The Radio Variability of Gamma-Ray Blazars

Thesis by

Joseph Lee Richards

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

California Institute of Technology

Pasadena, California

2012

(Defended October 11, 2011)

Acknowledgments

I must of course thank my family for their love, support, and so much more throughout this research. Thanks to my wife, Elina, for introducing me to astronomy, for providing the impetus to return to graduate school, and for tolerating several extra years in sunny California while I finished up. Our children, Jesse and Wendy, have provided a constant reminder that there are always things more important than research. To my parents, I owe thanks for teaching me the value of education and hard work, and for making any of this possible.

I am extremely grateful to my fellow student, Walter Max-Moerbeck, who performed far more than his share of observing, and for his patiently helping a physics student learn some astronomy. Thanks also to Matthew Stevenson, who, among many things, kept the monitoring program going at a critical juncture. Postdoctoral scholar Vaso Pavlidou has provided uncountable theoretical discussions as well as reliable advice and valuable CPU cycles, for which I am grateful. Thanks are also due to Tim Pearson for scientific guidance, always-thoughtful criticism, and technical support. Martin Shepherd's tireless dedication to precision and completeness has been inspirational, and any understanding I have of spherical trigonometry is owed to him. Erik Leitch helpfully laid out the fundamentals of calibrating 40 m data in his thesis, and his assistance debugging and extending his CMBPROG software was a great help. I would also like to thank Keith Grainge for many discussions about calibration and data filtering and for his unflappable good cheer. Mark Birkinshaw also provided very helpful suggestions on this topic, as well as fascinating stories about the 40 m. Many thanks to Dave Meier for the privilege of reading a manuscript of his forthcoming textbook on AGN physics. Thanks to Michael Seiffert for his support and mentorship during my work on the QUIET project. He taught me a great deal about radiometers, and provided very helpful advice about navigating in the academic world.

That the telescope functioned during this program is thanks to Russ Keeney. His dedication to the telescope and his understanding of its intricacies are awe inspiring. It has been an honor to work with him.

Tony Readhead, my adviser, has been a great teacher and mentor. His enthusiasm for astronomy is infectious, and his obvious love of knowledge in all its forms is inspiring. The support he provided for a graduate student facing the challenges of starting a family while carrying out his research was beyond expectations. It has been a wonderful experience working in his group, and I hope that we will continue to collaborate long into the future.

There are doubtless others who should be named here but have been omitted. My sincere thanks go to all those who have assisted me in this work.

Abstract

Since late 2007, we have regularly monitored over 1100 systematically selected blazars at 15 GHz using the Owens Valley Radio Observatory 40 m radio telescope. The number of sources in the program has grown to nearly 1600, including all the active galactic nuclei associated with *Fermi* Large Area Telescope (LAT) gamma-ray point source detections north of our declination limit of -20°. Here, we describe the first 42 months of this program, including the design and implementation of an automated data reduction pipeline and a MySQL database system for storing the reduced data and intermediate data products. Using the "intrinsic modulation index," a maximum-likelihood method, we estimate the variability amplitudes for 1413 sources from their radio light curves and compare the properties of physically defined subpopulations of the sample. We find that, among our preselected sample, gamma-ray–loud blazars detected by the LAT are significantly more variable at 15 GHz, attributable to a difference in variability between the gamma-ray–loud and gamma-ray–quiet flat spectrum radio quasars. The BL Lacertae objects in the samples do not show this division in variability amplitudes. In the first two years of our program, a 3σ-significant difference between variability amplitudes for sources at redshift $z \geq 1$ and for sources at $z < 1$ was found. This difference is found no longer to be significant in the full 42-month data set, particularly after we apply an analysis method to account for the effect of cosmological time dilation.

Contents

Figures

Tables

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 $\sim 10^{60}\ \mathrm{erg}$ 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 $\mathrm{km\,s^{-1}}$. Further from the disk lie the cooler narrow-line region clouds, where narrower emission lines (widths $\sim$ 500–1000 $\mathrm{km\,s^{-1}}$) 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} $\mathrm{W\,Hz^{-1}}$ (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.

Chapter 2

Telescope, Receiver, and Radiometry

In this chapter, we discuss the telescope, optics, and receiver used to carry out the blazar monitoring program. We also describe the radiometry and calibration procedures employed to make the measurements.

For a monitoring program of this size and cadence, making efficient use of telescope time is critically important. Although we were extremely fortunate to have use of the Owens Valley Radio Observatory (OVRO) 40 m telescope 100% of the time from the start of the program until mid-2011, and six days a week after that, our cadence of 1500 sources every three days requires careful planning. Fortunately, the majority of our sources are reasonably bright—more than 50 mJy at 15 GHz—so the sensitivity requirements of the program are relatively modest. This has allowed us to optimize for rapid observations and easily repeatable measurements rather than scrabbling for sensitivity at all costs. Still, a full understanding of the behavior of the telescope and receiver and careful measurement and calibration are essential.

2.1 The Hardware

We begin by introducing the OVRO 40 m Telescope, its optics, and the Ku-band receiver used for this program.

2.1.1 The OVRO 40 m Telescope

The OVRO "40 m" telescope is actually a 130-foot-diameter f/0.4 parabolic reflector with approximately 1.1 mm rms surface accuracy on an altitude-azimuth mount. The telescope is located on the floor of the Owens Valley near Big Pine, California, at $37^\circ 13' 53\farcs7$ N latitude, $118^\circ 16' 53\farcs83$ W longitude, and 1236 m elevation (Pearson 1999). Construction of the 40 m telescope was completed in 1966. The telescope has been used with several different receivers since then, at frequencies as high as 45 GHz, where the surface accuracy of the dish becomes a serious limit on antenna efficiency.

2.1.1.1 Telescope Control Systems

The 40 m telescope is controlled by a computer control system that provides a user interface, executes schedules, controls the drive system servos, and records radiometer output and housekeeping data for later analysis. From well before the inception of this monitoring program until August 2010, these services were provided by a control system running a Digital Equipment Corporation VAX microcomputer with user interface functions on another VAX microcomputer connected via local-area network. This control system will be henceforth referred to as the *VAX control system*. Although the VAX control system and hardware had performed admirably since its installation in the early 1990s, increasingly frequent hardware failures and concern about replacement components and maintenance led to the design and implementation of a replacement system.

On 11 August 2010, the VAX control system was permanently disconnected and a new control system designed and written by Martin C. Shepherd, henceforth the *MCS control system*, was put into operation. The MCS control system runs on a personal computer using a real-time variant of the Linux operating system. In addition to operating on more easily replaced hardware, the new control system makes use of the vastly increased capabilities of modern computer hardware to log data at a greatly increased rate and to provide a more sophisticated scheduling system.

Although the control system was replaced, the receiver, digitizer, and drive hardware were unchanged. As a result, in large part the observing methods were unaffected by this change. The most significant impact will be discussed in section 3.1 where we describe the software tools and reduction scripts, which were rewritten to work with the new data format.

2.1.1.2 Mount and Drive System

The 40 m telescope is mounted on an altitude-azimuth mount. Azimuth is measured from North through East with 0° at due North. The telescope can slew through 425°, with an overlap region in the northwest quadrant between -90° and $+335^\circ$. In elevation, the telescope can be pointed from 11.5° above the horizon to about 10° past zenith. In practice, however, the control system limits the maximum elevation to 90°, and observations are normally only made between 30° and 70° elevation to avoid excessive airmasses at low elevations and because the drive system has difficulty matching the sidereal rate near zenith. The telescope can be slewed at a maximum rate of about 15° per minute in both azimuth and elevation, but can only track a moving source at half this rate or less.

In a small azimuth range pointed due South (azimuth 180°), the telescope can be tilted down to 7° elevation. In this "service position," a ladder in one of the focus support legs enables access to the prime focus. This position is used for service, maintenance, and calibration procedures that require access to the receiver or optics. In particular, hot/cold load Y-factor measurements are performed in this position (see section 2.2.1.1).

Table 2.1. List of thermometers instrumenting the 40 m telescope

Label	Typ. Val.	Units	Description
T_{RX}	24	°C	Ambient section of receiver enclosure
T_{HEMT}	12	K	HEMT LNA temperature
$T_{70\ \mathrm{K}}$	66	K	70 K stage
$T_{15\ \mathrm{K}}$	11	K	15 K stage
T_{plate}	21	°C	Cold plate temperature
T_{switch}	80	K	Dicke switch temperature
T_{hot}	300	K	Hot load temperature during Y-factor measurement, approximate outdoor temperature otherwise
T_{backend}	26	°C	Receiver backend

2.1.1.3 Tilt and Temperature Monitoring

The 40 m telescope is equipped with two orthogonal tilt meters located in the teepee of the telescope in the alidade above the azimuth bearing. These are carefully aligned with the telescope axes and are referred to as the aft/forward and left/right meters. The purpose of these sensors is to monitor for tilt of the drive system relative to the topocentric coordinate frame due to gravity or wind. These tilts are in the range of up to a few arcminutes at most. The tilt sensor readings are used in the pointing model (see section 2.2.2.1) with a scale parameter to compensate for errors in their output calibration.

A set of thermometers at the prime focus monitor temperatures related to the receiver and its support electronics. Table 2.1 lists the thermometers and their purposes. These are primarily used to monitor for problems in the receiver, except for T_{hot}, which is used during the hot/cold load Y-factor measurement (see section 2.2.1.1).

2.1.1.4 Weather Measurements

The weather is an important consideration for our observations, most critically because in moderate winds the telescope cannot be accurately pointed, and high winds can even threaten the safety of the telescope. From the beginning of the monitoring program until the transition to the MCS control system, a simple digital weather station equipped with an anemometer reported the instantaneous wind speed to the control system. In 2009, a Columbia Weather Systems[1] Capricorn 2000EX weather station was installed. The data from the new weather station were logged for offline use when using the VAX control system. With the switch to the MCS control system, which uses the Capricorn 2000EX for real-time weather monitoring, the old weather station was retired. The Capricorn 2000EX performs instantaneous and peak gust wind measurements, as well as ambient temperature, precipitation, barometric pressure, and relative humidity measurements.

In high wind situations, wind loading on the telescope could exceed the power of the drive system, potentially leading to damage to the telescope. To prevent this, in high wind conditions, the observing program is suspended and the telescope is steered to the "stow position" at about 180° azimuth, 90° elevation. In this

[1] `http://www.columbiaweather.com`

position, the cross section of the telescope to the wind is minimized, so it is the safest orientation in high winds. With the original weather station, such a "wind stow" was triggered when the instantaneous wind speed exceeded 25 mph (11.2 $\mathrm{m\,s^{-1}}$). Using the Capricorn 2000EX, a wind stow is triggered by a 60 min peak gust speed of 25 mph (11.2 $\mathrm{m\,s^{-1}}$) or an instantaneous wind speed of 20 mph (8.9 $\mathrm{m\,s^{-1}}$. The wind stow is maintained for at least 60 min and until the instantaneous speed has fallen below 18 mph (8.0 $\mathrm{m\,s^{-1}}$) and the 60 min peak gust has fallen below 22 mph (9.8 $\mathrm{m\,s^{-1}}$).

In addition to triggering wind stows, the wind speed is also used to identify periods when the pointing of the telescope was degraded due to high winds. This is discussed in section 3.2.1.2.

2.1.2 Optics

At the prime focus of the 40 m telescope, two symmetric off-axis corrugated horn feeds are installed inside the receiver cryostat. Coupled to the parabolic reflector of the 40 m telescope, this produces a pair of approximately Gaussian beams with 157″ FWHM, separated by 12′.95 on the sky. We refer to these two beams, somewhat arbitrarily, as the "antenna" beam and the "reference" beam, or *ant* and *ref*. The beam separation is in the azimuthal direction, and because the beams are offset symmetrically from the optical axis, they are always located at equal elevations.

Figure 2.1 shows a schematic view of the optics and the waveguide section of the receiver before the low-noise amplifier. In this section, both the *ant* and *ref* signal paths are identical. After the feed, a dielectric waveguide polarizer selects left-hand circular polarization (LCP); because the received radiation is reflected from the telescope, this corresponds to right-hand circular polarization (RCP) on the sky (M. W. Hodges, personal communication; Moffet 1973).[2] As a result, linearly polarized sources of all orientations may be monitored in total intensity. The signal in each then passes through a circular-to-square waveguide transition, through a 30 dB directional coupler, and then into the Dicke switch. The Dicke switch common port and the directional couplers' ports each pass through a transition to a coaxial connector that connects to the rest of the receiver.

On the *ant* side, the directional coupler connects to the calibration noise diodes with a 30 dB reduction of the diode signal. The *ref* directional coupler is simply terminated and is included only to maintain symmetry between the two signal paths. The Dicke switch port connects to the HEMT low-noise amplifier. The calibration diodes and the receiver are discussed in section 2.1.3.

[2]Here, we adopt the Institute of Electrical and Electronics Engineers (IEEE) circular polarization convention that RCP corresponds to a clockwise temporal rotation (at a fixed position) of the electric vector from the point of view of the source, i.e., when looking in the direction of propagation.

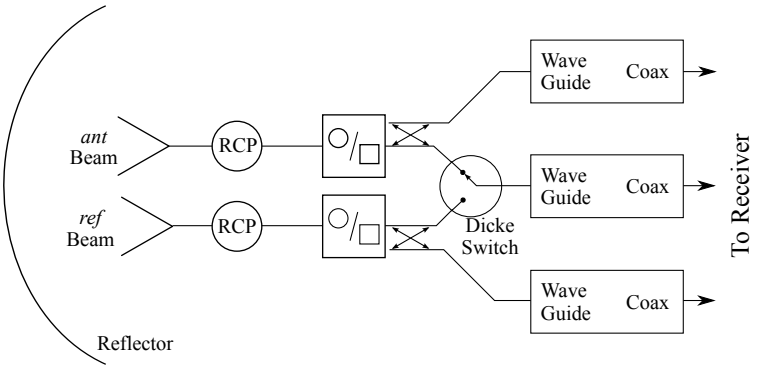

Figure 2.1. Optics and waveguide section block diagram.

2.1.2.1 Aperture Efficiency

The power received by a radio telescope with effective receiving area A_e sensitive to a single polarization is

$$P = \frac{1}{2} A_e \, S_\nu \, \Delta\nu, \tag{2.1}$$

where S_ν is the incident flux density and $\Delta\nu$ is the receiver bandwidth. If the antenna has a physical aperture area A_p, the aperture efficiency, η_a, is defined by

$$A_e = \eta_a \, A_p. \tag{2.2}$$

The aperture efficiency can be factored into contributions from a number of different causes. For example, if we combine the effects of feed illumination, spillover, and blockage into η_i, and quantify the effect of phase errors due to surface irregularities with η_p,

$$\eta_a = \eta_i \, \eta_p. \tag{2.3}$$

For an unresolved point source—a source of angular extent much smaller than the beam of the telescope—the specific intensity is effectively a delta function in angle. The response to a point source of flux density S_ν will then be

$$P = \frac{1}{2} A_e \, \Delta\nu \iint S_\nu \, \delta(\theta - \theta_0) \, \delta(\phi - \phi_0) \, B(\theta, \phi) \, d\Omega = \frac{1}{2} A_e \, \Delta\nu \, S_\nu \, B(\theta_0, \phi_0), \tag{2.4}$$

where $B(\theta, \phi)$ is the normalized antenna gain and (θ_0, ϕ_0) is the position of the source in the beam. Assuming this is centered, $B(0, 0) = 1$, so $P = A_e \, \Delta\nu \, S_\nu$. The antenna temperature, T_a, due to this point source is the temperature of a blackbody filling the aperture that gives the same response. The response to the blackbody is

$$P = \frac{1}{2} A_e \, \Delta\nu \iint I_\nu^{RJ} \, B(\theta, \phi) \, d\Omega = \frac{1}{2} A_e \, \Delta\nu \, I_\nu^{RJ} \, \Omega_a, \tag{2.5}$$

Table 2.2. Aperture efficiency measurement results

Date	η_a	Source(s)
2008-11-08	0.240	DR 21, NGC 7027, 3C 286
2009-03-11	0.255	3C 286
2009-05-19	0.247	3C 48
2009-07-07	0.261	3C 286
2009-08-10	0.275	3C 286
2009-09-10	0.264	3C 286
2009-11-11	0.269	DR 21, NGC 7027, 3C 48
2010-02-08	0.247	DR 21, NGC 7027, 3C 48
2010-04-26	0.260	3C 48
Mean	0.258 ± 0.004	

where $\Omega_a = \iint B(\theta, \phi)\, d\Omega$ is the beam solid angle. In the Rayleigh-Jeans limit ($h\,\nu \ll k_B\,T_a$), the specific intensity is proportional to temperature:

$$I_\nu^{RJ} = \frac{2\,k_B\,T_a}{\lambda^2}, \tag{2.6}$$

where k_B is Boltzmann's constant and λ is the wavelength. By equating the two detected powers from equations (2.4) and (2.5), we find

$$2\,k_B\,T_a = \frac{S_\nu\,\lambda^2}{\Omega_a}. \tag{2.7}$$

By the antenna theorem (e.g., Rohlfs & Wilson 2000),

$$A_e\,\Omega_a = \lambda^2, \tag{2.8}$$

so

$$2\,k_B\,T_a = S_\nu A_e = S_\nu\,\eta_a\,A_p. \tag{2.9}$$

We use equation (2.9) to measure the aperture efficiency of the 40 m telescope. Solving for η_a, we have

$$\eta_a = \frac{2\,k_B\,T_a}{A_p\,S_\nu} \tag{2.10}$$

where T_a is the measured antenna temperature for an unresolved point source of known flux density S_ν. In practice, we determine T_a by comparing the detected signal to a measurement of the CAL diode, whose equivalent noise temperature we know from the Y-factor tests described in section 2.2.1.1. This both converts the digitizer units (DU) to K and corrects for nonlinearity because the CAL diode measurement is affected by nearly the same amount of gain compression as the source measurement.

In table 2.2, we tabulate the measurements of η_a made during this program. Combining these and estimating uncertainty from the sample standard deviation, we find the aperture efficiency $\eta_a = (0.258 \pm 0.004)$. This relatively low aperture efficiency is due to deliberate underillumination of the dish by the feed—for monitoring observations of a large sample of objects aiming at flux density measurements repeatable to within

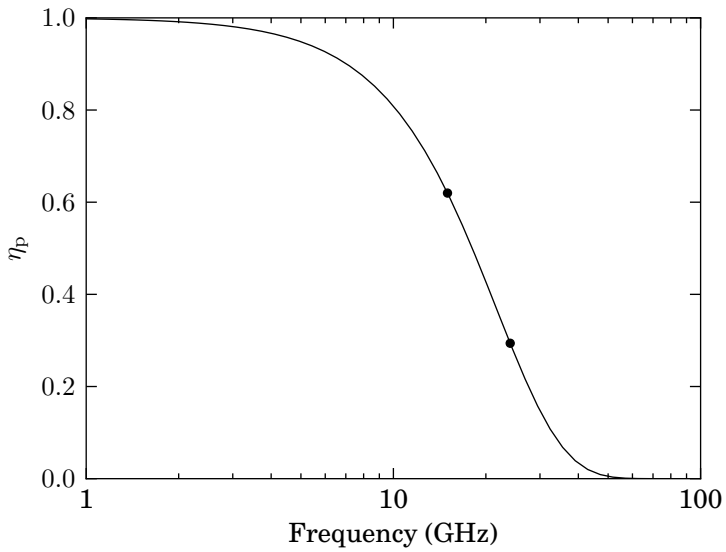

Figure 2.2. Predicted efficiency factor η_p calculated from equation (2.11). Points indicate 15 GHz (0.62) and 24 GHz (0.29) values.

a few percent we must consider the trade-off between aperture efficiency and pointing accuracy. Underillumination of the antenna increases the beamwidth and reduces susceptibility to pointing errors relative to more fully illuminating the antenna, in addition to reducing exposure to thermal noise from ground spillover. Experience has shown that we are operating at close to the optimum illumination for the most efficient use of the telescope at 15 GHz: increasing the aperture efficiency gains little because the thermal noise is already acceptably low for observing the objects in our monitoring sample.

2.1.2.2 Surface Accuracy

The surface of the 40 m telescope is composed of 852 individually adjustable panels, each with an surface accuracy of about 0.36 mm. After adjustment to match a parabolic figure at about 50° elevation, the total surface accuracy is about 1.1 $\mathrm{mm\ rms}$. The *Ruze formula* predicts the reduction of the aperture efficiency at frequency ν due to surface errors with $\mathrm{rms}\ \epsilon$ to be

$$\eta_p = e^{-(4\pi\,\epsilon\,\nu/c)^2}, \tag{2.11}$$

where c is the speed of light. Figure 2.2 shows the predicted values for η_p at various frequencies. At 15 GHz, $\eta_p = 0.62$. This accounts for a significant factor in the total aperture efficiency. To obtain the observed aperture efficiency $\eta_a = 0.258$, the illumination and blockage factors must amount to $\eta_i = \eta_a/\eta_p \approx 0.42$.

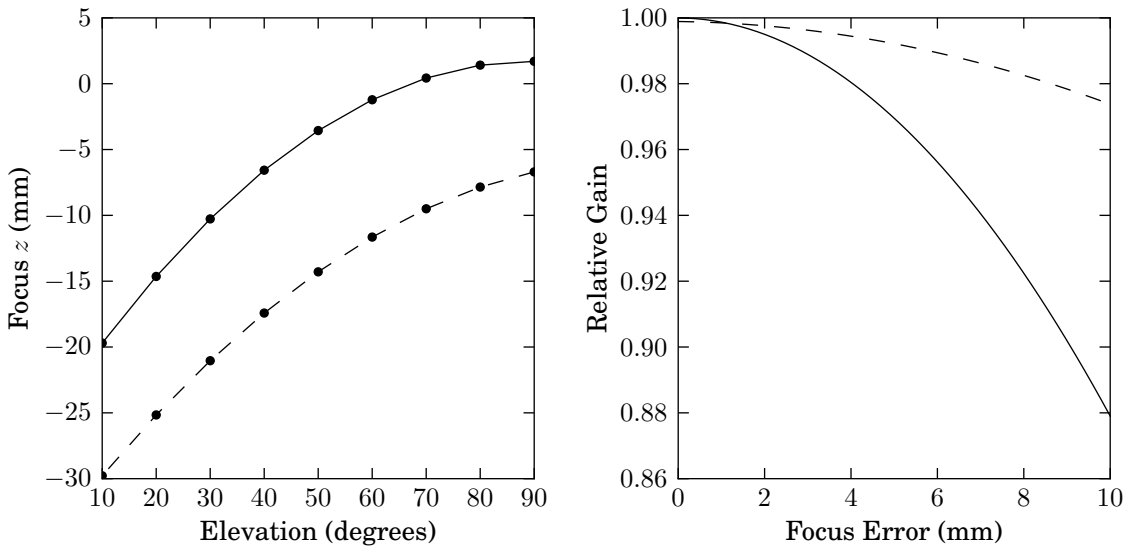

Figure 2.3. *Left:* Focus curve used during observations to predict optimum focus as a function of elevation. Points indicate the look-up table values, which are linearly interpolated as shown. The solid (dashed) line indicates the curve before (after) the shift in April 2010. *Right:* Relative gain that results from a focus error before (solid) and after (dashed) the shift. Measured flux densities are divided by the value of this curve to compensate for use of the simple focus curve rather than the more accurate model.

2.1.2.3 Antenna Gain and Focus

When the 40 m telescope moves in elevation, gravity deforms its surface, changing the antenna gain and focus location. The entire feed/receiver system can be moved along the optical axis to adjust the focus. The optimum focus position as a function of elevation is measured about once per year, but has not been found to vary significantly except when the receiver has been removed and reinstalled during maintenance. Due to thermal effects, the optimum focus also varies slightly between day and night operation and with the angle between the telescope structure and the Sun.

In normal operation, the focus is set before each observation procedure using a polynomial fit to the measured optimum focus as a function of elevation using a linearly interpolated look-up table. An example focus curve is shown in the left-hand panel of figure 2.3 with the plotted values given in table 2.3. The focus models do not change frequently. From 2008 until April 2010, measurements indicated there was no need to modify the focus curves. In April 2010, the receiver was removed from the telescope for maintenance and when reinstalled, an offset of nearly 1 cm was found in the optimum focus positions was observed, so a new focus model was determined.

During data calibration, a more complicated focus model that includes solar elongation and elevation is evaluated and a correction is applied to account for the focus error relative to that model. The ideal focus

Table 2.3. Focus curve values plotted in the left-hand panel of figure 2.3

Elevation (°)	z (Before) (mm)	z (After) (mm)
10	-19.70	-29.78
20	-14.64	-25.16
30	-10.27	-21.04
40	-6.57	-17.42
50	-3.56	-14.29
60	-1.22	-11.65
70	0.43	-9.50
80	1.41	-7.85
90	1.70	-6.70

Table 2.4. Polynomial coefficients for the focus models (before and after April 2010)

	Before			After		
n	a_n	b_n	c_n	a_n	b_n	c_n
0	7.02	—	—	-1.99	—	—
1	0.0102	-0.0677	-0.0355	-0.0241	-0.061	-0.0521
2	-0.00366	0.000209	0.000189	-0.00327	0.000156	0.000244

position, z, is given by

$$z = \sum_{n=0}^{N_a} a_n (90^\circ - \theta)^n + \sum_{n=1}^{N_b} b_n (90^\circ - \theta_\odot)^n + \sum_{n=1}^{N_c} c_n {\zeta_\odot}^n, \tag{2.12}$$

where θ is the source elevation, $\theta_\odot$ is the solar elevation, and $\zeta_\odot$ is the angular distance on the sky between the source and the sun (all measured in degrees) and a_n, b_n, and c_n are polynomial coefficients. These values are tabulated in table 2.4. The correction is calculated from a polynomial "focus miss" model, as shown in the right-hand panel in figure 2.3 with coefficients given in table 2.5. The focus miss model also changed significantly in April 2010, with the relative gain becoming much less sensitive to focus errors. This focus model and the parameters were developed and measured by Walter Max-Moerbeck by measuring the optimum focus position for point sources at many elevations and times of day.

Even with the focus adjustment and correction, the gain of the telescope is found to vary with elevation. This is due to reduced antenna gain as the reflector deforms under its own weight as it slews relative to the vertical. The surface of the reflector was set to provide an optimum parabolic surface at about 50° elevation (Pearson 1999). Figure 2.4 shows a 5th-order polynomial fit to the relative gain as a function of elevation, measured by tracking 3C 286 as it moved from about 20° to 80° elevation. The polynomial

Table 2.5. Polynomial coefficients of the focus miss curve (before and after April 2010)

n	c_n (before)	c_n (after)
0	1.0	0.999
1	-9.84×10^{-5}	-1.81×10^{-4}
2	-1.20×10^{-3}	-2.33×10^{-4}

Table 2.6. Gain curve polynomial coefficients (before and after April 2010)

n	c_n (before)		c_n (after)	
0	7.158	$\times 10^{-1}$	8.464	$\times 10^{-1}$
1	4.122	$\times 10^{-3}$	-1.210	$\times 10^{-2}$
2	7.457	$\times 10^{-4}$	1.744	$\times 10^{-3}$
3	-2.913	$\times 10^{-5}$	-6.030	$\times 10^{-5}$
4	3.752	$\times 10^{-7}$	8.270	$\times 10^{-7}$
5	-1.689	$\times 10^{-9}$	-4.106	$\times 10^{-9}$

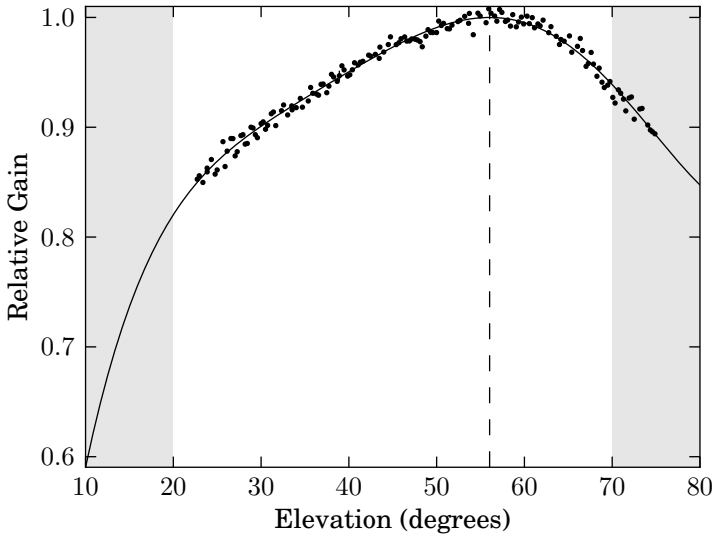

Figure 2.4. Example antenna gain curve plotting the relative peak gain as a function of elevation. Curve is a 5th-order polynomial fit to data collected from 3C 286 observations at a range of elevations on 09 March 2011. The coefficients of this polynomial are given in table 2.6 in the "After" column. Grey areas indicate regions where observations are not normally permitted and where the gain curve fit may be unreliable. Points indicate the data used for the fit, including measured uncertainties. The dashed line indicates the peak at 56.0° elevation.

coefficients for the gain curves used before and after the receiver was taken down and reinstalled in April 2010 are shown in table 2.6. Flux density measurements are corrected for this gain variation by dividing the observed flux density of a source by the value of the gain curve at the elevation of the observation. It is important to note that this correction is based on measurements of the response at the peak of the beam and is only appropriate for point-source observations like those that make up this observing program.

2.1.2.4 Beam Map

The telescope beam is characterized by the normalized power pattern $B(\theta, \phi)$, which gives the power response of the telescope to a uniform plane wave incident from direction (θ, ϕ) relative to the peak response $\max\{B\} = 1$. The two feeds at the 40 m prime focus project two symmetric beams on the sky, *ant* and *ref*,

Table 2.7. Properties of a few point sources suitable for beam mapping

Name	RA (J2000) ($^{\rm h}$ $^{\rm m}$ $^{\rm s}$)	Dec (J2000) ($^\circ$ $'$ $''$)	Size[a] ($''$)	Flux Density (Jy)
3C 48	01 37 41.3	+33 09 35	<1	1.72
3C 147	05 42 36.1	+49 51 07	<1	2.65
3C 286	13 31 08.3	+30 30 33	<5	3.44
3C 295	14 11 20.5	+52 12 10	4	1.61

Source: Flux densities and angular sizes are from Baars et al. (1977).
[a] Angular sizes specified at 1.4 GHz.

