Chapter 1

Introduction

Active galactic nuclei (AGN) house some of the most powerful particle accelerators in the universe. These objects are believed to consist of a supermassive black hole (SMBH) at the center of a galaxy. The SMBH supports a complex structure that in some cases can outshine the combined emission of all the stars in the galactic host. A great variety of AGN exist, a small fraction of which are bright radio sources. Explaining the difference between the radio-quiet and radio-loud AGN is a major outstanding question in this field. A major observational difference is that the structure of many (and perhaps all) radio-loud AGN includes a pair of collimated jet structures that are thought to give rise to many of the puzzling—and exciting—features of the AGN phenomenon, including the prodigious gamma-ray production in many of these sources.

A small fraction of radio-loud AGN happen to be aligned with their jet axis pointed very nearly toward Earth. Among these are the sources collectively known as *blazars*, which exhibit the most extreme behavior found in AGN. Blazars are broadband sources, emitting brightly over the entire electromagnetic spectrum (e.g., Krolik 1999). Many are bright gamma-ray emitters, with emission extending in some cases to the TeV regime (Punch et al. 1992). Furthermore, these sources are strongly variable in all bands, with significant variation on timescales ranging from many years down to a few minutes in some bands (e.g., Hughes et al. 1992; Aharonian et al. 2007). Blazars are some of the most significant sources of extragalactic high-energy emission. They were the most numerous sources identified in the the Third EGRET Catalog, the gamma-ray source catalog based on data from the Energetic Gamma Ray Experiment Telescope (EGRET; Hartman et al. 1999; Thompson et al. 1993) and have continued their dominance of the extragalactic gamma-ray sky in the *Fermi* era (Abdo et al. 2009a, 2010a, 2010b).

The launch of the *Fermi* Gamma-ray Space Telescope in June of 2008 provides an unprecedented opportunity for the systematic study of blazar jets (Atwood et al. 2009). Its Large Area Telescope (LAT) instrument is a pair-conversion telescope that is used to observe the sky at energies between 100 MeV and a few hundred GeV. It provides a large (2.4 sr) instantaneous field of view coupled with a set of precision trackers and calorimeters to permit accurate determination of the trajectory and energy of the incoming gamma ray responsible for the detected charged particle pair. During most of its mission, *Fermi* has been operated in sky-scanning survey mode. In this mode, the telescope rocks between $\pm 35^\circ$ of zenith on alternate orbits, scanning the whole sky with reasonably uniform sensitivity every two orbits (Abdo et al. 2010a). In its low-earth orbit at an altitude of about 565 km at 25.5° inclination, *Fermi* thus observes the entire sky with the LAT about every three hours (Atwood et al. 2009). The survey mode is occasionally interrupted, either because of a gamma-ray burst detected by the Gamma-ray Burst Monitor (GBM), the other instrument on board the *Fermi* satellite, or because of a planned pointed observation triggered by an extraordinary astronomical event. The LAT achieved an absolute efficiency of 73.5% (245.6 days) for its survey-mode operation during the first 11 months of science operations (Abdo et al. 2010a).

Although the LAT provides far better angular resolution than did EGRET, it still provides at best $0.6°$ (68% containment) resolution for single 1 GeV photons and a typical 95% position error of $10'$ for $\mathrm{TS} = 25$ point sources in the first-year catalog, although this depends on the gamma-ray spectral index (Abdo et al. 2010a). Association of point sources detected by the LAT relies on correlation of the point-source catalog derived from the gamma-ray observations with lists of candidate sources from radio AGN and blazar catalogs (Abdo et al. 2010b). Thus, even a revolutionary gamma-ray instrument like *Fermi* relies heavily on multiwavelength observations to produce its basic astronomical results.

However, the importance of radio observations runs far, far deeper than simply providing point source seed catalogs—in fact, the discovery of AGN is inextricably connected with the development of radio astronomy. As we will discuss below, the AGN and blazars that may produce the greatest fraction of the extragalactic gamma-ray emission are complex, intrinsically broadband emitters. Understanding these objects requires that their spectral energy distributions (SEDs) be measured at all frequencies, from radio through gamma rays. Furthermore, because these objects are often violently variable, simultaneous multiwavelength coverage is essential, and continuous, fast-cadence monitoring is extremely valuable.

