimes chew through cables). The power cable uses 1/2" conduit, while the ethernet and Firewire use 3/4" conduit; the larger diameter is needed to accommodate the connectors. The enclosure is weatherproof and contains two 120 mm fans with heavy-duty dust filters to provide cooling. The computers are only run at night since daytime temperatures in Arizona are often high enough to cause the computers to overheat.

The parts list for the VPM computers can be found on the VERITAS internal wiki[1] and is presented here for completeness with some modification to the formatting.

[1] `https://veritas.sao.arizona.edu/wiki/index.php/Pointing_Monitor_Maintenance`

- L-com 14x12x7 inch Weatherproof Fiberglass Enclosure (NB141207-10F)
- VIA EPIA-EN15000G CN700 1.5GHz Mini ITX Motherboard
- Kingston ValueRAM 1GB DDR2 533 (KVR533D2N4/1G)
- Sparkle FSP180-50PLA 180W ATX Power Supply
- SiliconDrive 1GB Compact Flash (SSD-C01GI-3100 or SSD-C01GI-3500)
- SYBA IDE to CompactFlash Adapter (SD-CF-IDE-DI)
- Unibrain Fireboard-Blue IEEE-1394 FireWire Adapter
- Scythe S-FLEX SFF21E 120mm Case Fan

The VPM computers run Linux off of a 1GB compact flash card, which are inexpensive and easy to replace. They are designed so that the components can all be replaced quickly and easily, including the motherboards. The components are all commercial, off-the-shelf parts and are not designed for operation in the extreme conditions of the Arizona Sonoran desert. Despite dust filters on the fan intakes, the fans eventually become clogged with dust and require replacement. A few of the power supplies have failed due to aging and exposure to the elements and have been replaced. The VPM computers sometimes become unresponsive and must be replaced from a stock of spare motherboards. Nevertheless, most of the components have proven to be remarkably durable, and the ease with which failed components can be replaced has allowed the VPM computers to successfully operate with little downtime since November 2006.

It should be noted that the VPM computers run Ubuntu Server Version 6.06 LTS, which is no longer supported. An attempt was made (by Dirk Pandel) to upgrade to Ubuntu 8.04 LTS several years ago, but it was found that the internal Linux clock would go out of sync very quickly (by as much as 5 minutes every hour) whenever the Firewire cameras

were capturing images. The VPM system requires that the computers keep accurate time because the image timestamps are used to correlate the VPM images with the the telescope tracking. The reason for the time lag was never determined, but is likely due to the direct memory access (DMA) transfer from the cameras causing the motherboard to miss some of the interrupt requests (IRQs) that advance the clock. This problem was not observed in Ubuntu 6.06, so it may be the result of a change to the Linux kernel. No further attempt has been made to upgrade the operating system of the VPM computers.

The last hardware component of the VPM system is a set of two LED modules which are securely mounted to the light cone plates, one on either side of the PMT camera. There are two red LEDs in each module for a total of four LEDs per telescope camera. The modules are connected to the preamplifier power supply using 6-position, 6-conductor RJ25 telephone cables (the same cable that is used for the current monitor boards). The focal plane camera images the four LEDs and uses their known positions on the light cone plate and their known separations to compute the location of the center of the PMT camera.

5.3 Pointing Monitor Online Software

The VPM online software is the software that runs on the VPM computers during normal VERITAS operations. It was written by Dirk Pandel at the University of Iowa with contributions from David Steele at the Adler Planetarium and many others. The main function of the online software is to capture images (at a rate of 0.5 Hz, or once every two seconds) from the two VPM cameras on each telescope, find the positions of the stars and LEDs in the images, and store the information from the images in the VERITAS database.

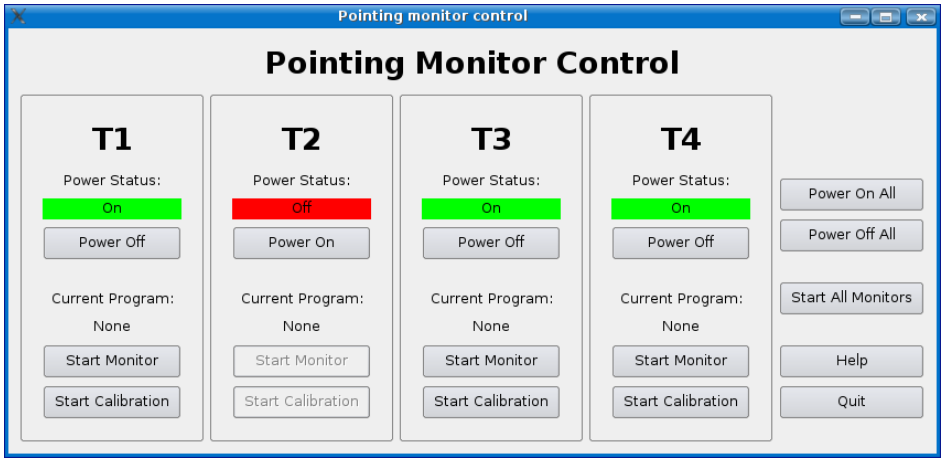

Figure 5.3. Screen capture of the pointing monitor control GUI, which is the main portal of the online software. This window is used to remotely power and shut down the VPM computers and to launch the Monitor and Calibration programs. Image credit: Dirk Pandel.

It is primarily written in C++ and uses the Qt 4 framework for its graphical user interface (GUI). Two libraries are required to capture images from the Firewire (IEEE 1394) cameras: libdc1394 version 1.2.2, and libraw1394 version 1.2.1. The VERITAS online software also includes rich functionality for interacting with the VERITAS SQL database. The VPM offline software, which post-processes data taken online, will not be discussed here.

When launched, the VPM software displays a control GUI (Figure 5.3). The control GUI is used to remotely power the VPM computers via wake-on-LAN and to shut them down at the end of each night. When powering up the VPM computers, the program polls the VPM computers over the network until it receives an acknowledgement, at which point it updates the VPM computer's power status. This process usually takes several seconds. The other function of the control GUI is to act as a portal to the other two VPM online software applications, the Monitor and Calibration programs.

The Monitor application is the workhorse of the online VPM software. It consists of a graphical front-end (shown in Figure 5.4) which shows observers live images from the

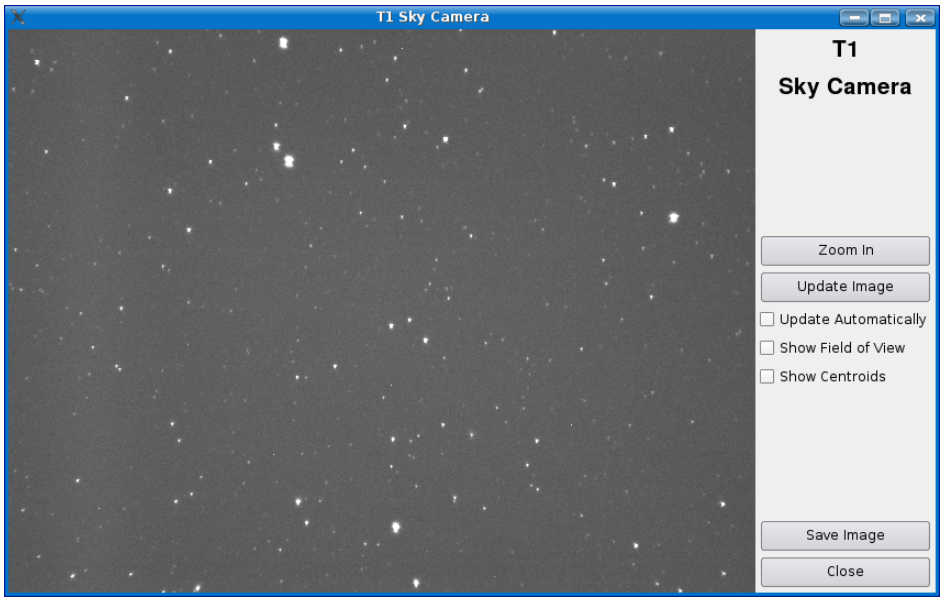

Figure 5.4. Screen capture of the Monitor application showing the view from the telescope 1 (T1) sky camera. An identical window is also available that shows the view from the focal plane camera. Image credit: Dirk Pandel.

focal plane and sky cameras from each telescope. These live images afford a quick visual confirmation of telescope pointing that is independent from the telescope tracking software. Other practical uses include allowing observers to see when clouds drift into the field of view of the telescopes and checking that the telescope camera shutters are open.

However, the main purpose of the Monitor program is to capture images from both of the cameras on each telescope every two seconds. The images from the focal plane cameras are processed to detect the locations of the four LEDs that are attached to the light cone plate inside each telescope's PMT camera (see the end of Section 5.2). The modified Julian date (MJD) and pixel coordinates (x,y) of the LEDs are written to the `tblPointing_Monitor_TelescopeX_LEDs` table in the `VERITAS` database, where the `X` in the table name is the telescope number running from 0-3 (i.e., the telescope number minus one, so that the telescope number follows the C-programming convention of numbering array

elements from zero). Each image from the sky cameras is processed to detect the 30 brightest stars in the image (typically 60-70 stars are visible in each image). The modified Julian date and the centroid pixel coordinates and brightness of the 30 brightest stars are all written to the `tblPointing_Monitor_TelescopeX_Centroids` table, which is also in the `VERITAS` database.

The final element of the online software is the Calibration program, which is used to perform VPM calibrations, point spread function (PSF) measurements, and telescope pointing (T-point) calibrations. Unlike the Monitor program, which does most of its work silently without user interaction, the Calibration program is driven by its graphical user interface (see Figure 5.5). Interactions with the user are limited to the ability to save images and the ability to adjust the exposure time (which controls the camera shutter speed). It captures images only when the user clicks the "Save images" button.

Both the Monitor and Calibration programs can centroid stars and LEDs in the images in real time and display their positions in the GUI. This can be seen in Figure 5.5, where the centroids of the four camera LEDs are indicated by blue crosses on the left and right sides of the image, and the centroid of an imaged star is shown by a red cross near the center of the image. The center of the PMT camera is also computed using the known positions of the LEDs on the light cone plate, and is shown with a blue cross in the center of the image. The Calibration program also computes the angular offset between the star centroid and the center of the camera.

VPM calibrations and PSF measurements are typically performed using a single set of measurements that are taken using the Calibration program. Detailed instructions on

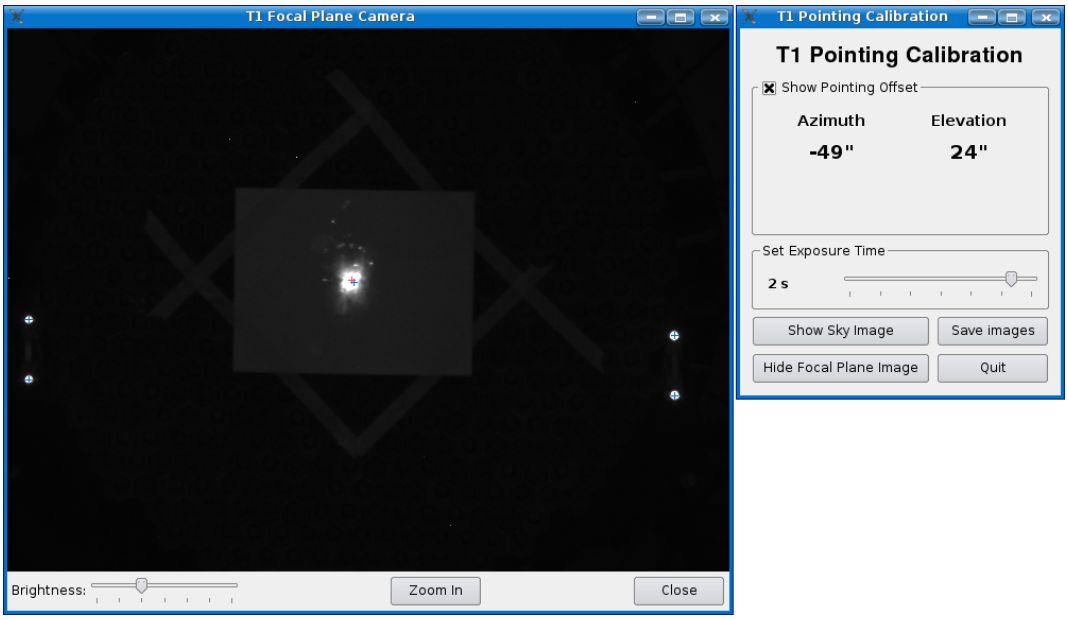

Figure 5.5. Screen capture of the Calibration application, which is used during PSF and T-point calibrations. This image is typical for a PSF calibration and shows a star imaged on a white screen in the center, as well as the four LEDs which are attached to the light cone plate on the camera. The blue crosses show the computed centroids of the LEDs as well as the location of the center of the telescope camera, which is computed from the LED locations. The red cross shows the centroid of the image of the star. Image credit: Dirk Pandel.

how these measurements are taken can be found on the VERITAS wiki page at `https://veritas.sao.arizona.edu/wiki/index.php/VPM_Calibration`.