Table 2.8. Results of fitting Gaussian components to the beam center scan in figure 2.5

Beam	Amplitude (DU)	ϕ ($'$)
ant	76.1	-0.24
ref	-82.2	12.73

with an angular separation Ψ. We decompose the beam response into separate terms for each beam, that is

$$B(\theta,\phi) = B_{\rm ant}(\theta, \phi - \frac{\Psi}{2}) + B_{\rm ref}(\theta, \phi + \frac{\Psi}{2}). \tag{2.13}$$

Because of the identical construction and symmetric placement of the feeds, the individual beams are expected to have very similar responses relative to their centers, with any deviations likely to be mirror-symmetric between the beams. Unless otherwise specified, when we describe properties "the beam," we refer to one of the two offset beams, $B_{\rm ant}$ or $B_{\rm ref}$.

The coordinates (θ, ϕ) are given relative to the optical axis of the telescope. Neglecting misalignment with the mount, the θ axis is the same as the elevation axis and the ϕ axis measures angle along the great circle on the sky that is tangent to the azimuth axis. These coordinates are properly measured in the spherical geometry of the sky. However, for the very small angular extent ($\lesssim 15'$) of the 40 m beam, we can safely treat the coordinates as Cartesian, remembering that the scale factor between ϕ and the mount's azimuth coordinate varies with elevation.

To measure the beam response, we use an unresolved astronomical source to sample the response of the telescope at various angular offsets. For the 40 m with a beam $\mathrm{FWHM} \approx 2\farcm6$, a source must have an angular size $\ll 1'$ to be unresolved and regarded as a point source. Table 2.7 lists the properties of a few suitable sources.

In figure 2.5, we show the result of 50 min continuous scanning in ϕ across a source (3C 295) positioned at the elevation center of the beams. Each scan spanned $\pm 2^\circ$ around the source, but only the center region is plotted. After removing the median background, we fitted an independent Gaussian profile with a fixed $157''$ FWHM and free amplitude and ϕ position to each beam. The fitting results are shown in table 2.8. The separation between the beams was thus measured to be $12\farcm97$. This is sufficiently close to the previous value of $\Psi = 12\farcm95$, reported by Bustos (2008), that we continue to quote the old value for continuity.

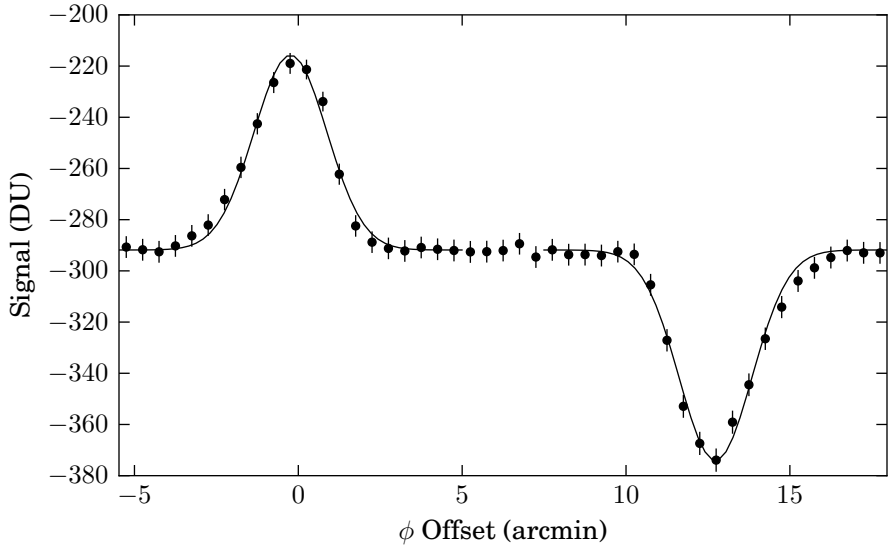

Figure 2.5. Binned switched signal (*ant* − *ref*) from 50 min of continuous azimuth scans through the elevation center of both *ant* and *ref* beams, measured on 31 Aug, 2011, using 3C 295. Lines are individual fixed-beamwidth Gaussian fits to the data spanned by the plotted line. The separation between the two beams is $12\farcm97$ on the sky.

In July 2011, we measured the beam response using 3C 286, 3C 295, and 3C 48. FLUX procedures were performed on a triangular grid with $45''$ spacing between centers. Because the beam changes as the telescope changes elevation, we restrict the measurements to elevations near 45°. Scans at constant elevation offsets from the source position were performed as sources rose or set through approximately 40°–50° elevation. To cover the region with radius $4'$ requires 109 pointings, or about 2 hours of observing, which we split into alternate elevation rows due to the elevation constraint. Several repetitions of this procedure are averaged to reduce noise, then contours are computed and plotted using the Matplotlib[3] `tricontour` routine to linearly interpolate between data points on the triangular grid.

Because FLUX procedures sample the source in both the *ant* and *ref* beams, this procedure measures the average of the two beams. This resulting average beam is shown in figure 2.6. To separate the beam responses, we also perform the same scans with an additional ϕ offset equal to the beam separation, Ψ.

Using this offset, a FLUX procedure now measures blank sky during the A and D segments and the source through the *ref* beam in the B and C segments. This, of course, reduces the on-source integration time by a factor of two, so results in a reduction in the gain of the FLUX procedure. Using an offset of $-\Psi$ similarly allows measurement of the *ant* beam alone. The separated beams, measured using 3C 48 are shown in figure 2.7. It appears that a slight inclination relative to the intended azimuthal separation is present. Elliptical Gaussian beams were fitted to the *ant* and *ref* beams individually and are plotted using the same (θ, ϕ) grid and contours in figure 2.8. Residuals from the fits are shown in figure 2.9 and the fit

[3] `http://matplotlib.scipy.org`

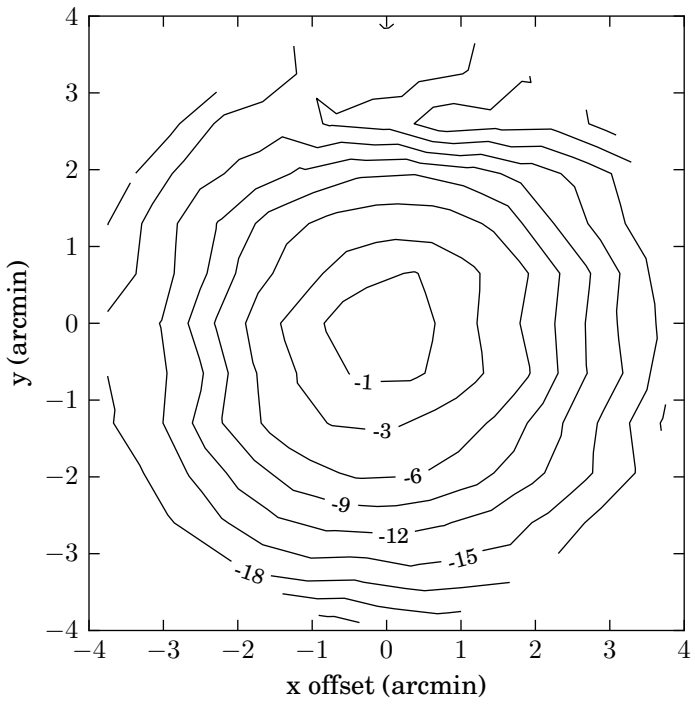

Figure 2.6. Normalized beam response for the average of the *ant* and *ref* beams, measured using 3C 286 and 3C 295. Contours are in dB relative to the peak. This depicts the effective beam for a FLUX procedure (neglecting the negative reference field lobes that are $\pm 12\farcm95$ away).

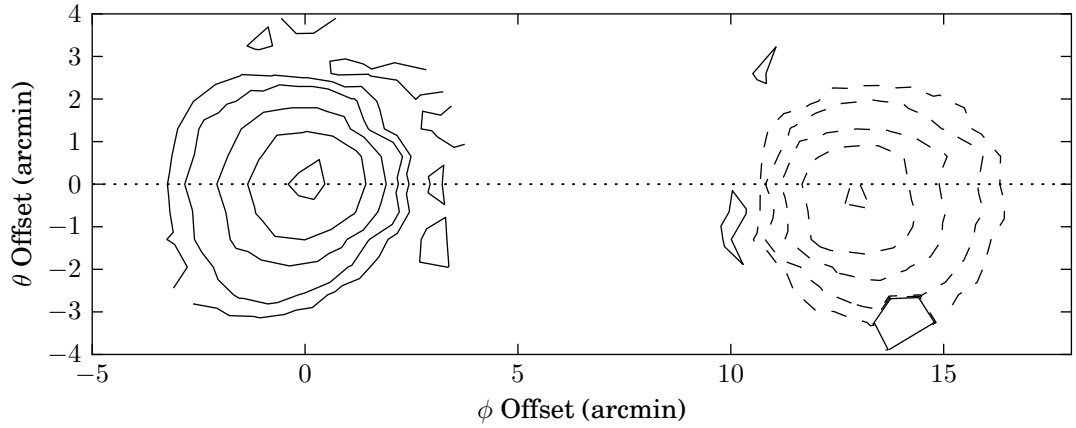

Figure 2.7. Normalized beam response for both *ant* (positive) and *ref* (negative) beams. Measured using 3C 48. Contours are 0, $\pm 5\%$, $\pm 10\%$, $\pm 25\%$, $\pm 50\%$, and $\pm 90\%$ of the *ant* peak, with dashed contours indicating negative values. The dotted line indicates pure azimuthal offset along which the beam separation is measured (see figure 2.5). Some inclination relative to this axis is apparent.

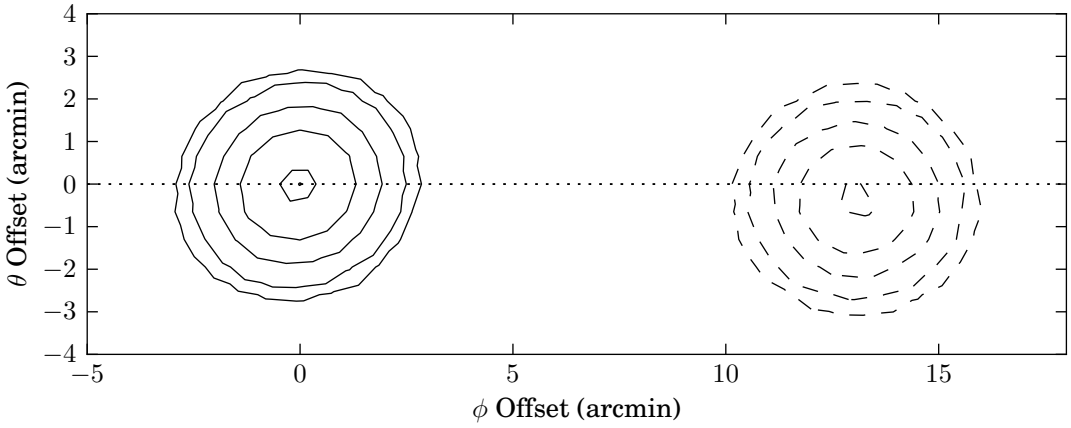

Figure 2.8. Elliptical Gaussian beam fits to the *ant* (positive) and *ref* (negative) beams. Contours are 0, $\pm 5\%$, $\pm 10\%$, $\pm 25\%$, $\pm 50\%$, and $\pm 90\%$ of the measured *ant* peak, matching the contours in figure 2.7.

Table 2.9. Elliptical Gaussian beam fit parameters

Parameter	*ant*	*ref*
Amplitude	1	0.99
Major FWHM	$3\farcm07$	$2\farcm88$
Minor FWHM	$2\farcm57$	$2\farcm69$
ϕ center	0	$13\farcm10$
θ center	0	$-0\farcm28$
Major axis inclination	$-57°$	$-60°$

Note: The Gaussian parameterization is that described in Leitch (1998).

parameters are shown in table 2.8. The results suggest that our nominal adopted beam FWHM of $157''$ ($2\farcm62$) is underestimated, but this does not affect the observations in this program.

2.1.3 Receiver

A block diagram of the receiver is shown in figure 2.10. The receiver operates in the Ku band with a center frequency of 15.0 GHz. The receiver noise temperature is about 30 K, and the typical system noise temperature including receiver, cosmic microwave background (CMB), atmospheric, and ground contributions is about 55 K.

The receiver front end consists of a cooled (T≈80 K), low-loss ferrite RF Dicke switch followed by a cryogenic (T≈13 K) HEMT low-noise amplifier. This is followed by additional room-temperature amplifiers, a 13.5–16.5 GHz band definition filter, and an electronically controlled attenuator used to adjust the overall gain of the receiver. The signal is detected directly using a square law detector diode. The detected signal is digitized with a 16-bit analog-to-digital converter and then recorded.

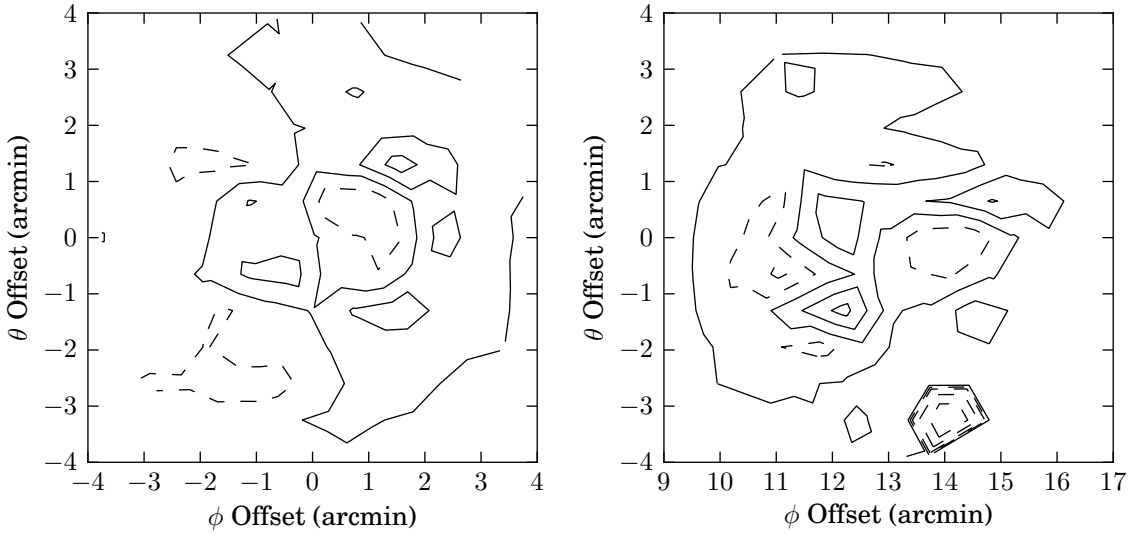

Figure 2.9. Residuals from fits to the beam map data shown in figure 2.7 for the *ant* (left) and *ref* (right). Contours show 0, $\pm 5\%$, $\pm 10\%$, and $\pm 25\%$ of the *ant* peak value.

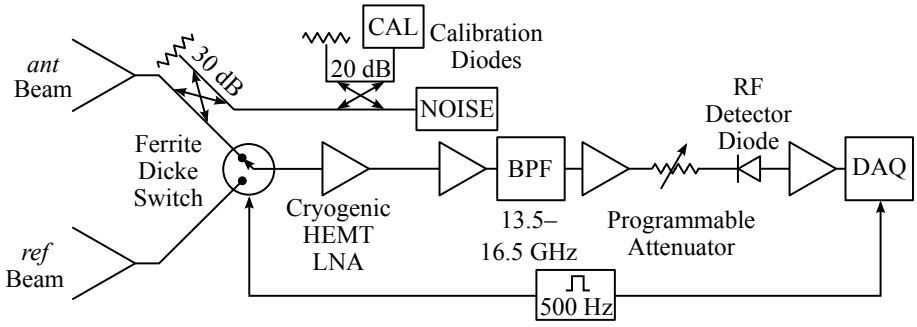

Figure 2.10. Block diagram of the Ku-band receiver.

2.1.3.1 Dicke Switching

In order to make the most efficient use of the telescope, a Dicke-switched dual-beam system is used (e.g., Rohlfs & Wilson 2000). The ferrite RF Dicke switch is switched at 500 Hz, alternately delivers the *ant* and *ref* beams to the receiver input. The radiometer output is integrated by an analog integrator circuit in each 1 ms half-period and then sampled. In software, the *ant* and *ref* samples are normally subtracted to produce the switched power, $\xi = P_{ant} - P_{refbeam}$. In some applications, such as when computing a nonlinearity correction, the average (or total) power $P = 0.5(P_{ant} + P_{ref})$ is required.

The most important benefit of Dicke switching is the removal of the large, slowly varying total power signal, which is made up of contributions from ground, atmosphere, and receiver thermal noise. Variations in the gain of the low noise amplifier cause variations in the large total power signal, and in addition the signals themselves vary slowly with time and with the position of the telescope. The resulting large variations in power limit the sensitivity of the receiving system, as discussed in section 2.1.3.3. Ground spillover, like gain variations, contributes directly to the system noise, but the effect is difficult to quantify due to the complexity of the far sidelobes of the telescope beam. Dicke switching removes or reduces these large slowly varying signals.

A second benefit of Dicke switching is the reduction of noise due to the rapidly varying atmosphere above the telescope. With a beam separation of $12\farcm95$, and for a water vapor scale height of 1.5 km, 75% of the total mass of water vapor seen by the telescope lies in the overlapping portions of the two beams. This fraction does not change substantially with scale height, dropping only to 72% (69%) for a water vapor scale height of 2 km (2.5 km). So Dicke switching reduces the effects of the varying atmosphere by about a factor of four.

A third benefit of Dicke switching is that the on-off measurement of the source against the reference allows the flux density of the source to be measured in a single pointing. This is much faster than the alternative strategy of scanning a single beam across the source. Additionally, because the source is near the peak of the beam response for the entire integration, the effective sensitivity is greater for the same integration time. This is at the cost of more stringent pointing requirements, since a mispointing will reduce the apparent brightness of the source. More details of the flux density measurement procedure are provided in section 2.2.2.2 below.

2.1.3.2 Bandwidth

The output of the receiver in response to a narrowband input signal varies depending on the frequency of that input. In general, this is a complicated function with peaks and valleys. However, the response is normally approximately zero except in some range of frequencies around the nominal frequency of the receiver. The width of this range is characterized by the *bandwidth* of the receiver. Qualitatively, the meaning of bandwidth is clear. However, there are several quantitative definitions, each useful for different calculations. Three common definitions are the half-power bandwidth, $\Delta\nu$, the noise bandwidth, $\Delta\nu_{\rm noise}$, and the radiometer

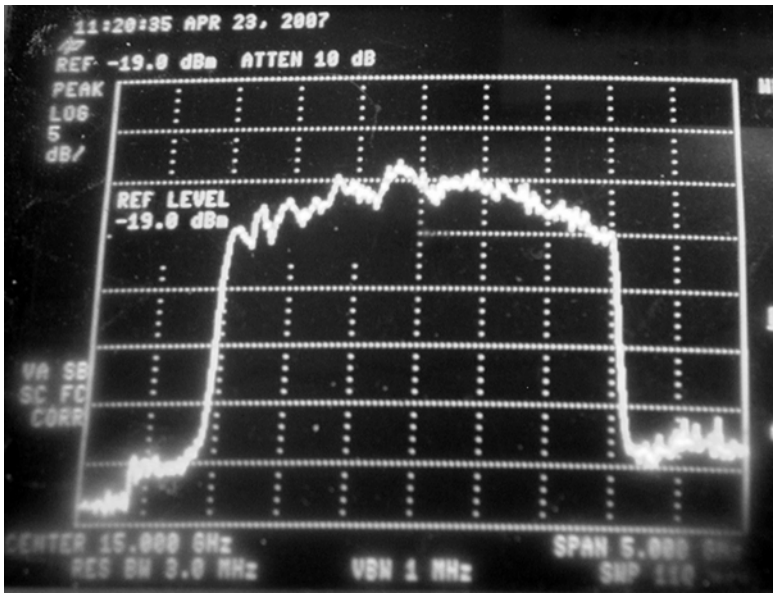

Figure 2.11. Photograph of spectrum analyzer sweep of the receiver response. Center frequency is 15 GHz, frequency span is 5 GHz (i.e., 500 MHz per division). Vertical scale is 5 dB per division. Reproduced from Bustos (2008).

reception bandwidth, $\Delta\nu_{\rm rec}$. For reasonably flat frequency responses, these bandwidths are of similar magnitude. In the literature, the nomenclature for these definitions varies, so care must be taken to determine what definition a particular author is using. Here, we adopt the convention used by Evans & McLeish (1977).

To compute these bandwidths, we begin with a spectrum analyzer trace of the receiver response. Figure 2.11 shows this response on a semilogarithmic plot, reproduced here from Bustos (2008). To work quantitatively, a piecewise linear approximation to the curve was estimated in the pass band between 13.5 and 16.5 GHz. This estimate is shown in figure 2.12 and the estimated values are tabulated in table 2.10.

Half-power bandwidth. The half-power, or 3 dB, bandwidth is the difference between the frequencies at which the receiver's power response is half that of the peak. If the ripple in the response is greater than 3 dB, the lowest and highest 3 dB points are used. This is the simplest bandwidth to measure and is frequently implied when a specific bandwidth definition is not given. Using the approximate bandpass data plotted in figure 2.12, the 40 m Ku-band receiver has a half-power bandwidth $\Delta\nu = 1.5$ GHz between 14.3 and 15.8 GHz.

Noise bandwidth. The noise bandwidth is the bandwidth of a hypothetical receiver with perfectly flat response, the same peak gain, and the same response to a wideband white noise input as the receiver in question. That is,

$$\Delta\nu_{\rm noise} \equiv \int_0^\infty \frac{G(\nu)}{G_{\rm max}} d\nu. \tag{2.14}$$

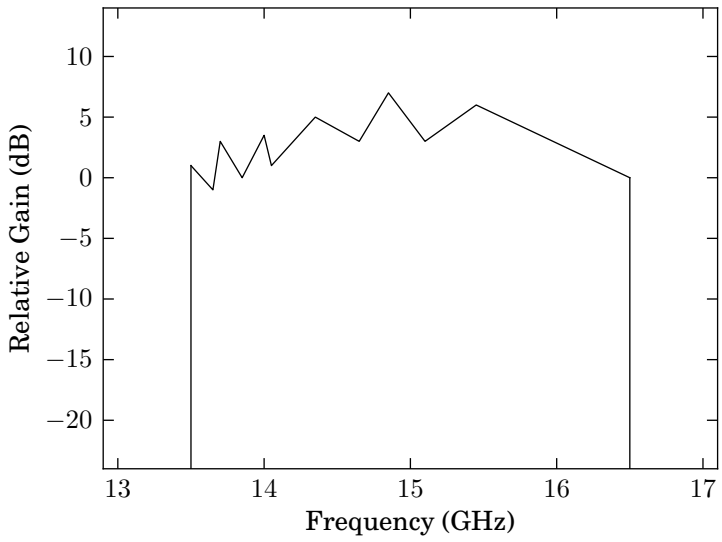

Figure 2.12. Piecewise linear approximation to the spectrum analyzer response shown in figure 2.11. Numerical values are listed in table 2.10. Attenuation is assumed infinite outside the 13.5–16.5 GHz band.

Table 2.10. Segment endpoints for the piecewise linear approximation to the measured receiver gain

ν (GHz)	G_{dB} (dB)
13.50	1.0
13.65	−1.0
13.70	3.0
13.85	0.0
14.00	3.5
14.05	1.0
14.35	5.0
14.65	3.0
14.85	7.0
15.10	3.0
15.45	6.0
16.50	0.0

Note: The approximation linearly connects these points in the semilogarithmic gain-frequency plane.

The noise bandwidth of the postdetection filter is a particularly important application of this definition. As discussed in section 2.1.3.3, it quantifies the postdetection circuit's contribution to the radiometer sensitivity. In this section, however, we compute the noise bandwidth of the receiver.

To compute the noise bandwidth from our approximate response, we must first convert the log-linear piecewise approximation into a linear-linear model. Normalizing for unity peak gain, the result is shown in the solid lines in figure 2.13. Note that our piecewise-linear model was expressed in dB, so it becomes a piecewise-exponential model rather than the piecewise-linear model shown in dashed lines. The distinction between the two is small, but we would slightly overestimate bandwidths by using the linear interpolation.

Although we could integrate numerically, it is straightforward to evaluate the integral analytically for one exponential segment of our model. For later use, we will evaluate the integral of an arbitrary nonzero power of the gain, p. If the kth segment connects (ν_0, G_0) with (ν_1, G_1), we have

$$I_k^{(p)} = \int_{\nu_0}^{\nu_1} \left(G_0 \, 10^{m(\nu-\nu_0)} \right)^p d\nu, \tag{2.15}$$

where m is the semilogarithmic slope of the segment. This has the solution

$$I_k^{(p)} = \frac{(\nu_1 - \nu_0)\,(G_1^p - G_0^p)}{p\,(\ln G_1 - \ln G_0)}. \tag{2.16}$$

Applying this to the data from table 2.10 gives $\Delta\nu_{\mathrm{noise}} = \sum_k I_k^{(1)}/G_{\max} = 1.37$ GHz.

As a test of the sensitivity of the result to errors in reading data points from figure 2.11, 10^4 perturbed piecewise models were generated by adding a random offset to each of the gain values in table 2.10 and recalculating $\Delta\nu_{\mathrm{noise}}$ with the perturbed model. Each offset was chosen from a uniform distribution between -1 and 1 dB. The mean value for this test was 1.38 GHz with a standard deviation of 0.07 GHz. We therefore quote $\Delta\nu_{\mathrm{noise}} = (1.4 \pm 0.1)$ GHz.

Radiometer reception bandwidth. The radiometer reception bandwidth $\Delta\nu_{\mathrm{rec}}$ is the bandwidth used to characterize the predetection radio-frequency bandpass of the receiver when calculating the sensitivity through the radiometer equation described in section 2.1.3.3. It is defined as

$$\Delta\nu_{\mathrm{rec}} = \frac{\left[\int_0^\infty G(\nu)d\nu\right]^2}{\int_0^\infty G^2(\nu)d\nu}. \tag{2.17}$$

Using the integral result from equation (2.16) and the data in table 2.10, we find

$$\Delta\nu_{\mathrm{rec}} = \frac{\left(\sum_k I_k^{(1)}\right)^2}{\sum_k I_k^{(2)}} = 2.57 \text{ GHz}. \tag{2.18}$$

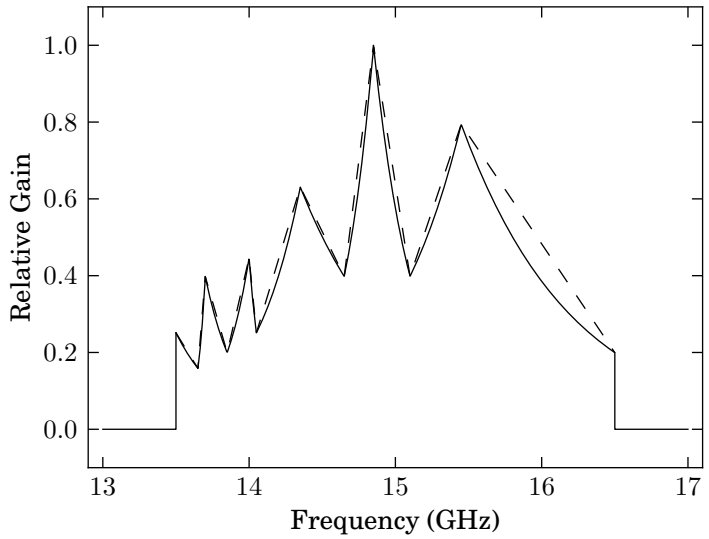

Figure 2.13. Comparison of the linear interpolation (dashed lines) with exponential interpolation (solid lines) in a linear plot of the receiver gain. Values have been normalized to unity at the peak response.

For most radiometers, the reception bandwidth is somewhat larger than the half-power or noise bandwidths, so it is not surprising that $\Delta\nu_{\rm rec}/\Delta\nu \approx 1.7$ (e.g., Evans & McLeish 1977). To evaluate the uncertainty, we use the same random perturbation method described for $\Delta\nu_{\rm noise}$, and find a mean value of 2.57 GHz and a standard deviation of 0.06 GHz. Thus, $\Delta\nu_{\rm rec} = (2.6 \pm 0.1)$ GHz. In section 2.1.3.3, we compare this to the observed thermal noise of the receiver.

2.1.3.3 Sensitivity

A very simple model of a direct detection radiometer, shown in figure 2.14, consists of a radio-frequency amplifier, a detector, and a postdetection filter. As derived in, e.g., Evans & McLeish (1977), the *radiometer equation* relates the $\rm rms$ variation in the output signal, ΔT, to the system temperature at the input, $T_{\rm SYS}$:

$$\frac{\Delta T}{T_{\rm SYS}} = K \left(\frac{2\Delta\nu_{\rm noise}}{\Delta\nu_{\rm rec}} \right)^{1/2}. \tag{2.19}$$

Here $\Delta\nu_{\rm noise}$ is the noise bandwidth of the postdetection filter and $\Delta\nu_{\rm rec}$ is the radiometer reception bandwidth of the amplifier, and the detector has been assumed to be a square-law device. The constant K is a factor that can be used to generalize this result to other receiver architectures. If the postdetection filter is a boxcar integrator of integration time τ, the noise bandwidth $\Delta\nu_{\rm noise} = (2\tau)^{-1}$. This gives the more familiar form

$$\frac{\Delta T}{T_{\rm SYS}} = \frac{K}{\sqrt{\Delta\nu_{\rm rec}\,\tau}}. \tag{2.20}$$

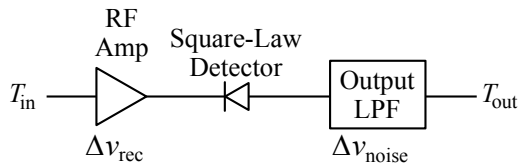

Figure 2.14. Simple radiometer model.