To address this need, in late 2007, we began the project described by this thesis: a fast-cadence 15 GHz radio monitoring program using the 40 m Telescope at the Owens Valley Radio Observatory (OVRO). The ultimate goal of this monitoring program is to combine our 15 GHz radio data with gamma-ray light curves to demonstrate or rule out the presence of physically significant correlations. If such a connection can be established, time lags between physically connected features in the light curves could be used to identify the location in the jet where the gamma-ray emission occurred relative to the radio emission. Using very long baseline interferometry (VLBI), the location of radio emission within the jet can often be directly resolved, and in many cases moving emission features can be tracked over multiple observation epochs (e.g., Kellermann et al. 2004), so the relative position could then be used to determine the location of gamma-ray generation. This would be a valuable constraint on models of the emission processes and AGN structure. However, establishing the physical significance of apparent correlations in blazar light curves is a challenge. The presence of frequent, apparently random outbursts in both the radio and gamma-ray bands makes coincidental correlations likely, particularly in short or intermittently sampled light curves. In this thesis, we focus on an important prerequisite to such detailed correlation studies: firmly establishing that a significant intrinsic connection between the emission at the two frequencies exists.

1.1 Historical Background

Although many AGN are intrinsically broadband emitters, spanning the entire electromagnetic spectrum, the discovery of their nature and much of the history of AGN research is deeply intertwined with the development of radio astronomy. Although sources now known to be AGN were first detected in the optical and recognized as peculiar, the connection of these sources with radio-loud objects was instrumental in understanding their physical characteristics.

Although astronomical source catalogs already contained many sources now known as AGN, it was in the early 1900s that detailed evidence for their physical properties began to be compiled, beginning with observations of emission lines in nebulae, suggesting the extremely high velocities (e.g., Slipher 1917). The observational history of AGN jets began shortly thereafter, with the description of an unusual ray-like structure in an observation of M87 (Curtis 1918). In 1943, six extragalactic nebulae were singled out because of the presence of extremely broad optical emission lines, indicating the presence material at unusually high velocities (Seyfert 1943). These *Seyfert galaxies*, later subdivided into Type 1 and Type 2 classifications (e.g., Weedman 1977), exhibit emission linewidths of up to $10^4\ \mathrm{km\,s^{-1}}$. Seyfert 1 galaxies are characterized by broad permitted emission lines and narrower ($\sim$500–1000 $\mathrm{km\,s^{-1}}$) forbidden lines, while in Seyfert 2 galaxies, both the permitted and forbidden lines are narrow. This can be explained if a region producing the narrow lines is visible in both types of Seyfert galaxies and the Seyfert 1 galaxies also contain a visible region producing the broad lines. This broad-line region is obscured by dust in Seyfert 2 galaxies, but it is clear that the region is present because, e.g., a Seyfert 1–like spectrum can be detected in the highly polarized emission from scattered nuclear radiation that avoids the obscuring material (e.g., Antonucci & Miller 1985). The majority of Seyfert galaxies are spiral in morphology, including five of the six original Seyfert galaxies (Weedman 1977; Seyfert 1943).

In the early days of radio astronomy, most radio sources appeared as unresolved "radio stars," but the limited resolution available made connection of these detections with optical counterparts very difficult. By 1950, optical counterparts for only seven of the 67 known radio sources had been suggested, and even these were somewhat tentative (Baade & Minkowski 1954b). In 1954, improved radio positions obtained through interferometry enabled the unambiguous identification of optical counterparts to the radio sources Cassiopeia A, Puppis A, and Cygnus A. The first two were identified as galactic objects, while a large redshift ($z = 0.056$) was found for Cygnus A, establishing it as an extragalactic radio source (Baade & Minkowski 1954a). Its cosmological distance and radio brightness imply a high synchrotron luminosity ($5.7 \times 10^{44}\ \mathrm{erg\,s^{-1}}$), the production of which requires a total energy of $\sim 10^{60}$ erg in particles and magnetic field (Burbidge 1959). Almost simultaneously with the optical identification, Jennison & Das Gupta (1953) showed that Cygnus A consists of two distinct radio components that straddle the optical counterpart, a morphology that proved ubiquitous as radio interferometers improved and more radio sources were resolved (e.g., Maltby & Moffet 1962).

Continuing improvements in interferometers and in handling source confusion effects (Mills & Slee 1957) led to the publication of the Third Cambridge catalog (3C) in 1959 and its revision, 3CR, three years later (Edge et al. 1959; Bennett 1962). Finding optical counterparts for these sources was difficult until a breakthrough when Schmidt (1963) identified 3C 273 with a peculiar starlike optical counterpart at $z = 0.158$. This, and the identification of similar counterparts (termed *quasars*) to radio sources that followed shortly thereafter, led to extensive optical searches for quasars based on the properties of these counterparts, with the expectation that these objects would be similar to their radio-selected counterparts. It turned out, however, that 90%–95% of the optically selected quasars were radio quiet, though not radio silent, and that the original radio-loud quasars were only a small fraction of the population (e.g., Begelman et al. 1984).

The Markarian (MRK) catalog is one of the most important early optically selected AGN catalogs (Markarian 1967; Markaryan et al. 1981). Sources in this catalog were selected based on their excess ultraviolet (UV) continuum emission in order to find Seyfert-like galaxies. Most UV-selected MRK sources were found to produce emission lines and about 10% of the sources were found to be Seyfert galaxies (Weedman 1977).