5.4 Pointing Monitor Performance

The performance of the VPM system depends primarily on two factors: the ability to determine the pointing of the sky camera via accurate astrometry, and the ability to relate the pointing of the sky camera to the center of the telescope's focal plane. The former is related to the imaging capabilities of the VPM sky cameras and the quality of the astrometry package used to determine the pointing of the sky camera in celestial coordinates (right ascension and declination). The latter depends on the calibration of the VPM system

and misalignments within the VERITAS telescopes. We will discuss both of these factors in the order listed, followed by a brief discussion of lower order effects which can modify the performance of the VPM system in certain conditions.

The resolution of the VPM sky cameras is 17.74 arcsec per pixel, a quantity which is determined by the focal length of the lens that is used (75 mm) and the size of the pixels on the CCD sensor. The point spread function of the sky camera optics is such that stars are spread over several pixels. This is advantageous because it allows stars to be centroided with sub-pixel accuracy. A reasonable estimate of the accuracy of the stellar positions in a sky camera image is $\sim$5 arcsec. The pointing of the sky camera is determined by Astrometry.net (Lang et al., 2010), which takes the positions of the 30 brightest stars in each sky camera image and computes the celestial coordinates the sky camera is pointed at. It has been established by Grube et al. (2011) that the uncertainty in the Astrometry.net solution is dominated by the uncertainties of the stellar positions in the sky images, and that the total uncertainty in the sky camera pointing is $\sim$5 arcsec.

Most of the uncertainty in the VPM pointing solutions comes from bending in the optical support structure (OSS; see Section 3.3). The monthly VPM calibration accounts for bending in the OSS by finding a functional relationship between the normal vector of the sky camera to the normal vector of the telescope's PMT camera. This calibration is performed approximately monthly at the beginning of each new dark run (dark runs are blocks of observing days where the moon is not fully illuminated; they typically start three nights after a full moon). Grube et al. (2011) found that the change in the monthly calibration equates to a $\sim$20 arcsec uncertainty in each telescope's pointing. The variations in the VPM

calibration appear to be the inversely related to the stiffness of the optical support structures of the telescopes, and may be influenced by factors such as average ambient temperature and wind loading.

One source of uncertainty that is not currently accounted for in the VPM software is atmospheric refraction. VPM calibrations are performed at elevations between 20-90 degrees (i.e., zenith angles from 0-70 degrees), so in principle the VPM corrections should be accurate over this range of zenith angles. However, at large zenith angles (greater than 60 degrees), Cherenkov light from high altitude extensive air showers is refracted 60-70% as much as stellar light (Bernlöhr, 2000). Since stellar light is used for the VPM calibration, this can cause an additional uncertainty in the gamma-ray source localization. At a zenith angle of 60 degrees, stellar light is refracted by approximately 100 arcsec (Auer & Standish, 2000). Since Cherenkov light is only refracted 60-70% as much, this will lead to an additional uncertainty of 30-40 arcsec in the source localization for observations made at large zenith angles.

CHAPTER 6

AUTOCOLLIMATOR ALIGNMENT SYSTEM

6.1 Global Alignment System Overview

This chapter will discuss an autocollimator alignment system which will be used as part of the alignment system for a prototype 12 meter diameter Schwarzschild-Couder telescope (SCT). This telescope represents the United States' contribution to the Cherenkov Telescope Array (CTA; the US contribution is known as CTA-US), which is the next generation of imaging atmospheric Cherenkov telescopes (IACTs). The Schwarzschild-Couder design is a recent innovation (Vassiliev et al., 2007) and differs most noticeably from traditional IACT optical designs in its use of two reflecting surfaces instead of one (see Figure 6.1). The use of two reflectors makes the task of aligning the telescope considerably more difficult since it requires that the primary and secondary mirrors are independently aligned and that the primary and secondary mirrors and the camera are aligned with each other.

The global alignment system works in concert with the panel-to-panel alignment system (P2PAS), which is being developed by Columbia University to measure the relative displacements of each mirror segment relative to its neighbors. This system rigidly mounts small, low-powered lasers on each mirror segment which shine onto cameras fixed to adjacent panels. The path of the laser beam is enshrouded within a flexible enclosure to prevent external light from reaching the cameras. As the panels shift, the cameras image the movement of the laser's spot. By combining multiple sets of lasers and cameras, the relative alignment of adjacent panels can be completely constrained. The alignment systems cannot measure the absolute alignment of the mirrors. Absolute alignment will require utilizing bright stars

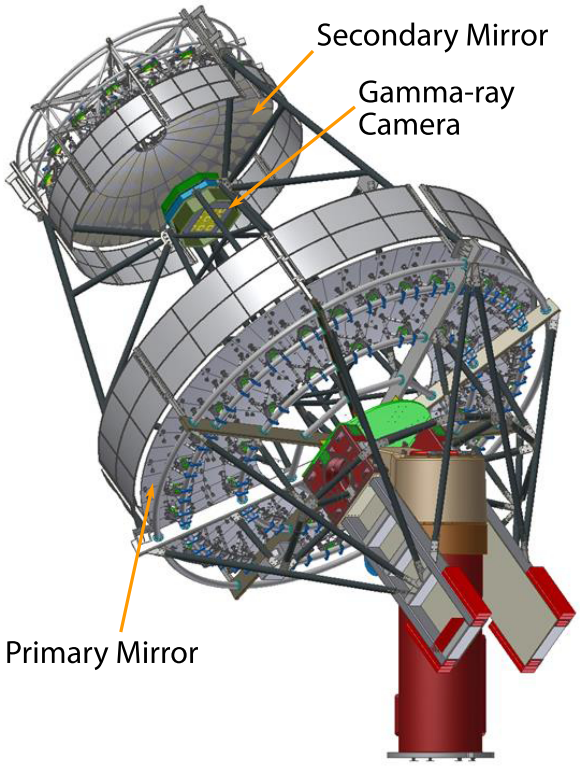

Figure 6.1. Rendering of a CAD model of the prototype CTA-US Schwarzschild-Couder telescope. Image credit: Victor Guarino, Argonne National Laboratory.

to initially align the mirrors to obtain the best image quality. After initial alignment, the P2PAS system will maintain the proper alignment of each of the mirror panels. The alignment of the primary and secondary mirrors and the telescope camera will be handled by the global alignment system.

The global alignment system is tasked with determining the positions of a handful of mirror segments in the primary reflector with respect to the center of the secondary reflector, and vice versa. Only one mirror segment in each of the primary and secondary mirrors needs to be positioned in space; the P2PAS system can set the alignment of the rest of the mirrors once one is positioned correctly. Small errors in the relative positions of each panel will accumulate as the distance away from a single aligned panel increases, so the

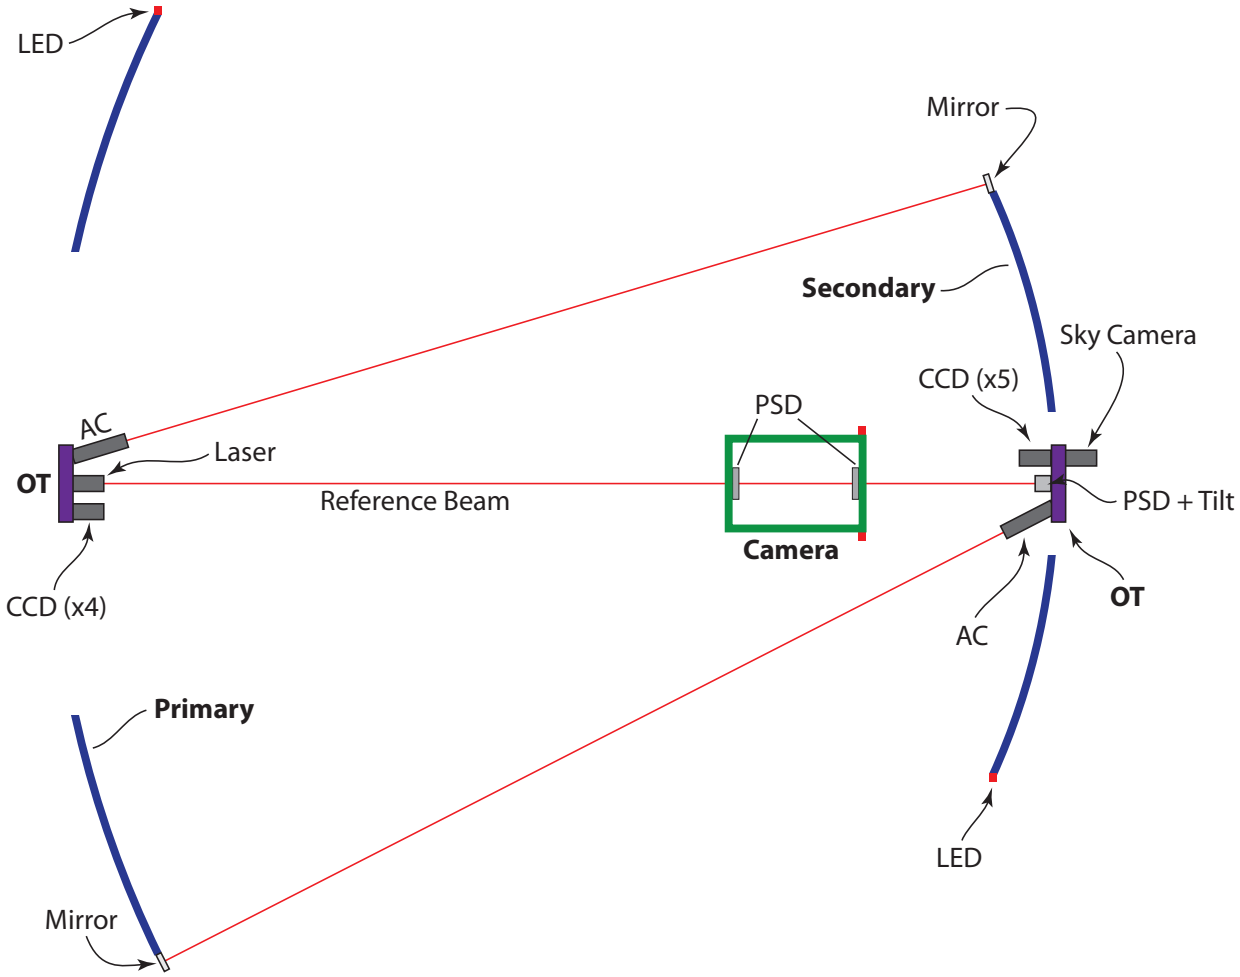

Figure 6.2. Diagram of the CTA-US global alignment system. Iowa is contributing the autocollimators and associated mirrors, CCDs (including the sky camera), and LEDs. The optical tables, PSDs, and the reference-beam laser (which defines the optical axis) will be contributed by UCLA.

global alignment system will be used to check the cumulative errors in the P2PAS alignment. To accomplish this, the global alignment system should measure the positions of at least two panels in each of the primary and secondary mirrors, although the positions of more mirror panels may be measured to provide redundancy.