The radiometer equation quantifies the *minimum* achievable noise level with a given receiver and integration time. In practice, additional sources of noise will result in a a higher level. For example, if the gain of the receiver varies significantly over the integration time, output changes due to gain fluctuations will be indistinguishable from changes due to input variation. As a result, gain fluctuations reduce the sensitivity of the receiver. If such fluctuations have rms amplitude ΔG, as shown in Rohlfs & Wilson (2000) this gives

$$\frac{\Delta T}{T_{\rm SYS}} = \sqrt{\frac{K^2}{\Delta\nu_{\rm rec}\,\tau} + \left(\frac{\Delta G}{G}\right)^2}. \tag{2.21}$$

Other noise sources will similarly add in quadrature on the right-hand side.

Equation (2.20) with $K = 1$ is valid for a total power radiometer. For a Dicke-switched receiver, $K = 2$. This can be understood as two factors of $\sqrt{2}$: one because only half the integration time τ is spent on-source and another because the source and reference integrations are subtracted, combining their independent noises in quadrature. It is important to note that τ is taken to be the full integration time including *both* Dicke switch states. Although Dicke switching appears to increase the noise level by a factor of two through the K factor, in practice it usually greatly reduces the radiometer noise level by eliminating much of the $\Delta G/G$ factor in equation (2.21).

We now compare the observed noise level to that expected from the radiometer equation. The simplest comparison results from observations of blank sky at zenith, where the input system temperature should be its most stable. In figure 2.15, we plot the first two seconds of the radiometer output measured in DU for the individual beams and the difference between the two beams. These data were collected when the telescope was pointed at zenith for a one-hour period on 17 September 2011.

Using the *ant* and *ref* beam data separately, the receiver acts as a total power radiometer ($K = 1$) with $\tau = 1\ \mathrm{ms}$ per sample. If we compute the difference *ant* $-$ *ref*, the output corresponds instead to a Dicke switched receiver ($K = 2$) with $\tau = 2\ \mathrm{ms}$ per sample. In table 2.11, we use the results from the hour-long data set to estimate the per-sample rms noise and compare this to the results of the radiometer equation. For convenience, we measure $T_{\rm SYS}$ and ΔT in DU rather than converting to K. In this case, $T_{\rm SYS}$ is the average level of the radiometer input—for the switched case, this is the average of the level for the two beams.

Both the *ant* and *ref* data give consistent results. This demonstrates that the two signal paths have nearly equal $T_{\rm SYS}$ and bandpass contributions from the sections of the signal path not in common, i.e., skyward of

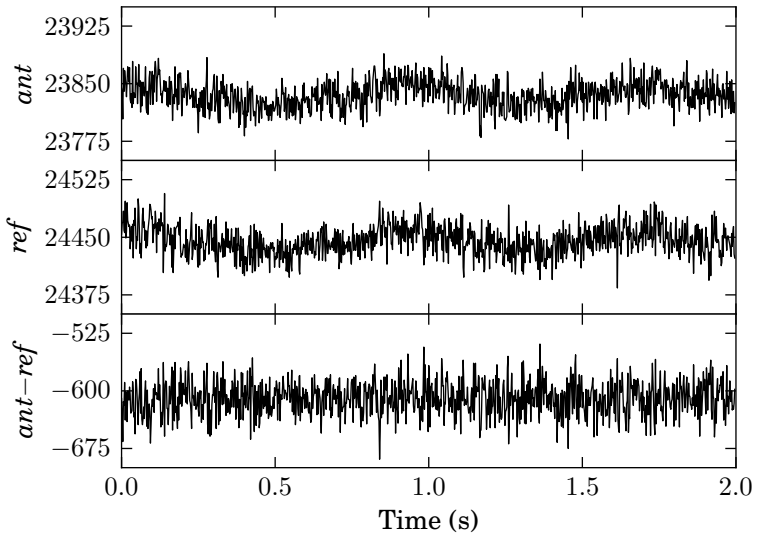

Figure 2.15. Uncalibrated full-rate (500 Hz) Dicke-switched radiometer samples (in DU) collected while pointed at zenith on 17 September 2010. The differenced signal is much flatter, indicating a reduction in slow fluctuations due to atmospheric or receiver variations or other common-mode noise sources. This is at the cost of a $\sqrt{2}$ increase in the white noise level, visible as an increase in the high-frequency scatter in the bottom panel.

Table 2.11. Calculation of receiver sensitivity and comparison with the radiometer equation

Input	K	τ (ms)	$T_{\rm SYS}$ (DU)	ΔT (DU)	Expected ΔT (DU)	$\Delta G/G$
ant	1	1	23,844.2	24.6	14.8	8.2×10^{-4}
ref	1	1	24,453.4	25.1	15.2	8.2×10^{-4}
ant − *ref*	2	2	24,148.8	22.9	21.2	3.6×10^{-4}

Note: Summary of data and calculations comparing the noise in one hour of blank-sky data (the first two seconds are shown in figure 2.15) to the radiometer equation. $K = 1$ for the total power radiometer mode and $K = 2$ for the Dicke-switched mode. Expected ΔT is computed from equation (2.20) and $\Delta G/G$ is computed from equation (2.21) assuming any excess noise results from gain fluctuations.

the Dicke switch. Thus, the difference in DC signal levels between the *ant* and *ref* signals is mostly due to a gain mismatch between the two signal paths rather than a source of excess noise in one.

The switched data match the radiometer equation more closely than do those for the individual channels. In fact, although the predicted ΔT is greater for the Dicke switched mode (21.2 DU versus 14.8 DU or 15.2 DU), the switching eliminates enough excess noise due to atmospheric and receiver gain fluctuations that the measured rms is lower (22.9 DU versus 24.6 DU or 25.1 DU). Trading a factor of two increase in the theoretical noise level (i.e., letting $K = 2$ in equation (2.21)) for the elimination of gain fluctuations through Dicke switching actually lowered the measured noise level.

2.1.3.4 Gain Fluctuations and *1/f* Noise

So far in the sensitivity discussion, we have implicitly assumed that fluctuations are spectrally white—that they are constant in amplitude over all frequencies. This leads to the result, embodied in the radiometer equation (equation (2.20)) that increasing the integration time will reduce the uncertainty in the measurement. This is generally a reasonable assumption for short times, but often breaks down for long integrations. The reason is the presence of noise processes with amplitudes that increase at lower frequencies. These processes are characterized by a power spectral density, $\Phi(f) \propto 1/f^{\alpha}$ for some $\alpha \approx 1$, and are frequently referred to as "red," "pink," or simply "$1/f$" noise. Such noise processes are ubiquitous in nature, and have been observed in systems ranging from turbulence scale distributions in lakes to gain fluctuations in semiconductor devices such as amplifiers and detector diodes (e.g., Schmid 2007; van der Ziel 1988). These latter phenomena affect radiometer sensitivity through the gain fluctuation term in equation (2.21). Note that in this section, we refer to frequency as f rather than ν, both to match the literature on this topic, and to avoid confusion between radio frequency, ν and frequencies in the postdetection signal, f.

On 27 October 2010, 27 min of data were collected while the telescope was pointed at blank sky. Figure 2.16 shows the power spectral density of the average of the *ant* and *ref* signals during 10 min of this period. The $1/f$ behavior is evident, as is contamination due to mains power at 60 Hz and harmonics and a signal from the cryogenic compressor cycling at about 1 Hz and harmonics. The white noise limit at high frequencies is estimated from data, neglecting the narrowband contamination, to be $0.53\ \mathrm{DU^2\,Hz^{-1}}$. The "knee frequency," f_{knee}, is the frequency at which the $1/f$ noise component equals the white noise, leading to a total noise that is double the white noise floor. This occurs at about $f_{\mathrm{knee}} = 17$ Hz.

In figure 2.17, we show the power spectral density for the differenced data, *ant* − *ref*, during the same 10 min period. The white noise level is found to be $2.07\ \mathrm{DU^2\,Hz^{-1}}$, or a factor of about four higher, as expected when comparing the variance of a difference to that of an average. There is no evidence for the onset of $1/f$ noise in this plot—it has been reduced tremendously by the differencing. Although in the ideal case *all* $1/f$ noise due to sources in the common signal path would be eliminated, in practice there is some residual effect due to imbalances between the inputs in the two switch states.

In figure 2.18, we plot the power spectral density of the difference *ant* − *ref* over the entire 27 min period, computed to a much lower minimum frequency. Because this increases the number of data points in the plot enormously, we have downsampled by a factor of 30 to reduce the number of points. The white noise level from these data is found to be $2.10\ \mathrm{DU^2\,Hz^{-1}}$ in agreement with figure 2.17. It appears that residual $1/f$ noise is becoming significant at very low frequencies, with an estimated $f_{\mathrm{knee}} = 5$ mHz.

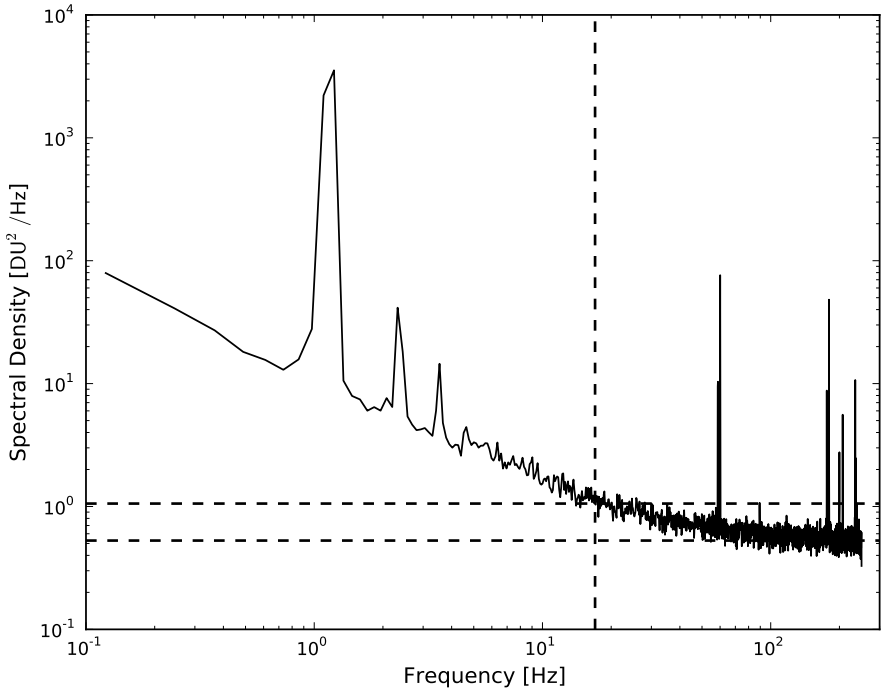

Figure 2.16. Power spectral density of 10 min of averaged *ant* and *ref* samples, illustrating clear $1/f$-type behavior at low frequencies. The horizontal lines indicate $1\times$ and $2\times$ the white noise level ($0.53\ \mathrm{DU}^2\ \mathrm{Hz}^{-1}$). The vertical line indicates the estimated knee frequency, $f_{\mathrm{knee}} \approx 17$ Hz.

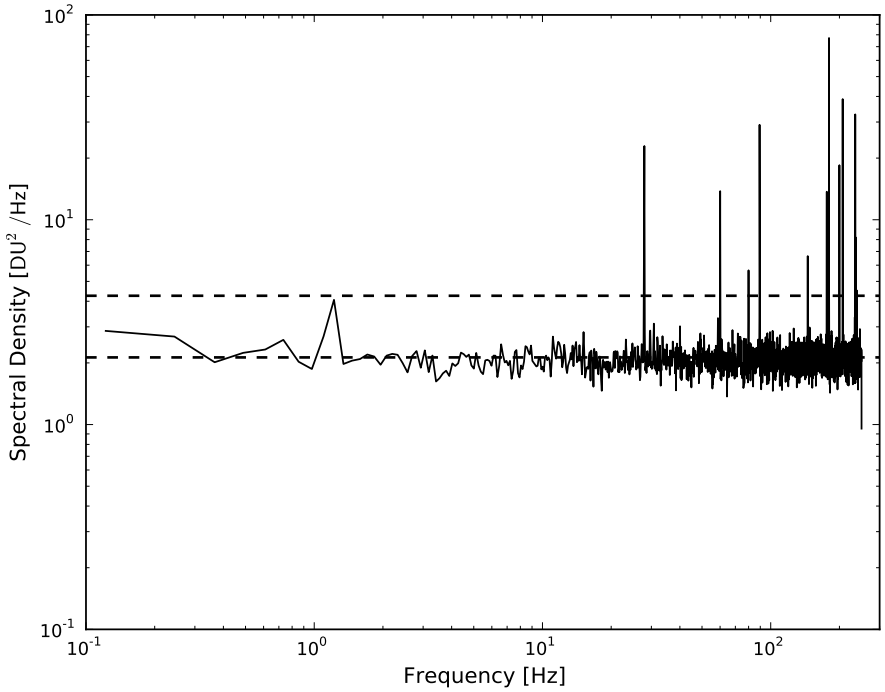

Figure 2.17. Power spectral density of the same 10 min of data shown in figure 2.16, now computed from the difference signal *ant* − *ref*. The reduction in $1/f$ noise is evident. The horizontal lines indicate $1\times$ and $2\times$ the white noise level ($2.07\ \mathrm{DU}^2\ \mathrm{Hz}^{-1}$).

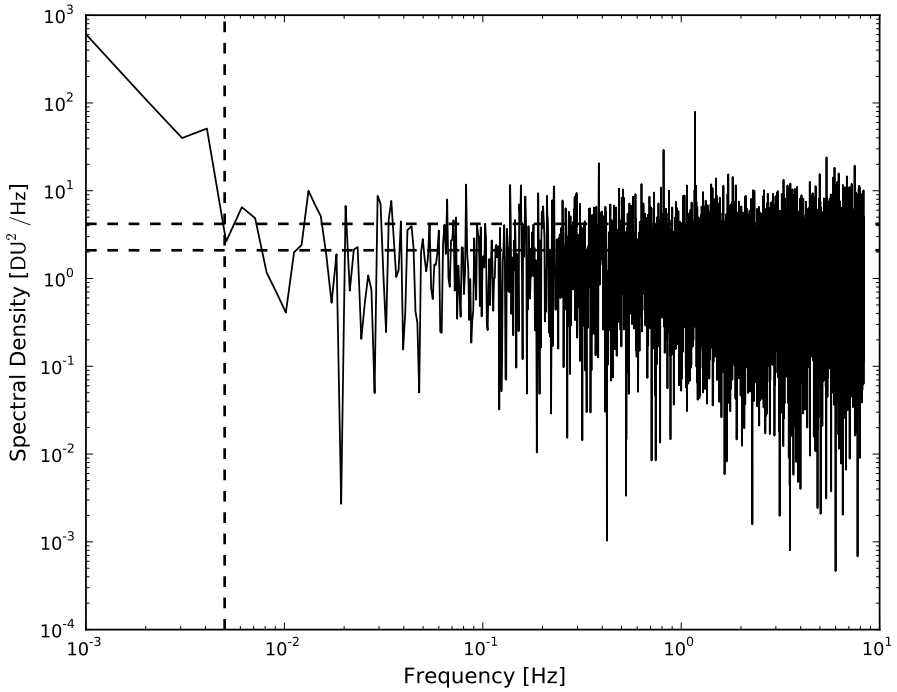

Figure 2.18. Power spectral density of about 27 min of data that overlaps the 10 min interval shown in figures 2.16 and 2.17. The spectrum is computed with finer resolution to accurately measure very low frequencies. To keep the number of data points manageable, the data were down-sampled by a factor of 30 and averaged. The horizontal lines indicate $1\times$ and $2\times$ the white noise level ($2.10\ \mathrm{DU}^2\ \mathrm{Hz}^{-1}$). The vertical line indicates the estimated knee frequency, $f_{\mathrm{knee}} \approx 5\ \mathrm{mHz}$. The "bushy" appearance of the plot at high frequencies is a visual artifact of the logarithmic binning—there is no actual increase in the white noise level at high frequencies.

The radiometer equation can be used to estimate the reception bandwidth of a receiver from the white noise level in the power spectral density of the output. Given a white noise power spectral density Φ_0 in units of power per Hz, we use $\Delta\nu_{\rm noise} = 1$ Hz in equation (2.19) and find

$$\Delta\nu_{\rm rec} = \frac{2\,K^2\,{T_{\rm SYS}}^2}{\Phi_0}. \tag{2.22}$$

In place of $T_{\rm SYS}$, we insert the average level of the input signal in DU. The average value of the input during this test was 2.5265×10^4 DU. We find $\Delta\nu_{\rm rec} = 2.41$ GHz using the averaged data ($K = 1$ and $\Phi_0 = 0.53\ {\rm DU}^2\,{\rm Hz}^{-1}$) and $\Delta\nu_{\rm rec} = 2.47$ GHz from the differenced data ($K = 2$ and $\Phi_0 = 2.07\ {\rm DU}^2\,{\rm Hz}^{-1}$). This agrees reasonably well with the value $\Delta\nu_{\rm rec} = (2.6 \pm 0.1)$ GHz we computed from the receiver bandpass in section 2.1.3.2.

2.1.3.5 Calibration Diodes

A pair of calibrated noise diodes, referred to as the NOISE and CAL diodes, are connected to the main beam input via directional couplers to the Dicke switch. At their outputs, these noise diodes provide an excess noise ratio of (31 ± 1) dB from 12–18 GHz with compensation to maintain output stability with temperature. The outputs of the noise diodes are reduced in amplitude without introducing excessive thermal noise by connecting them to the *ant* signal chain through directional couplers as shown in figure 2.10.

These calibration diodes provide two equivalent noise temperatures for calibration. The NOISE diode provides a noise temperature comparable to the system temperature and the CAL diode provides a noise temperature comparable to the antenna temperature of the astronomical sources we are observing. The equivalent noise temperatures of the NOISE and CAL diodes at the receiver input are about 67.3 K and 1 K—see figures 2.23 and 2.25.

The temperature stability of the NOISE and CAL diodes was measured using a calibrated continuous-wave RF power meter to measure the output level as the temperature was raised from room temperature (near 300 K) to about 325 K using a hot air gun. The diode under test was removed from the receiver and its metal case was bolted to a thick aluminum plate. The hot air gun was applied to the back side of the plate away from the diode and a few seconds were allowed for the diode to equilibrate with its case temperature before the output of the power meter was recorded. The case temperature was measured using an infrared thermometer (specified accuracy ± 0.2 K) aimed at a piece of tape with high infrared emissivity that was securely attached to the diode case. After the temperature was raised to the maximum tested, the diode was allowed to cool back to room temperature. Measurements were made both during heating and cooling.

Both diodes were tested, but the more accurate absolute power reading was only recorded for the CAL diode. The results are plotted in figure 2.19, showing variation of about $8 \times 10^{-4}\ {\rm dB} \cdot {\rm K}^{-1}$ assuming a linear model. The NOISE diode exhibited nearly identical relative output measurements, so we believe its temperature stability to be similar.

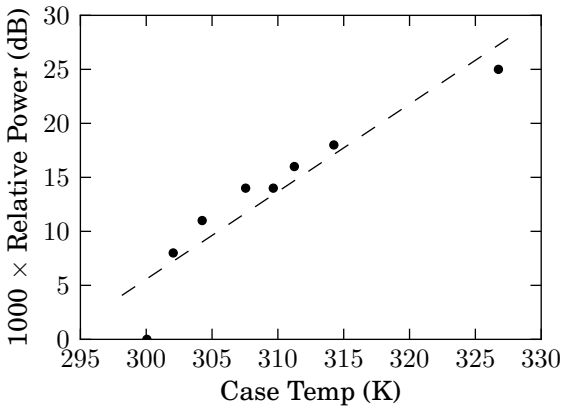

Figure 2.19. Change in absolute CAL diode power output versus its case temperature. The dashed line is a linear fit with slope 8×10^{-4} dB $\cdot$ K^{-1}.

Table 2.12. Results of a stability test of CAL diode output versus diode case temperature

Case Temp (°C)	Relative (dB)	Absolute (dBm)
26.9	0.00	−40.457
36.5	0.01	−40.443
41.1	0.02	−40.439
53.6	0.02	−40.432
38.1	0.01	−40.441
34.4	0.01	−40.443
31.1	0.01	−40.446
28.9	0.01	−40.449

Table 2.13. Results of a stability test of NOISE diode output versus diode case temperature

Case Temp (°C)	Relative (dB)
27.1	0.00
28.6	0.00
31.7	0.00
35.6	0.01
36.1	0.01
39.7	0.01
44.4	0.01
48.4	0.02
45.0	0.02
33.7	0.00
28.8	−0.01

Note: During this test, the absolute output was not recorded.

Although the calibration diodes' physical temperatures are not regulated or monitored directly, they are located inside the warm section of the receiver enclosure. This is regulated at an air temperature of approximately 300 K with rms fluctuations of about 1.1 K. In 2009, during hot weather when the ambient air temperature reached $\gtrsim 310\ K$, excursions up to about 307 K inside the enclosure were observed. During winter months occasional brief decreases as low as 288 K occurred. Using the temperature coefficient we measured, this corresponds to a 0.015 dB (0.3%) change in the output level from one extreme to the other if the diodes vary over the full range of temperatures in the enclosure. If we fit a quadratic instead of linear model to the data in figure 2.19, we predict a 0.031 dB (0.7%), still well less than 1%.

Even this level of temperature-induced output variation is unlikely to occur on short timescales. Temperature swings this large occur between seasons, not diurnally. We therefore need not worry about variations in calibration diode output due to temperature changes in the course of a day—we may treat their outputs as constant to the level of precision we require.

2.1.3.6 Beam Isolation

The Dicke switch is a ferrite switch that couples one of two feed-side ports to the common port connected to the radiometer. The device is controlled electronically by reversing the magnetic field domains in a ferrite. Such an electronic device is the only practical microwave RF switch capable of continuous 500 Hz switching for years. A trade-off must be made in the isolation between the *ant* and *ref* ports—the isolation of a ferrite switch is much lower than a mechanical switch, but the latter would wear out in a very short time. Furthermore, devices of this sort designed for cryogenic operation were not available at the time the receiver was constructed (and are now extremely rare in any form). Thus, the Dicke switch is a device intended for room temperature operation that was found to behave reasonably at temperatures down to about 70 K (Leitch 1998). As the temperature is lowered below this, the isolation drops to unacceptable levels.

By observing the response in the *ref* channel when the NOISE diode (in the *ant* channel), the isolation of the Dicke switch is found to be about 15 dB. The isolation need not be symmetric, but Leitch (1998) found the two directions to match to about 0.5 dB. Because the isolation is nearly symmetric, the effect of isolation is to reduce the effective gain for a flux density measurement measured as described in section 2.2.2.2. Because we observe our astronomical calibration sources with the same procedures as our program sources, therefore, an explicit isolation correction is not required.

2.1.3.7 Nonlinearity

Although amplifiers and detectors are designed to exhibit linear response, deviation from linearity at some level is inevitable. The most important nonlinear effect for our receiver is gain compression, a reduction in gain as the output level increases. This effect is demonstrated in figure 2.20, where the difference between the ideal linear response (solid line) and the actual response (dashed line) is due to compression.

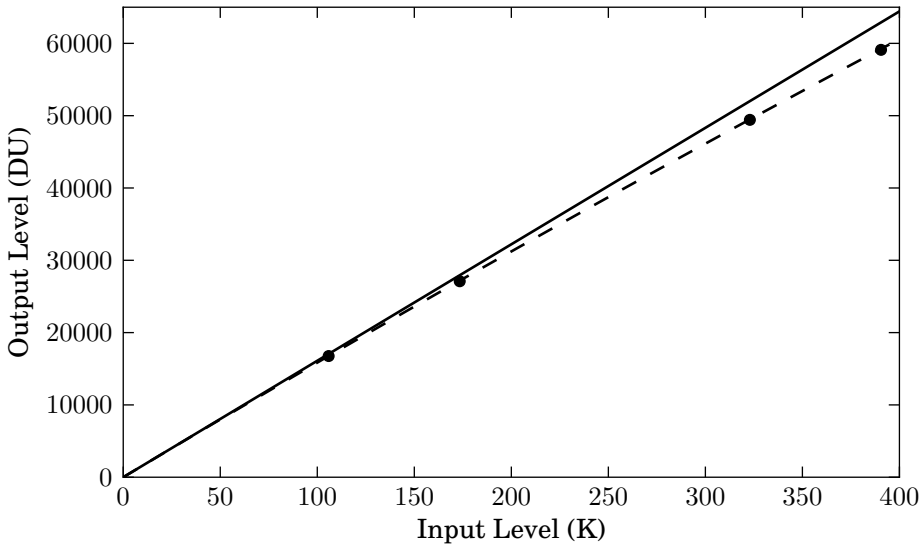

Figure 2.20. One trial of a hot/cold Y-factor measurement demonstrates the presence of gain compression in the receiver. The four input levels, here measured in K, correspond to the LN2 cold load (assumed to be 77.4 K) and the ambient-temperature hot load (measured to be 294.4 K), each with and without the NOISE diode enabled (equivalent noise temperature measured to be 67.6 K). The measured receiver temperature from this trial, 28.5 K, is also included. The solid line is a constant gain of 161 $\mathrm{DU} \cdot \mathrm{K}^{-1}$ with no compression. The dashed line through the data points is the same gain, but including compression with $b = -9.824 \times 10^{-7}\ \mathrm{DU}^{-1}$.

Modeling the exact nonlinear behavior of a receiver system is complicated, but for our needs, a simple empirical model is sufficient. We use the same nonlinearity model developed in Herbig (1994) and also described in Leitch (1998), where deviation from linear response is characterized by a single parameter, b. The model simply adds a gain factor proportional to the output level to the constant component of the gain,

$$y = G\ (1 + b\, y)\, x, \tag{2.23}$$

where x and y are the input and output signals, respectively, and G is the nominal gain measured well below where compression effects are significant. Here, b is given in inverse units of the output signal. This model is reasonable for responses that exhibit slight compression, when $|b| \ll y^{-1}$. Figure 2.21 demonstrates the effect of various b values. For a system with gain compression, $b < 0$. When $b > 0$, the model response describes a gain expansion.

Equation (2.23) describes the response of the system for any input. For many measurements, we are interested in the small-signal gain. To find this, we solve for y as a function of x, giving

$$y = \frac{G\, x}{1 - b\, G^{-1}\, x}. \tag{2.24}$$

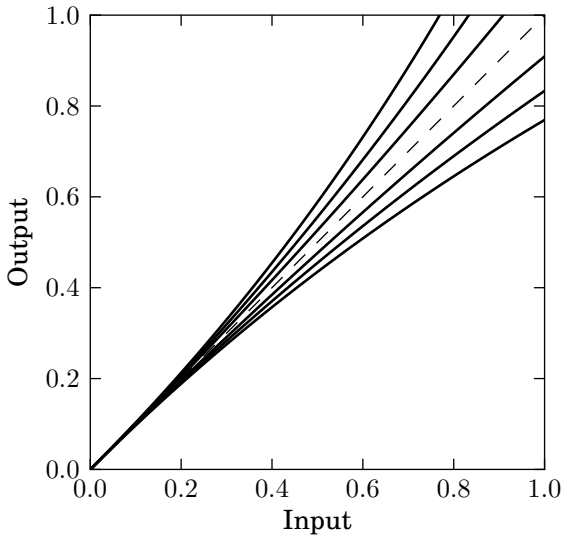

Figure 2.21. Behavior of the compression model given by equation (2.23) for $b = 0$ (dashed line) and -0.3, -0.2, -0.1, 0.1, 0.2, 0.3 (solid lines, bottom to top).

The small-signal gain near an input operating point x_0 with corresponding output y_0 is then given by the derivative at that point,

$$\left.\frac{dy}{dx}\right|_{x_0} = \frac{G}{\left(1 - b\,G^{-1}\,x_0\right)^2} = G\,\left(1 + b\,y_0\right)^2. \tag{2.25}$$

This is the correction that would apply to a differential measurement, such as when Dicke switching is employed.

This compression model is convenient because the correction can be applied given the output level and a measured value for b. For a measured output y, the corrected input level is simply

$$x = \frac{G^{-1}\,y}{1 + b\,y}. \tag{2.26}$$

For a differential measurement Δy, the input difference is given by

$$\Delta x = \frac{G^{-1}\,\Delta y}{\left(1 + b\,y\right)^2}. \tag{2.27}$$

Our receiver is constructed of several components that each individually exhibit compression. As described in Leitch (1998), the total compression of the system will depend on the signal level presented to each component and will change if the attenuation between components changes. Most significantly, the nonlinearity parameter must be measured for each setting of the programmable attenuator shown in figure 2.10. During normal observations, the attenuator is very rarely adjusted, so we simply measure the overall b parameter for the operating attenuation. However, during hot/cold calibration procedures, a higher attenuation level is used, so b must be measured separately for these procedures.

2.2 Radiometry Techniques

In order to make use of the hardware described in the previous section, its performance characteristics must be measured and understood and measurement techniques must be specified. In this section, we first describe several calibration procedures that are performed to measure and monitor the receiver performance. We then describe the measurement procedures that are used as the basis for the monitoring program.

2.2.1 Calibration and Diagnostic Procedures

Regular diagnostic tests must be performed to ensure that the telescope and receiver are performing as expected and that the data can be properly calibrated. Fortunately, the 40 m telescope and receiver have been very stable and reliable, simplifying these tasks a great deal. Calibrations have been performed regularly, on a schedule that has evolved as we have better understood the stability and requirements of the instrument.

From 2007 until November 2008, receiver calibrations were performed occasionally but without a regular schedule. Beginning in November 2008, approximately monthly calibrations were performed. In December, 2009, the schedule was reduced to semimonthly, which has proven adequate. These calibrations include Y-factor measurements to characterize receiver temperatures, sky dips to measure atmospheric optical depth and to determine the ground spillover, calibration diode effective temperature measurements, and observations of calibration sources to measure the aperture efficiency. We describe these methods here (see section 2.1.2.1 for the aperture efficiency measurement procedure).

2.2.1.1 Measuring Receiver and Calibration Diode Noise Temperatures

We use a single calibration procedure to determine the receiver noise temperature and to find the equivalent noise temperature of the calibration diodes. The basis for our measurement is the Y-factor method (e.g., Evans & McLeish 1977; Rohlfs & Wilson 2000).