With the release of the 3C/3CR radio catalogs, statistical approaches to exploring the populations of extragalactic radio sources became important. Studies of the number densities of various source populations as a function of flux density using the N–S and $\langle V/V_{\rm max}\rangle$ methods were executed to test the steady-state cosmological model and to look for cosmological evolution of the spatial density of sources. Among 3CR quasars, Schmidt (1968) found that there was evidence for cosmological evolution of the source counts that was compatible with that reported for extragalactic radio sources (Scott & Ryle 1961). It was found that quasars and radio galaxies were far more numerous in the cosmological past, peaking at $z \approx 2$ with a comoving density $\gtrsim$1000 times the local ($z = 0$) density (e.g., Begelman et al. 1984)

Although quasars were first discovered in radio, the first reports of variability came from optical observations (Matthews & Sandage 1963). The first radio variability in quasars was reported in 3C 273, 3C 279, and 3C 345 (Dent 1965) and in a Seyfert galaxy in NGC 1275 (Dent 1966). In 1968, a radio source was identified with the "variable star," BL Lacertae (BL Lac; Schmitt 1968; MacLeod & Andrew 1968). A faint nebulosity was found in the optical image of BL Lac, and its spectrum was similar to that of then-known extragalactic radio objects. Other similar objects were soon identified, and noted for their quasar-like spectra, aside from a lack of spectral lines (e.g., Blake 1970). As we will discuss below, these BL Lac objects are one of the subclasses of the class known as blazars. The other blazar subclass is now known as flat-spectrum radio quasars (FSRQs). FSRQs mostly coincide with the class of radio-loud quasars known as optically violently variable (OVV) quasars. These comprise only about 10% of the radio-loud quasar population.

1.2 Blazar Structure and Emission

Much of the difficulty in understanding the nature of AGN results from their strongly anisotropic physical structure. In AGN with jets, the intrinsic anisotropy is enhanced by relativistic beaming. The most dramatic observational effects are found in blazars, where Doppler boosting effects are strongest due to the small viewing angles relative to the jet axis. In this section, we briefly outline the generally accepted model of the physical structure of these sources and the mechanisms thought to be responsible for their observed emission. In the ensuing discussion and in this thesis, we restrict our discussions to radio-loud AGN, which make up a minority (no more than 15%–20%) of AGN (e.g., Kellermann et al. 1989).

1.2.1 Structure

The generally accepted model for the AGN structure is similar to that described by Lynden-Bell (1969), who first suggested that black holes ("dead quasars") could be responsible for the AGN phenomena. The model consists of an SMBH ($M_{\rm BH} \sim 10^5$–$10^9 M_\odot$) at the center of a host galaxy. The SMBH is surrounded by a hot accretion disk consisting of material falling in from the host galaxy. A dusty cloud, frequently depicted as a torus, surrounds the accretion disk, blocking direct observation of the accretion disk from some viewing angles. Near the accretion disk, hot clouds form the broad-line region, responsible for the production of emission lines with equivalent widths of up to $10^4\ \rm km\,s^{-1}$. Further from the disk lie the cooler narrow-line region clouds, where narrower emission lines (widths ~ 500–$1000\ \rm km\,s^{-1}$) are produced. In most radio-loud AGN, a pair of axial relativistic jets are produced by the SMBH with various observed morphologies that are thought to depend on the intrinsic strength of the jet, the characteristics of the surrounding medium, and the angle of the line of sight relative to the jet axis.

1.2.2 Emission Processes

The broadband emission from blazars is characterized by an SED with two broad components. The first, spanning the radio through ultraviolet or soft X-ray band with a peak typically in the infrared (IR) through UV, is widely accepted to result from synchrotron emission from a nonthermal population of ultrarelativistic electrons spiraling in a magnetic field. The emission mechanism producing in the second component, which typically spans the X-ray through gamma-ray bands, is not as well understood. The two common approaches are divided between leptonic and hadronic models (e.g., Böttcher 2007).

Synchrotron emission results from the acceleration of highly relativistic electrons (and/or positrons) in a magnetic field (e.g., Schott 1912; Rybicki & Lightman 1979). The spectrum of synchrotron radiation is characterized by two regimes, the optically thick low-frequency regime and the optically thin high-frequency regime. The turnover between these two regimes typically lies in the radio band. In the optically thick regime, the spectrum is described by a spectral index $\alpha = 2.5$, using the convention $S_\nu \propto \nu^\alpha$. Above the turnover frequency, the spectrum is optically thin and the spectral index $\alpha = (p+1)/2$ where p is the exponent of an assumed power law distribution of particle energies, $N(E) \propto E^p$. The optically thin spectral index in AGN and blazars is typically around $\alpha = -0.7$. At still higher energies, radiative cooling rapidly depletes the population of emitting particles, resulting in a steepening of the spectrum.