The global alignment system will continuously monitor the positions (three linear axes and three axes of rotation) of the camera and primary and secondary mirrors relative

to one another. These measurements must be made at a distance since the primary and secondary are ∼9 meters apart and the center of the secondary is approximately 2 meters away from the gamma-ray camera focal plane. Figure 6.2 shows a schematic of the global alignment system. The system relies on a large number of high-resolution grayscale cameras, which are discussed in detail in Section 6.3. We will refer to these cameras as CCDs (even though they use CMOS image sensors) to disambiguate them from the gamma-ray camera. It also uses several semi-transparent position sensitive detectors (PSDs). The main elements of the global alignment system are:

1. CCDs that image LEDs mounted to mirror panels
2. CCDs that image LEDs mounted to the gamma-ray camera
3. CCD-based autocollimators to measure the angular offsets between the optical tables and mirror panels
4. A laser to generate a reference beam which defines the optical axis
5. PSDs to measure the position of the gamma-ray camera
6. A PSD and tilt sensor to measure the relative alignment of the optical tables

The design of the global alignment system did not initially include autocollimators; these were added later in response to the alignment requirements for the prototype SCT becoming more stringent. Table 6.1 lists the current alignment requirements for the prototype SCT. Figure 6.3 shows the outline of a mirror panel with the definitions of the parameters in Table 6.1, which are the six degrees of freedom of each mirror panel. The autocollimator was added to meet the very strict requirements for $\Delta\theta_x$ and $\Delta\theta_y$ of ± 0.025 mrad or ± 5.1

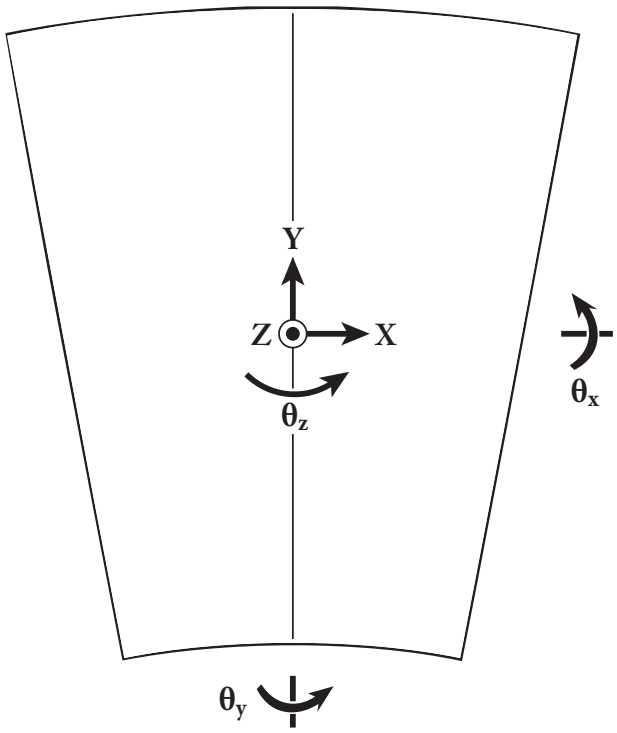

Figure 6.3. Diagram showing the six degrees of freedom (three linear, three rotation) for each mirror panel.

arcsec. In fact, $\Delta\theta_x$ and $\Delta\theta_y$ are the only parameters that the autocollimators are capable of measuring.

6.2 Autocollimator Principle of Operation

The operational principle behind an autocollimator is very simple. A light source illuminates a target object, which is typically either a crosshair reticle or a pinhole. This

Table 6.1. CTA-US global alignment system requirements

Parameter	Primary $\leftrightarrow$ Secondary	Secondary $\leftrightarrow$ Camera	Comments
$\Delta x, \Delta y$	$\pm$0.25 mm	$\pm$0.25 mm	perpendicular to optical axis
Δz	$\pm$8 mm	$\pm$1.2 mm	along optical axis
$\Delta\theta_x, \Delta\theta_y$	$\pm$0.025 mrad	$\pm$0.025 mrad	relative to optical axis
$\Delta\theta_z$			no requirements

target is located at the focal point of an objective lens. An eyepiece or digital camera sensor is likewise located at the focal point of the same objective lens. Locating both objects at the focal point of the objective lens requires using a beam splitter, which is typically either a thin piece of glass or a pellicle (a thin plastic membrane with a coating to make it partially reflective). Because the illuminated target is located at the focal point of the objective lens, the diverging rays from the target will emerge from the lens collimated (i.e., parallel with one another). The collimated light emitted by the autocollimator is reflected by a mirror or other reflective surface, and provided the reflective surface is flat, the reflected rays will remain collimated. The collimated rays then enter the autocollimator through the objective lens and are focused onto a digital image sensor (in traditional autocollimators, the reflected light formed an image on another crosshair reticle, and an eyepiece was used to visually compare the offset between the eyepiece crosshair and the image of the illuminated, reflected crosshair). This is shown schematically in Figure 6.4.

If the plane of the mirror is orthogonal to the optical axis, the reflected rays will form an image at the center of the image sensor (assuming that the optical axis passes through the exact center of the image sensor, which is unlikely in practice). If, however, the mirror is tilted at an angle α with respect to the optical axis, then the rays will be reflected at an angle of 2α with respect to the optical axis. This is because the angle of incidence between the incoming rays and the mirror's surface normal is α, so the angle of reflection *relative to the mirror normal* is also α and the sum of the two is 2α. The displacement between where the image would form if the mirror were orthogonal to the optical axis and where the image

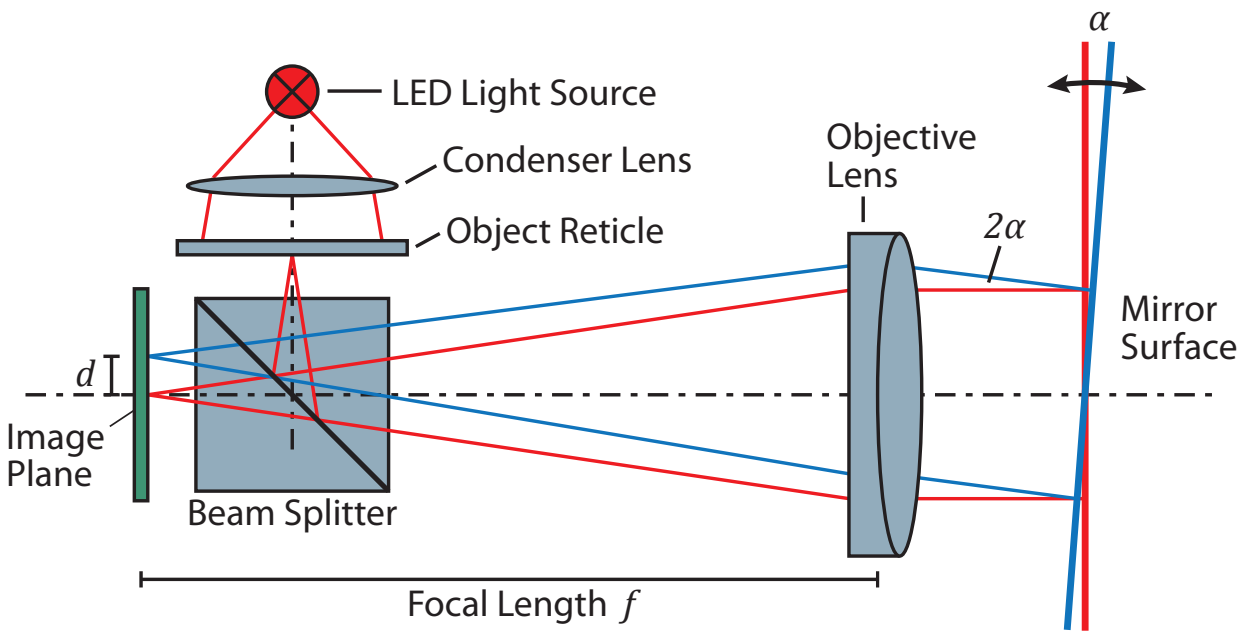

Figure 6.4. Diagram showing the parts and operational principle of a typical digital autocollimator. The red rays illustrate the scenario where the mirror (shown as a thick red line) is perpendicular to the optical axis; the blue rays show the path taken by rays reflected by a mirror (shown as a thick blue line) which is inclined at an angle α relative to the optical axis. The reflected rays make an angle of 2α with respect to the optical axis. This diagram is modeled after a figure by TRIOPTICS GmbH, a manufacturer of digital autocollimators; the original figure can be found at `http://www.trioptics.com/knowledge-base/optical-angle-measurement`.

from the tilted mirror actually forms is given by the simple formula

$$d = 2\alpha f, \tag{6.1}$$

where d is the displacement of the spot on the image sensor, f is the focal length of the autocollimator's objective lens, and α is the angular offset of the mirror. Note that the distance between the mirror and the autocollimator does not appear in Equation (6.1) – an autocollimator can measure the angle of the reflecting surface irrespective of the distance to the reflector. In fact, this is one of the more useful features of an autocollimator.

The distance to the reflecting surface does impact the range of angles that can be

measured by an autocollimator. If we consider the light emitted by the autocollimator to be a bundle of rays, then in order for the autocollimator to measure the tilt of the reflector, at least some of the collimated rays in the reflected bundle must pass through the autocollimator's objective lens. Assuming that the size of the lens is equal to the entrance aperture of the autocollimator (i.e., the lens tube holding the objective does not extend much past the lens), then the maximum angle at which rays will enter an autocollimator with an aperture of diameter A is given by

$$\tan(2\alpha) = \frac{A}{D} \quad \rightarrow \quad \alpha_{max} = \frac{1}{2}\arctan\left(\frac{A}{D}\right), \tag{6.2}$$

where D is the distance between the autocollimator and the reflector, and α_{max} is the maximum tilt of the reflector. For a 1" (25.4 mm) diameter aperture autocollimator measuring the tilt of a mirror 8 meters away, the maximum measurable angle is 0.09° or 327 seconds of arc. Moreover, since α_{max} is the angle at which rays cease to enter the autocollimator aperture, the reflected image observed by the autocollimator should get dimmer as $\alpha \rightarrow \alpha_{max}$, and therefore the effective angular range might be somewhat smaller.

We have designed an instrument which operates using a somewhat different principle. A schematic of our instrument is shown in Figure 6.5. Our instrument uses an LED light source with an aspheric condenser lens to collect the light from the wide-angle LED beam. The back of the aspheric condenser lens is ground, which turns it into a diffuser, which enhances the uniformity of the LED light. This light passes through a ~ 1 mm diameter iris that is located at the focal point of an achromatic doublet, which serves as the objective lens. Due to the finite size of the iris, it is impossible to perfectly collimate the light emitted from it (Diaz & Knobl, 2012). The full angle of the beam emerging from the autocollimator

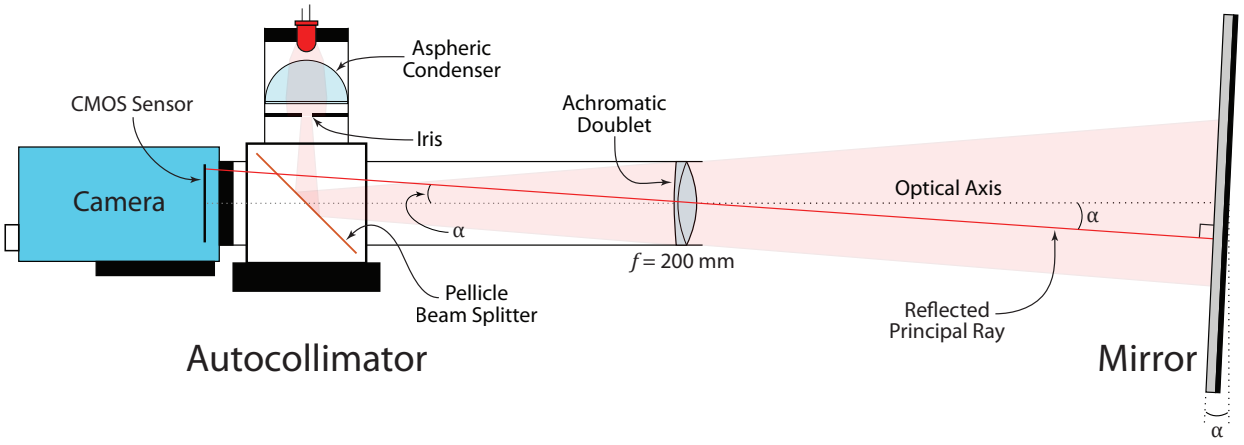

Figure 6.5. Diagram showing the parts and operational principle of our autocollimator. The finite diameter of the iris means that the light emerging from the achromatic doublet is not truly collimated. Note that the principle ray of the reflected spot makes an angle of α (instead of 2α) with the optical axis. This increases the angular range of the instrument at the expense of losing half of its sensitivity.

is given by

$$\theta = \frac{D_s}{f}, \tag{6.3}$$

where D_s is the diameter of the source (the iris diameter), and f is the focal length of the objective lens. For our instrument, the half-angle of the beam emerging from the objective lens is 0.143°; this angle can be increased by opening up the iris. Because of this, the principal ray that determines the center of the imaged spot on the autocollimator image sensor is reflected normal to the mirror surface, and the ray therefore makes an angle of α with respect to the optical axis instead of 2α. The displacement of the image on the image sensor is therefore not given by Equation (6.1), but instead is

$$d = \alpha f. \tag{6.4}$$

The sensitivity of the instrument is thus reduced by half at the benefit of an increased angular range.