To measure the receiver noise temperature using the Y-factor method, two input loads of different known brightness temperatures are needed. The Y-factor is defined as

$$Y \equiv \frac{P_H}{P_C}, \tag{2.28}$$

where P_H and P_C are the receiver output levels when presented with hot and cold loads at temperatures T_H and T_C, respectively. For a receiver with a noise temperature $T_{\rm RX}$ and an input temperature T,

$$P = G\left(T + T_{\rm RX}\right). \tag{2.29}$$

Using the two loads and solving, we find

$$T_{\rm RX} = \frac{T_H - Y\,T_C}{Y - 1}. \tag{2.30}$$

The uncertainty in this measurement is then given by

$$\sigma^2_{T_{\mathrm{RX}}} = \left(\frac{1}{Y-1}\right)^2 \sigma^2_{T_H} + \left(\frac{Y}{Y-1}\right)^2 \sigma^2_{T_C} + \left(\frac{T_H - T_C}{(Y-1)^2}\right)^2 \sigma^2_Y, \tag{2.31}$$

where σ^2_x is the mean-squared uncertainty in the value of x. Note that in equation (2.29), the receiver response was assumed to be linear with no offset. In practice, neither assumption is valid, so corrections must be applied.

The procedure we follow when performing a hot/cold load test comprises a series of five measurements performed with the receiver on the telescope:

1. fill the feed aperture with the hot load and measure P_H;
2. enable the NOISE diode and measure P_{H+N};
3. disable the NOISE diode, disconnect the radiometer input from the DAQ back end and measure P_Z;
4. remove the hot load, fill the feed aperture with the cold load, and measure P_C; and
5. enable the NOISE diode and measure P_{C+N}.

This procedure is repeated several times. The results are averaged and their scatter is used to estimate the uncertainty in the result.

Both the hot and cold loads are constructed from blocks of radio-absorbent foam mounted with a reflective metal backing. The hot load is kept at the ambient temperature and is instrumented with a thermometer to measure T_H. The cold load is enclosed in a radio-transparent insulating box that is filled with liquid nitrogen (LN2), so $T_C \approx 77$ K. Care must be taken to minimize uncertainty in T_H and T_C, e.g., by giving the cold load time to equilibrate with its LN2 bath, by keeping the hot load in the shade to avoid direct solar heating, and by working quickly to minimize changes in ambient temperature and avoid excessive LN2 boil-off. Even with these precautions, it is difficult to precisely estimate the effective load brightness temperature, which depends both on the physical temperature of the LN2 and the effect of reflections due to imperfect matching and absorption by the absorbing foam. Based on experience, the effective 2 cm brightness temperature for a well-matched box load is likely within 1–2 K of the physical temperature (D. P. Woody and J. W. Lamb, personal communication). We adopt a value of $T_C = 77.4$ K for the cold load, which may be a slight underestimate of its true brightness temperature. For the typical value $Y \approx 3$ in our tests, each 1 K underestimate in T_C would result in a $\sim$1.5 K overestimate of T_{RX}.

The hot load brightness temperature is similarly uncertain due to reflection and imperfections in the load, as well as the possibility that the thermometer does not reflect the physical temperature of the load due to thermal gradients in the foam. It is more difficult to estimate the uncertainty because reflections from the load may be terminated on the ground ($T \approx 300$ K), on the sky ($T \approx 10$ K, including the CMB contribution), or within the cryostat. In any case, the temperature of the terminating material is likely to be similar to or less

than the measured temperature of the hot load, so the measured T_H is probably a slight overestimate of the true brightness temperature. Each 1 K overestimate in T_H increases the measured $T_{\rm RX}$ by ~ 0.5 K for our typical value $Y \approx 3$. Thus, the receiver temperatures we report are perhaps slightly conservative.

Given a single trial, the data are reduced as follows. First, we estimate the nonlinearity parameter b_{HC} in equation (2.23) applicable for the measurement. Because the ~ 300 K temperature of the hot load dictates a higher programmable attenuator setting for this test than is used for normal observing, this measurement of b_{HC} cannot be used to correct other data. We find b_{HC} by comparing the change in receiver output due to the NOISE diode when measured against the hot load ($\Delta_H \equiv P_{H+N} - P_H$) versus when measured against the cold load ($\Delta_C \equiv P_{C+N} - P_C$). The parameter b_{HC} is given by

$$b_{HC} = \frac{2\left(\sqrt{\Delta_C} - \sqrt{\Delta_H}\right)}{\Sigma_C\sqrt{\Delta_H} - \Sigma_H\sqrt{\Delta_C}}, \tag{2.32}$$

where $\Sigma_H = P_H + P_{H+N} - 2P_Z$ and $\Sigma_C = P_C + P_{C+N} - 2P_Z$ are the average offset-corrected output levels.

Next, using equation (2.26) we correct our measured values for nonlinearity and offset using the measured value for b_{HC}. For P_H, this is

$$P'_H = \frac{P_H - P_Z}{1 + b_{HC}\left(P_H - P_Z\right)}, \tag{2.33}$$

and similar for P_C, etc. Now using the corrected P'_H and P'_C values to compute Y, equation (2.30) yields $T_{\rm RX}$.

The receiver gain (in $\mathrm{DU} \cdot \mathrm{K}^{-1}$) at this attenuator setting can be calculated as

$$G_{HC} = \frac{P'_H - P'_C}{T_H - T_C}. \tag{2.34}$$

This can be used to convert the measured noise diode response from DU into its equivalent noise temperature in K:

$$T_{\rm NOISE} = G_{HC}^{-1} \times \left(P'_{H+N} - P'_H\right). \tag{2.35}$$

Although b_{HC} and G_{HC} are only applicable to the attenuator setting used for the hot/cold load test, $T_{\rm NOISE}$ is a property of the NOISE diode output and does not depend on the attenuator. The equivalent temperature of the CAL diode, $T_{\rm CAL}$, is determined from the ratio of its response to that of the NOISE diode. Because the response of the CAL diode is much smaller than that of the NOISE diode, this comparison is performed at the normal operating attenuator setting rather than at the hot/cold load test setting.

Results from the Y-factor tests and the related measurements are shown in figures 2.22 through 2.27. In figure 2.22, the large $T_{\rm RX}$ outlier in February 2010 resulted from a hardware problem in the receiver that was corrected shortly after the measurement.

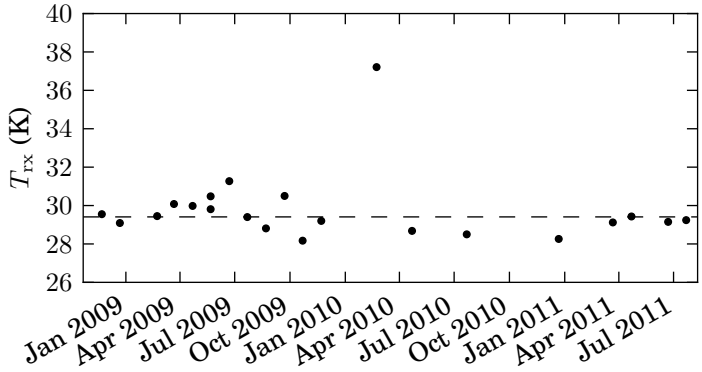

Figure 2.22. Receiver temperature measured between November 2008 and July 2011. Error bars estimated on individual measurements, typically about 0.15 K, are smaller than the plotted points. Dashed line indicates the mean, $T_{\rm RX} = (29.4 \pm 0.2)$ K, excluding the outlier in February 2010 and using the sample standard deviation to estimate the uncertainty.

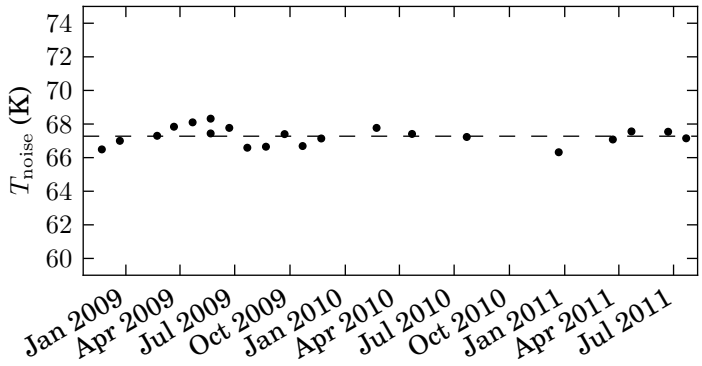

Figure 2.23. NOISE diode equivalent noise temperature measured between November 2008 and July 2011. Error bars estimated on individual measurements, typically about 0.05 K, are smaller than the plotted points. Dashed line indicates the mean, $T_{\rm NOISE} = (67.3 \pm 0.1)$ K, where the uncertainty is estimated from the sample standard deviation.

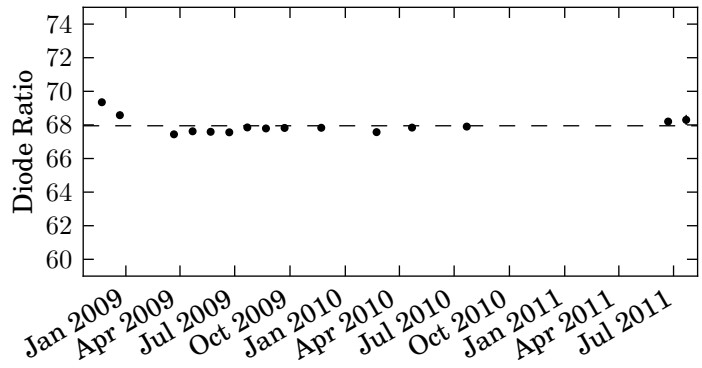

Figure 2.24. Nonlinearity-corrected NOISE to CAL output ratio measured between November 2008 and July 2011. Error bars estimated on individual measurements, typically about 0.06, are mostly smaller than the data points. Dashed line indicates the mean ratio, (68.0 ± 0.1), where the uncertainty is estimated from the sample standard deviation.

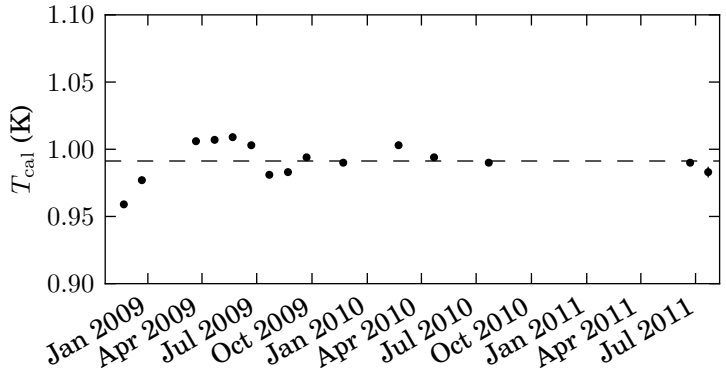

Figure 2.25. CAL diode equivalent noise temperature measured between November 2008 and July 2011. Error bars estimated on individual measurements, typically about 0.001 K, are mostly smaller than the data points. Dashed line indicates the mean, $T_{\rm CAL} = (0.991 \pm 0.003)$ K, where the uncertainty is estimated from the sample standard deviation.

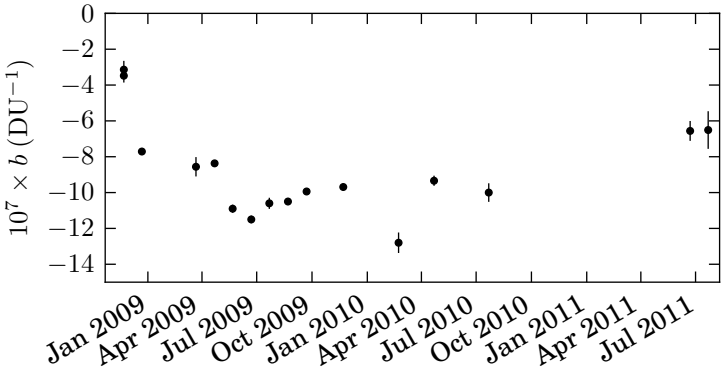

Figure 2.26. Nonlinearity parameter b measured during skydip tests between November 2008 and July 2011. The uncertainty is typically about 10^{-8} and many error bars are smaller than the data points. The 2008 data were measured with the programmable attenuator set to 5, leading to a smaller $|b|$ on those dates. The decrease in $|b|$ in the June and July 2011 may result from a change in receiver gain following maintenance or from a change in the measurement procedure.

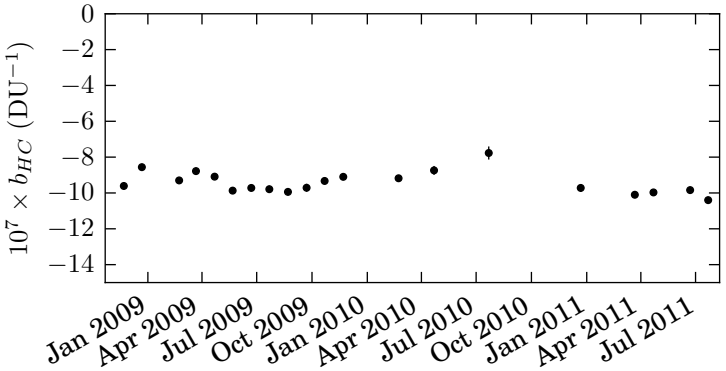

Figure 2.27. Nonlinearity parameter $b_{\rm HC}$ measured during hot/cold load tests between November 2008 and July 2011. The uncertainty is typically about 10^{-8} and many error bars are smaller than the data points.

2.2.1.2 Measuring Atmospheric and Ground Pickup

In order to measure the emission due to the atmosphere, we perform a "sky dip" measurement. That is, we measure the receiver output as we dip the telescope from near zenith to near the horizon. At lower elevations, the path length through the atmosphere—the air mass—increases. The air mass at the moderate elevations at which we observe is well approximated by assuming a plane-parallel atmosphere, where the air mass $d_{\rm atm}$ at zenith angle z in atmospheric depths is

$$d_{\rm atm} \approx \sec z. \tag{2.36}$$

If we assume that the absorption coefficient α of the atmosphere is uniform, then the atmospheric optical depth is simply

$$\tau = \int_0^{d_{\rm atm}} \alpha \, dy = \alpha \, d_{\rm atm}. \tag{2.37}$$

When observing an astronomical source with a brightness temperature T_s, the detected signal will then be

$$\begin{aligned} T &= (T_s + T_{\rm CMB})\, e^{-\tau} + T_{\rm ATM}\, \left(1 - e^{-\tau}\right) + T_{\rm GND} + T_{\rm RX} && (2.38) \\ &\approx (T_s + T_{\rm CMB})\, (1 - \alpha \, d_{\rm atm}) + T_{\rm ATM}\, \alpha \, d_{\rm atm} + T_{\rm GND} + T_{\rm RX}, && (2.39) \end{aligned}$$

where $T_{\rm ATM}$ is the atmospheric temperature, assumed constant, $T_{\rm GND}$ is the temperature due to ground pickup, and $T_{\rm RX}$ is the receiver noise temperature. The latter approximation applies if $\tau \ll 1$. We are normally justified in assuming $T_{\rm GND}$ is constant over short times because the antenna sidelobes that couple it to the receiver are extended in solid angle, so it measures the average temperature over a large area of the ground. Exceptions to this occur, for example, during periods when the ambient temperature is changing rapidly, when localized radio-frequency interference sources are present, or when the telescope is tipped to low enough elevation that near-in sidelobes intersect the ground. The latter can occur at relatively high elevations when pointed toward nearby mountains, but as these are less than 10° elevation, this is not a concern for normal observing.

Figure 2.28 shows the results for a single sky dip performed on 27 April, 2011. The radiometer data were corrected for nonlinearity, then converted to temperatures by comparing their response to NOISE diode measurements just before the sky dip. In panel (c), the dashed line is the fit

$$T_{\rm ant} = (5.2\, d_{\rm atm} + 47.3)\ {\rm K}. \tag{2.40}$$

Comparing this to equation (2.39) with $T_s = 0$, we conclude $\alpha\, (T_{\rm ATM} - T_{\rm CMB}) = 5.2$ K and $T_{\rm GND} + T_{\rm RX} + T_{\rm CMB} = 47.3$ K. Although the temperature of the atmosphere varies, in the troposphere where most of the water vapor responsible for atmospheric radio absorption and emission is located, temperatures are typically

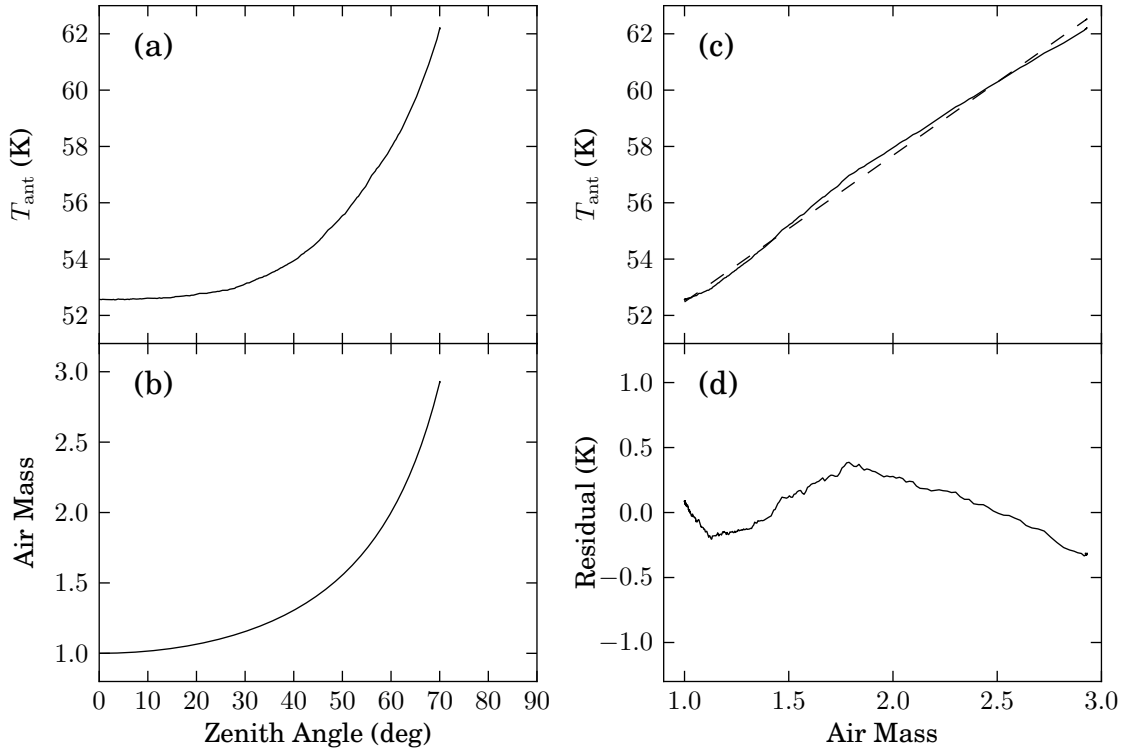

Figure 2.28. Downward portion of a sky dip (27 April, 2011). (a) Calibrated radiometer output as a function of zenith angle. (b) Approximate air mass as a function of zenith angle, computed from equation (2.36). (c) Radiometer output, now plotted against air mass. The dashed line is a linear fit to the data with slope 5.2 K and intercept 47.3 K. (d) Residual after removing the linear fit from the data, showing small nonrandom deviation from the linear model.

230–300 K (e.g, NOAA 2010)), so 270 K is a reasonable estimate of $T_{\rm ATM}$. This gives $\tau_0 \approx 0.02$ as the atmospheric optical depth at zenith, justifying the small-τ approximation in equation (2.39).

The residuals from the fit, shown in panel (d) of figure 2.28, systematically deviate from the linear model—much of this structure is consistent from one sky dip to the next. This most likely results from changes in the ground pickup term as the telescope slews. In particular, the sharp increase near zenith occurs when the telescope feed is looking into the ground, perhaps coupling around the edge of the dish or detecting emission *through* the dish, much of which is perforated rather than solid.

A sky dip procedure is also used to measure the nonlinearity parameter, b, of the receiver, defined by equation (2.23). The calibration diodes produce a constant input noise temperature increment. By measuring the resulting receiver output increment at several background levels, and fitting to the nonlinearity model, b can be estimated. By measuring the calibration diode response at several elevations, the varying atmospheric signal that results from variations in τ with elevation provides this varying background level. Normally the NOISE diode, with a temperature increment of about 70 K, is used for this measurement. The fainter CAL diode signal ($\sim$1 K) is not strong enough for a reliable measurement of the nonlinearity effect for typical background variations.

2.2.2 Observation Procedures

2.2.2.1 Pointing

Pointing the 40 m telescope accurately is a crucial requirement for this observing program. As discussed in section 2.2.2.2, flux densities are measured using an on-off sky switching, which requires that the 40 m telescope be able to reliably position the center of the beam pattern on the target. Although the beam width of $\theta_{\rm FWHM} \approx 2\farcm6$ is relatively broad, even a $15''$ pointing error will result in a 2.5% reduction in gain.

Pointing model. The 40 m telescope is equipped with encoders on the azimuth and elevation shafts. These readings and the readings from the tilt sensors (see section 2.1.1.3) are combined in a pointing model that generates encoder azimuth and zenith angle offsets based on the requested position on the sky. The pointing model has 9 terms for the azimuth angle correction and 5 terms for the zenith angle correction,

$$\begin{aligned} \Delta\phi_{model} &= A_1 \sin\theta + A_2 + A_3 \sin\phi\cos\theta \\ &+ A_4 \cos\phi\cos\theta + A_5 \cos\theta + A_6 \sin\phi\sin\theta \\ &+ A_7 \cos\phi\sin\theta + A_8 \sin(4\theta) + A_9 T_{LR}\cos\theta, \end{aligned} \tag{2.41}$$

$$\Delta\theta_{model} = Z_1 + Z_2\sin\theta - Z_3\cos\phi + Z_4\sin\phi + Z_5 T_{AF}. \tag{2.42}$$

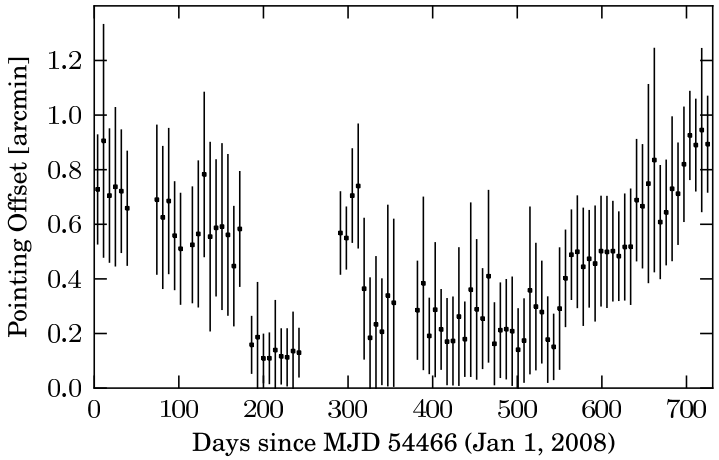

Figure 2.29. Residual error between the pointing model and the actual requested position, plotted in week-long bins for 2008 and 2009. The plotted data and errors are the weekly means and standard deviations of the pointing offsets measured by the pointing calibrations.

Here, ϕ and θ are the requested azimuth and zenith angles, $\Delta\phi_{model}$ and $\Delta\theta_{model}$ are the pointing model corrections for the azimuth and zenith angles, A_i and Z_i are the pointing model coefficients, and T_{AF} and T_{LR} are the aft-forward and left-right tilt meter readings.

The pointing model is determined by performing POINT procedures, described below, at a range of mount coordinates that spans its range. Each POINT produces a measurement of the offset between the encoder readings and the actual position of the telescope beam on the sky. The pointing model coefficients are fit using a least-squares procedure to minimize the error between the model prediction and the measured offsets. The rms residual between the pointing model and the offsets is typically about $15''$ on the sky. This quantifies our ability to point the telescope blindly at a desired position.

We have found that the pointing model terms drift slowly with time. Figure 2.29 shows the residual offset between the pointing model and the actual requested position for 2008 and 2009. The sharp steps in the average offset correspond to adjustments in the pointing model. Before 2010, we adjusted the pointing model two to three times per year to minimize the scatter in the offset and maintain an average offset less than about $0.'5$ to ensure accurate pointing. The data for the fit were collected in a two day period when the telescope was dedicated to measuring the pointing model.

In 2010, we changed our procedure and eliminated the dedicated pointing model measurements. Instead, the POINT procedures executed as part of the monitoring program were collected and the residuals were used to fit pointing model updates as necessary. This eliminates the engineering time overhead of the dedicated measurements and ensures that the model is fit to data collected at mount positions corresponding to those spanned by our program.

Pointing offsets. In addition to the pointing model correction, at least once per hour we measure the pointing offset between a bright pointing calibrator and the model prediction. This measures the effect of wind and thermal loading. In early 2009, we determined that these pointing offsets have the accuracy we require only at separations up to about 30° from the position where the pointing offset was measured. Because of this effect, after MJD 54906 (16 March 2009), care was taken when scheduling to ensure that flux density measurements were always made at separations of less than 15° from the pointing offset measurement. Prior to this, no such limit was in place. We have discarded flux densities measured with a separation of more than 30°.

POINT procedure. Before December 2010, pointing offsets were measured using the POINT procedure, which was implemented in both the VAX and MCS control systems. This procedure operated by performing 3-point cross-scans of the calibrator in both azimuth and zenith angle and fitting a fixed-width Gaussian beam profile to each axis to determine the position of the peak. A pointing offset measurement is considered invalid if its signal-to-noise ratio is less than two, or if the offset indicates that the peak was outside the span of the cross-scan, $\pm\theta_{\mathrm{FWHM}}/2$. When using the POINT procedure, several iterations are normally attempted, moving the cross-scan center by up to $\theta_{\mathrm{FWHM}}/2$ after each attempt. This allows offsets less than the θ_{FWHM} to be measured reliably.

POINT2D procedure. In December 2010, the POINT2D procedure was introduced in the MCS control system as a more efficient procedure for measuring pointing offsets. As suggested by its name, the POINT2D procedure measures the two dimensions of the sky offset at once, rather than first scanning in azimuth and then in elevation. In this procedure, measurements are performed at the expected peak position of the beam and 12 other positions in a pair of hexagons centered on the expected peak at $0.4\,\theta_{\mathrm{FWHM}}$ and $0.7\,\theta_{\mathrm{FWHM}}$ radii. A final point at an azimuth offset about $8'$ from the expected peak is measured to sample the sky background. The pattern of measurement positions is shown in figure 2.30. Using a least-squares fit to the results, the offsets in both directions can be measured simultaneously. This procedure has been found to be more robust and, importantly, faster than the original POINT procedure.

2.2.2.2 Beam Switching and Flux Density Measurements

While Dicke switching does much to reduce the large error terms due to the atmosphere, the ground, and gain fluctuations in the receiver, it does not remove linear drifts in any of these quantities and the situation can be further improved by beam switching. Beam switching in azimuth is optimum because by maintaining a constant elevation we minimize changes to the atmospheric and ground spillover signals and thereby maximize their cancellation. We therefore adopt the same "double switching" technique introduced in Readhead et al. (1989) and also discussed in Myers et al. (1997), Leitch (1998) and Angelakis et al. (2009). In this method, in addition to the dual-beam Dicke switching, we also switch the target between the two beams and hence

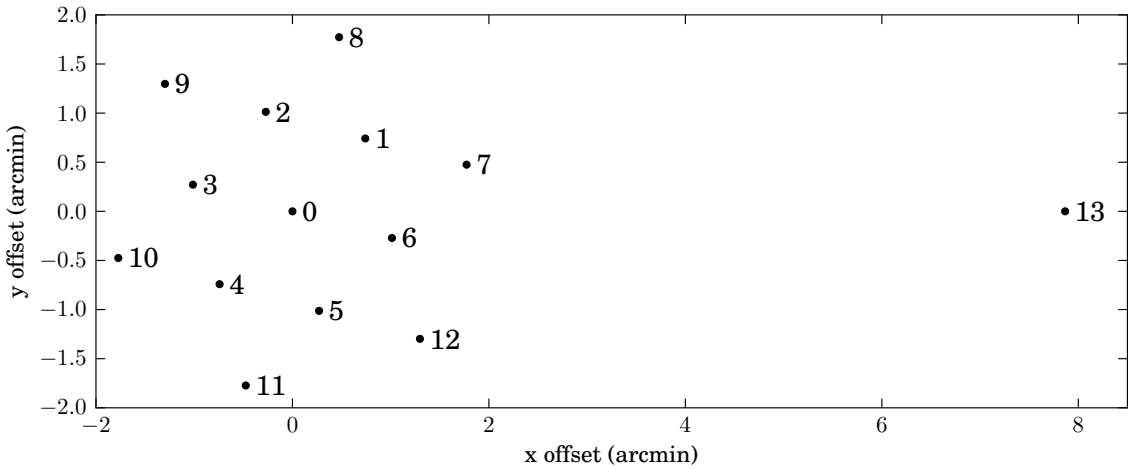

Figure 2.30. Pattern of offsets for a POINT2D procedure. The expected position of the source is at 0.

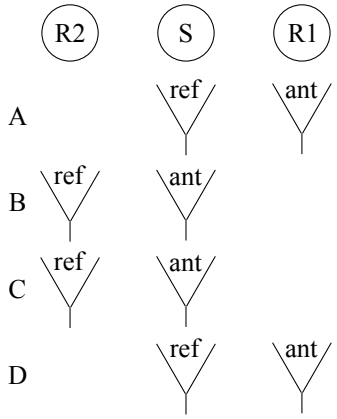

Figure 2.31. Schematic illustration of the FLUX procedure. The source (S) and reference fields (R1 and R2) alternate between the *ant* and *ref* in the four segments of the procedure.

remove both the constant term and any linear drifts in the power from these unwanted components of the signal.