In leptonic models of the high-energy component, the emission is ascribed to inverse Compton scattering of photons by electrons and/or positrons. The same population of ultrarelativistic leptons responsible for the low-frequency synchrotron emission can up-scatter seed photons into the high-energy regime. In the synchrotron self-Compton (SSC) models, the seed photons are from the same synchrotron emission produced within the jet (e.g., Konigl 1981; Marscher & Gear 1985; Maraschi et al. 1992). An alternative possibility is that the seed photons originate outside the jet. In these external Compton (EC) models, seed photons may originate, e.g., in the accretion disk (either encountered directly by the jet, or first reprocessed by the broad-line region or other structures), from jet synchrotron emission reflected by other structures, or from other emission sources in the AGN structure (e.g., Dermer et al. 1992; Sikora et al. 1994; Blandford & Levinson 1995; Dermer et al. 1997; Błażejowski et al. 2000).

In the hadronic jet models, protons are accelerated to relativistic energies where pion production can occur, leading to pair cascades (e.g., Mannheim 1993). In these models, which require strong ($\gtrsim 10$ G) magnetic fields, the emission results from pionic emission from the primaries as well as inverse Compton emission from the secondaries. Additional complications, such as synchrotron emission from the primary protons and secondary muons and mesons, must be considered as well (e.g., Böttcher 2007, and references therein).

1.2.3 Jets and Beaming

Jets are highly collimated, radio-bright outflows found to be extending from near the central region of radio-loud AGN. The jet phenomenon is among the first unusual AGN features to be reported, although about half a century would elapse between the first observation and a clear picture of the physical origin of the jet (Curtis 1918). These structures are found on parsec scales using high-resolution VLBI imaging, as well as on longer kiloparsec to megaparsec scales in larger-scale observations. Evidence for alignment of the jet structures between the parsec-scale nuclear jets and the kilo- to megaparsec-scale extended jets or radio lobes has been found in many sources (e.g., Kellermann et al. 1975; Readhead et al. 1978a), requiring a mechanism or mechanisms for collimation that can operate over these extended distances. The alignment is not always present, with bright, compact objects more frequently showing curvature in their small-scale jets (Readhead et al. 1978b).

The composition, acceleration, and collimation mechanisms of these jets are not well understood (e.g., Fragile 2008, for a recent review). From observations of “hot spots” in the lobes of Cygnus A, it was found that a continuous source of energetic electrons was needed to maintain emission over the necessary timescales (Hargrave & Ryle 1974). The first theoretical mechanism for such a continuous injection was proposed by Blandford & Rees (1974), who also suggested a collimation mechanism based on the de Laval nozzle. This is now known to be a likely cause of recollimation on kiloparsec scales, but not of the initial collimation, which, as revealed by VLBI, clearly occurs on subparsec scales. Other viable early theories for AGN jet power sources include Blandford & Znajek (1977), which uses energy extracted from a rotating black hole threaded by magnetic fields supported by currents in the inner accretion disk, and Blandford & Payne (1982) which proposes a magnetohydrodynamic wind mechanism to extract energy from the accretion disk. These theories, which suggest magnetic jet collimation, remain promising explanations for the basic jet mechanisms.

The relativistic nature of the material in the jet is revealed in several ways. Multiepoch VLBI studies have found components traveling down the jet with *apparent* superluminal velocities in many sources. These were first observed using model fits to VLBI visibilities (Cohen et al. 1971; Whitney et al. 1971; Moffet et al. 1972) but became widely accepted after VLBI mapping methods improved (e.g., Pearson et al. 1981; Kellermann et al. 2007). This phenomenon, predicted by Rees (1966), is indicative of relativistic (but, of course, *actually* subluminal) motion of the radiating material toward the observer. Apparent velocities of $\beta_{\rm app} \sim 30$ or higher have been observed, corresponding to $\gamma = (1 - \beta^2)^{-1/2} \gtrsim 30$, where $\beta = v/c$ is the actual component speed (Cohen et al. 2007).

Further evidence that the material in the jets is relativistic comes from the rapid variability observed in many sources. Rees (1966) showed that rapid variability can be explained by expansion of the emitting region toward the observer. Using light travel time arguments, the variability timescale can be used to constrain the size of the emission region, which can then be used to compute the necessary brightness temperature, T_B, to produce the observed emission. As was shown in Readhead (1994), brightness temperatures in synchrotron emission regions greater than about 5×10^{10} K are unlikely to persist because an enormous departure from energy equipartition between the particles and magnetic field would be required. Higher brightness temperatures are frequently found in blazar sources. For example, in Abdo et al. (2009c), we used OVRO 40 m data for the narrow-line Seyfert 1 galaxy PMN J0948+0022 to estimate $T_B \approx 2 \times 10^{13}$ K. Because relativistic beaming will enhance the apparent brightness temperature, by postulating that the emission region is beamed with a Doppler factor $\delta = \gamma^{-1}\,(1 - \beta\cos\theta)^{-1} \gtrsim 7$ where θ is the angle to the line of sight, we reduce the necessary brightness temperature below the equipartition limit.