For our application, a larger angular range is preferable since the telescope mirrors are independently adjustable, and it is not yet clear how much they will move. Moreover, by using a high resolution image sensor (5 megapixels), we are able to maintain sufficient angular sensitivity. We present the results of testing our instrument in the lab in Section 6.5, where we find that our instrument performs well over an angular range of $\pm 0.126^\circ$. This exceeds the angular range available to a typical digital autocollimator.

6.3 Hardware

In this section, we will describe the components of each autocollimator and how they are assembled. The single most expensive part of the autocollimator, and the only part that users interact with, is the camera. All of the "CCDs" in the global alignment system are The Imaging Source DMK 23GP031 5 megapixel monochrome GigE industrial cameras. These cameras (an example of which is shown in Figure 6.6) utilize an 1/2.5" format Aptina MT9P031 CMOS image sensor with a maximum resolution of 2592 x 1944 pixels and a pixel size of 2.2 x 2.2 microns. The pixel depth can be set to either 8- or 12-bits. The rolling shutter is adjustable between 50 μs – 30 s. It can be powered via an external 12V DC power supply using a customer-supplied harness with a Hirose HR10A-7P-6S circular connector, or via power-over-Ethernet (POE).

The cameras offer simple remote configuration over Ethernet using a tool which is a part of The Imaging Source's Linux Repository, which can be downloaded from GitHub (`https://github.com/TheImagingSource/tiscamera`). The tool is called *camera-ip-conf* and can be found in `tiscamera/tools/camera-ip-conf`. Figure 6.7 shows a screen shot of

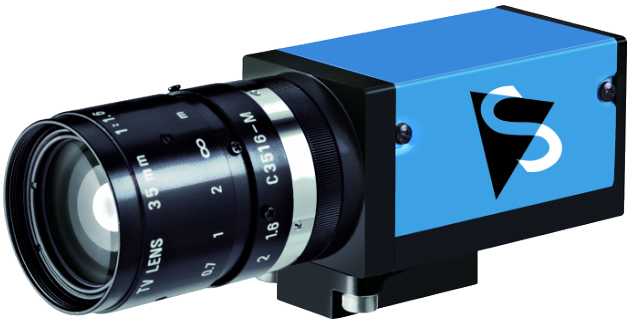

Figure 6.6. A stock image of The Imaging Source's DMK 23GP031 GigE monochrome industrial camera. The camera is shown here with a detachable C-mount lens.

camera-ip-conf. This tool is used to change the network configuration of the cameras (which use DHCP by default) and to set the user-defined name, which can be used by the Aravis library (discussed in Section 6.4) to address specific cameras on the network. Note that a label with the serial number is attached to the top of each camera for easy identification. The *camera-ip-conf* tool can also be used to install camera firmware updates.

The light source for the autocollimator is a simple epoxy-encapsulated LED. These LEDs are procured from ThorLabs, but they are quite generic and could probably be sourced from elsewhere for pennies apiece. They have a peak wavelength of 635 nm with a spectral FWHM of 10 nm. At 20 mA they have an optical output of 4.0 mW, which is not especially bright. Their opening half-angle is rather wide at 20°, which necessitates the use of an aspheric condensing lens to collect the light. An integral diffuser in the condensing lens helps increase the uniformity of the LED's light; without it, it is actually possible to see the structure of the LED in the output beam from the autocollimator. Light from the LED and the aspheric condensing lens shines through a circular aperature that is formed by an adjustable, ring-actuated iris diaphragm.

Along with the camera, the other major element of the autocollimator optics is its

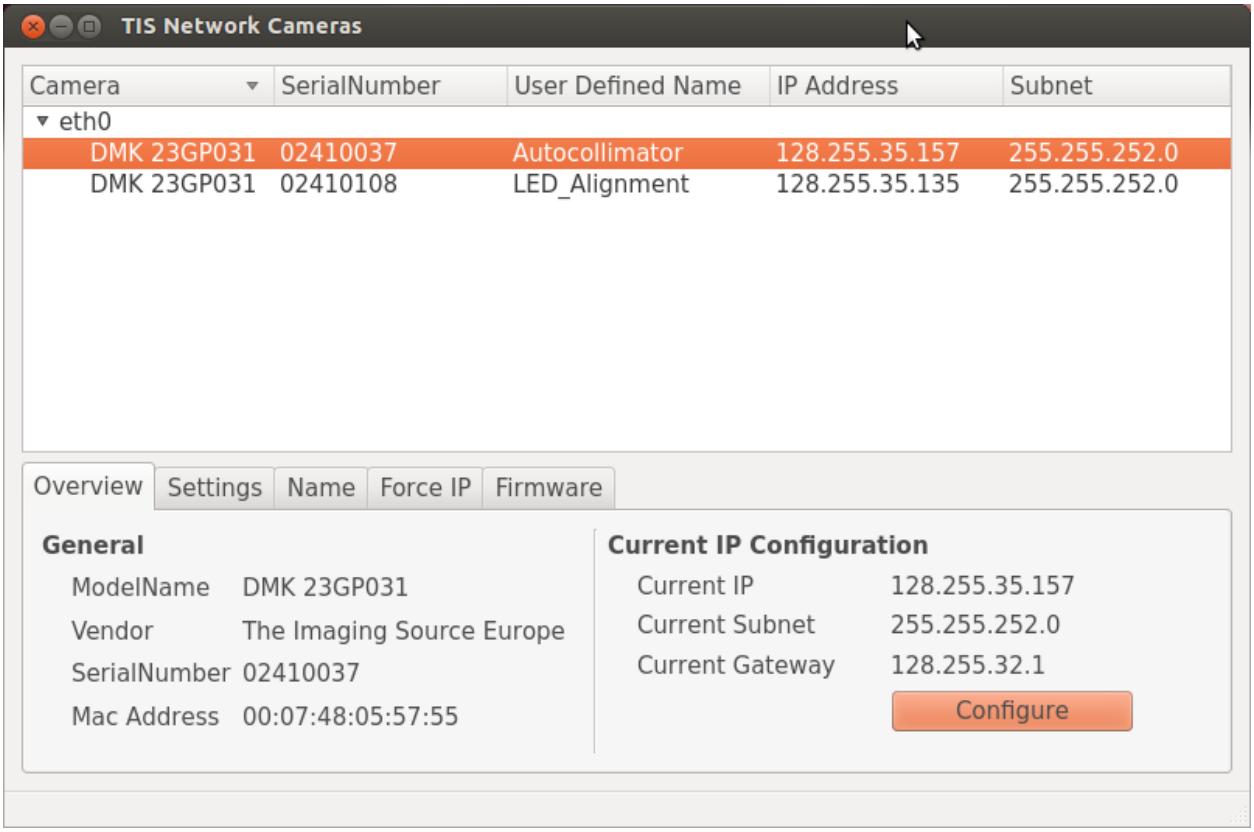

Figure 6.7. Screen capture of the vendor-supplied *camera-ip-conf* tool that is used to name The Imaging Source cameras. This utility can also be used to change the network configuration of the cameras and update their firmware.

lens. The autocollimators use 1” diameter, 200 mm focal length achromatic doublet lenses. The use of an achromatic doublet is not strictly necessary in this application since the spectral peak of the LED is fairly narrow, making chromatic aberration unlikely. The diameter of the lens is the result of an early choice to use 1” diameter lens tubes for the autocollimator’s construction and was not determined by optical considerations. By contrast, a long focal length was deliberately chosen since the focal length determines the resolution of the autocollimator (see Section 6.2 for details), as well as the field of view. The 5 megapixel CMOS image sensor on the camera, combined with the 200 mm focal length lens, yields a field of view of 1.63 x 1.23 degrees.

The LED light source (specifically, the iris diaphragm) and the CMOS image sensor are both positioned so that they are at the focus of the lens. Moreover, they are both on the same side of the lens. This requires the use of a beam splitter (see Figure 6.5). The autocollimator uses a pellicle beam splitter with a 50:50 reflection to transmission ratio. A pellicle is a thin (2 micron thickness), transparent plastic membrane which is coated to make it partially reflective. The thinness of the pellicle eliminates offsets in the optical path and image-doubling from secondary reflections. The pellicle is mounted at 45° inside an aluminum cube with four SM1-threaded circular apertures. The camera, lens, and LED assembly each occupy an aperture of the cube. The exact positions are not especially important except that the backing ring of the pellicle should be facing away from the lens to avoid occulting the image. One aperture is not used and is sealed against external light and contaminants.

The last part we will discuss is the second most expensive part of the autocollimator: its adjustable mount. As shown in Figure 6.2, the autocollimators are mounted to optical tables (OTs) at the center of the primary and secondary mirrors. The beams from the autocollimators are reflected by 2" (50.8 mm) diameter mirrors some 9 meters away which are rigidly mounted to the backs of several mirror panels. With a field of view of 1.63 x 1.23 degrees, or, at a distance of 9 meters, 25.7 x 19.2 cm, the 2" mirror will occupy about a quarter of the field of view. Since the mirror panels have some freedom in how they move, some degree of adjustability in the autocollimator mount is required. The Foveal FM0_2a fine-adjustment 2-axis camera mount was selected for its compact size and locking adjusters, which are necessary for long-term stability.

The remainder of the parts in the autocollimator will not be discussed in detail since