FLUX procedure. The FLUX procedure implements double-switching and is the measurement procedure by which we collect the flux densities for this monitoring program. Figure 2.31 illustrates the procedure schematically. Each FLUX procedure encompasses four Dicke-switched integration periods of length τ, labeled A, B, C, and D. The FLUX measurement is executed as follows. First, the *ref* beam is positioned on the source and integrated, yielding the power difference $\xi_{\rm A}$. Then the *ant* beam is positioned on the source and integrated to produce $\xi_{\rm B}$. Next, $\xi_{\rm C}$ is measured with *ant* still on-source, and finally the *ref* beam is again positioned on the source for a final integration, $\xi_{\rm D}$. Thus we spend a total time of 4τ actually integrating on the source for each flux density measurement. Of course, slewing and settling times have to be allowed for at the beginning of the A, B and D integrations. In this program, $\tau = 8$ s is used and with slewing overhead,

the total time for a FLUX procedure is about 1 min, yielding an on-source efficiency of about 50% for the FLUX procedure.

From the A, B, C, and D integrations, the corresponding flux density is given by

$$S_{15} = \frac{\kappa}{4}(\xi_{\rm B} + \xi_{\rm C} - \xi_{\rm A} - \xi_{\rm D}), \tag{2.43}$$

where κ is the calibration factor required to turn $\rm DU$into Jy, and the $\rm rms$ error is given by

$$\sigma_{15} = \frac{\kappa}{4}\sqrt{{\sigma_A}^2 + {\sigma_B}^2 + {\sigma_C}^2 + {\sigma_D}^2}. \tag{2.44}$$

The calibration factor consists of a relative calibration factor that is computed for each measurement (section 3.2.2.1) and an flux density calibration factor (section 3.2.2.3).

The four integrations also contain interesting information about the stability of the instrument and, more importantly, the atmosphere, during the observations. For each flux density measurement, we therefore also compute two other quantities—one that we call the "switched power," ψ, given by

$$\psi = \frac{\kappa}{4}\left(\xi_{\rm B} + \xi_{\rm D} - \xi_{\rm A} - \xi_{\rm C}\right), \tag{2.45}$$

and the other that we call the "switched difference," μ, given by

$$\mu = \frac{\kappa}{4}\left(\xi_{\rm C} + \xi_{\rm D} - \xi_{\rm A} - \xi_{\rm B}\right). \tag{2.46}$$

Both ψ and μ should be zero in the absence of gain or atmospheric drifts so we use these as a way of estimating such variations in our error model (section 3.2.3) and to reject badly contaminated measurements (section 3.2.1.6). The uncertainties in ψ and μ are clearly given by equation (2.44).

A third combination of the four integrations cancels the contribution of the source field. We call this the "source-nulled flux," and it is defined as

$$S_{\rm null} = \frac{\kappa}{4}\left(\xi_{\rm A} + \xi_{\rm B} + \xi_{\rm C} + \xi_{\rm D}\right). \tag{2.47}$$

In $S_{\rm null}$, the source field is canceled, leaving behind a signal proportional to the difference in brightness between the two reference fields. This property makes it a good detector for astronomical confusion, of which solar interference is an example. We do not use this combination during normal observing or data reduction, but we use it in section 4.4 to detect solar interference.

In practice, the $1\ \rm ms$ samples recorded by the data acquisition system are accumulated into $1\ \rm s$ averages for the *ant* and *ref* beams. The uncertainties ${\sigma_i}^2$ are measured from the sample variances in the $1\ \rm s$ averages rather than from the $1\ \rm ms$ samples directly.

Impact of contamination. Although we will show in section 4.5 that bright radio sources at high galactic latitudes are rare enough that very few source or reference fields will be affected by confusion, it is important to understand the effect of such contaminants. A contaminating source near enough the intended target to appear in the main on-source beam will introduce an artificially high flux density measurement. Fortunately, the average beam for the on-source field (figure 2.6) is nearly circularly symmetric, so if the 40 m pointing is repeatable, such contamination will occur with a constant antenna gain in every observation regardless of the parallactic angle. As a result, although the flux density measurements may be spuriously offset, no or very little artificial variability will be inferred.

Confusion from sources found in the off-source reference beams will reduce the measured flux density for a source. Fortunately, because each reference field is only integrated half as long as the source field, the amplitude of the contamination will be reduced by a factor of two. However, contamination in the reference fields is more likely to result in false variability because the beams are no longer symmetric under parallactic angle rotations. We discuss the effects of confusion further in section 4.5.

2.2.2.3 Calibration Diode Measurements

The output of the CAL and NOISE diodes is measured via the CAL procedure, which is very similar in operation to the FLUX procedure. If a source is being tracked, the telescope first slews to an offset position so the antenna beams are (most likely) positioned on blank sky near the target. A four-integration A, B, C, D procedure is executed, but no slew is required. Instead, during the B and C integrations, the NOISE or CAL diode is enabled. The output of the CAL procedure is then given by

$$S_{\mathrm{CAL}} = \frac{1}{2}\left(\xi_{\mathrm{B}} + \xi_{\mathrm{C}} - \xi_{\mathrm{A}} - \xi_{\mathrm{D}}\right), \tag{2.48}$$

and the rms error by

$$\sigma_{\mathrm{CAL}} = \frac{1}{2}\sqrt{\sigma_A^2 + \sigma_B^2 + \sigma_C^2 + \sigma_D^2}, \tag{2.49}$$

with similar expressions for S_{NOISE} and σ_{NOISE}. Note that the normalization factor is $1/2$ rather than $1/4$ as for the FLUX—this is because the diode is only active for two of the four integrations of the CAL procedure, whereas one beam is on-source for all four integrations in the FLUXPROC procedure.

Because the offset position at which the CAL measurement is executed is near the position of the source being tracked, the background level of the radiometer due to receiver noise, CMB, atmospheric, and ground contamination should be similar to that of a subsequent FLUX procedure. This is convenient because it leads to a similar level of gain compression for a CAL diode measurement as for a FLUX measurement of the source. Thus, unless the source is extraordinarily bright, calibrating a measured source flux density by comparing it to a nearby CAL will implicitly correct for nonlinearity. This is not true for the NOISE diode—the NOISE diode produces so large a signal that it induces a significant change in compression. Neither the CAL diode nor typical program sources are bright enough to induce such changes.

2.2.2.4 Other Observation Procedures

Two other procedures are used for calibration purposes. The ZERO procedure is used to measure the output of the detector when the radiometer is disconnected. This is achieved via an RF switch that can connect the detector diode input to a matched termination instead of the radiometer output. This procedure is used to distinguish the amplitude of the detected signal due to power incident on the radiometer from the offset due to the back-end itself. Such an offset measurement is not required for the differenced output, ξ, since any such offset would be canceled by the subtraction.

The AVERAGE procedure was used in the VAX control system to integrate the switched, ξ, and total power, P, signals from the radiometer without performing sky switching through the FLUX procedure. Because the MCS control system constantly stores the raw output samples, an explicit procedure for this is no longer required.

Chapter 3

Data Reduction Pipeline

In this chapter, we describe the data reduction pipeline that is used to convert the raw data logged by the telescope control system into light curves. This includes the software tools, the data filtering and calibration steps implemented by those tools, and the database system used to store the results of this pipeline for later scientific analysis. A major goal in the the design of the pipeline was to make it as automatic as possible, consistent with producing reliable, high-quality output data. This is important both due to the size of the monitoring program and to avoid opportunities for bias to enter the data set.

The reduction pipeline is organized into three levels of reduction, depicted in figure 3.1. This division is intended to avoid needless computational expense by storing intermediate data products at convenient points in the calibration process so that minor changes to the calibration procedures do not require a complete end-to-end recalibration. Additionally, the pipeline is designed to provide enough logging to ensure that a reduced data set can be reproduced in the future if needed, even if the standard parameters have been changed in the meantime. This also allows multiple reductions with different parameters to be stored in the database independently without requiring redundant copies of the intermediate data products.

The pipeline begins with a log file (VAX control system) or data archive (MCS control system) which contains the raw data stored by the control system. The *low-level* reduction script operates on these data to perform basic calibration steps that are unlikely to be changed often. Essentially, this reduction step consists of obtaining the results of the FLUX, CAL, and POINT procedures and filtering out those procedures that are completely unusable. In the current implementation, the CAL procedures are used to perform relative calibration as described in section 3.2.2.1 below, although this step properly belongs in the high-level calibration. The low-level scripts differ markedly in implementation between the VAX and MCS control systems, but both perform the same steps using essentially the same algorithms. The output of the low-level reduction is written into the reduction database in a common set of tables for both control systems. Further reduction steps are control-system agnostic—this is the only pipeline layer with different implementations for the two control systems.

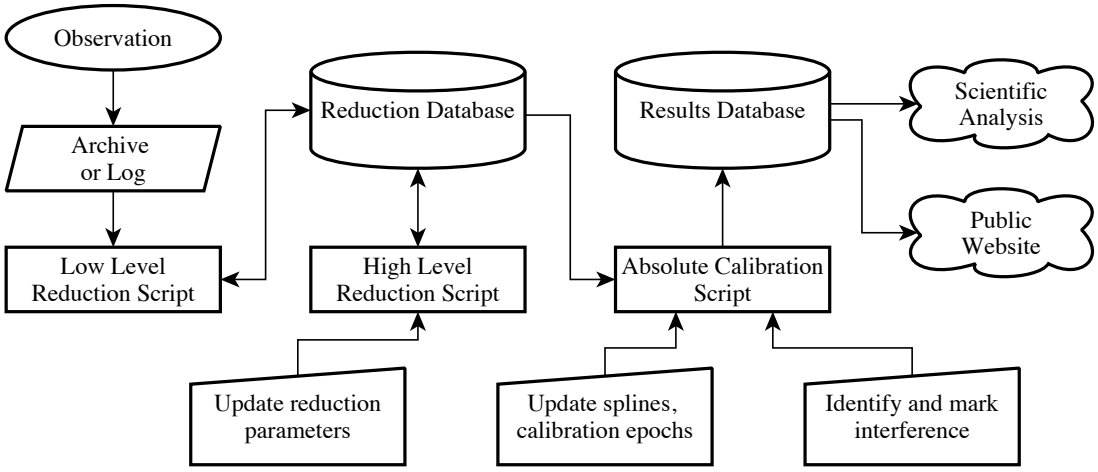

Figure 3.1. Overview of the data reduction pipeline.

The second stage in the pipeline is the *high-level* reduction script. In this stage, most of the flux density calibration steps are applied, including the gain curve and focus correction (section 2.1.2.3), and the error model (section 3.2.3). The output of the high-level reduction stage is written back to the reduction database.

Finally, the *flux density calibration* script convert the results of the high-level reduction into the fully calibrated end product of the pipeline. This script applies the data flags to remove unreliable measurements (section 3.2.1), uses the astronomical calibrator (3C 286) to provide calibration of the flux density scale, and applies the calibrator variation spline (section 3.2.2.2) and nonthermal error correction factor (section 3.2.3.2) to the data. The output of this script is a set of comma-separated-value (CSV) files containing the fully calibrated flux density measurements for each source. These files are then used to update the results database and for scientific analysis.

The storage requirements for a monitoring program of this size present some challenges. The volume of data, while not enormous by modern standards, is sufficiently large that careful organization is needed. This is especially true for a long-term monitoring program like this. Because the data will be reduced on approximately a monthly basis, an initial investment in proper organization of data and code yields a repeated benefit. To achieve this goal, two database systems—one for storing the intermediate data reduction products and metadata describing the reduction process and one to contain the processed data—have been developed.

This chapter is organized as follows. In section 3.1, we introduce the software tools used to implement this pipeline. In section 3.2, we describe in detail the steps that make up the reduction pipeline and discuss their implementation. Finally, in section 3.3, we describe the organization of the database that stores the results of the pipeline and makes them available for analysis and to the outside world. Additional details about the software and database structure may be found in appendixes A and B.

3.1 Software Tools

There is certainly no shortage of numerical software and programming languages suitable for data analysis and reduction. For this project, we have chosen to use the Python[1] programming language. Python is an object-oriented, bytecode-interpreted language that is suitable for either scripted or interactive use. It is supported on a wide variety of operating systems including Linux, Apple OSX, and Microsoft Windows, making it convenient for use in a mixed computing environment. The SciPy[2] and NumPy[3] libraries provide a wide variety of scientific and numerical methods that make Python especially appropriate for scientific computing. High-quality plotting suitable for either interactive use or publication plots is available using the Matplotlib[4] library (Hunter 2007). A wide variety of Python library modules are available, making it convenient to interface with other software. Additionally, Python is increasingly popular for astronomical data analysis, with projects such the National Radio Astronomy Observatory's (NRAO) Common Astronomy Software Applications (CASA) using it to provide a convenient user interface.[5]

Regardless of the choice of programming language and environment, a library of functions to implement the particular data handling and calibration routines is needed. We have implemented most of these in a Python module named *Arcreduce*, described in section 3.1.2.2 below. A set of utility routines needed by the reduction scripts and by Arcreduce, called *py40m*, has also been developed. The py40m module contains routines for interacting with the database, performing the actual calibration computations, and includes a number of convenience routines for, e.g., converting between date formats.

As described in section 2.1.1.1, data in this program were recorded using two different telescope control systems. The VAX and MCS control systems differ greatly in their capabilities, reflecting the enormous increase in computing power and data storage capability available to the newer system. The data interface hardware used for both systems is the same, with radiometer sampling and telescope monitoring data recorded through the Universal Back-End (UBE). In particular, both systems receive radiometer samples at 1 kHz, normally Dicke switched to provide an effective 500 Hz switched sampling rate. The VAX control system, due to its limited bus speed, normally did not record the full stream of samples, instead integrating these in software and storing a summarized result of the procedure to disk for analysis. Between scheduled procedures, no radiometer data were recorded. The MCS control system, on the other hand, permanently records every 1 ms sample, regardless of whether an active observation has been commanded. A human-readable log including on-line integrations for quick debugging is also provided, but the off-line analysis software is responsible for converting the raw samples into procedure results for science data. The MCS control system also stores continuous records of telescope monitoring data, including its pointing, thermometer readings, weather station outputs, etc., whereas the VAX control system only stored relatively low-rate snapshots of these status data.

[1] `http://www.python.org`
[2] `http://scipy.org`
[3] `http://numpy.scipy.org`
[4] `http://matplotlib.sourceforge.net`
[5] `http://casa.nrao.edu/`

The VAX control system stored approximately 1 MB of data per day. The MCS control system stores a bit more than 1 GB per day—more than a three order of magnitude increase. The massive increase in available data has many effects. It enables much finer study of the telescope's behavior for debugging, and is particularly better at identifying rare and unpredictable errors because all the data are available at all times, without the need to enable a debugging mode at the right moment. The increased rate also increases the amount of storage by a factor of 1000. Although advances in storage technology have made this a negligible cost increase, the data management requirements (e.g., arranging to back up the data regularly) are non-trivial. Last, and most relevant to this chapter, the two control systems require very different reduction software. In particular, the VAX control system performed the step of translating radiometer samples into procedure results, whereas the MCS control system leaves this to the reduction software layer.

3.1.1 CMBPROG and the VAX Control System

For the first two years of the program, the telescope was controlled using the VAX control system. In this section, we briefly describe the VAX control system and CMBPROG, the analysis tool used to process the data.

3.1.1.1 VAX Control System

The VAX control system was used to control the 40 m telescope from the early 1990s until August, 2010. A full description of the control system is available in Pearson (1999), here we briefly summarize that description. The telescope hardware was controlled by a single-board DEC MicroVAX *satellite computer* located in the pedestal of the 40 m, called OV40M. OV40M was equipped with analog-to-digital and digital-to-analog interfaces as well as a parallel data bus that enabled it to communicate with and control the telescope hardware. One component of the control software stack, called the *satellite program* or SAT, executed on OV40M. This program included processes for handling servo controls, radiometer data collection, logging, and command and response processing for interacting with the other layers in the software stack. For reliability, OV40M was not equipped with a hard disk. Instead, it booted via Ethernet from a boot image on another MicroVAX, OVCM, which was located in the control building.

In addition to hosting the boot image for the satellite computer, OVCM also executed the two other programs in the control program software stack: the *command and control program* (CCP) and the *terminal interface program* (TIP). The TIP was a terminal-based user interface that allowed commands to be sent to the telescope, schedules to be queued, and the status of the telescope to be monitored. Multiple copies of the TIP could be executed simultaneously, enabling several users to monitor the telescope. The CCP acted as an intermediate layer between the TIP and the SAT programs, processing and serializing commands from multiple TIPs, dispatching responses to the appropriate TIPs, processing schedules, and writing data to the output logs.

A simple, reasonably general scripting language was implemented in the CCP. This included control commands, looping constructs, and subroutine handling. An example observing script is shown in Listings 3.1 and 3.2.

3.1.1.2 Overview of CMBPROG

To analyze the data stored in the log files from the VAX control system, the CMBPROG program was developed by Erik Leitch using a framework written by Martin Shepherd for the DIFMAP package (Leitch 1998; Shepherd 1997). CMBPROG provides a parser to load the data from the log files generated by the CCP, a graphical user interface for plotting and interactively editing the data, and facilities for automated analysis and filtering of the data.

Data within CMBPROG are organized in a series of data structures called *Procedures*, each representing the time series of results of the executions of a particular radiometry or observation-related procedure. Each Procedure contains a set of *Members* containing the actual data recorded by the CCP. For example, the output of the execution of a FLUX measurement is stored in the FLUX Procedure. The actual measured flux densities of a series of flux density measurements is stored in the flux Member, which is referred to as FLUX.flux.

Procedure results can be filtered based on various criteria, e.g., to restrict the results to a particular source. Search criteria are implemented via a data structure called a *ferret*.[6] Several ferrets can exist in parallel, allowing the results for different selection criteria to be handled separately. The ferret also provides a mechanism for reconciling the time bases of different procedures. The results of one procedure may be *referenced* onto the time base of another within the context of a ferret. Several referencing algorithms are implemented. Referencing can be used to associate, e.g., the most recent previous wind measurement with a POINT procedure with that POINT, or the average of all CAL procedure flux measurements within a two-hour window around a FLUX procedure with that FLUX.

A wide variety of data filtering and calibration procedures are provided by CMBPROG. CMBPROG is best suited for interactive use, but does provide a scripting language for batch processing. An excerpt from the data reduction script is given in Listing 3.3. CMBPROG can export processed data in a log file format similar to the log files written by the CCP, or by writing the contents of selected Members to a text file.

3.1.1.3 Retiring CMBPROG

Although CMBPROG is a powerful and useful tool for data reduction, because it consists of a monolithic compiled C program, it is somewhat cumbersome to add features. Also, many of its plotting and numerical computing features are now available as libraries to the general-purpose Python programming language, reducing the need for a custom package. In addition, only a relatively small subset of the functionality provided by CMBPROG is actually used for the data reduction pipeline. For these reasons, we chose not to

[6]The name “ferret” is derived from the phrase “to ferret out,” meaning to search tenaciously for and find something.

```
subfile doflux
!====================================
!      CGRABS Flux measurement Schedule
!      04jan10_0_0.sch
!====================================
  anwait
  weather
  setfocus
  setfocus
  flux samples=8 reps=1
  tag
endsubfile

subfile dopoint
  anwait
  setfocus
  setfocus
  flux samples=8 reps=1
  setfocus
  setfocus
  record tilt
  poff 0.0 0.0
  setfocus
  setfocus
  point samples=10 reps=1
  record tilt
  setfocus
  setfocus
  point samples=10 reps=1
  record tilt
  setfocus
  setfocus
  point samples=10 reps=1
  record offsets
  zero
  weather
  setfocus
  setfocus
  flux samples=8 reps=1
  cal reps=1 diode=cal
  cal reps=1 diode=noise
endsubfile

subfile calib
  anwait
  setfocus
  setfocus
  flux samples=10 reps=1
  setfocus
  setfocus
  record tilt
  poff 0.0 0.0
  record tilt
  setfocus
  setfocus
  point samples=8
  record tilt
  setfocus
  setfocus
  point samples=8
  record tilt
  setfocus
  setfocus
  point samples=8
  record offsets
  setfocus
  setfocus
  ave
  weather
  setfocus
  setfocus
  flux samples=10 reps=1
  cal reps=1 diode=cal
  cal reps=1 diode=noise
  zero
endsubfile
```

Listing 3.1. Example VAX control system schedule. This excerpt defines the compound procedures for performing a FLUX, CAL, or POINT procedure.

```
subfile region_120
obs/nowait J2331-1556
sch/sub/nowait/lst dopoint start=21:34:29 stop=01:25:17
obs/nowait J2336-1451
sch/sub/nowait/lst doflux start=21:34:29 stop=01:40:20
endsubfile

subfile region_119
obs/nowait J0116-1136
sch/sub/nowait/lst dopoint start=22:49:45 stop=03:40:45
obs/nowait J0111-1317
sch/sub/nowait/lst doflux start=22:54:46 stop=03:25:42
obs/nowait J0110-0741
sch/sub/nowait/lst doflux start=22:24:40 stop=04:00:50
obs/nowait J0110-0415
sch/sub/nowait/lst doflux start=22:04:35 stop=04:15:53
obs/nowait J0127-0821
sch/sub/nowait/lst doflux start=22:44:44 stop=04:10:52
obs/nowait J0125-0005
sch/sub/nowait/lst doflux start=22:04:35 stop=04:45:59
obs/nowait J0141-0202
sch/sub/nowait/lst doflux start=22:29:41 stop=04:56:01
obs/nowait J0141-0928
sch/sub/nowait/lst doflux start=23:04:48 stop=04:20:54
obs/nowait J0140-1532
sch/sub/nowait/lst doflux start=23:39:55 stop=03:40:45
obs/nowait J0132-1654
sch/sub/nowait/lst doflux start=23:44:56 stop=03:20:41
endsubfile
```

Listing 3.2. Example VAX control system schedule. Here, the subfiles shown in Listing 3.1 are called to perform the observations for a series of regions. Two such regions are shown here.

```
!! chop low-snr points (using a random flag)
clip point.snr, 0, 2, inc, swp

!! flag failed points
setintd 4800
pointflag flux
useflags -8


!! cut out saturated / abnormal fluxes
!! This is a fairly low threshold -- not truly "saturation", but
!! grabs some periods of anomalously high TP also.
clip flux.atp, 10000, 50000, exc, flux
clip flux.btp, 10000, 50000, exc, flux
clip flux.ctp, 10000, 50000, exc, flux
clip flux.dtp, 10000, 50000, exc, flux

!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!! Now clean up the calibration diode before we use it
!!
@/home/bigjoe/caltech/glast/pipeline/cmbprog/calcuts.cmb

!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!! Now let's apply the calibration diode to put various
!! epochs on equal scales before doing any outlier filters.
!!

!! Create a ferret with cals for all sources
diode CAL
caltag ANT
make cal, allcal

!! do the calibration with the CAL diode
diode CAL
setcalt box, false
setcalb 7200, 2
setcald 0, 0.1
setcaln allcal
! tcal_ant in Jy (divide by 2 since it's only on half the time in a CAL)
telpar "tcal_ant", 8.33 / 2
calib flux, flux

setintt prev, false, true
setintd 3600, 0.1
reference flux.time, setfocus.cz, focus

!! Add some aliases
source ""
make pcor, allpcor
make poffsets, allpoff
make point, allpoint
```

Listing 3.3. Main CMBPROG calibration code, executed after loading the data from the log files and before dumping the data for each source into text files for import into Python.

extend CMBPROG to support data from the MCS control system. Instead, we have reproduced the essential elements of CMBPROG as a Python module and used this as the basis for reducing data after the control system upgrade.

For data recorded under the old control system, however, we continue to use CMBPROG for reduction. Because this is done in the low-level reduction stage, these early data are rarely re-reduced, so the CMBPROG stage is executed only occasionally. When changes that affect the low-level reduction are made, care must be taken to ensure that the CMBPROG scripts are consistent with any changes to the Python scripts that implement the calibration procedures for the newer data (see section 3.1.2).

3.1.1.4 Interface between CMBPROG and Python

Although CMBPROG was not designed to be executed as part of another program, it is possible to execute it from within a Python script by issuing commands over a pipe to its standard input. For this purpose, we implemented a set of routines in `py40m.cmbprog` to execute and control an instance of CMBPROG. Within the low-level reduction script, this interface is used to load the raw data from the telescope log files and execute the CMBPROG calibration scripts. The output for each source is then written to a series of text files and loaded into Python. The data from the text files are then stored in the reduction database for further processing.

3.1.2 Python, Arcreduce, and the MCS Control System

We now turn our attention to the newer MCS control system and the software tools for reducing the data collected with it. We first give an overview of the new control system hardware and software, then describe the Arcreduce Python module that is used to manage those data.

3.1.2.1 MCS Control System

A full description of the MCS control system architecture, focused on the implementation of the drive system, is given in Shepherd (2011a). An online manual for the control system, including a description of the scripting language, user interface, and other details, may be found at Shepherd (2011b). In this section, we give a brief overview of the control system hardware and software.

Like the VAX control system, the MCS control system divides the critical telescope control tasks and user interface tasks between separate networked computers. This architecture is shown in figure 3.2. The *control computer* is responsible for managing the telescope operation, including all time-critical operations. It runs a real-time variant of the Linux operating system and is equipped with analog and digital input/output cards connected to the telescope hardware. The *observing computer* runs an ordinary (not real-time) Linux distribution. Communication with the control computer is managed by the control program, *telcontrol*. The user interface for queuing schedules, executing commands, and monitoring the telescope and radiometer

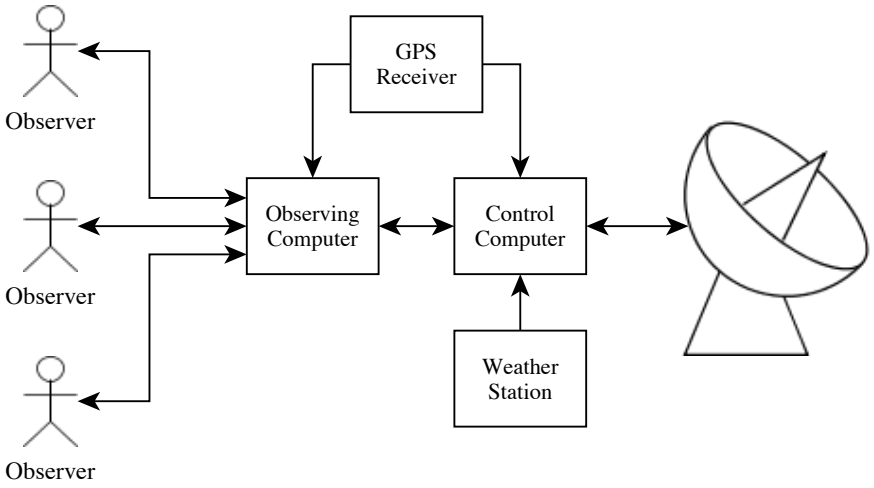

Figure 3.2. Schematic overview of the MCS control system architecture. The global positioning system (GPS) receiver provides a pulse-per-second signal to keep the clocks on the observing and control computers precisely synchronized with UTC.

status is provided by the *telviewer* program. To access the interface, observers connect to the observing computer via a remote desktop connection and execute instances of the telviewer program.

The MCS control system supports a schedule scripting language that is more general than that provided by the VAX control system. The MCS control system provides script-level access to low-level telescope control commands. Functions to implement radiometry procedures are implemented as script-level functions, rather than having hard-coded implementations defined in the control program. Listings 3.4 through 3.6 show a few examples of excerpts from MCS control system scripts.

3.1.2.2 The Arcreduce Python Module

We have implemented a Python module, Arcreduce, which contains the classes and methods needed for reducing MCS control system data. The `ArchiveReader` class is the engine for converting the low-level radiometer samples into procedure results. The replacement pieces of CMBPROG were implemented as the `CalManager` class and its related methods and classes. The Arcreduce module is described in detail in appendix A, here we give a brief overview of the software.