The importance of relativistic beaming in explaining the observed characteristics of radio-loud AGN was put forth by Blandford & Königl (1979), who proposed that the radio emission in these objects originates in a collimated relativistic jet. Beaming introduces complications in observational studies of relativistic jets. The continuum emission is strongly beamed along the jet axis, introducing strong observational selection effects. Because components beamed toward an observer are enhanced while those beamed away are diminished, the apparent morphology of a beamed source often does not directly reflect its actual structure. Strong boosting of the continuum synchrotron emission from the jet also frequently swamps optical line emission, making it difficult or even impossible to obtain a redshift for the source.

1.2.4 AGN Unification

Unification refers to the identification of observationally different classes of AGN as intrinsically similar structures viewed under different conditions. A good review of the unification efforts that were led by the radio astronomy community in the late 1970s and early 1980s is found in Begelman et al. (1984). A relatively recent review of unification with emphasis on later multiwavelength developments is given in Urry & Padovani (1995). Attempts to unify the various classes began in the late 1970s. A very early model was developed by Readhead et al. (1978b), who suggested that radio galaxies and quasars could be the same type of object simply viewed from different angles relative to the jet axis. They pointed out that the relativistic beaming effects that follow from this suggestion simply explained the observed superluminal motions in quasars, the larger bends between the inner and outer jet structures observed in the two classes, the difference in sizes of objects in the two classes, and the finding that quasars are more numerous at higher redshifts. In Blandford & Rees (1978), it was similarly suggested that BL Lac objects (BL Lacs) (as well as OVV quasars, now generally included in the FSRQ class) were radio galaxies viewed along their jet axis. A similar unification proposal was put forth by Orr & Browne (1982), who instead suggested that flat-spectrum and steep-spectrum quasars were identical objects with different observed properties due to the line of sight effect.

The Readhead et al. (1978b) idea has been widely adopted in the present unification paradigm: blazars are understood to correspond to radio galaxies viewed along the jet axis (e.g., Barthel 1989). Radio galaxies are divided into two classes, the low-luminosity Fanaroff-Riley type I (FR I) galaxies and the high-luminosity FR II galaxies. The two classes were originally divided at a 178 MHz luminosity of 2×10^{25} W Hz^{-1} (Fanaroff & Riley 1974), although the luminosity threshold was later shown to vary strongly with optical luminosity, likely due to deceleration of weaker jets by the interstellar medium in larger, more luminous host galaxies (Ledlow & Owen 1996). In addition to the luminosity distinction, the two classes differ morphologically. FR I galaxies are brightest at the core with dimmer lobes, whereas FR II galaxies show prominent edge-brightened radio lobes. This morphological difference is probably due to the FR II galaxies containing a more powerful, faster jet, which can penetrate the interstellar medium of the host galaxy without being significantly disrupted.

1.2.5 Blazars

Blazars are widely understood to be the beamed counterparts to the radio galaxies. Unfortunately, AGN taxonomy is complicated by rather frequent changes in terminology, in some cases to reconcile conventions from different branches of astronomy, and in some reflecting a shift in the physical understanding of the sources. The original definition of the blazar class was rather informal, so there is some variation in use of the term. In this thesis, we adopt a simple division of blazars into two classes: FSRQs and BL Lacs. This division is consistent with the most common modern usage, in particular with the *Fermi* publications (e.g. Abdo et al. 2009a, 2010a, 2010b). We adopt the definitions of FSRQs and BL Lacs based on optical emission lines specified in Healey et al. (2008). The optical spectra of FSRQs are dominated by strong, broad emission lines. BL Lacs, on the other hand, are characterized by optical spectra dominated by continuum emission, with emission lines absent or weak, with emission-line equivalent widths of < 5 Å. As a result, determining the redshift of BL Lac objects is often difficult, and the redshifts for BL Lac samples are often less than 50%

complete. Under the present unification paradigm, FSRQs are normally associated with the high-luminosity FR II radio galaxies, while BL Lacs are associated with the weaker FR I galaxies.