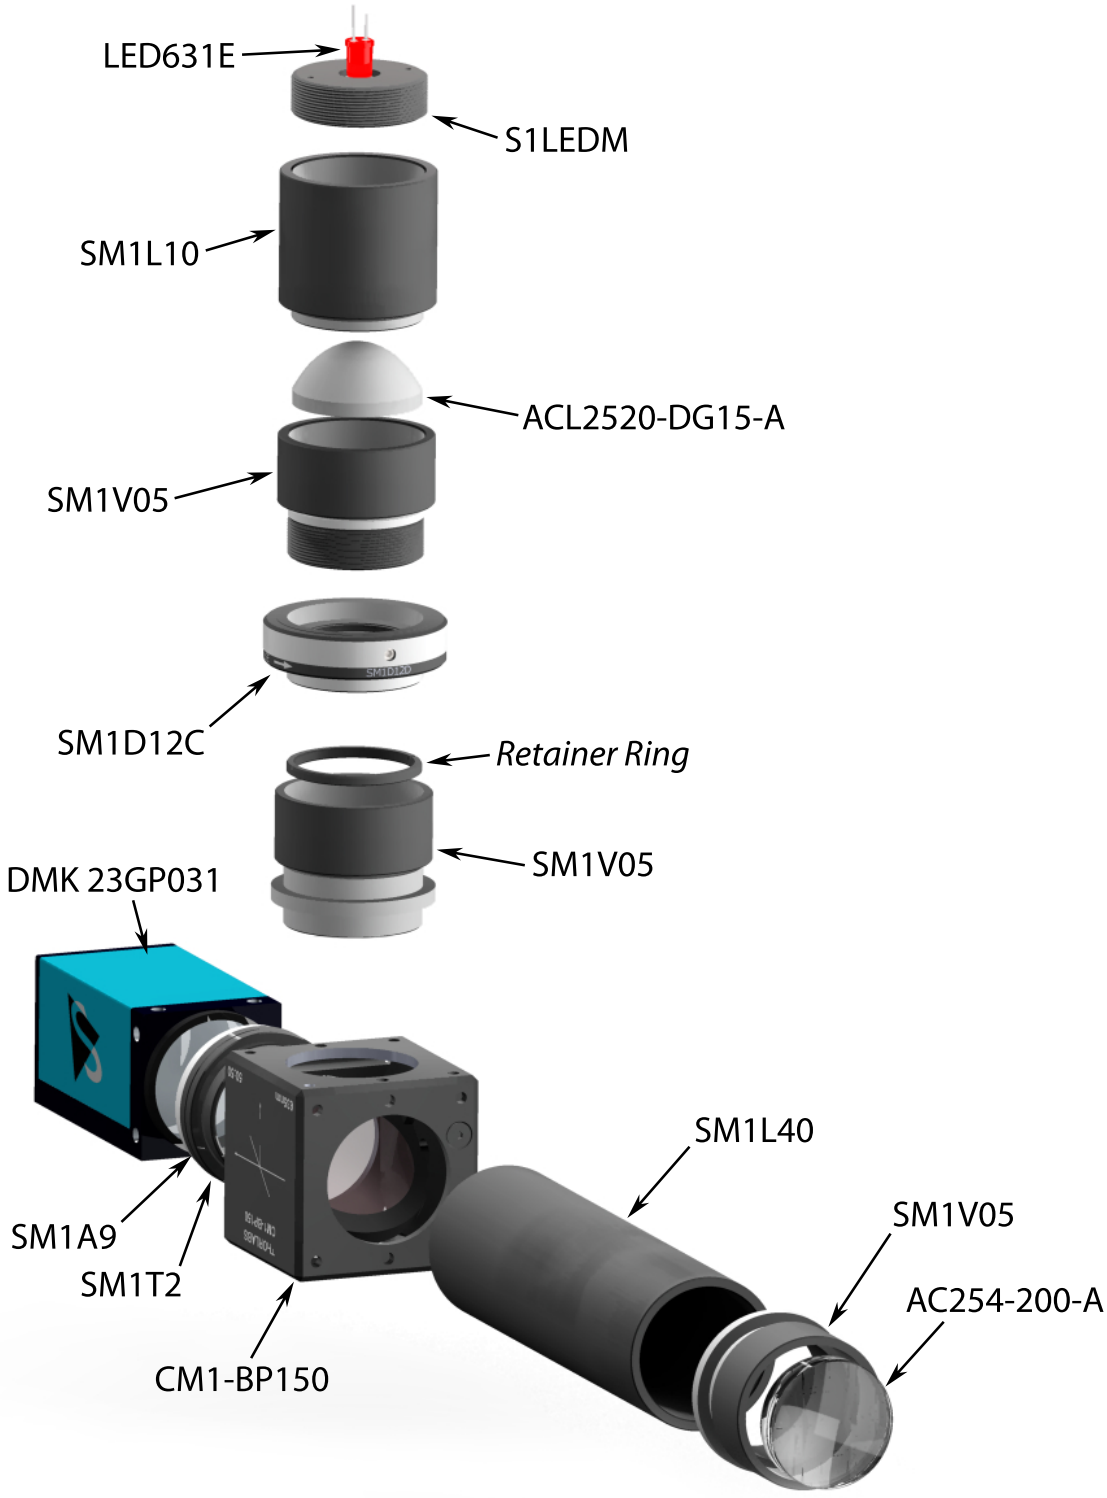

Figure 6.8. Render of exploded autocollimator assembly with parts labeled. See Table 6.2 for descriptions of the parts. The Foveal FM0_2a mount is not shown. Retainer rings are included with the SM1 lens tubes. The retainer ring shown in the image is not in the correct location: two rings are necessary to hold each of the lenses in place.

Table 6.2. Autocollimator bill of materials

Supplier	Description	Item #	Quantity
The Imaging Source	Camera	DMK 23GP031	1
Foveal Mounts	Camera Mount with Custom Drilled Holes	FM0_2a_CDH	1
ThorLabs	Cube-mounted Pellicle Beam Splitter	CM1-BP150	1
ThorLabs	Red LED, 635 nm, T-1 3/4 package	LED631E	1
ThorLabs	SM1 LED Mount	S1LEDM	1
ThorLabs	SM1 Lens Tube – 1" Length	SM1L10	1
ThorLabs	Aspheric Condenser Lens & Diffuser - 20 mm	ACL2520-DG15-A	1
ThorLabs	SM1 Adjustable Lens Tube – 0.5" Length	SM1V05	3
ThorLabs	SM1 Ring-Actuated Iris Diaphragm	SM1D12C	1
ThorLabs	C-Mount to Internal SM1 Adapter	SM1A9	1
ThorLabs	SM1 Externally-Threaded Coupler	SM1T2	1
ThorLabs	SM1 Lens Tube – 4" Length	SM1L40	1
ThorLabs	Achromatic Doublet Lens – 200 mm	AC254-200-A	1

these parts are mostly structural and do not influence the operation of the instrument. An exploded assembly drawing with all of the components labeled can be seen in Figure 6.8, and a bill of materials containing the remaining parts of the autocollimator can be found in Table 6.2.

6.4 Software

In this section, we will primarily discuss the calibration software for the autocollimator. The calibration software is distinct from the final production software that will be used during actual telescope operations, although it is expected that there will be some overlap between the two.

The calibration software was developed to facilitate laboratory testing of the autocollimator. Operation of the autocollimator requires a computer and software to display images from the GigE camera. Preliminary testing was carried out using The Imaging Source's IC Capture software, which displays images from the camera and allows for simple adjustments

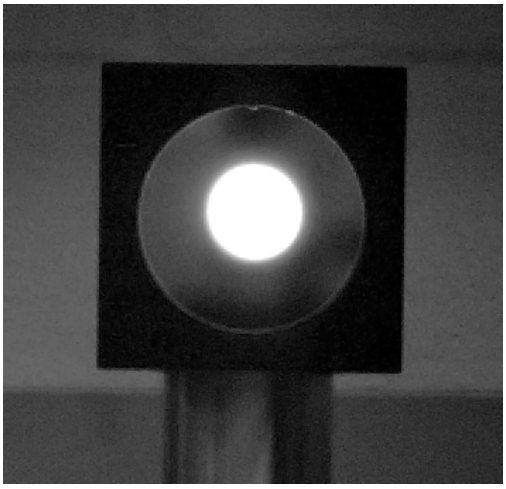

Figure 6.9. An overexposed image from the autocollimator camera shows the typical view of an autocollimator during lab testing. The large white spot is the image of the autocollimator LED (which would be red in a color image) reflected by a 2" diameter mirror. The mirror is held in place by a ThorLabs C1507 mirror mount which is clamped to a 1.5" diameter steel post, both of which can be seen in the image.

to the camera's settings (exposure, gain, resolution, etc.). However, the IC Capture software runs only on Windows machines, and since it is not specific to an autocollimator, it contains no functionality to make measurements using the images.

The calibration software is written in C++11 and uses Qt 5.1 for the graphical user interface. Communication with the gigabit cameras is accomplished using the Aravis library, which is developed by Emmanual Pacaud of the GNOME Project. Aravis is an open source library that is designed to interface with gigabit Ethernet-protocol (GigE) industrial cameras which adhere to the Genicam standard. The calibration software uses version 0.3.2 of Aravis; version 0.3.7 is already available at the time of this writing, and new versions (sometimes including significant changes) are released every 2-3 months. Because GigE cameras are still fairly new, Aravis is currently the only stable software library for interfacing with them.

Figure 6.9 shows a typical autocollimator image that was captured by the IC Capture

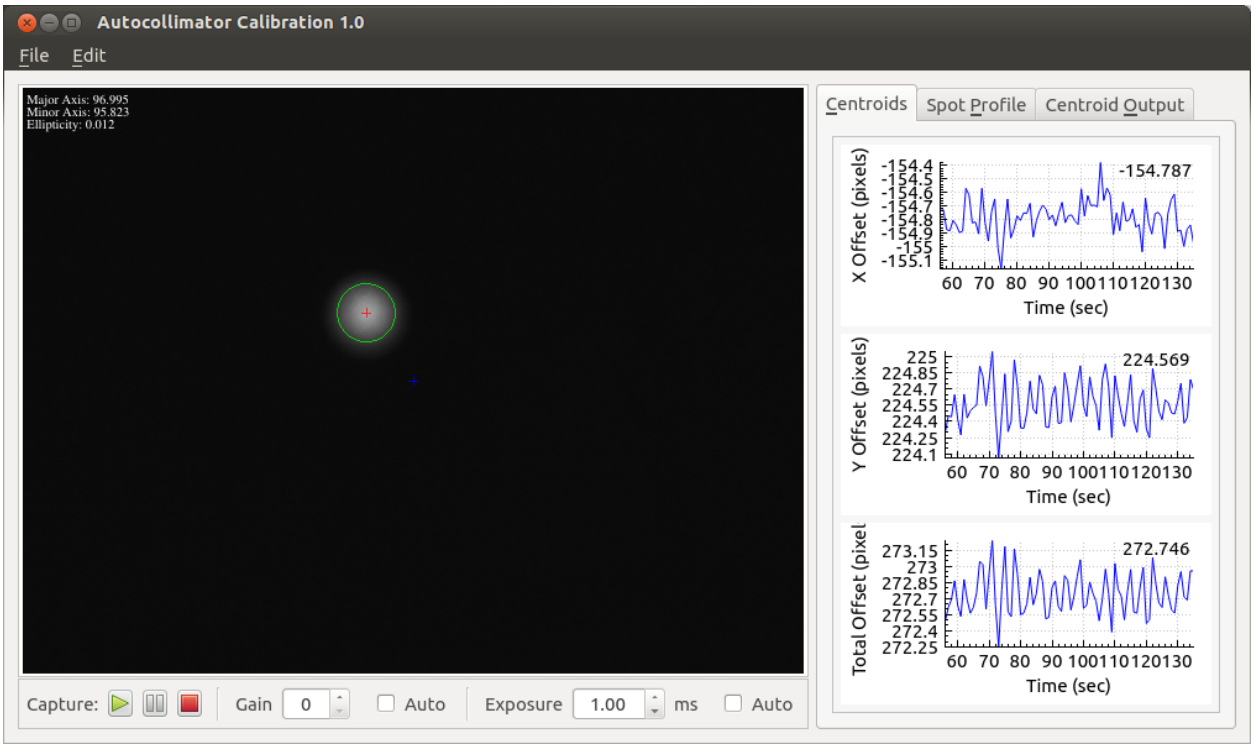

Figure 6.10. Screen capture of the autocollimator calibration software. Most of the window is occupied by the image from the autocollimator, which contains the reflected image of the autocollimator's LED. A red cross marks the center of the LED spot, and a blue cross marks the image center. On the right side of the window are three graphs showing the offset of the LED spot centroid from the center of the image: the top graph shows the number of pixels in the horizontal x direction, the middle the vertical y direction, and the bottom the total ($\sqrt{x^2 + y^2}$) offset.

software. The main feature in the image is the large white spot, which is an over-exposed (∼50 ms exposure) image of reflected light from an autocollimator's LED. The 2" diameter mirror and mirror mount used to reflect the light are also visible in the image. In a properly exposed image (a 4 ms exposure is typical), the LED spot size is much smaller and has the structure of a 2D Gaussian, and background elements such as the mirror mount are not easily visible.

The autocollimator calibration software is designed to measure the position of the LED spot in camera coordinates, measured as the number of pixels from the center of the

image. It expects to find an image similar to the one shown in Figure 6.9 containing a single reflected spot with proper exposure. The software requires the exposure to be set manually—the automatic exposure feature does not work well and is disabled by default—but it is imbued with an empirically-determined appropriate exposure time and warns if the image is over-exposed (i.e., if any pixel has the maximum attainable value, which would be 255 for an 8-bit image).

Images are captured at an adjustable cadence of once a second, and the LED spot centroid is computed for each frame. The distance between the spot centroid and the center of the image frame is printed into a text box (which can be saved to a file) and graphed. The graphs show the total distance between the spot centroid as well as the projections along the x and y axes. The projection plots are particularly useful for the initial alignment of the mirror mounts since the y axes of the camera and mirror mounts are intentionally aligned vertically in lab coordinates.