The `ArchiveReader` class provides a flexible engine for processing data from the MCS control system data archives. Figure 3.3 shows a schematic view of the architecture of the `ArchiveReader` engine. The user first instantiates an `ArchiveReader` object to manage the data decoding process, specifying the start and end dates for the reduction. The `ArchiveReader` opens a connection to the data archive using the readarc module.[7] The user then instantiates decoder objects and connects them to the `ArchiveReader`. Each decoder class is a subclass of the `GenericDecoder` class, which defines the interface between the

[7] `http://www.astro.caltech.edu/~mcs/ovro/40m/help/readarc.html`

```
# Arrange for the telescope to be stowed when this schedule ends.
at_end {
  stow
}
# Get the definitions of the standard procedures.
import "$TCS_SCHED_DIR/schedlib.sch"
# Load fermi-specific procedure definitions.
import "$TCS_SCHED_DIR/fermilib.sch"
#-----------------------------------------------------------------------
# Define a set of regions on the sky. Each parameters of each region
# are as follows:
#  String name,          # The name of the region.
#  Double start,         # The earliest LST (hours) to observe this region.
#  Double stop,          # The latest LST (hours) to observe this region.
#  Source psrc,          # The source to peak up the pointing on.
#  Double point_dt,      # The pointing integration time.
#  listof Source srcs    # The list of sources whose fluxes are to be measured.
#-----------------------------------------------------------------------
listof SourceRegion regions = {
  {"region_103", 23:30:29, 00:29:30, J2323-0317, 0:0:10,
   {PKS2320-021, J2335-0131, J2337-0230, J2301-0158, BBJ2247+0000, J2247+0310,
       J2257+0243}
  },
  {"region_120", 23:52:29, 00:52:01, J2331-1556, 0:0:10,
   {J2336-1451}
  }
}

# Start the observations, passing the following function the list of regions
# to be observed, followed by the allowed elevation range.

observe_fermi_sources $regions, 30, 70
```

Listing 3.4. Example MCS control system schedule. This is an example of a program observing schedule. It defines a number of regions and their pointing calibrator sources. In a full schedule, many more regions would be defined.

```
command flux(Count reps, Double integ_dt, Double idle_dt, Double xoff,
             Boolean both) {
  Double flux_a = 0.0
  Double flux_b = 0.0
  Double flux_c = 0.0
  Double flux_d = 0.0
  Double sdev_a = 0.0
  Double sdev_b = 0.0
  Double sdev_c = 0.0
  Double sdev_d = 0.0

  Double dx = $xoff
  TrackingOffsets old = $tracking_offsets()
  Boolean ok = false
  at_end {
    if(!$ok) {
      mark one, f0+f1    # Use feature markers 0 and 1 to indicate failure.
    }
    label "none"
    offset x=$old.x    # Restore the original X offset.
    print "Procedure FLUX ended"
  }
  print "Procedure FLUX starting"
  catch $signaled(source_set) | !$drives_enabled() | $signaled(suspended) |
      $aborted() {
    while($iteration < $reps) {
      label "flux:A"       # Perform the measurement at the +ve off-source
          position.
      flux_integ $integ_dt, $idle_dt, $old.x + $dx, true, true, $flux_a, $sdev_a
      label "flux:B"      # Perform the first on-source measurement.
      flux_integ $integ_dt, $idle_dt, $old.x, true, false, $flux_b, $sdev_b
      label "flux:C"      # Perform a second on-source measurement.
      flux_integ $integ_dt, 0.0, $old.x, false, false, $flux_c, $sdev_c
      # In two-sided mode, toggle the sign of the offset.
      if($both) {
        dx = -$dx
      }
      label "flux:D"      # Perform the second off-source measurement.
      flux_integ $integ_dt, $idle_dt, $old.x + $dx, true, false, $flux_d,
          $sdev_d
      print "Fluxes A=", $format_double(".4f",$flux_a),
                    " B=", $format_double(".4f",$flux_b),
                    " C=", $format_double(".4f",$flux_c),
                    " D=", $format_double(".4f",$flux_d)
      print "Flux (B+C-A-D)/2 = ", $format_double(".4f",($flux_b + $flux_c -
          $flux_a - $flux_d)/2.0), " +/- ", $format_double(".4g",$sqrt($sdev_a*
          $sdev_a + $sdev_b*$sdev_b + $sdev_c*$sdev_c + $sdev_d*$sdev_d))
      mark one, f0      # Indicate that this iteration of the procedure was
          successful.
    }
    ok = true    # Indicate that the procedure ran to completion.
  }
}
```

Listing 3.5. MCS control system schedule for performing a FLUX procedure.

```
#-----------------------------------------------------------------------
# This command is part of the flux() command. It moves the telescope
# to a given absolute X offset, and if this takes less than
# $idle_dt, waits until $idle_dt hours have passed. It then performs
# an integration of $integ_dt hours.
#-----------------------------------------------------------------------
command flux_integ(Double integ_dt, Double idle_dt, Double xoff, Boolean move,
                   Boolean refocus, Double flux, Double sdev) {
  # Change the tracking offset?
  if($move) {
    offset x=$xoff
    # Wait for longest of the acquisition time and the idle time.
    until $acquired(source) & $elapsed >= $idle_dt
  }
  if($refocus) {
    focus
    until $acquired(focus)
  }
  # Integrate for the specified length of time.
  request_flux $receiver, switched, $integ_dt
  until $acquired(flux) | $elapsed >= $integ_dt+0:0:5

  # Get the returned flux.
  flux = $requested_flux.flux
  sdev = $requested_flux.sdev
}
```

Listing 3.6. MCS control system definition of the `flux_integ()` routine used by the FLUX procedure. In the VAX control system, functions at this low level were defined within the control program, rather than implemented through schedule scripts.

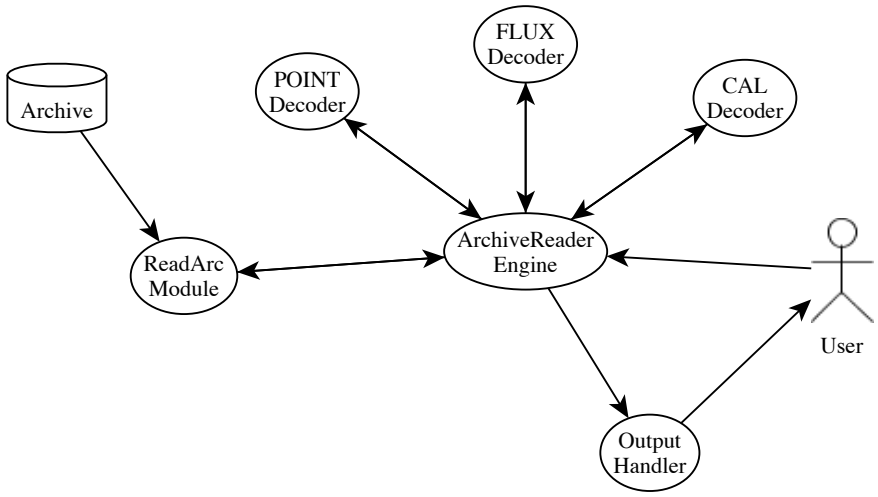

Figure 3.3. Architecture of the `ArchiveReader` engine. The user or reduction script instantiates and interacts with an `ArchiveReader` object. This object uses the readarc module to obtain raw data from the archive. These data are passed frame-by-frame to a set of decoder objects, each of which is responsible for identifying and processing a single type of radiometry procedure. When the decoder finishes decoding a procedure, it passes the result to the `ArchiveReader`, which notifies one or more output handler objects of the result.

decoder and the `ArchiveReader`. The class hierarchy for the decoder objects is shown in figure 3.4. Finally, the user instantiates and connects one or more output handler objects to the `ArchiveReader`.

Data in the archive are organized into named registers, which are stored in data frames, one frame per second. For data collected at a higher rate, such as the radiometer samples (1 kHz) or the mount encoder readings (10 Hz), array registers allow multiple data points to be stored in a single frame. The `ArchiveReader` must specify to the readarc module which registers are to be read from the archive. By default, it provides the real-time clock time stamp for each frame, the current source name register, and the frame label (a free-form string register that is used to identify the currently active radiometry procedure). When a decoder is first connected, it passes a list of the registers for which it requires data to the `ArchiveReader`, which then configures the readarc module to read those registers if they are not already active.

Once the `ArchiveReader`, decoders, and output handlers have been set up, the user or reduction script repeatedly calls the `handle_frame()` method until it indicates that the end of the requested date period (or the end of the data archive) has been reached. For each frame, which corresponds to 1 s of data, the `ArchiveReader` reads the data from the archive using readarc, then notifies each decoder object of the new data. The decoder objects each implement a state machine to identify and process the data for a radiometry procedure. When the end of a procedure is encountered, the decoder returns the results of the procedure to the `ArchiveReader` object, which in turn passes this result to any connected output handlers.

Although several output handlers are implemented, the `CallbackOutputHandler` is normally used during data reduction. When a procedure is completed, this output handler calls a user-specified function

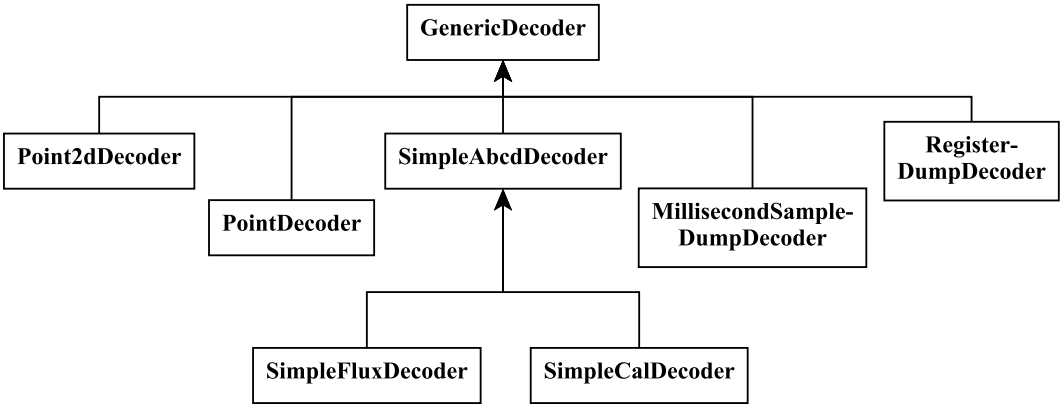

Figure 3.4. Class hierarchy for procedure decoders in the ArchiveReader module. The `SimpleAbcdDecoder` provides generic decoding of radiometry procedures that execute an A-B-C-D double-switching pattern. The `RegisterDumpDecoder` and `MillisecondSampleDumpDecoder` are "phony" decoders—these are used for dumping raw 1 s or 1 ms samples to a file, rather than for decoding procedures.

with the the results of the procedure. This is normally used with the `CalManager` class, which provides a callback function that stores the results of each procedure in the appropriate `Procedure` object, described in section A.1.

3.1.2.3 Reducing Radiometry Procedure Data

The MCS control system provides access to a much greater amount of observation data than did the VAX control system, but this additional detail comes at a cost: decoding the results of an observation requires a much more sophisticated analysis. Whereas the VAX control system software processed and reported the results of each radiometry procedure, requiring only parsing the output logs to determine the results, the MCS control system stores the raw radiometer samples, which must be processed. This greater complexity is undoubtedly worthwhile, however, as it permits more careful scrutiny of the telescope performance, and enables improvements in the procedure reduction algorithm to be applied retroactively because all the data, not just the summarized final result, are available.

The `ArchiveReader` class supports decoding three radiometry procedures: FLUX, CAL, and POINT. The first two procedures are very similar, so these are implemented using a single set of logic with very minor adjustments. The POINT procedure is quite different, and must be handled in both the original and "2D" implementations, as described in section 2.2.2. The decoders for these procedures currently use the 1 s averaged radiometer data available in the data archive rather than working with the full-rate millisecond samples.

FLUX procedure. As described in section 2.2.2, the FLUX procedure consists of four integrations, labeled ξ_{A}, ξ_{B}, ξ_{C}, and ξ_{D}. The control system labels the 1 s data frames with the strings "FLUX:A" through

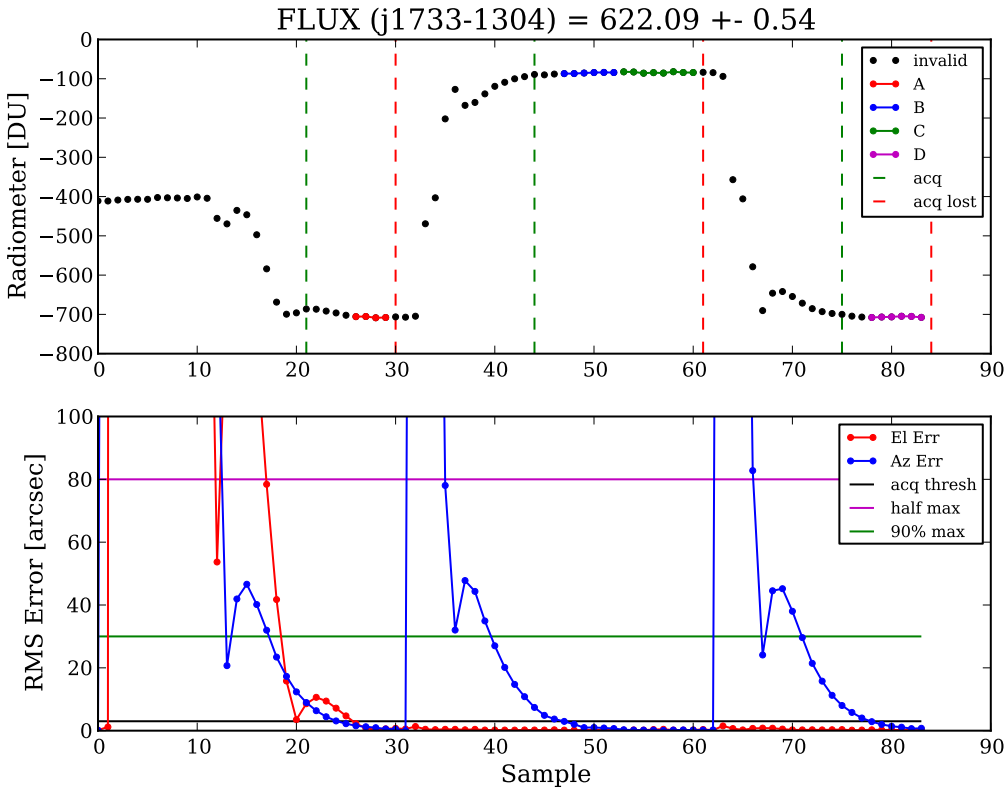

Figure 3.5. Illustration of the FLUX procedure decoder. *Top:* Dots show the uncalibrated 1 s average difference between the *ant* and *ref* inputs. The valid samples that are integrated for the A, B, C, and D segments of the procedure are indicated. The vertical dashed lines indicate where the control system signaled that the source being observed had been acquired or when acquisition was lost. *Bottom:* Red and blue data show the elevation and azimuth pointing error. The horizontal lines indicate the offset from the nominal Gaussian beam profile corresponding to the half power and 90% power points, and the acquisition threshold used by the software decoder.

"FLUX:D" to enable the decoder to identify the starts and ends of FLUX procedures as the archive is processed. The *tracker* thread within the control system produces an *acquired* flag that indicates when the tracking servo has first acquired the target source. Relying solely on this flag is not an adequate measure to reject samples where tracking errors are significant. In the example shown in figure 3.5, just after the acquisition signal is raised (indicated by the first vertical line in the upper plot), an elevation tracking glitch occurs, resulting in a clear change in the radiometer output. To ensure that only data collected with accurate pointing are included in the integrations, the decoder identifies the last continuous run of samples with tracking errors less than $3''$ in each segment and integrates only these samples. Using the selected samples in each segment, the mean and standard deviation are computed, and these are combined to produce the FLUX result.

The control system records an estimate of the result of the FLUX procedure in a human-readable log file as the data are collected. This is useful for debugging, but because it uses a less cautious decoder, its results

are somewhat less accurate than those from the `ArchiveReader` decoder. A comparison between the results in the log file and those obtained from `ArchiveReader` finds that the flux density measurements reported in the log file are systematically lower than those computed by `ArchiveReader`. The median difference is about $0.9 \times \sigma_{15}$, a rather large disagreement. This is due to the inclusion of less accurately pointed data in the online integrations.

CAL procedure. The CAL procedure is nearly identical to the FLUX procedure so the same decoder described above is used for the CAL procedure as well. Because the CAL procedure requires a slew only before the A segment (when it slews off-source), the tests for accurate tracking do not typically discard data except during the A segment.

POINT procedure. The decoding requirements for the POINT procedure are different than for the FLUX and CAL procedures. In addition to the obvious difference in the data collected, the result of a POINT procedure is required in real time—the pointing offset must be computed as the data are taken so that it can be used to point the telescope accurately. The control system performs this calculation in real time and, following a successful POINT, indicates the result by moving the telescope to the measured offset and raising a flag. If a POINT fails, this flag is not raised, thus indicating the failure to the decoder. During later analysis, the `ArchiveReader` decoder has no need to repeat the fitting procedure. Instead, the result indicated by the control system is simply adopted.

The `ArchiveReader` decoder does, nonetheless, integrate the samples within the segments of the POINT (or POINT2D) procedure. An example of the samples collected during a POINT2D procedure is shown in figure 3.6. Using the integrated averages and estimated uncertainties together with the fitted position indicated by the control system, the decoder fits a fixed-beamwidth Gaussian beam profile to the results to determine the flux density of the source and to compute the signal-to-noise ratio. These results are used to filter bad pointings out of the data stream.

3.2 Reducing the Data

We now turn to the specific procedures used for reducing the data. These steps are implemented using the software tools—CMBPROG and ArcReduce—described in the previous section. In this section, we begin with the reduced procedure results, assuming that ArcReduce has already processed the raw samples for the MCS control system data. In this section, we organize the calibration steps into groups by function, first discussing data editing and filtering, then flux density calibration, and finally error estimation. In the pipeline, for convenience these functions are not necessarily performed in that order, but the end result is the same as if they were.

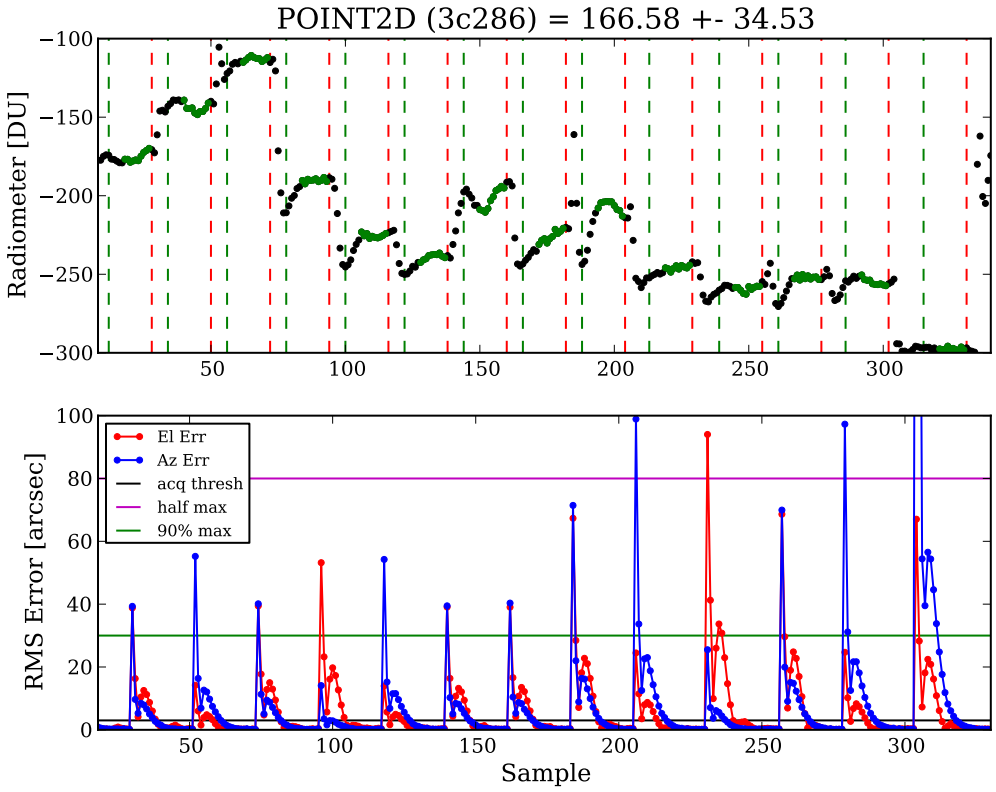

Figure 3.6. Illustration of the POINT2D procedure. *Top:* Dots show the uncalibrated 1 s average difference between the *ant* and *ref* inputs with the valid samples used for each integration indicated. Vertical lines indicate acquisition and loss of acquisition at the start and end of each segment. The segments are arranged in order as numbered in figure 2.30. *Bottom:* Azimuth and elevation tracking errors during the procedure.

3.2.1 Data Editing and Filtering

We first describe the data editing and filtering used to remove unreliable data. Parts of this procedure are performed in the low-level reduction scripts, but most filters are applied as part of the final flux density calibration procedure. This is done so that the suspect data remain in the database for later examination to evaluate whether the filters being applied are working as expected. Data filtered out early in the calibration process are inconvenient to evaluate later, so this is only done for data so obviously corrupt that access to them is unlikely to be needed often. This includes data lost due to pointing offset measurement and calibration diode reading failures.

Editing and removal of corrupted data is performed using both automated and manual filters. Where possible, we have used automated filters both to avoid laborious examination of data and to eliminate opportunities for bias to enter the pipeline. Care is taken to avoid flagging data based on criteria that could induce bias. For example, no flagging based on the flux density of a measurement is implemented.

Data to be edited or filtered are marked with *flags*. The flux density calibration script accepts a list of flags to either require or reject for each data point. Some of the flags require parameters—these are specified as reduction parameters which are stored as described in section 3.3.4 . In each section, we specify the flag name and the names of any parameters that affect it. The database tables that define the flags are described in section 3.3.4 and a full list of the flags is given in section B.2.

3.2.1.1 Date Interval Cuts

A list of date intervals for which all data should be ignored is maintained. This is the mechanism by which manual data editing is implemented. In addition, telescope calibration and maintenance periods, observations other than the monitoring program, and any other periods during which the data should not be included in our results are indicated with an entry in the list. The reduction database permits multiple sets of flagged date intervals to be stored simultaneously. The date interval cuts are enabled by excluding the `dateflag` flag and the set of flagged intervals is specified using the `DATEFLAG_SET_NAME` parameter.

In a few cases, anomalous calibrator source readings are dropped using a date interval cut. This was done frequently with 3C 161, which sometimes experiences anomalous pointing failures (see section 4.1.3 for further discussion). This sort of data editing is only used when there is very strong evidence for a pointing failure or other data contamination. Most frequently these edits are used for calibration sources that are expected to exhibit stable flux densities.

3.2.1.2 Wind, Sun, Moon, and Zenith Angle Cuts

Under high winds there is a systematic reduction in observed flux densities due to mispointing and poor tracking. Observations made when the instantaneous wind speed exceeds 15 mph ($6.7\ \mathrm{m\,s^{-1}}$) are discarded. A FLUX procedure is rejected if the wind speed exceeded the threshold either during that FLUX procedure

or during the POINT procedure that was used to obtain the pointing offset for that FLUX. To protect the telescope a "wind watchdog" program stows the telescope to a safe position pointing near zenith. For more details, see section 2.1.1.4. If wind data are not available, due to a weather station or logging failure, data are discarded.[8] Under the VAX control system, wind measurements were recorded with each POINT procedure, or at least about once per hour. Under the MCS control system, the weather station outputs are logged every second, but only one reading per minute is used by the reduction pipeline for wind speed flagging. To reject data flagged due to high wind, the `wind` flag must be enabled and the wind threshold in mph must be specified through the `WIND_THRESH` reduction parameter.

Observations at zenith angles less than 20° are discarded because the telescope is unable to track fast enough in azimuth to match the sidereal rate near zenith. The scheduling algorithm avoids scheduling sources for observation at these zenith angles, so few observations are lost. Observations at solar or lunar elongations less than 10° are also discarded. The scheduler does not avoid these areas of the sky so a small number of observations are lost. The zenith angle cut is enabled using the `za_limit` flag with a upper and lower zenith angle thresholds set with the `ZA_LIMIT_MIN` and `ZA_LIMIT_MAX` reduction parameters, which are normally set to 20° and 90°, respectively. The solar and lunar elongation limits are enabled via the `sun_angle` and `moon_angle` flags, with the thresholds set by the `SUN_ANGLE_THRESH` and `MOON_ANGLE_THRESH` parameters.

3.2.1.3 Pointing and Calibration Failures

An observation is rejected if a pointing offset was not obtained within the prior 4800 s, or if the pointing offset measurement immediately preceding the observation failed. FLUX procedures affected by these pointing failures are dropped during the low-level calibration process. Occasional scheduling errors resulted in observations without adequately measured pointing offsets. These observations are discarded. These flags are enabled by requiring the flag and by excluding the `pointing_model_only` and `pointtest` flags. During 2009 and 2010, FLUX observations of the pointing calibrators were performed just prior to measuring the pointing offset. These were used to evaluate the pointing performance and are excluded from our normal data set. These "point test" FLUX observations are identified by the presence of a POINT on the same source as a FLUX, within a short time period specified by the `POINT_TEST_MAX_DELAY` parameter. Finally, the `pdist` flag will identify FLUX procedures for removal if they were observed more than `PDIST_THRESH` degrees on the sky away from the preceding POINT measurement.

An observation is rejected if fewer than two reliable calibration procedures using the CAL diode were successfully executed within a two-hour interval centered on the time of the observation, or if the difference between the largest and smallest CAL diode measurement within that interval differ by more than 10%. FLUX procedures affected by this are dropped during the low-level calibration process.

[8] In the data published with Richards et al. (2011), data were accepted when wind data were not available. This change has caused a few data points to be rejected that were accepted in that published data set.

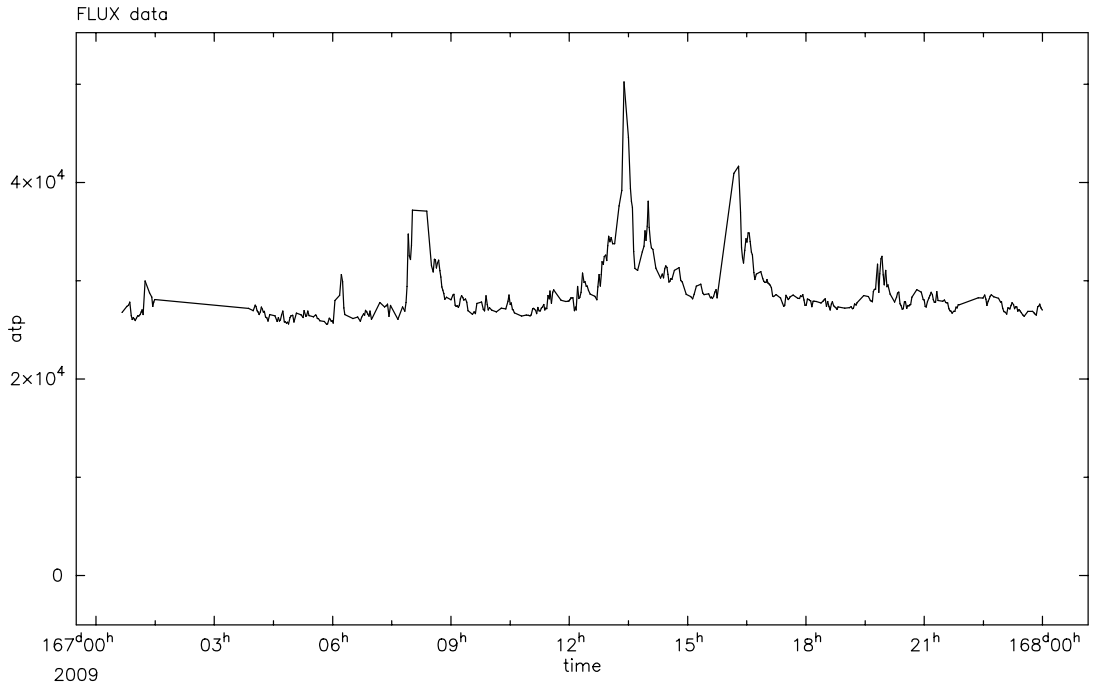

Figure 3.7. Average total power during the A segments of the FLUX procedures recorded on 16 June 2009. The strong spikes resulted from interference from inclement weather.

3.2.1.4 Saturation and Total Power Anomalies

The total power varies depending on the attenuator setting, receiver gain fluctuations, atmospheric conditions, and the observed zenith angle. Observations that indicate saturation or other total power anomalies are rejected. Heavy cloud cover or precipitation often causes large fluctuations in total power. Such periods are identified by inspection of the total power time series and manually discarded. In figure 3.7, the average P for flux procedures executed during a thunderstorm is shown, illustrating the sort of behavior that is selected for removal. These manual flags are implemented as date interval flags as described above.

3.2.1.5 Measured Uncertainty Cuts

We reject flux density measurements with anomalously large measured uncertainties, σ_{15}, defined by equation (2.44). However, a straightforward cut at a fixed value or a fixed multiple of the expected thermal uncertainty introduces a bias against larger flux densities. This occurs because there are contributions to the measured error that are proportional to the flux density of the target radio source, such as telescope tracking errors. We therefore apply a flux density-dependent threshold and discard flux density measurements for which

$$\sigma_{15} > \zeta\sqrt{1 + (\rho \cdot S_{15})^2}. \tag{3.1}$$

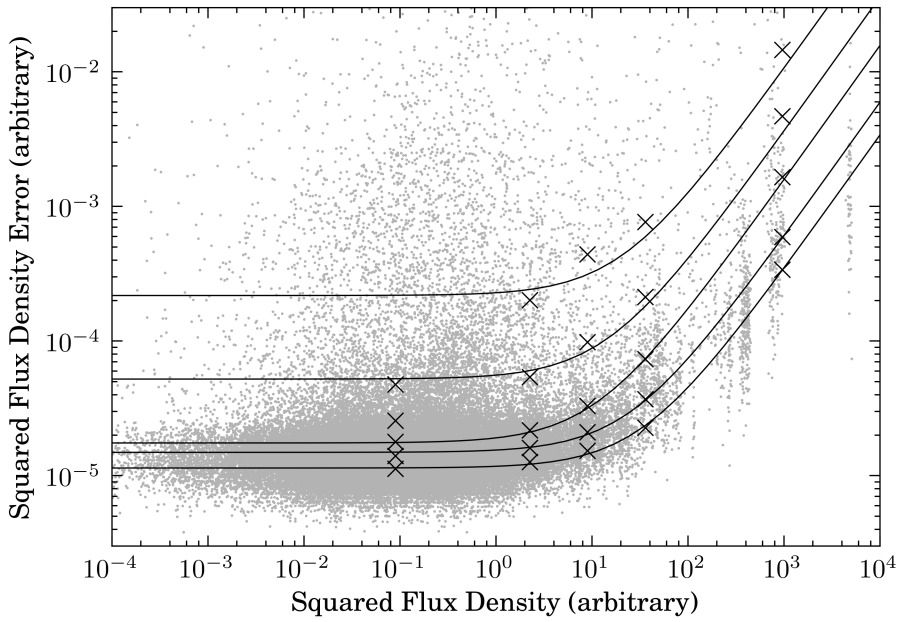

Figure 3.8. Logarithmic plot of σ_{15}^2 versus S_{15}^2 for the data used to fit ρ in the measured uncertainty filter. Grey dots are the individual data, black crosses are the 25th, 50th, 75th, 90th, and 95th percentile values in the bins indicated in table 3.1, and the solid lines are linear fits to the black crosses for each percentile. The flux density values are plotted in approximate Jy units.

Table 3.1. Flux density bins used for fitting the measured uncertainty filter parameters

Bin (Jy)	Center (Jy)	# Data	Threshold (Jy)	Percent Rejected
0.2–0.4	0.3	19498	0.0208	98.4
1–2	1.5	6771	0.0217	97.1
2–4	3.0	2574	0.0242	96.3
4–8	6.0	1230	0.0324	95.9
22–40	31	230	0130	97.0

Note: The threshold is computed at the center of each bin for this study.

To determine the optimal values of ζ and ρ, we examined the data collected between March and October, 2009. In figure 3.8, we plot the square of the measured uncertainty for each flux density as a function of the square of that flux density in grey. The black crosses indicate the 25th, 50th, 75th, 90th, and 95th percentile of the data in each of the five flux density bins listed in table 3.1. The solid lines show linear fits to the data. Although there is peculiar behavior at low flux densities, the high-flux density behavior is similar in all the fits. We used the 50th percentile (i.e., median) fit to determine the optimal value $\rho = 0.200$. Finally, $\zeta = 0.0208$ Jy was set to discard as many obviously bad flux density measurements as possible while maintaining selectivity. In table 3.1, we tabulate the percentage of data rejected in each bin, which is reasonably insensitive to the flux density of the bin. Overall, about 2% of the data are eliminated by this filter.