There is growing evidence that this appealingly simple unification is not accurate, or at least is not the complete picture. At least as early as Blandford & Rees (1978), reservations as to whether exceptional, high-luminosity BL Lac objects such as AO 0235+164 (known as J0238+1636 in our sample) should be grouped with more typical low-luminosity BL Lacs like BL Lacertae itself. The peak of the synchrotron SED component has been used as a convenient index for categorizing blazars. Abdo et al. (2010c) suggested a simple scheme dividing BL Lac objects into high synchrotron peak (HSP), intermediate synchrotron peak (ISP), and low synchrotron peak (LSP) classes, and finds common properties among FSRQ sources and LSP BL Lacs. The *blazar sequence* (Fossati et al. 1998; Ghisellini et al. 1998) suggested blazars exhibited a continuous trend of decreasing bolometric luminosity with increasing synchrotron peak frequency, with FSRQs populating the high luminosity/low-peak region, moving toward BL Lacs as the luminosity decreased and the peak frequency increased. However, the existence of low-luminosity FSRQs and high-luminosity BL Lacs is difficult to explain in this picture. Recently, Meyer et al. (2011) suggested a modified scheme that eschews optical classifications, instead dividing blazars into strong jet (FR II-like) and weak jet (FR I-like). In this picture, most FSRQs and LSP BL Lac objects fit into the strong class, while HSP and ISP BL Lacs fall into the weak class.

It is clear from observations that blazars are a major source of extragalactic gamma-ray emission. However, the exact location of the gamma-ray emission region and its proximity to the central black hole remain subjects of debate. Two possible models of the GeV emission region are that this emission comes from a *gamma-sphere* close to the base of the jet (Blandford & Levinson 1995), or that it comes from the same shocked regions that are responsible for the radio emission seen in VLBI observations much further out in the jet (Jorstad et al. 2001a). If the former model is correct then the gamma-ray observations might well provide evidence of the initial collimation mechanism.

An observational difficulty is that, except in a few cases (e.g., M87), radio observations, which provide the most detailed images of active galaxies, only probe the relativistic jets down to the point at which the jets become optically thick at a point some light-weeks or light-months from the site of the original collimation. Higher-frequency observations are needed to probe deeper into the jets, although interstellar scintillation observations do in some cases reveal the presence of radio emission features in some AGN that are $\sim$ 5–50 μas in extent (Kedziora-Chudczer et al. 1997; Dennett-Thorpe & de Bruyn 2000; Jauncey et al. 2000; Rickett et al. 2002, 2006; Lovell et al. 2008), which can be very persistent (Macquart & de Bruyn 2007). These mysterious, very high brightness temperature features are by no means understood, and are certainly of great interest. At optical wavelengths, rapid swings in the polarization position angle have been used to tie together flux density variations at TeV energies and variations at millimeter wavelengths (Marscher et al. 2008). At very high energies of hundreds of GeV to TeV, very rapid variations down to timescales of minutes have been observed by the HESS, MAGIC and VERITAS instruments (e.g., Aharonian et al. 2007, 2009;

Acciari et al. 2009, 2010). Full three-dimensional (non-axisymmetric) magnetohydrodynamic relativistic simulations are now being carried out that enable detailed interpretation of the observations over the whole electromagnetic spectrum (e.g., McKinney & Blandford 2009; Penna et al. 2010).

1.3 The Radio-Gamma Connection

Many of the proposed models for blazar jet emission predict correlation between the emission at radio and gamma-ray wavelengths. Certainly if the synchrotron-inverse Compton explanation for the blazar SED components is correct, we would expect this to be the case. A major aim of the radio observation program described in this work is to establish whether a significant intrinsic correlation between radio and gamma-ray emission from blazars is present. If such a correlation exists, then it may be possible to use cross-correlation or other techniques to measure time lags between common features in the light curves for objects measured in the radio and gamma-ray bands. The presence or absence of such detailed correlations will provide useful constraints on the models for AGN structure and emission processes.

Since the EGRET era, many attempts have been made to establish whether a significant correlation exists between the radio and gamma-ray emission from AGN. Beyond providing information about the emission processes, such a correlation is of interest, e.g., to determine the contribution of unresolved blazars to the extragalactic gamma-ray background. In some cases, evidence for a correlation was reported (e.g., Stecker et al. 1993; Padovani et al. 1993; Salamon & Stecker 1994). However, when the impact of redshift and truncation bias of nonsimultaneous observations were included, these correlations were found not to be statistically significant and the significance of apparent correlations was shown to be frequently overestimated (Mücke et al. 1997). In Taylor et al. (2007), it was reported that EGRET gamma-ray and VLBI radio flux densities did not strongly correlate among weaker radio sources.

Similar studies have continued with the availability of *Fermi* data. Kovalev et al. (2009) reported that among the bright gamma-ray sources detected by the LAT in its first three months of operation (Abdo et al. 2009a), the gamma-ray and quasi-simultaneous (within a few months) compact 15 GHz radio fluxes exhibited a correlation, and that the jets were preferentially in an active state in the epoch near their gamma-ray detection. Mahony et al. (2010) and Ghirlanda et al. (2010) found a similar flux-flux correlation among the first-year *Fermi* catalog sources using archival data from the Australia Telescope 20-GHz survey, which was conducted from 2004 to 2008. It is interesting (and encouraging) that signs of a connection continue to be found, however these studies have used limited sample sizes and/or nonsimultaneous data. Additionally, their correlation studies have not fully addressed effects like those pointed out in Mücke et al. (1997) that can lead to a serious overestimation of the significance of apparent flux-flux correlations.