The images that are captured by the autocollimator software can be saved in a variety of standard image file formats, including bitmap (BMP), portable network graphics (PNG), and JPEG. The calibration software also supports export to the flexible image transport system (FITS) image format via the CCfits library. One major advantage of the FITS format over the other image formats is that it easily accommodates 12-bit grayscale images (the autocollimator camera can capture images with either 8- or 12-bit pixel depth). FITS is also lossless and allows for an image header containing customizable keys which can store meta-data that is relevant to the specific application. FITS is thus the default export format for the calibration software.

6.5 Performance

The performance of the autocollimator was evaluated in a laboratory setting. Figure 6.11 is an image of the autocollimator in the lab on a ThorLabs adjustable 2-axis stage. This stage was affixed to the laboratory workbench with clamps and used to point the autocollimator towards a 2" diameter front-aluminized mirror. This mirror was generally positioned 8 meters away from the autocollimator to match the approximate distance expected to separate the autocollimator and the mirror on the prototype CTA-US telescope. The mirror was contained in a Newport U200-AC kinematic optical mount which was mounted via optical post to a breadboard, which in turn was clamped to another laboratory workbench (see Figure 6.12). We attached Newport SM-13 Vernier micrometers to the mount which have 10 micron graduations and 1 micron sensitivity with the Vernier scale. The distance between the micrometers and the pivot point is 60 mm; therefore 1 micron of linear motion translates to an angular tilt of $\arctan(1\,\mu\text{m}/60\,\text{mm}) = 16.7$ microradians or 3.44 arc seconds.

The performance of the autocollimator was evaluated by adjusting the angular tilt of the mirror mount by some known quantity. Most often the 10 micron graduations on the micrometer were used (corresponding to 34.4 arcsec adjustments) because they were easier to read; the Vernier scale was used to check the precision of the setting so that the error was typically less than about 3.5 arcsec. The autocollimator was then used to measure the tilt of the mirror from 8 meters away. For each mirror tilt, the x, y coordinates of the autocollimator spot centroids were collected for about ten seconds at a rate of 1 Hz. The exposure was set to 0.04 seconds, so the sampling rate could be much faster. The mean x and y pixel positions were then used to compute the angle of the mirror. We only adjusted the tilt of the mirror

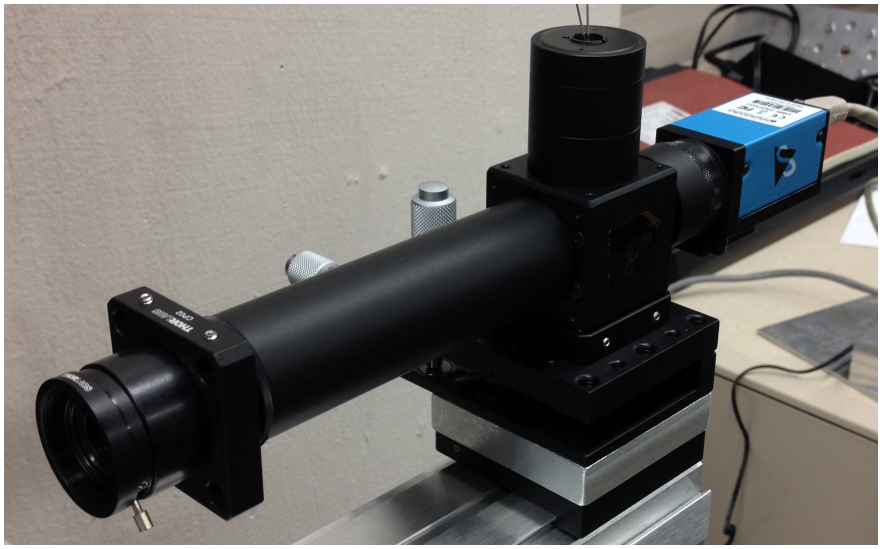

Figure 6.11. Photo of the autocollimator in the lab during testing. In this photo, the autocollimator is attached to a 2 meter long aluminum beam, but typically it was simply clamped to the workbench. The square object near the objective lens is for mounting ThorLabs 30 mm cage mount components, and was only used during early testing.

along one rotation axis at a time, and the stability of the mirror mount and the rotational alignment of the autocollimator were good enough that the other axis never varied by more than about 1 pixel. Using two sets of measurements, manually adjusting the angle of the mirror in between, the angular difference along the rotation axis could be computed using the formula

$$\Delta\alpha = \frac{(2.2 \text{ pixels}/\mu\text{m})\,\Delta d}{f} \text{ radians}, \tag{6.5}$$

where 2.2 pixels/μm is the pixel size, Δd is the difference in pixel means along either the x or y axis, f is the focal length of the objective lens in μm, and $\Delta\alpha$ is the difference between the two angle tilts as measured by the autocollimator.

The above formula is generally used in practice since the autocollimator acts as an angle comparator. However, we wished to show that the autocollimator would measure the correct angle throughout the range of rotations along each axis in which the spot was still

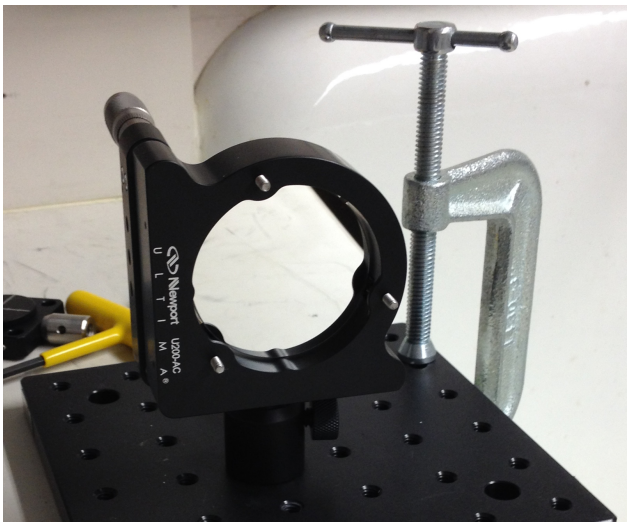

Figure 6.12. Photo of a 2" diameter mirror housed in a Newport U200-AC kinematic mirror mount. This assembly was used to test the autocollimator in the lab. Micrometers were used in place of adjustment screws in the mirror mount to increase adjustment precision.

entirely visible on the mirror. The angle of the mirror mount was generally recorded in terms of micrometer graduations, starting from zero at one edge of the mirror and increasing until the other edge of the mirror was reached. The x, y pixel values of the spot centroid were recorded in terms of distance away from the image center. After multiplying the micrometer graduations by 16.7 microradians to get the mirror rotation relative to some arbitrary starting point, and applying Equation (6.5) to the pixel values to obtain the mirror rotation relative to the center of the autocollimator image, we fit a line to the data and adjusted the y-intercept so that zero on both axes corresponds to perpendicularity (along the rotation axis) between the mirror plane and the autocollimator's line of sight.

The results are shown in Figure 6.13. The top row of plots in Figure 6.13 shows that the autocollimator's response is linear across the entire rotation range of the mirror (about ± 2200 microradians or ± 0.126 degrees) for both laboratory-frame rotation axes, horizontal and vertical. For horizontal rotation, the fitted slope is -0.998, while for vertical rotation

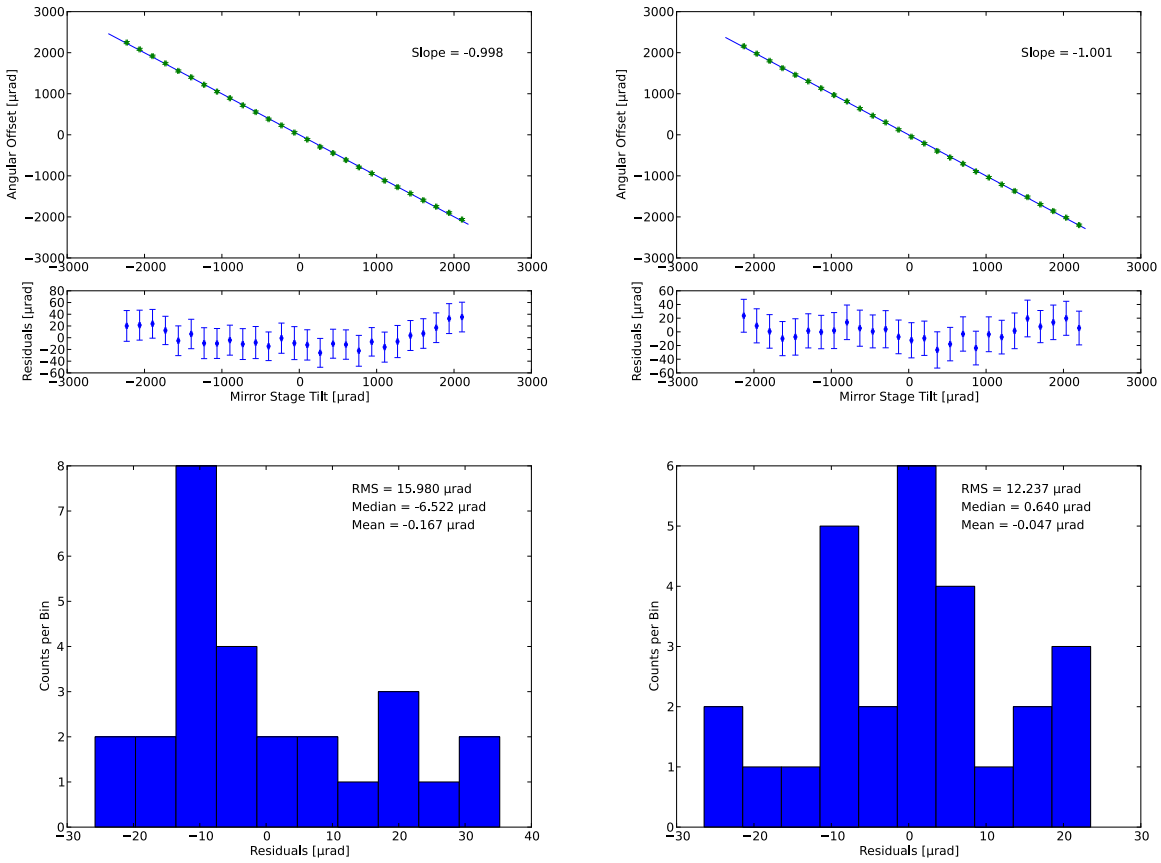

Figure 6.13. Plots demonstrating the performance of the autocollimator. The left column shows horizontal rotation in the lab frame, while the right column shows vertical rotation. In the top row, the mirror's angular offset as measured by the autocollimator is plotted versus the tilt angle measured with the mirror's kinematic mount. These plots are fitted so that zero angular offset corresponds to perpendicularity between the mirror and the autocollimator's optical axis (with respect to the rotation axis being used). The slopes of the fitted lines are shown in the top plots, and both are very close to unity in magnitude. The bottom row contains histograms of the residuals. The RMS spread of the residuals along both axes is less than 16 microradians, or 3.3 arc seconds.

it is -1.001. The minus sign results from a mismatch between the direction in which the micrometers were turned, and the assignment of an increasing angle to the right/upward direction in the autocollimator image; the *magnitude* of the slope is what matters. A unity magnitude slope means that the angle measured by the autocollimator is equal to the angle set by the mirror mount micrometer. The bottom row of plots shows the distribution of

residuals (differences between the measured autocollimator angles and the fitted line). The root mean square (RMS) spreads of these distributions are both less than 16 microradians, or 3.3 arcsec. The angle subtended by a single pixel is 2.27 arcsec, so this error corresponds to about 1.5 pixels. We note that the RMS error is within the tolerances for $\Delta\theta_x$ and $\Delta\theta_y$ outlined in Table 6.1.

CONCLUSION

None of the proposed quantum gravity theories require Lorentz invariance to be broken, although many of them allow it. Many studies have been conducted to search for Lorentz invariance violation (LIV) using astrophysical observations, and they have been successful to the extent that their published limits strongly disfavor the possibility of dispersion with linear energy dependence (linear LIV) for a quantum gravity energy scale less than the Planck energy (Abdo et al., 2009; Vasileiou et al., 2013). Dispersion with a quadratic energy dependence (quadratic LIV) still remains a real possibility, but only if the quantum gravity energy scale is greater than 10^{10} GeV. A successful detection of LIV, whether by astrophysical time-of-flight tests or some other method, would represent a tremendous advance in the field of physics.