3.2.1.6 Switched Difference Cuts

We also use the switched difference μ, defined by equation (2.46), to determine whether flux density measurements might be contaminated by systematic errors. The expected value of μ is 0, provided that the ground spillover and atmospheric noise in the *ant* and *ref* beams are identical. Pointing and tracking errors again give flux density-dependent contributions to μ, so to avoid bias against brighter radio sources we flag points where

$$\left|\frac{\mu}{\sigma_{15}}\right| > \beta \cdot \frac{(\mu_0 + \rho_s \cdot S_{15})}{\sqrt{1 + (\rho_t \cdot S_{15})^2}}. \tag{3.2}$$

The optimum values of the parameters ($\beta = 5$, $\mu_0 = 1.148$, $\rho_s = 0.0682$, and $\rho_t = 0.0243$) are determined from the data. Figure 3.9 shows the data that were used to determine the parameters. The switched difference values for FLUX procedures between October 2008 and October 2009 are plotted in grey. These were binned into eight bins (in Jy, 0–0.1, 0.1–0.3, 0.3–1, 1–3, 3–10, 10–25, 25–40, 60–80; there were no data between 40 and 60 Jy) and the median value in each bin was computed; these are shown by the black points. The parameters in equation (3.2) were fitted to the binned data with $\beta = 1$, giving the other parameters. The value $\beta = 5$ was chosen to drop as many suspect data points as possible while maintaining the selectivity of the filter.

This filter discards about 2% of flux densities. Because the tracking performance of the telescope should not change, this procedure is expected to give consistent results across epochs for a fixed set of parameters. This expectation appears to be correct—approximately equal fractions of the data are dropped from each epoch. Among bright ($S_{15} \geq 10$ Jy) sources in the 2008–2009 period, 1.4% of data were flagged by this filter. This is comparable to the 1.7% overall flagging rate for the same period, which indicates that we have successfully removed the bias against bright flux densities.

3.2.2 Flux Density Calibration

The data recorded by the control system is measured in the arbitrary units that come from the 16-bit analog-to-digital converter, referred to as DU. In this section, we describe the procedure for determining the calibration factor to convert DU into physical units. This calibration factor changes with time due to, e.g., fluctuations in the receiver gain, intentional changes in the programmable attenuator setting, or inadvertent but unavoidable changes in the loss or mismatch in the signal path between the antenna feed and the receiver when components are disassembled and reassembled for maintenance. Our fundamental calibration strategy is to use the stable CAL diode output to account for short-term variations within the receiver, then periodically determining the ratio between the CAL diode level and an astronomical calibrator (3C 286) to account for longer-term variations, including those outside the receiver.

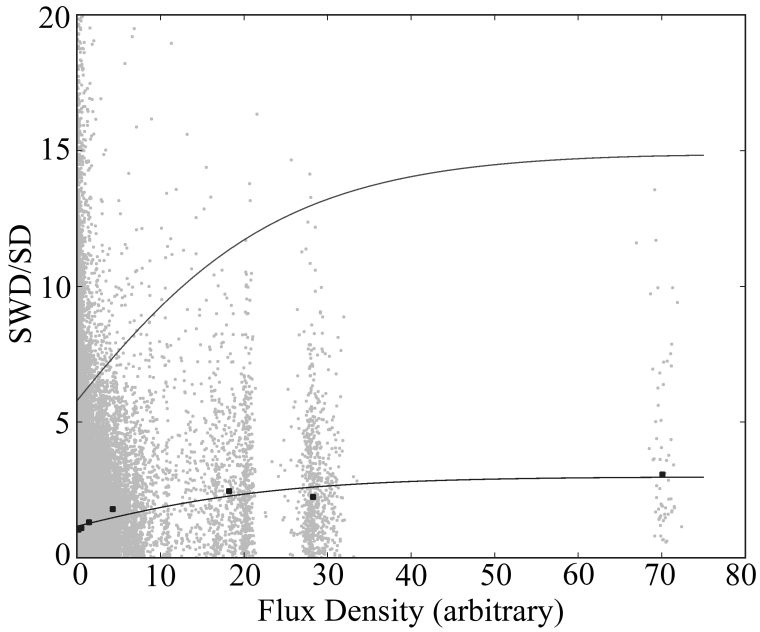

Figure 3.9. Plot of the switched difference, μ normalized by σ_{15}, versus S_{15} for FLUX procedures between October 2008 and October 2009. Grey points are the individual data points, black points are the binned medians. The lower black line is a fit of equation (3.2) to the black data points with $\beta = 1$. The upper black line corresponds to the same parameters except with $\beta = 5$, corresponding to the adopted threshold. The flux density values are plotted in approximate Jy units.

3.2.2.1 Relative Calibration

To correct for slow gain fluctuations of the receiver, we first divide each flux density measurement by a calibration factor measured using the small noise diode CAL. A measurement of the strength of the CAL diode is made after each pointing observation, and no less than once per hour. Gain fluctuations that affect the Dicke-switched data are rather slow, so the calibration factor is averaged over a two-hour window, centered on the time of the flux density measurement. If there are fewer than two good measurements of the strength of the CAL diode in that window or if the measurements disagree by more than 10%, then the flux density observation is discarded.

Due to gravitational deformation of the telescope structure and the increase in atmospheric attenuation with airmass, the effective antenna gain varies substantially with zenith angle. We model this variation with a polynomial gain curve and scale flux density measurements to remove the effect, as described in section 2.1.2.3. Day-to-day changes in atmospheric opacity are found to vary with $<1\%$ rms, so these variations are accounted for within our error model and no correction is applied.

Additionally, the optimal axial focus position varies with zenith angle, as well as solar zenith angle and elongation. During observations, the focus position is set using a polynomial model of the zenith angle variation, and a correction is applied during calibration using a more complete model that accounts for solar zenith angle and elongation. This procedure is described in more detail in section 2.1.2.3.

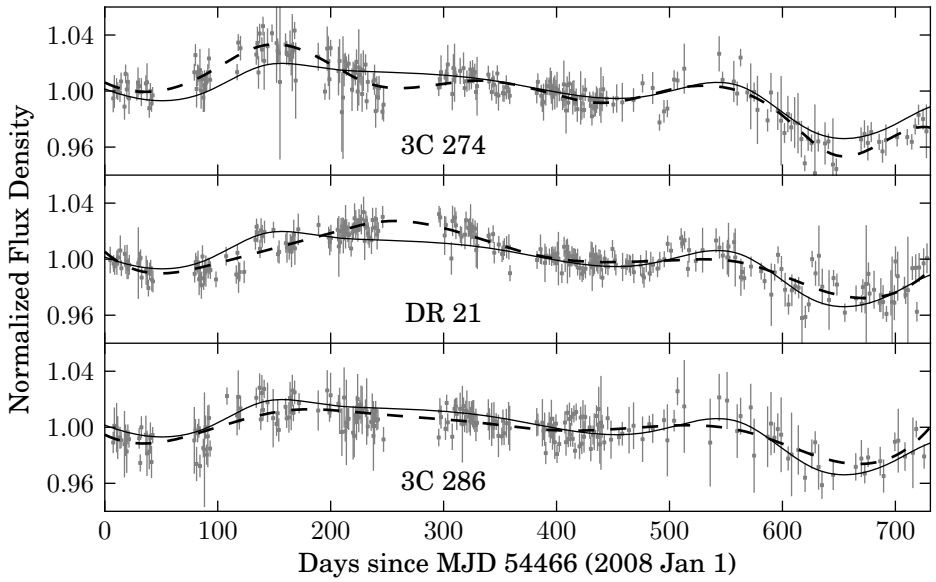

Figure 3.10. Normalized flux densities for 3C 274 (top), DR 21 (center), and 3C 286 (bottom) after outlier removal. Each light curve is normalized by its median. The solid line in each plot is the spline fit to the combined data. The dashed line is a similar spline fit to the data for the individual source.

The combined effect of these corrections is a factor, κ_{rel}, that is computed for each flux density measurement. This is one component of the overall κ introduced in equation (2.43).

3.2.2.2 Long-Term Trends in 3C 286, 3C 274, and DR 21

After carrying out the above editing and calibration steps we returned to the residual 1%–2% long-term (about 6-month) variations in the light curves for stable-flux-density calibration sources. We chose 3C 286, 3C 274, and DR 21 for this study because they are well-known to be stable on timescales of many years. The fractional variations in flux density of these objects are shown in figure 3.10 and are clearly correlated, indicating the presence of an unidentified source of multiplicative systematic error. For each of these sources, we removed 2σ outliers in a 100-day sliding window and normalized the resulting data by the median flux density. We then combined the data for all three sources and fitted a cubic spline to the result.

We apply the corresponding correction to all light curves in our program by dividing each flux density by the value of this spline. Figure 3.11 shows the residuals for the three fitted sources after dividing out the spline fit. The 1% residual variation that remains is the level of systematic uncertainty after correction for this long-term trend.

Each time we perform a reduction, it is necessary to check whether these calibration splines need to be updated. This is a bit challenging because it can be difficult to determine whether a few points above or below the mean flux density for a stable source merely represent a few statistical deviations or actually indicate a new trend. We normally assume the former, only extending the spline fits when there is clear evidence for a correlated trend among the calibrator sources. We extend the spline curves by refitting a cubic spline to

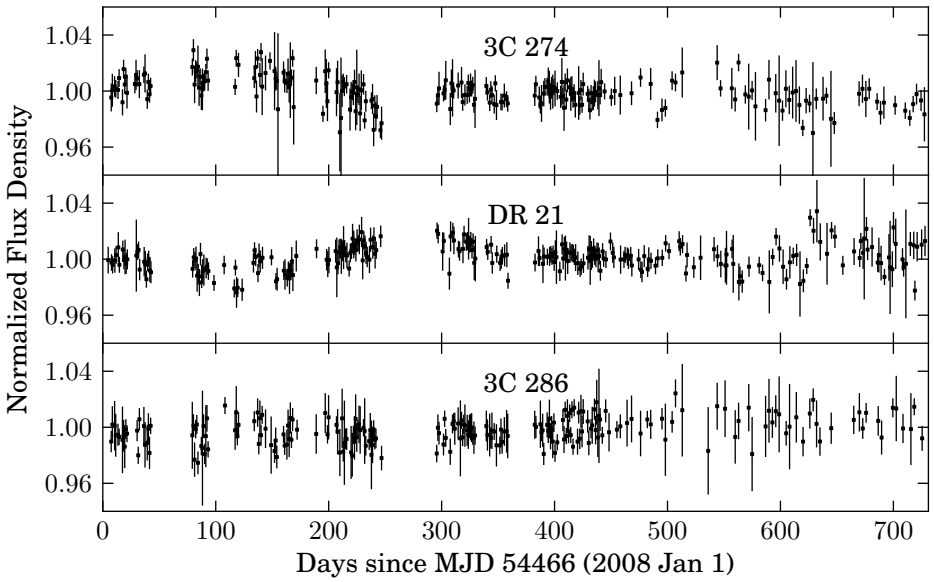

Figure 3.11. Normalized flux densities for 3C 274 (top), DR 21 (center), and 3C 286 (bottom) after dividing by the spline fit to remove long-term systematic trends.

Table 3.2. Calibrator spline epochs

Epoch	Start		End	
	MJD	Date	MJD	Date
1	54466	01 Jan 2008	55197	01 Jan 2010
2	55197	01 Jan 2010	55624	04 Mar 2011
3	55624	04 Mar 2011	55728	16 Jun 2011

the entire light curve for the calibrators, adjusting the number of knots in the spline to match the end of the previously determined spline as well as possible. To avoid changing the values of old data, however, we continue to use the old spline for the older data. The boundaries between spline fits are stored using a set of *calibrator spline epochs* that are updated as necessary. These epochs are tabulated in table 3.2.

3.2.2.3 Flux Density Calibration

We divide our observation period into *flux density calibration epochs* characterized by a consistent ratio between the calibration diode and feed horn inputs to the receiver. This ratio might change if, for example, the signal path is disconnected and reconnected for maintenance, resulting in a slight change in loss along one path. Within a single epoch, the ratio of the CAL diode signal to a stable astronomical source should therefore be constant. Table 3.3 lists the epochs we have used in our analysis. Flux density calibration[9] is applied to each epoch separately.

[9]N.B., the script that implements the flux density calibration steps is named the "absolute calibration" script, although we do not actually perform absolute calibration.

Table 3.3. Flux density calibration epochs

Epoch	Start		End	
	MJD	Date	MJD	Date
1	54466	01 Jan 2008	54753	14 Oct 2008
2	54753	14 Oct 2008	54762	23 Oct 2008
3	54762	23 Oct 2008	55420	12 Aug 2010
4	55420	12 Aug 2010	55482	13 Oct 2010
5	55482	13 Oct 2010	55511	11 Nov 2010
6	55511	11 Nov 2010	55543	13 Dec 2010
7	55543	13 Dec 2010	—	—

For each epoch, a calibration factor is determined from regular observations of the primary calibrator, 3C 286. We adopt the spectral model and coefficients from Baars et al. (1977). At our 15 GHz center frequency, this yields 3.44 Jy, with a quoted absolute uncertainty of about 5%. The calibration factor for epoch i, κ_i, is the ratio of the adopted flux density for the calibrator to the weighted mean of the observations:

$$\kappa_i = \frac{3.44 \text{ Jy}}{\left(\sum S'_{15} \cdot {\sigma'_{15}}^{-2}\right) / \left(\sum {\sigma'_{15}}^{-2}\right)}, \tag{3.3}$$

where S'_{15} and σ'_{15} denote the flux densities for the calibrator with only the relative calibration applied.

The total calibration factor for a flux density measurement in equation (2.43) is then $\kappa = \kappa_{rel} \cdot \kappa_i$, and reflects both relative and flux density calibration. Comparing our calibrated flux densities for 3C 48, 3C 161, and DR 21 with the Baars et al. (1977) values, we find a scale error of $(-0.8 \pm 4.1)\%$. This is probably a conservative estimate (i.e., an overestimate) of the scale error because some of this disagreement may result from variation in the sources. After correcting for different flux density scales, cross-checks of our calibration against 14.6 GHz observations of a number of common sources observed with the Effelsberg 100 m telescope through the F-GAMMA project confirm the overall accuracy of our flux density scale.

3.2.3 Uncertainties in Individual Flux Density Measurements

In a perfect observing system with no sources of systematic error the uncertainties in the flux density measurements would be given by the thermal noise on each observation. In practice there are many sources of systematic error, including the effects of weather and the atmosphere, mispointing due to wind, and focus errors. Many of these are correctly identified and accounted for in the automatic and manual editing and calibration described in the preceding sections. However, even after flux density measurements affected by these problems are filtered out, there remain many observations that are significantly affected by systematic errors. Such systematic errors can lead to significant errors in the measurement that are not reflected in the thermal noise of the observation and can give rise to bad flux density measurements with small thermal errors. This leads to *outliers* in the light curves, i.e., points which do not lie close to the level determined from interpolation of adjacent observations and which have small errors. The task of identifying and eliminating or

Table 3.4. Error model parameters

Parameter	Pointing Calibrator	Normal Source	
		Early	Late
ϵ	0.0057	0.0200	0.0135
η	3.173	3.173	3.173

Note: The "early" and "late" periods are before and after 16 March 2009 (MJD 54906), respectively.

allowing for the wide variety of systematic errors leading to such outliers is challenging and time-consuming. Great care must be taken not to assume that the behavior of the source is known, and hence to eliminate a real and potentially extremely interesting flux density variation.

3.2.3.1 Error Model

We first apply an error model to determine the uncertainty of each flux density measurement:

$$\sigma^2_{\rm total} = \sigma^2_{15} + (\epsilon \cdot S_{15})^2 + (\eta \cdot \psi)^2 \ , \tag{3.4}$$

which is an extension of the model described in Angelakis et al. (2009). The first term represents the measured scatter during the flux density measurement. This includes thermal noise, rapid atmospheric fluctuations, and other random errors. The second term adds an uncertainty proportional to the flux density of the source. This term allows for pointing and tracking errors, variations in atmospheric opacity, and other effects that have a multiplicative effect on the measured flux density. In the third term ψ is the switched power, defined by equation (2.45). This term takes account of systematic effects that cause the A-B segment of the flux density measurement to differ from the C-D segment, such as a pointing offset between the A and D segments, or some rapidly varying weather conditions.

The error model is defined by the two parameters, ϵ and η, whose values must be determined from the observations. Because ϵ describes the error contribution due to pointing errors, its value depends on whether a source is used as a pointing calibrator. Furthermore, for non-pointing calibrator sources, ϵ is found to differ between the scheduling algorithms used before and after 16 March 2009 (MJD 54906). The parameter, η, is found to be adequately described by a single value for all sources and all epochs. The adopted values are given in table 3.4.

For pointing calibrator sources, both ϵ and η were estimated simultaneously using the stable flux calibrators 3C 286, 3C 48, 3C 161, and DR 21. Due to systematic errors, these sources and other stable-flux density calibrators show long-term variations of 1%–2% so we fitted a 7th-order polynomial to remove this trend from each source, then computed the residual standard deviation, median flux density, the rms, and mean ψ for each source, then used these to fit the error model parameters.

To determine the error model parameter ϵ for ordinary sources, we selected 100 sources that exhibited little variation or slow, low-amplitude variations in flux density, between the start of our program and 05 Au-

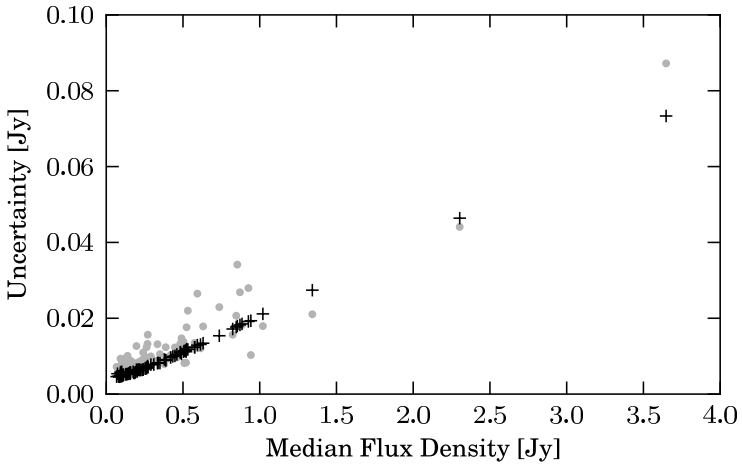

Figure 3.12. Residual standard deviation (grey points) and fitted ϵ-only error model values (black crosses) for ordinary sources in the early (MJD $<$ 54906) period. The fit in the late period is similar. A single high-flux density data point was omitted to limit the scale.

gust 2009 (MJD 55048). This interval was split into two periods, "early" and "late," at MJD 54906 and this procedure was separately applied to each period. For each light curve, we fitted and removed a second-order polynomial trend, then iteratively removed outlier data points with residuals greater than three standard deviations. We repeated the fitting and outlier removal until no further outliers were removed and we discarded any source with fewer than 10 remaining data points (retaining 94 and 88 sources in the early and late periods, respectively). From the surviving points in each light curve, we computed the median and the rms flux densities, and the standard deviation of the residuals. We then fitted equation (3.4) to these data, omitting the η term. The data and the error model results for the early period are shown in figure 3.12. We then adopted the same value of η for these sources as was determined for pointing calibrator sources.

3.2.3.2 Scaling of the Nonthermal Error

The uncertainty model we have introduced combines two qualitatively different components. The first component is that directly obtained during the flux measurement, σ_{15}, which represents random errors such as thermal noise and rapid atmospheric fluctuations. The second, quantified by the ϵ and η parameters, is introduced to take into account other, flux-density–dependent effects. Thus far, we have assumed these to be source independent. However, many sources exhibit coherent long-term variations with random scatter about those that is clearly smaller than what would be expected as a result of the quoted errors. This indicates that the assumption of source-independent ϵ and η is invalid and has resulted in an overestimation of the actual uncertainty in some cases. We now describe a correction that is applied to the errors to reduce this effect. The method for computing correction factors described here was developed and implemented by Walter Max-Moerbeck based on ideas and requirements that we developed together.

To correct these constant scale factors on a source-by-source basis, we begin by fitting a cubic spline to the light curve, adjusting the number of knots to achieve a residual χ^2 per degree of freedom near to one. Due to the large number of sources and the requirement of an uniform and consistent method, an automated method was developed for this procedure. We begin by removing any extreme outliers using a low-order cubic spline fit, rejecting the data in the top 5% of the absolute residuals. Next, for each number N_k in a reasonable range (typically 1 to 80), we fit a cubic spline with N_k knots to the remaining data.[10] Not all the fits are acceptable—some cases have obviously correlated residuals or a large departure from normality. Acceptable fits are selected using two statistical tests: Lilliefors' test for normality (Lilliefors 1967) and the runs test for randomness (e.g., Wall & Jenkins 2003).[11] Only the fits for which both null hypotheses cannot be rejected at the 10^{-3} level are considered acceptable. For each acceptable fit, a scale factor that makes the χ^2 per degree of freedom equal to one is calculated. Among the scale factors for all the acceptable fits, the median scale factor is selected as the final correction. The value of the scale factor is not very sensitive to the exact number of polynomial sections. A typical example of the behavior of the scale factor is shown in figure 3.13.

Using the correction factor, we rescale only the nonthermal part of the errors (the S_{15} and ψ terms in equation (3.4)). This correction is only applied to those sources for which the resulting correction factor is smaller than one (i.e., the rescaling would result in *smaller* errors). The latter choice was made for two reasons. First, a correction factor larger than one simply indicates that the spline fit cannot provide an adequate description of the data. This may result from a light curve more variable than can be fit by spline with a given number of knots, so such a correction could mask real variability. Only the reverse is cause for concern—when the spline fit is too good a fit, given the quoted errors. Second, this choice ensures a smooth transition between scaled and nonscaled errors, as the transition point (correction factor equal to one) is equivalent to no error scaling.

The error scale factors were first computed only for the CGRaBS sample using the two-year data set, then were computed for all sources using the 42-month data set. For the CGRaBS sources, the error scale factors changed somewhat between the two data sets. In figure 3.14, the histogram of the differences between the 42-month error scale factor and the two-year error scale factor for each source is plotted for the sources that had a scale factor less than one (i.e., had a correction applied) in either interval. The mean change is an increase by 0.134.

[10] We use the MATLAB Spline Toolbox function spap2, which automatically selects the positions of the knots for the spline.

[11] We have used the implementation of both tests that are part of the MATLAB Statistics Toolbox.

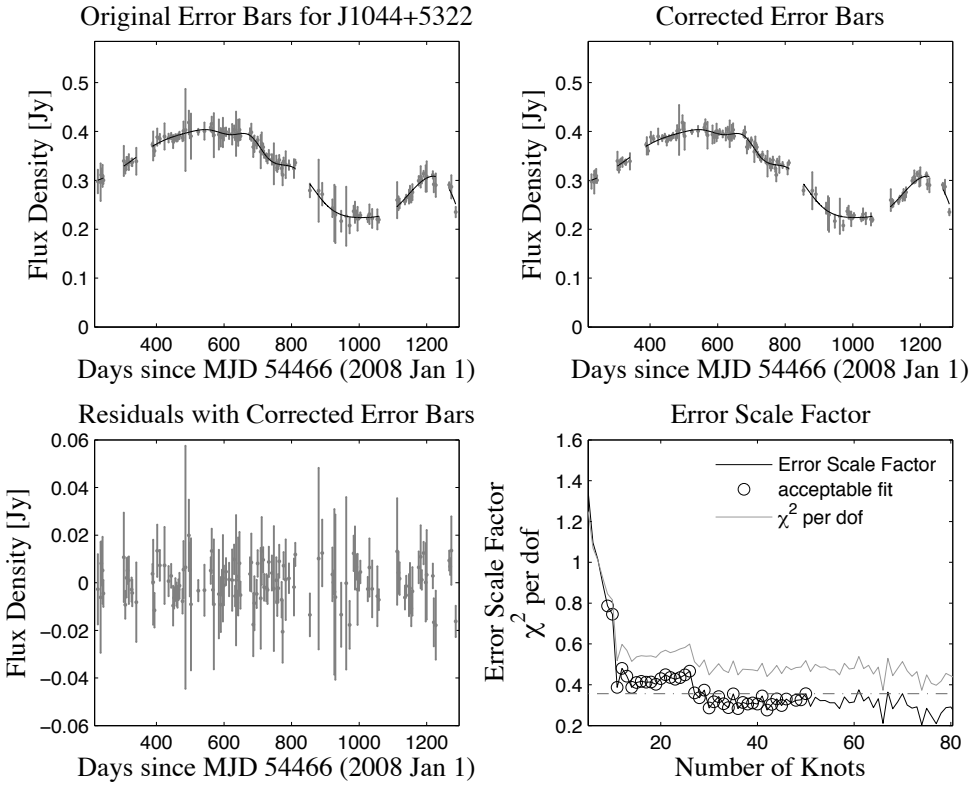

Figure 3.13. Example of the error scale factor correction using data for J1044+5322. The two upper panels show the light curve with the original (left) and corrected (right) error bars (grey points) and a typical spline fit (black line). The bottom left panel shows the residuals from the spline fit using the corrected error bars. In the bottom right panel, the χ^2 per degrees of freedom (solid grey line) and correction factor (solid black line) are shown, with black circles marking the correction factors for fits that pass the acceptance tests, and a dashed line showing the adopted correction factor for the source (0.356).

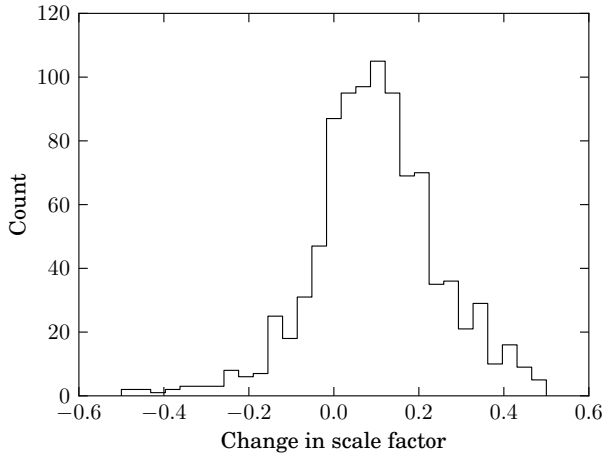

Figure 3.14. Histogram of the change in error scale factors for CGRaBS computed using the 42-month versus the two-month data sets. Only sources for which one or both error scale factors were less than one are included.

3.3 Storing and Retrieving the Data

We now turn our attention to the method by which the data for the program are stored. This is accomplished using a MySQL[12] database. Two such databases are used by this program: the *reduction* database, which is used by the data reduction pipeline, and the *results* database, which is used to store the fully reduced data. In this section, we will primarily discuss the reduction database. We begin this section by discussing the design principles that guided the design of the database schema, then describe the most important components of the schema. Finally, we will briefly describe the results database.

In this section, we present diagrams of several of the major parts of the reduction database. Diagrams of the remaining tables in the database are found in appendix B. Tables listing the contents of the domain tables are also found in the appendix.

3.3.1 Database Design Principles

A schema is, informally, the set of constraints that define the tables that store the data within the database. MySQL implements a relational database (at least approximately—there are some technical differences between an SQL database and a true relational database), so each table within the database represents a relation. That is, each table represents true values of some true-or-false statement, called a predicate. For example, our `Source` table is defined by the predicate, "An astronomical source with right ascension α, declination δ, redshift z, etc., is part of the monitoring program." The schema, then, is a technical definition of the predicates that define the tables in the database. The relational database was a groundbreaking development—it represented the first database to be rigorously defined and analyzed in mathematical terms. However, in this section, we will restrict our discussion to informal terms, and in some cases describe behavior that is specific to MySQL. In most cases, other SQL databases will have similar behavior.

3.3.1.1 Normal Forms

Although there is great freedom to design database schema, a well-known set of conventions, or *normal forms*, has been developed. By organizing a schema according to the normal forms, redundancies and inefficiencies in the database structure are avoided. We have consciously attempted to apply the first three normal forms to the reduction database schema.

The *first normal form*, or 1NF, requires that the columns in a table not contain any repeating groups. For example, because astronomical sources are frequently known by many names or numbers according to different catalogs, it is tempting to define a table of sources with columns, "Name 1," "Name 2," etc. This violates 1NF because each of these columns represents the same sort of information. To comply with this, we must restrict ourselves to nonrepeating data columns. We could comply with this normal form in a source table by storing multiple copies of each source row, one for each name.

[12] `http://www.mysql.com/`

This is obviously inefficient—the data for each source is repeated many times, both requiring extra space and extra effort if any of those data need to be updated. The *second normal form*, or 2NF, addresses this. It requires that rows within a table not contain this sort of redundant information. Thus, rather than repeating the data for each source for each of its names, we must instead create a separate table of source names. Each entry in this table should contain a single source name and a pointer to the row in the source table to which that name applies. The pointer is implemented using a *key* (or *candidate key*) in the source table—a key is a subset of the data in the row that uniquely identifies that row in the table. Each table must have one key defined to be the *primary key*.

Although in principle a key can be constructed from the actual data in the table, in our database, we use the convention of assigning each row a unique integer identifier known as a *surrogate key*. These keys are identified by a "_ID" suffix in the column name. In our example, the reference between the source name table and the source table would then be implemented by storing the surrogate key from the source table in source name with each name for that source. This reference to a key in another table is called a *foreign key*.

The *third normal form*, or 3NF, requires that all fields within a row must have a functional dependence on the key for that row. That is, all the data within the row must actually be tied to the entity described by that row. In section 3.3.2 below, we give an example of the application of 3NF. Several other normal forms exist, but a full discussion is beyond the scope of this section. We have not explicitly attempted to normal forms beyond 3NF to our database schema. In a few cases, we have consciously broken even the first normal form. This has been done to simplify the database schema. For a database of modest size, like ours, this is generally acceptable.