In the *Fermi* LAT Collaboration paper Ackermann et al. (2011), we carried out a systematic study of the connection between radio and gamma-ray emission from the AGN detected by *Fermi* in its first year of operation. This study used concurrent radio data from this OVRO 40 m program for the 199 sources that were part of our sample during the full *Fermi* observation period and also used archival 8 GHz data for a study covering all the *Fermi* sources. A statistically significant ($p < 10^{-7}$) correlation between the radio and gamma-ray energy fluxes was found with the archival data. Using the OVRO 15 GHz radio data, it was found that using concurrent data improves the significance of the correlation, reinforcing the importance of using nearly simultaneous data when performing multiwavelength studies of variable sources. In this work, a surrogate data technique described in Pavlidou et al. (2012, *submitted*) was applied to account for selection effects and redshift biases, ensuring that robust significance estimates were obtained.

The importance of contemporaneous measurements is also supported by recent 5 GHz VLBI studies. In Linford et al. (2011), a marginal correlation was found between *Fermi* gamma-ray flux and radio flux density. In a follow-up using contemporaneous data, evidence for a strong correlation was found (Linford et al. 2012). Thus, it seems clear that comparing archival observations is of limited value in sources as strongly variable as blazars.

1.4 The OVRO 40 m Monitoring Program

The testing of models of the location, structure, and radiative properties of the gamma-ray emission region in blazars requires, in addition to the *Fermi* observations, supporting broadband observations of likely gamma-ray sources in various activity states. Such multiwavelength efforts can occur in two modes:

1. regular monitoring of a preselected, statistically complete sample of likely gamma-ray–bright objects, independent of their gamma-ray activity state; and

2. intensive observations of archetypal objects or objects exhibiting unusual behavior.

The blazar monitoring program we discuss here is focused on the first mode. In anticipation of the unique opportunities offered by the *Fermi* LAT sky monitoring at gamma-ray energies, in late 2007 we began the biweekly 15 GHz monitoring of a large sample of blazars expected to be gamma-ray emitters. We also apply our observations in studies of the second mode through LAT multiwavelength campaigns for flaring sources (e.g., Abdo et al. 2009c; *Fermi*-LAT Collaboration et al. 2010)) and through collaboration with the F-GAMMA project, a complementary effort representing the second mode, focused on radio and submillimeter spectral monitoring of about 60 prominent sources (Angelakis et al. 2010; Fuhrmann et al. 2007). In this thesis, however, we will focus only on studies in the first mode.

The initial OVRO 40 m monitoring sample included a uniformly preselected sample of 1158 blazars. As described in section 4.1.1, this sample was selected based on EGRET-era results to represent blazars likely to be detected by the LAT. Since the beginning of *Fermi* science operations in August 2008, we have expanded the sample to include AGN and blazar sources associated with LAT gamma-ray detections. The monitoring sample currently contains nearly 1600 sources, each observed twice per week.

This sample is statistically well defined and large enough to allow for statistical analyses and comparisons of subsamples. In addition, as the 40 m telescope is dedicated full time to this project, the cadence is high enough to allow sampling of the radio light curves on timescales comparable with those typically achieved by the LAT for bright gamma-ray blazars, and in this sense the 40 m and the LAT are ideally matched. This combination of sample size and cadence is unprecedented. Other long-term AGN and blazar radio monitoring programs have been carried out using single-dish (e.g., Aller et al. 1999; Teräsranta et al. 2004; Fuhrmann et al. 2007) and interferometric (e.g., Kellermann et al. 2004; Jorstad et al. 2001a, 2001b; Lister & Homan 2005; Ojha et al. 2010), but the OVRO 40 m program is unique due to its large number of sources and fast cadence.

1.4.1 Impact of OVRO 40 m Data

Data from this program, in combination with *Fermi* observations, will allow us to derive the radio and radio/gamma-ray observational properties of the blazar population, including

- the radio variability properties of the blazar population, their dependence on redshift, spectral classification, luminosity, and gamma-ray activity;
- any differences between the radio properties of gamma-ray–loud blazars and blazars with similar radio luminosity which have not been detected by the LAT;
- the properties (e.g., the significance of correlation and the length and sign of any time delays) of cross-correlations between radio and gamma-ray flares of gamma-ray–loud blazars; and
- the combination of radio properties, if one exists, that can predict the apparent gamma-ray luminosity of a blazar (which, in turn, could be used to derive blazar gamma-ray luminosity functions from radio luminosity functions).