In the first part of this work, we searched for energy-dependent dispersion resulting from LIV in observations of the blazar Markarian 421 made by VERITAS. Only four blazars have been successfully used to constrain LIV: PKS 2155-304, PG 1553+113, Markarian 501, and Markarian 421. The LIV limit using Markarian 421 was published by the Whipple collaboration (Whipple was the predecessor of VERITAS) and was the first study of its kind to use blazars; it was published 16 years ago. The enhanced sensitivity of VERITAS and the exceptionally bright flare occurring on Februrary 17, 2010 presented an excellent opportunity to advance the field by updating the original Whipple limit, which only constrained linear LIV (see Table 1.1). Furthermore, at least one segment of the data showed gamma-ray fluxes variable on minute time scales; variability is necessary for a measurement of dispersion in the data, and the minute time scale variability suggested that the resulting limit would be

competitive with published limits from other instruments similar to VERITAS (imaging air Cherenkov telescopes, or IACTs).

Unfortunately, the structure of the flare light curves was not ideal for searching for dispersion. Due to the way IACTs take data (wobble-mode or false-source tracking), our 5-hour long data set was chopped up into a dozen 20-30 minute observations with minute-long gaps in between. Dispersed photons could be lost in these observing gaps, so we had to be careful in handling them. Moreover, because the flaring activity occurred on top of a constant background flux, we had to be especially careful to avoid adding sharp boundaries to our data since the dispersion estimation algorithms would see these as particularly variable parts of the data and produce incorrect results. Although it is possible to use data containing flares on top of a constant flux (Albert et al., 2008), the flare needs to be large enough to compensate for the loss of sensitivity resulting from the constant flux. We were not able to obtain useful LIV constraints with the available data.

Nonetheless, we successfully implemented and tested two very different kinds of dispersion estimation algorithms and showed that they could accurately determine the dispersion present in simulated data. The PairView algorithm was implemented according to the prescription given in the literature. The dispersion cancellation algorithm we use differs somewhat from the prescription given by Scargle et al. (2008) because the original algorithm was developed for use with space-bourne observatories with better energy resolution and less background. Our results show that the PairView algorithm provides better sensitivity when determining the dispersion in a data set consisting of a single (Gaussian-shaped) flare with no background. However, we found that our dispersion cancellation algorithm provided

superior performance with our non-ideal data, although its sensitivity was still insufficient to provide a viable result. We also made suggestions for how future observations of blazar flares can be made in ways that maximize their potential for yielding a useful LIV limit.

In the second part of this work, we described two alignment and positioning systems used by IACTs. We first described the pointing monitor system currently used by VERITAS to accurately determine where the telescopes are pointed on the celestial sphere. The Davies-Cotton optics used by VERITAS allow for simple alignment of the faceted reflector with the gamma-ray camera, and allow the pointing of the telescopes to be determined with an accuracy of 25 arcsec with only two CCD cameras. The more complex Schwarzschild-Couder optics planned for the medium-sized telescopes in the U.S. contribution to the Cherenkov Telescope Array (CTA-US) require a much more sophisticated alignment system, since in this case there are two reflectors to align with the gamma-ray camera and with each other.

The focus of this work is on an autocollimator alignment system, which is one element of the global alignment system that will be used in a prototype Schwarzschild-Couder telescope (SCT) for CTA-US. This prototype telescope is currently under construction at the Fred Lawrence Whipple Observatory in southern Tucson. We provide a detailed description of the autocollimator hardware as well as the custom-made software that was used to test it in the lab. We found that our autocollimator can successfully measure the angle (with respect to the line of sight) of a mirror located 8 meters away with a precision better than 5 arcsec. Moreover, the configuration of our autocollimator accommodates measurement of angular offsets between the mirror and the line of sight up to 0.126 degrees when used with a 2" diameter mirror, which should be sufficient for use with the prototype SCT.

REFERENCES

Aad, G., Abajyan, T., Abbott, B., et al. 2012, Physics Letters B, 716, 1

Abdo, A. A., Ackermann, M., Arimoto, M., et al. 2009, Science, 323, 1688

Abdo, A. A., Ackermann, M., Ajello, M., et al. 2009, Nature, 462, 331

Abdo, A. A., Ackermann, M., Ajello, M., et al. 2011, ApJ, 736, 131

Abramowski, A., Acero, F., Aharonian, F., et al. 2011, Astroparticle Physics, 34, 738

Abramowski, A., Aharonian, F., Benkhali, F. A., et al. 2015, The Astrophysical Journal, 802, 65

Acciari, V. A., Beilicke, M., Blaylock, G., et al. 2008, ApJ, 679, 1427

Ade, P. A. R., Aghanim, N., Armitage-Caplan, C., et al. 2014, A&A, 571, A16

Agudo, I., Jorstad, S. G., Marscher, A. P., et al. 2011, The Astrophysical Journal Letters, 726, L13

Aharonian, F., Akhperjanian, A. G., Bazer-Bachi, A. R., et al. 2007, The Astrophysical Journal Letters, 664, L71

Aharonian, F., Akhperjanian, A. G., Barres de Almeida, U., et al. 2008, Phys. Rev. Lett., 101, 170402

Albert, J., Aliu, E., Anderhub, H., et al. 2007a, The Astrophysical Journal, 663, 125

Albert, J., Aliu, E., Anderhub, H., et al. 2007b, The Astrophysical Journal, 669, 862

Albert, J., Aliu, E., Anderhub, H., et al. 2008, Physics Letters B, 668, 253

Aleksandrov, E. B. 1965, Soviet Ast., 9, 519

Aliu, E., Arlen, T., Aune, T., et al. 2011, Science, 334, 69

Allen, S. W., Dunn, R. J. H., Fabian, A. C., Taylor, G. B., & Reynolds, C. S. 2006, MNRAS, 372, 21

Alväger, T., Farley, F., Kjellman, J., & Wallin, L. 1964, Physics Letters, 12, 260

Amelino-Camelia, G., Ellis, J., Mavromatos, N. E., & Nanopoulos, D. V. 1997, International Journal of Modern Physics A, 12, 607

Amelino-Camelia, G., Ellis, J., Mavromatos, N. E., Nanopoulos, D. V., & Sarkar, S. 1998, Nature, 393, 763

Angel, J. R. P., & Stockman, H. S. 1980, ARA&A, 18, 321

Antonucci, R. 1993, ARA&A, 31, 473

Auer, L. H., & Standish, E. M. 2000, The Astronomical Journal, 119, 2472

Baade, W., & Minkowski, R. 1954, ApJ, 119, 215

Babcock, G. C., & Bergman, T. G. 1964, J. Opt. Soc. Am., 54, 147

Bagchi, J., Vivek, M., Vikram, V., et al. 2014, The Astrophysical Journal, 788, 174

Benbow, W., & Collaboration, H. 2005, AIP Conference Proceedings, 745, 611

Beringer, J., Arguin, J. F., Barnett, R. M., et al. 2012, Phys. Rev. D, 86, 010001

Bernlöhr, K. 2000, Astroparticle Physics, 12, 255

Biller, S. D., Breslin, A. C., Buckley, J., et al. 1999, Phys. Rev. Lett., 83, 2108

Blandford, R. D. 1976, MNRAS, 176, 465

Blandford, R. D., & Königl, A. 1979, ApJ, 232, 34

Blandford, R. D., & Payne, D. G. 1982, MNRAS, 199, 883

Blandford, R. D., & Rees, M. J. 1974, MNRAS, 169, 395

Blandford, R. D., & Rees, M. J. 1978, Physica Scripta, 17, 265

Blandford, R. D., & Znajek, R. L. 1977, MNRAS, 179, 433

Bloom, S. D., & Marscher, A. P. 1993, AIP Conference Proceedings, 280, 578

Boggs, S. E., Wunderer, C. B., Hurley, K., & Coburn, W. 2004, ApJ, 611, L77

Bojowald, M., Morales-Técotl, H. A., & Sahlmann, H. 2005, Phys. Rev. D, 71, 084012

Bolmont, J., Jacholkowska, A., Atteia, J.-L., Piron, F., & Pizzichini, G. 2008, The Astrophysical Journal, 676, 532

Bondi, M., Marchã, M. J. M., Dallacasa, D., & Stanghellini, C. 2001, MNRAS, 325, 1109

Brecher, K. 1977, Phys. Rev. Lett., 39, 1051

Bretz, T. 2005, International Cosmic Ray Conference, 4, 315

Bridle, A. H., & Perley, R. A. 1984, ARA&A, 22, 319

Buckley, J. H., Dowkontt, P., Kosack, K., & Rebillot, P. F. 2003, International Cosmic Ray Conference, 5, 2827

Cannon, R. D., Penston, M. V., & Brett, R. A. 1971, MNRAS, 152, 79

Čerenkov, P. A. 1934, Doklady Akademii Nauk SSSR, 2, 451

Čerenkov, P. A. 1937, Phys. Rev., 52, 378

Chatrchyan, S., Khachatryan, V., Sirunyan, A., et al. 2012, Physics Letters B, 716, 30

Cheng, T. 2010, Oxford Master Series in Physics, Vol. 11, Relativity, Gravitation and Cosmology: A Basic Introduction, 2nd edn. (Oxford University Press)

Cogan, P. 2008, International Cosmic Ray Conference, 3, 1385

Colladay, D., & Kostelecký, V. A. 1997, Phys. Rev. D, 55, 6760

Curtis, H. D. 1918, Publications of the Lick Observatory, 13, 9

Daniel, M. K. 2008, International Cosmic Ray Conference, 3, 1325

Daum, A., Hermann, G., Heß, M., et al. 1997, Astroparticle Physics, 8, 1

Davies, J. M., & Cotton, E. S. 1957, Solar Energy, 1, 16, the Proceedings of the Solar Furnace Symposium

de Almeida, U. B., & Daniel, M. 2012, Astroparticle Physics, 35, 850

de Naurois, M., & Rolland, L. 2009, Astroparticle Physics, 32, 231

Dermer, C. D., & Schlickeiser, R. 1993, ApJ, 416, 458

Diaz, E., & Knobl, M. 2012, Photonik, 36

Doeleman, S. S., Fish, V. L., Schenck, D. E., et al. 2012, Science, 338, 355

Dondi, L., & Ghisellini, G. 1995, MNRAS, 273, 583

DuPuy, D., Schmitt, J., McClure, R., van den Bergh, S., & Racine, R. 1969, ApJ, 156, L135

Efron, B. 1979, The Annals of Statistics, 7, 1

Efron, B., & Tibshirani, R. J. 1993, An Introduction to the Bootstrap, Chapman & Hall/CRC Monographs on Statistics & Applied Probability (New York, NY: Chapman & Hall)

Einstein, A. 1905, Annalen der Physik, 322, 891

Einstein, A. 1908, Jahrbuch der Radioaktivität und Elektronik, 4, 411

Einstein, A. 1915, Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften (Berlin), Seite 844-847., 844

Einstein, A. 1916, Annalen der Physik, 354, 769

Ellis, J., Farakos, K., Mavromatos, N. E., Mitsou, V. A., & Nanopoulos, D. V. 2000, ApJ, 535, 139

Ellis, J., Mavromatos, N., Nanopoulos, D., & Sakharov, A. 2003, A&A, 402, 409

Ellis, J., Mavromatos, N., Nanopoulos, D., Sakharov, A., & Sarkisyan, E. 2006, Astroparticle Physics, 25, 402

Ellis, J., Mavromatos, N., Nanopoulos, D., Sakharov, A., & Sarkisyan, E. 2008, Astroparticle Physics, 29, 158

Fazio, G. G., Helmken, H. F., Rieke, G. H., & Weekes, T. C. 1968, Canadian Journal of Physics, 46, S451

Fegan, D. J. 1997, Journal of Physics G: Nuclear and Particle Physics, 23, 1013

Fegan, S., & Vassiliev, V. 2005, internal VERITAS Memo

Fichtel, C. E. 1994, ApJS, 90, 917

Filippas, T. A., & Fox, J. G. 1964, Phys. Rev., 135, B1071

Finnegan, G. 2011, ArXiv e-prints, arXiv:1111.0121

Fossati, G., Maraschi, L., Celotti, A., Comastri, A., & Ghisellini, G. 1998, MNRAS, 299, 433