Such a systematic study of radio and radio/gamma-ray population properties should allow us to address a series of long-standing questions on the physical properties of blazar jets, including the location, structure, and radiative properties of the gamma-ray emission region, and the collimation, composition, particle acceleration, and emission mechanisms in blazar jets.

1.4.2 Statistical Considerations

Broadly stated, the studies described in this thesis concern the correlation of data sets with the goal of identifying causal connections and ultimately demonstrating a physical mechanism responsible for that correlation. Assessing the actual significance of a result detected in this manner is challenging. Standard methods for quantifying the statistical significance of a correlation are based on assumptions that can easily be violated. In some cases, such as those addressed in Mücke et al. (1997) and Pavlidou et al. (2012, *submitted*), these result from selection effects or hidden correlations within the data themselves. However, in some cases the scientific process itself introduces problems.

Publication bias, or the *file drawer effect* is one such effect (e.g., Scargle 2000). This occurs when the ability to publish the results of a study depend on the results of that study. This feedback typically leads to publication of only the most significant results—studies that find significance of less than $p = 0.05$ are instead put away in the "file drawer" and forgotten. As a result, the publication record comprises a strongly biased, incomplete sample of the actual results.

Studies that are designed to generate hypotheses rather than to test a specific, preconceived relationship are particularly likely to generate false results (Ioannidis 2005). This results because the number of hypotheses tested is large and, since such tests are likely not independent, difficult to quantify. If a few dozen hypotheses are tested en route to publication, 3σ events will occur by chance in more than 10% of such studies. Coupled with the file drawer effect, naive calculations will grossly overestimate the statistical significance of such results. This problem is actually worsened when multiple groups work independently on similar problems unless great care is taken to account for the number of nonsignificant trials in the unpublished results of the various groups.

Avoiding fallacious statistical conclusions due to these types of effects can be difficult because their impact is even difficult to quantify. In this work, we have adopted several practices to limit our exposure to these problems and to avoid contaminating the publication record through the file drawer effect. First, we have attempted to limit the number of uncounted hypothesis trials. For example, rather than blindly testing every possible population pairing we have only evaluated correlations between populations drawn from complete, well-defined parent samples, divided into subpopulations according to criteria that correspond to plausible physical distinctions. Second, we have attempted to fully disclose the results of all such comparisons, whether or not a significant result was obtained. Through these steps, we limit our own exposure to rare chance correlations, and we minimize our own contribution to the file drawer effect.

Additionally, in the likelihood analyses discussed in this thesis and presented in full detail in Richards et al. (2011) we have been careful to account for uncertainties in our data. Without a proper analysis of uncertainties, it is difficult or impossible to determine the significance of a result. Furthermore, the application of standard significance estimators relies on the validity of their assumptions. The uncertainties encountered in monitoring data are frequently non-Gaussian and are rarely fully independent. Although in this thesis, we do assume Gaussian error distributions, our methods can easily admit more accurate models in future studies. Our group is also working to develop robust Monte Carlo and data scrambling methods for accurately estimating significances in data that do not admit a simple analytical characterization. Examples of these methods are given in Max-Moerbeck et al. (2010) and the forthcoming Pavlidou et al. (2012, *submitted*), which describes the significance estimation method used in Ackermann et al. (2011).

1.5 Overview of Thesis

The remainder of this thesis is organized as follows. In chapter 2, we examine and discuss the telescope, receiver, and radiometry procedures used to carry out the monitoring program. This includes the contributions of the various elements to the sensitivity and performance and the methods used to reduce interference and to achieve accurate and stable flux density calibration. In chapter 3, we describe the data reduction pipeline, including an overview of the software, a description of the editing, filtering, and calibration steps implemented by that software, and a detailed discussion of the database system used to store the results and intermediate data products. Chapter 4 describes the observing program, explaining the source selection criteria and the resulting properties of our samples. Some basic results from the observing program are presented in this chapter as well. In chapter 5, we examine the properties of the radio light curves we obtained from the monitoring program using a likelihood method to calculate the intrinsic variability amplitude for each source. We then compare the variability amplitudes between physically defined subpopulations of our samples.

The first three of the four appendixes discuss further details of the monitoring program that were not necessary for the discussion in the main text. Appendix A contains a user's guide to Arcreduce, the Python data reduction module developed for the reduction pipeline. Appendix B contains documentation of parts of the database system that were not discussed in chapter 3. Appendix C contains a full list of the program sources with their coordinates, redshifts, classifications, and a summary of their flux density monitoring and variability results.

In appendix D we discuss another project, the Q/U Imaging ExperimenT (QUIET). I was a part of this program during the first three years of my doctoral study and contributed to the electronics hardware design and testing, to the characterization of the polarimeter modules, and to assembly and preparation of the telescope for deployment. In appendix D, we first discuss several of my QUIET-related projects, then present the text of the submitted paper that describes the first results from the program.