Franceschini, A., Rodighiero, G., & Vaccari, M. 2008, A&A, 487, 837

Gaidos, J. A., Akerlof, C. W., Biller, S., et al. 1996, Nature, 383, 319

Gambini, R., & Pullin, J. 1999, Phys. Rev. D, 59, 124021

Garofalo, D., Evans, D. A., & Sambruna, R. M. 2010, Monthly Notices of the Royal Astronomical Society, 406, 975

Gear, W. K., Stevens, J. A., Hughes, D. H., et al. 1994, MNRAS, 267, 167

Gehrels, N. 1986, ApJ, 303, 336

Ghisellini, G., & Madau, P. 1996, MNRAS, 280, 67

Gibbs, K., Criswell, S., Falcone, A., et al. 2003, International Cosmic Ray Conference, 5, 2823

Giommi, P., Barr, P., Pollock, A. M. T., Garilli, B., & Maccagni, D. 1990, ApJ, 356, 432

Gorelik, G. E. 1992, in Einstein Studies, Vol. 3, Studies in the History of General Relativity, ed. J. Eisenstaedt & A. J. Kox (Boston: Birkhäuser), 364–379

Gould, R. J., & Schréder, G. P. 1967, Phys. Rev., 155, 1408

Greenstein, J. L., & Matthews, T. A. 1963, AJ, 68, 279

Grube, J., Pandel, D., Gall, D., et al. 2011, internal VERITAS Memo

Hawking, S. 1978, Nuclear Physics B, 144, 349

Hawking, S., Page, D., & Pope, C. 1979, Physics Letters B, 86, 175

Hawking, S., Page, D., & Pope, C. 1980, Nuclear Physics B, 170, 283, volume {B170} [FSI] No. 3 to follow in Approximately Two Months

Hays, E. 2007, 30th International Cosmic Ray Conference, Merida, Mexico, arXiv:0710.2288

Heinz, S., & Begelman, M. C. 2000, The Astrophysical Journal, 535, 104

Heitler, W. 1954, The Quantum Theory of Radiation, Dover Books on Physics and Chemistry (Dover Publications)

Hillas, A. M. 1985, International Cosmic Ray Conference, 3, 445

Hoffmeister, C. 1929, Astronomische Nachrichten, 236, 233

Hofmann, W., Jung, I., Konopelko, A., et al. 1999, Astroparticle Physics, 12, 135

Holder, J., Atkins, R., Badran, H., et al. 2006, Astroparticle Physics, 25, 391

Holder, J., Acciari, V. A., Aliu, E., et al. 2008, in American Institute of Physics Conference Series, Vol. 1085, American Institute of Physics Conference Series, ed. F. A. Aharonian, W. Hofmann, & F. Rieger, 657–660

Jacob, U., & Piran, T. 2008, J. Cosmology Astropart. Phys., 1, 31

Jelley, J. 1958, Čerenkov Radiation and its Applications (Pergamon Press)

Jelley, J. V. 1966, Nature, 211, 472

Jones, T. W., O'dell, S. L., & Stein, W. A. 1974, ApJ, 188, 353

Kaaret, P. 1999, A&A, 345, L32

Kellermann, K. I., & Pauliny-Toth, I. I. K. 1969, ApJ, 155, L71

Kieda, D. 2011, International Cosmic Ray Conference, 9, 14

Konopelko, A., Aharonian, F., Hemberger, M., et al. 1999, Journal of Physics G Nuclear Physics, 25, 1989

Kostelecký, V. A., & Mewes, M. 2009, Phys. Rev. D, 80, 015020

Kostelecký, V. A., & Russell, N. 2011, Reviews of Modern Physics, 83, 11

Kostelecký, V. A., & Samuel, S. 1989, Phys. Rev. D, 39, 683

Kouveliotou, C., Meegan, C. A., Fishman, G. J., et al. 1993, ApJ, 413, L101

Krawczynski, H., Carter-Lewis, D., Duke, C., et al. 2006, Astroparticle Physics, 25, 380

Krennrich, F., Badran, H. M., Bond, I. H., et al. 2001, The Astrophysical Journal Letters, 560, L45

Krennrich, F., Bond, I., Boyle, P., et al. 2004, New Astronomy Reviews, 48, 345

Krolik, J. H., & Begelman, M. C. 1988, ApJ, 329, 702

Lamon, R., Produit, N., & Steiner, F. 2008, General Relativity and Gravitation, 40, 1731

Lang, D., Hogg, D. W., Mierle, K., Blanton, M., & Roweis, S. 2010, The Astronomical Journal, 139, 1782

LeBohec, S., Duke, C., & Jordan, P. 2005, Astroparticle Physics, 24, 26

LeBohec, S., Degrange, B., Punch, M., et al. 1998, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 416, 425

Ledden, J. E., O'Dell, S. L., Stein, W. A., & Wisniewski, W. Z. 1981, ApJ, 243, 47

Lewis, D. 1990, Experimental Astronomy, 1, 213

Li, T.-P., & Ma, Y.-Q. 1983, ApJ, 272, 317

Lide, D. R. 2005, in CRC Handbook of Chemistry and Physics, Internet Version (Boca Raton, FL: CRC Press), 224

Lin, Y. C., Bertsch, D. L., Chiang, J., et al. 1992, ApJ, 401, L61

Lovell, B., Whipple, F. L., & Solomon, L. H. 1964, Nature, 202, 377

MacLeod, J. M., & Andrew, B. H. 1968, Astrophys. Lett., 1, 243

MacLeod, J. M., Andrew, B. H., Medd, W. J., & Olsen, E. T. 1971, Astrophys. Lett., 9, 19

Mannheim, K. 1993, A&A, 269, 67

Mannheim, K., & Biermann, P. L. 1992, A&A, 253, L21

Maraschi, L., Ghisellini, G., & Celotti, A. 1992, ApJ, 397, L5

Marscher, A. P., Jorstad, S. G., D/'Arcangelo, F. D., et al. 2008, Nature, 452, 966

Marscher, A. P., Jorstad, S. G., Larionov, V. M., et al. 2010, The Astrophysical Journal Letters, 710, L126

Martínez, M., & Errando, M. 2009, Astroparticle Physics, 31, 226

Martínez, M. R., & Piran, T. 2006, Journal of Cosmology and Astroparticle Physics, 2006, 006

Massaro, E., Giommi, P., Leto, C., et al. 2009, A&A, 495, 691

Massaro, F., D'Abrusco, R., Ajello, M., Grindlay, J. E., & Smith, H. A. 2011, ApJ, 740, L48

Matthews, J. 2005, Astroparticle Physics, 22, 387

Matthews, T. A., & Sandage, A. R. 1963, ApJ, 138, 30

Michelson, A. A., Pease, F. G., & Pearson, F. 1935, ApJ, 82, 26

Misner, C., Thorne, K., & Wheeler, J. 1973, Gravitation (W. H. Freeman)

Misner, C. W., & Wheeler, J. A. 1957, Annals of Physics, 2, 525

Mushotzky, R. F., Boldt, E. A., Holt, S. S., et al. 1978, ApJ, 226, L65

Nakamura, K., & Group, P. D. 2010, Journal of Physics G: Nuclear and Particle Physics, 37, 075021

Oke, J. B. 1963, Nature, 197, 1040

Oke, J. B. 1967, ApJ, 147, 901

Oke, J. B., & Gunn, J. E. 1974, ApJ, 189, L5

Oke, J. B., Neugebauer, G., & Becklin, E. E. 1969, ApJ, 156, L41

Otte, N. 2011, International Cosmic Ray Conference, 7, 256

Padovani, P., & Giommi, P. 1995, ApJ, 444, 567

Parsons, R., & Hinton, J. 2014, Astroparticle Physics, 56, 26

Pérez, I. 2011, European Journal of Physics, 32, 993

Perley, R. A., Dreher, J. W., & Cowan, J. J. 1984, ApJ, 285, L35

Petry, D., Bradbury, S. M., Konopelko, A., et al. 1996, A&A, 311, L13

Planck, M. 1900, Annalen der Physik, 306, 69

Punch, M., Akerlof, C. W., Cawley, M. F., et al. 1992, Nature, 358, 477

Rees, M. J. 1966, Nature, 211, 468

Rees, M. J. 1978, MNRAS, 184, 61P

Sambruna, R. M., Maraschi, L., & Urry, C. M. 1996, ApJ, 463, 444

Scargle, J. D., Norris, J. P., & Bonnell, J. T. 2008, ApJ, 673, 972

Schaefer, B. E. 1999, Phys. Rev. Lett., 82, 4964

Schmidt, M. 1963, Nature, 197, 1040

Schmitt, J. L. 1968, Nature, 218, 663

Semakin, N. K. 1954, Peremennye Zvezdy, 10, 238

Shannon, C. E., & Weaver, W. 1949, The Mathematical Theory of Communication (Urbana, IL: University of Illinois Press)

Shapiro, S. L., & Elliot, J. L. 1974, Nature, 250, 111

Shayduk, M., Hengstebeck, T., Kalekin, O., Pavel, N. A., & Schweizer, T. 2005, International Cosmic Ray Conference, 5, 223

Sikora, M., Begelman, M. C., & Rees, M. J. 1994, ApJ, 421, 153

Stecker, F. W. 2003, Astroparticle Physics, 20, 85

Strittmatter, P. A., Serkowski, K., Carswell, R., et al. 1972, ApJ, 175, L7

Swanenburg, B. N., Bennett, K., Bignami, G. F., et al. 1978, Nature, 275, 298

Tamm, I., & Frank, I. 1937, Doklady Akademii Nauk SSSR, 14, 107

Tanimori, T., Sakurazawa, K., Dazeley, S. A., et al. 1998, ApJ, 492, L33

Tasson, J. D. 2014, Reports on Progress in Physics, 77, 062901

Tavecchio, F., Becerra-Gonzalez, J., Ghisellini, G., et al. 2011, A&A, 534, A86

Tavecchio, F., Maraschi, L., & Ghisellini, G. 1998, The Astrophysical Journal, 509, 608

Urry, C. M. 1996, in Astronomical Society of the Pacific Conference Series, Vol. 110, Blazar Continuum Variability, ed. H. R. Miller, J. R. Webb, & J. C. Noble, 391

Urry, C. M. 1998, Advances in Space Research, 21, 89, space based astronomy: ISO, AGN, radiopulsars and the sun

Urry, C. M., Mushotzky, R. F., Kondo, Y., Hackney, K. R. H., & Hackney, R. L. 1982, ApJ, 261, 12

Urry, C. M., & Padovani, P. 1995, Publications of the Astronomical Society of the Pacific, 107, 803

van Paradijs, J., Groot, P. J., Galama, T., et al. 1997, Nature, 386, 686

Vasileiou, V., Jacholkowska, A., Piron, F., et al. 2013, Phys. Rev. D, 87, 122001

Vassiliev, V., Fegan, S., & Brousseau, P. 2007, Astroparticle Physics, 28, 10

Visvanathan, N. 1969, ApJ, 155, L133

Vlahakis, N., & Königl, A. 2004, The Astrophysical Journal, 605, 656

Völk, H., & Bernlöhr, K. 2009, Experimental Astronomy, 25, 173

Vovk, I., & Babić, A. 2015, A&A, 578, A92

Wagner, R. 2009, AIP Conference Proceedings, 1112, 187

Weekes, T. C. 1988, Physics Reports, 160, 1

Weekes, T. C. 1992, Space Sci. Rev., 59, 315

Weekes, T. C. 2003, Very High Energy Gamma-Ray Astronomy, Series in Astronomy and Astrophysics (CRC Press)

Weinstein, A. 2008, International Cosmic Ray Conference, 3, 1539

Wheeler, J. A. 1957, Annals of Physics, 2, 604

Winston, R. 1970, Journal of the Optical Society of America, 60, 245

Worrall, D. M., Boldt, E. A., Holt, S. S., Mushotzky, R. F., & Serlemitsos, P. J. 1981, ApJ, 243, 53

Zahejsky, J., & Kolesnikov, V. 1966, Nature, 212, 1227

Zitzer, B. 2013, Proceedings of the 33rd International Cosmic Ray Conference (ICRC2013), Rio de Janeiro (Brazil), arXiv:1307.8360