3.3.1.2 Table Indexes

Searching for an entry in an unsorted table is an $\mathcal{O}(N)$ operation—a table with N entries requires time proportional to N. When executing a query that joins several large tables in the database, the search space is the Cartesian product of those tables, with a total number of rows equal to the product of the number in each table. As a result, N can easily grow to be astronomically large, resulting in extremely slow response from the database. Modern databases like MySQL attempt to organize queries to avoid such enormous searches, but in many cases the database designer must provide additional assistance by defining *indexes* on the tables.

An index is a precomputed ordering of the rows in a table using the data values in that row. In MySQL, searching a table using an index is often a constant-time ($\mathcal{O}(1)$) operation, compared to the linear-time ($\mathcal{O}(N)$) full-table search. The response time to queries can be improved enormously by appropriately defining indexes on large tables. Adding an index to a table does incur some costs, however. While an index speeds up data retrieval, it requires that the index be recomputed whenever data are added to the table. Thus inserting data into the database is generally slowed by each index. Additionally, each index increases the disk space used by the database table.

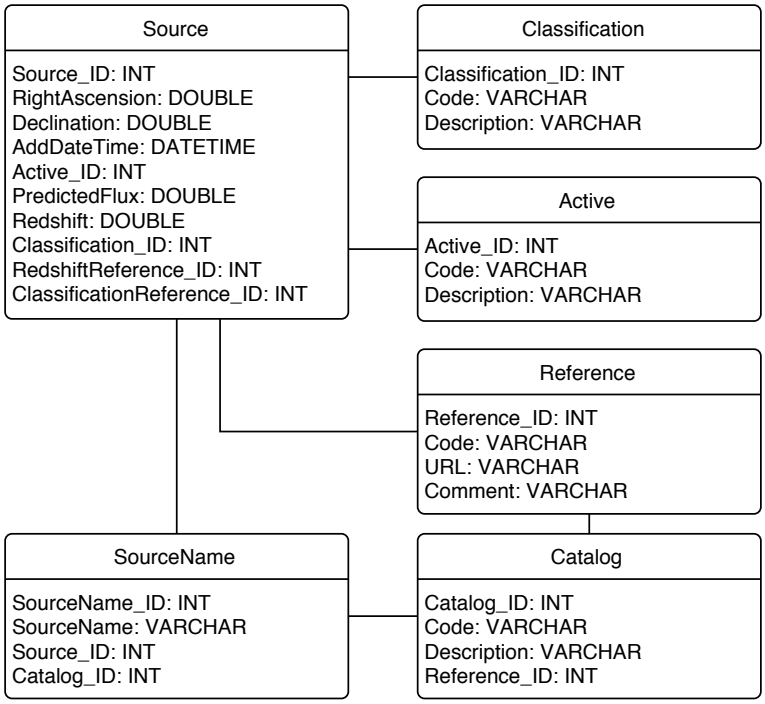

Figure 3.15. Diagram illustrating the database tables used for storing source data.

An index is defined by an ordered list of columns from the table. Indexes can be defined as *unique*, which will prevent more than one row with the same value in each index column from being added to the table. Every table has at least one unique index—the primary key. In the reduction and results databases, we have used indexes both to improve the speed of data retrieval and to enforce uniqueness constraints, for example to prevent accidentally inserting multiple copies of a FLUX result.

3.3.2 Representing Sources

In figure 3.15, a diagram illustrating the database tables relevant to storing source names, coordinates, and other data is shown. Each box in the diagram represents a table in the database with the name given in the header. The entries below the header give the names and data types of the columns in that table. The first column listed for each table is the surrogate primary key for that table. Lines between boxes indicate a foreign key reference between the two tables, which can be identified by a column in the referring table with the same name as the primary key of the table to which it refers.

Rows in the `Source` table represent astronomical sources. For each source, the J2000 right ascension and declination are stored in degrees. The `PredictedFlux` column contains the recent typical 15 GHz flux density measured for the source, which is used by the scheduler to choose suitably bright pointing calibrator sources. The redshift and optical classifications are also stored, as is the date at which the source was added to the database.

In keeping with the 1NF and 2NF rules described above, names for the sources are stored in a separate `SourceName` table which associates one or more names with each source. The `Catalog` table defines the names of catalogs of sources, and a unique index prevents more than one source name within a single catalog to be created (although different catalogs can use overlapping source names).

The `Active` and `Classification` tables (and also the `Catalog` and `Reference` tables) are examples of *domain tables*. These define the domain of values that are valid for entries in a column in another table. In this case, the `Active` table defines options for specifying whether a source is actively being observed as part of the monitoring program, and the `Classification` table defines the range of classifications that can be assigned to a source. Use of a domain table rather than, say, a free-form `VARCHAR` string field, ensures that entries in a table use a consistent set of values.

The `Reference` table stores a list of literature references that are used to track references for the source catalogs as well as for the redshift and classifications stored for each source. The `Reference` table stores a short name for the reference, plus a comment and a full reference to the cited material (the `URL` column). This is an example of applying 3NF: we could have opted to store both the short name for the reference and the full reference with each redshift in the `Source` table. However, including the full reference would violate 3NF because that reference is a property of the cited paper, not of the astronomical source. By separating these data into the Reference table, we are able to update the `URL` field for all the references simply by changing one value in one row of the `Reference` table rather than replacing every instance in the `Source` table. In addition, we can reuse the `Reference` table entries to provide values for the source classifications in the `Source` table and for the `Catalog` table.

There is a minor violation of 2NF in the `Source` table: the `ClassificationReference` and `RedshiftReference` columns are technically repeated instances of the same type of information. These should properly be stored through an additional table linking `Source` table entries to the `Reference` table with a reference type (“redshift” or “classification”) for each link. However, we decided that the overhead of adding this additional table was not warranted for the minor benefit of strict 2NF compliance.

3.3.3 Storing Observation Data

In figure 3.16, a diagram showing the tables relevant to storing observation data from the monitoring program is shown. The `Reduction` table has foreign key references to tables not shown here. It is discussed further in section 3.3.4.

3.3.3.1 FLUX and POINT Procedures

Two tables, `PointProcedure` and partially calibrated `FluxProcedure` store the POINT and FLUX procedures that result from the low-level reduction script. The data stored in these tables are largely based on the data that were reported for each procedure by the VAX control system. Some additional data, such as the

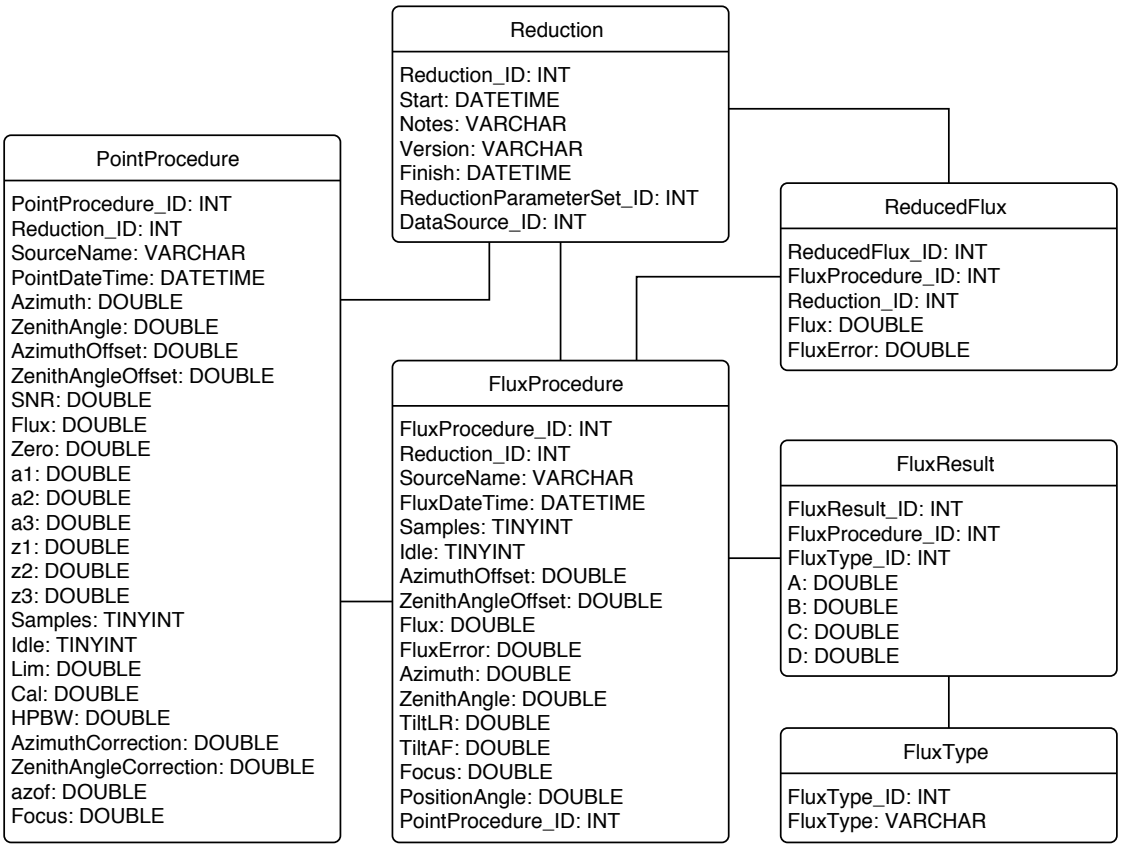

Figure 3.16. Diagram illustrating the database tables used for storing observation data.

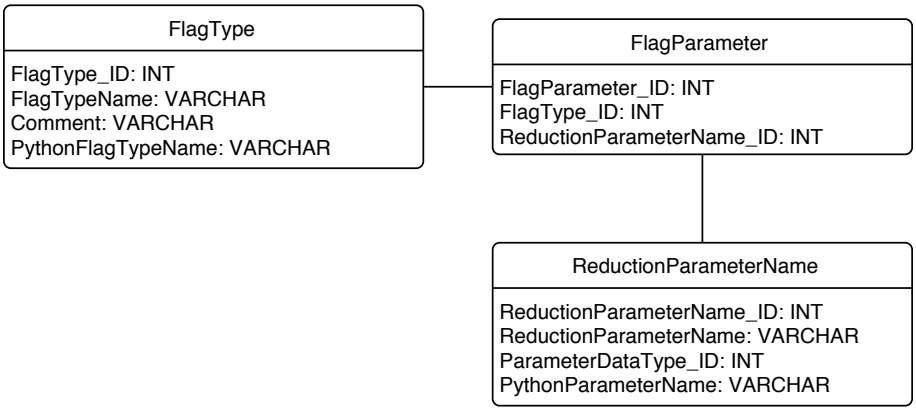

Figure 3.17. Diagram illustrating the database tables used for storing flag types and parameter definitions.

tilt meter readings and focus position, are stored in these tables for convenience. Each `FluxProcedure` entry refers to a `PointProcedure` entry to indicate which POINT procedure was used to compute the pointing offset in effect during the FLUX. For FLUX procedures executed with no measured pointing offset, this reference is null.

The `SourceName` entry in these tables contains whatever source name appeared in the log file or archive as the name of the target for that observation procedure. It is *not* a reference to a source name in the `SourceName` field. Resolving which program source is referred to by this name is part of the flux density calibration script.

Each FLUX procedure reports two sets of integrations for the A, B, C, and D segments. One is the switched average for each segment, the other is the total power. In keeping with 1NF and 2NF, these are stored as separate entries in the `FluxResult` table to avoid repeated encoding of the same type of data in the `FluxProcedure` table.

The `ReducedFlux` table stores the results of the high-level calibration script for each FLUX. Because there is a one-to-one correspondence between `FluxProcedure` and `ReducedFlux` rows, the latter refers back to the former rather than than store a duplicate of the the data from that table. Only the `Flux` and `FluxError` values need to be stored (and the `Reduction_ID` field will refer to a different `Reduction`).

3.3.3.2 Flag Tables

Diagrams of the tables that contain the data editing flag types and their parameters are shown in figure 3.17. The parameter values are actually stored in the `ReductionParameter` table described in section 3.3.4. The `FlagParameter` table connects each flag with the correct `ReductionParameterName`.

Another set of tables are used to store the date intervals to be flagged. Diagrams of these tables are shown in figure 3.18. Each entry in the `FlaggedDate` table describes one interval to be flagged. A comment is stored with each entry for logging purposes, and each is assigned a `FlaggedDateType` that classifies the

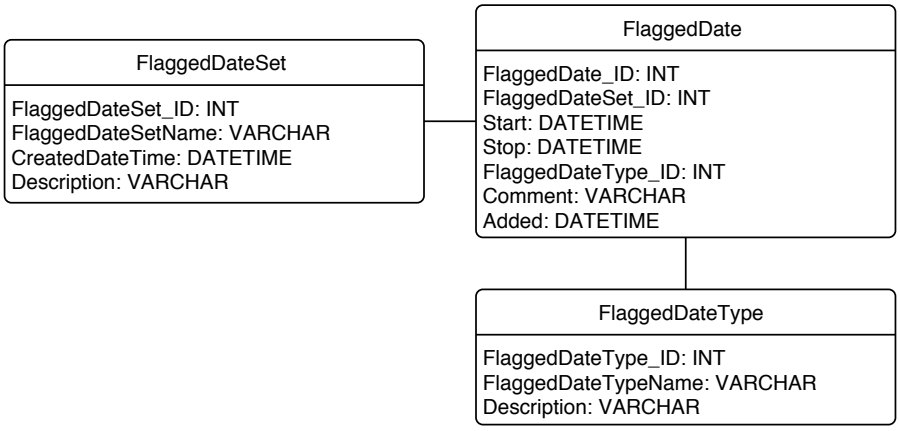

Figure 3.18. Diagram illustrating the database tables used for storing date interval flags.

reason for the flagging. The `FlaggedDateSet` table is used to define independent sets of `FlaggedDate` entries so that separate reductions can be stored in the database simultaneously without requiring them to share common date interval flags.

3.3.4 Managing Data Reduction

In figure 3.19, a diagram of the tables used to manage the data reduction process is shown. The `Reduction` table stores a record of every execution of a low-level or high-level reduction script. Flux density calibrations are not stored because these do not modify the data in the reduction database. When a reduction is begun, a new entry in the `Reduction` database is created with the `Start` field filled in with the UTC date and time of creation. Upon completion of the reduction process, the `Finish` field is filled in.

The low-level reduction script reduces data in units of one day. Upon completion of a low-level reduction, the dates that were processed are noted in the `ReducedData` table. Each date is associated with the `Reduction` during which it was processed, and the type of reduction performed (i.e., whether it was processed using the VAX control system scripts, the MCS control system scripts, or was copied from an existing low-level reduction) is noted using the `ReductionType` table. Using these tables, the reduction scripts can automatically determine what time periods need to be reduced.

Each `Reduction` is given a version and the name of a set of reduction parameters. The version is stored as a free-form string in the `Version` field of the table. There is no domain table for this, the user must keep track of the version names outside of the database. By convention, low-level reductions are simply numbered, while high-level reductions are given numbers with the prefix “HL” (e.g., “HL15”). The data reduced for this thesis were given the special version “JLR01” which is based on low-level data from version 15.

The parameter list is stored as a reference to the `ReductionParameterSet` table. Entries in the `ReductionParameter` table define the set of parameters and their values that belong to each set. The names and data types for the reduction parameters are specified by the `ReductionParameterName` and

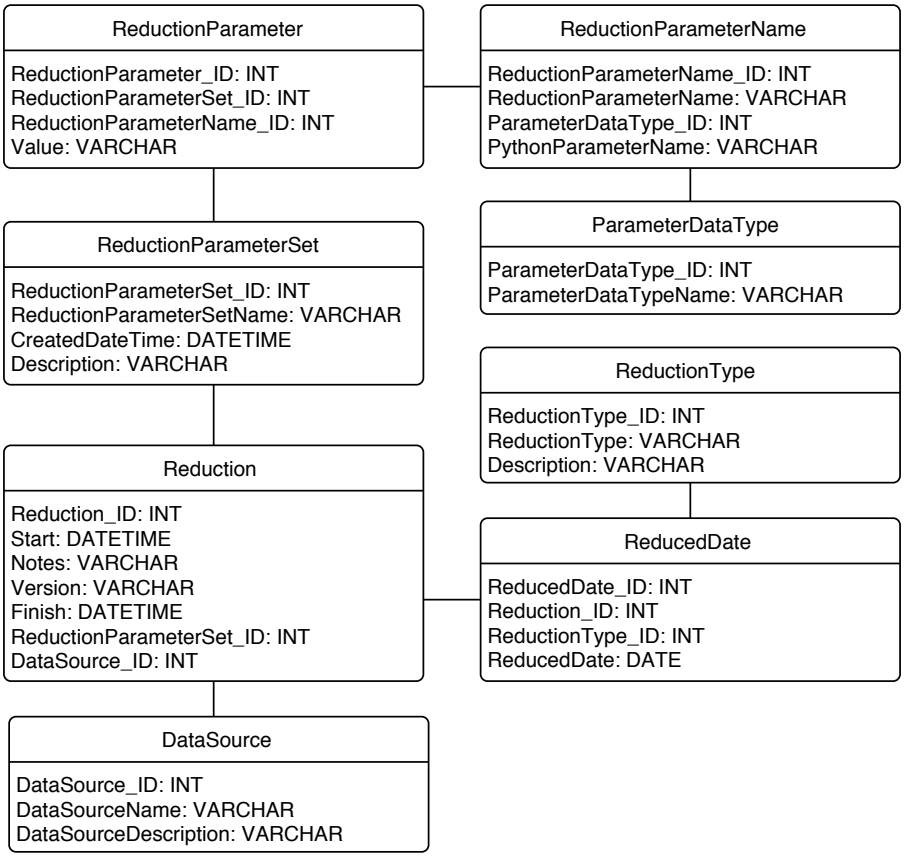

Figure 3.19. Diagram illustrating the database tables used for storing reduction data.

`ReductionParameterType` tables. Within the database, reduction parameter values of all types are stored as `VARCHAR` strings. When a Python script reads these from the database, the value is parsed according to its `ReductionParameterType` value. Multivalued parameters are, by convention, stored as comma-separated string parameters.

3.3.5 The Results Database

As of the time of writing, our 15 GHz light curve data for the CGRaBS sample are made publicly available through the web within a few months of observation.[13] This web page is powered by another MySQL database, the *results database*, into which our fully reduced data are inserted. This database schema, scripts to create and populate it, and the HTML and PHP code for the web interface were developed by Matthew A. Stevenson. Starting from his work, I have made a number of small changes to the scripts and database schema and added indexes to the database to improve performance. The source table schema described in section 3.3.2 is based on the design from the results database.

[13] `http://www.astro.caltech.edu/ovroblazars`

Chapter 4

Observing Program

In this chapter, we discuss the source samples that make up our monitoring program and examine the characteristics of the populations that make up those samples. We then consider the effects of confusion due to radio sources outside the program that lie near the program sources. Finally, we present the basic results of the observations.

In addition to the samples of blazars discussed in this chapter, a few other small samples are included in the 40 m monitoring program. These include (i) any objects not already included in our sample that are being studied in the F-GAMMA (Angelakis et al. 2010) or VERITAS programs (Weekes et al. 2002); (ii) a variety of galactic objects, such as microquasars and cataclysmic variables; and (iii) a few bright radio galaxies that show interesting jet properties. We are continually adding sources of interest to our monitoring sample, so that as of this publication, the sample comprises over 1550 objects that are monitored twice weekly. We will not discuss these other samples further, however.

4.1 Source Selection

Choosing an appropriate sample of sources for monitoring is, obviously, a critical element of a campaign such as this. In order to draw robust statistical conclusions that we may confidently extrapolate to the parent population, the sample should be complete with respect to physical characteristics that could affect those conclusions. When completeness is not possible, care must be taken to understand thoroughly the impact of selection effects before claiming that a result is physically significant. AGN and blazar samples are especially sensitive to selection effects due to the effects of Doppler beaming (e.g., Lister & Marscher 1997). With this in mind, the selection of the core sample for our monitoring program was driven by three considerations. First, since we are interested in the detailed study of the radio variability properties of the blazar population and the dependence of these properties on other observables such as redshift, the sample should be large enough to divide into subsamples (e.g., in redshift or luminosity bins) with each subsample large enough to permit statistical characterization.

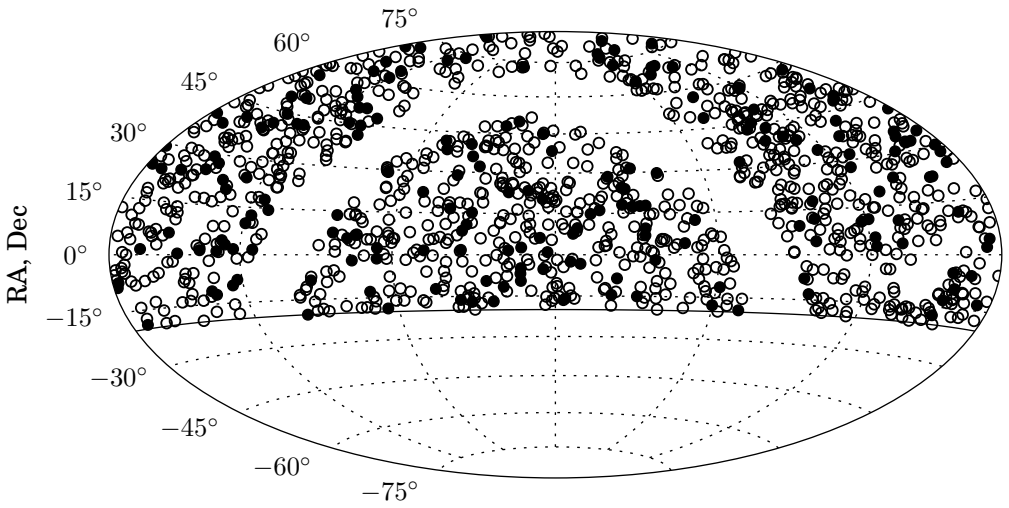

Figure 4.1. Positions of the CGRaBS sources in our program in equatorial coordinates. Filled circles indicate CGRaBS that are also in the 1LAC sample. The solid line marks the -20° declination limit of our program.

Second, to allow for the evaluation of the confidence level of any correlations or variable dependencies identified in our data through Monte Carlo simulations, and the generalization of our findings to the blazar population, the sample should be well-defined statistically, using uniform and easily reproducible criteria. Simply choosing bright, well-known, easily observable sources does not suffice for robust statistical study.

Finally, one of the major goals of our monitoring program is the cross-correlation of 15 GHz data with *Fermi* gamma-ray data, including cross-correlation of light curves in the two bands. For this reason we would like our sample to include a large number of gamma-ray–loud blazars. On the other hand, we would also like to be able to address the question of why some blazars are gamma-ray loud while other blazars, with apparently similar properties, are not. For this reason we would like our sample to be *preselected*—before *Fermi* data bias our understanding of what constitutes a likely gamma-ray–loud blazar—and, ideally, to include a comparable number of blazars which are not gamma-ray loud.

4.1.1 CGRaBS

The initial sample for our program was drawn from the Candidate Gamma-Ray Blazar Survey (CGRaBS Healey et al. 2008). The CGRaBS blazars in this survey satisfy all of the requirements above. These sources were selected from a flat-spectrum parent sample (complete to 65 mJy flux density at 4.8 GHz and radio spectral index $\alpha > -0.5$ where $S \propto \nu^\alpha$) by a well-defined figure-of-merit criterion based on radio spectral index, 8.4 GHz radio flux density, and X-ray photon flux from the ROSAT All Sky Survey, to resemble blazars that were detected by the Energetic Gamma Ray Experiment Telescope (EGRET, the precursor of *Fermi*-LAT). The CGRaBS sample is a total of 1625 active galactic nuclei (AGN) over the whole sky outside a $\pm 10^\circ$ band around the galactic plane. This sample was compiled before the launch of the *Fermi* and was expected to contain a large fraction of the extragalactic sources that would be detected by *Fermi*-LAT.

Table 4.1. Source counts in the CGRaBS and *Fermi* 1LAC samples

	Total		CGRaBS Subset	
Sample	All Decl.	Decl. $> -20^\circ$	All Decl.	Decl. $> -20^\circ$
CGRaBS	1625	1158	—	—
1LAC	709	545	291	221
1LAC Clean	599	454	263	199

Note: The 1LAC Clean sample above declination -20° is used for the population studies in this work, but is normally identified simply as 1LAC.

The initial sample for this monitoring program was the CGRaBS objects above declination -20°, a total of 1158 sources. The sky positions of these sources are shown in figure 4.1. These sources have been continuously monitored since the inception of our program, with publication-quality data available since 01 January 2008. Although, as we will see, CGRaBS proved only to be moderately successful at preselecting gamma-ray–loud blazars, we have kept all 1158 CGRaBS in our monitoring program (those that have been detected by *Fermi*-LAT are shown as filled circles in figure 4.1). This provides us a uniformly selected sample that can be used to study the differences between gamma-ray–loud and gamma-ray–quiet sources.

4.1.2 *Fermi*-Detected Sources: The 1LAC Sample

The *Fermi*-LAT collaboration published the LAT first-year point source catalog based on the first 11 months of science operations (1FGL Abdo et al. 2010a). Based on this catalog, a catalog of high-latitude blazar and AGN associations was assembled (1LAC Abdo et al. 2010b). Within the 1LAC catalog, a source was considered *clean* if it had a high association probability ($P > 80\%$), was the only source associated with the corresponding 1FGL gamma-ray source, and was not flagged during LAT analysis due to a problem or anomaly. In this work, we consider a blazar to be gamma-ray loud if it was associated with a blazar in the clean subset of the 1LAC catalog. The numbers of sources in the 1LAC catalog and its clean and CGRaBS subsets are tabulated in table 4.1. Figure 4.2 shows the positions of the 1LAC sources in our declination range in a Hammer equal-area projection. Henceforth, unless otherwise stated, when we refer to the 1LAC sample, we mean the clean subset of the 1LAC catalog above declination -20°.

CGRaBS sources made up 44% of the clean associations in the first-year *Fermi* AGN catalog. This number is thus far smaller than anticipated; in the full 1LAC clean sample only $\sim 16\%$ of the CGRaBS sources were detected, and a large number of blazars not in CGRaBS have been detected. This suggests that the CGRaBS (EGRET-like) blazar sample is substantially different from that seen in the early *Fermi* mission. This finding represents a unique opportunity to investigate why gamma-ray activity is found only in certain blazars, and for this reason we retain in our monitoring program all of the blazars in our original core sample even if they have not yet been detected by the LAT. However, in order to optimize the potential for studies of the cross-correlation between radio and gamma-ray light curves, we have since *added* (and we continue to add) to our monitoring program all new LAT-detected blazars north of -20° declination.

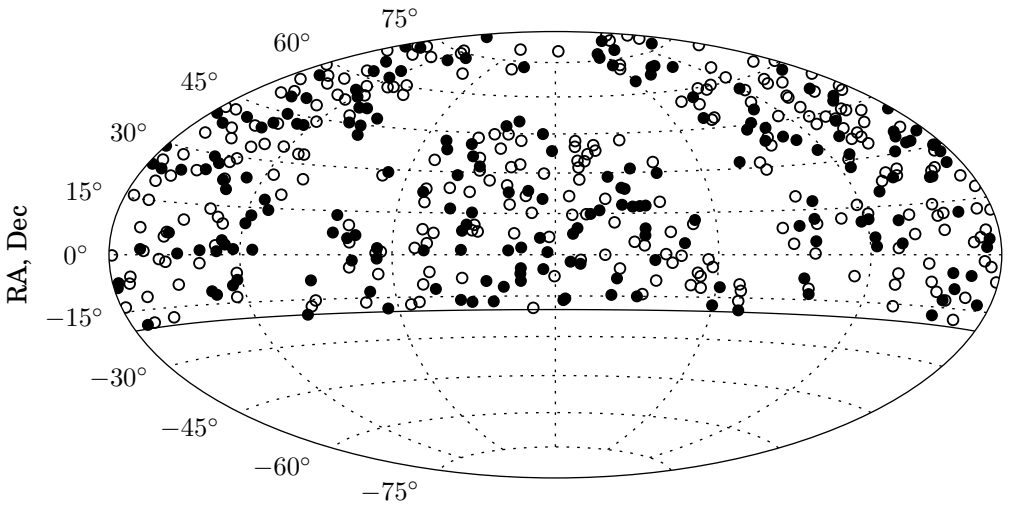

Figure 4.2. Positions of the 1LAC sources in our program in equatorial coordinates. Filled circles indicate 1LAC sources that are also in the CGRaBS sample. The solid line marks the $-20°$ declination limit of our program.

Table 4.2. Usage of calibration sources in this program

Use	Calibration Sources
Flux Density Scale	3C 286
Scale Error Est.	3C 48, 3C 161, DR 21
Cal. Spline Fit	3C 286, 3C 274, DR 21
Error Model	3C 286, 3C 48, 3C 161, DR 21
Beam Mapping	3C 286, 3C 48, 3C 295
Gain Curve	3C 286

It is important to note that 1LAC, which was derived from the 1FGL catalog, is not a true flux-limited sample. Its 1FGL parent sample is instead complete to a "test statistic" (TS) limit. The TS, defined as twice the log-likelihood difference between models with and without the presence of a point source, is a measure of the significance of a point source identified in the gamma-ray data (Mattox et al. 1996). Because diffuse emission and point source densities are not equal over the sky, the TS for the detection of a point source at a given gamma-ray flux can vary. The 1LAC sample was then generated from the association of gamma-ray point sources with radio sources using several parent radio source catalogs. This further complicates the explicit characterization of the sample selection criteria. However, in the work we present here, we do not rely on the properties of a flux-limited gamma-ray sample, so this distinction does not create a problem.

4.1.3 Calibration Sources

In addition to our blazar samples, several bright, stable sources are included in our program to provide flux density calibration and to monitor instrumental variability. These are the primary calibrator, 3C 286, plus 3C 48, 3C 161, 3C 274, and DR 21. Table 4.2 specifies how each of these sources is used as a calibrator in our program.

3C 286. Our primary astronomical flux density reference is 3C 286. It has been observed to produce a stable radio flux density over several decades (Ott et al. 1994), and is very widely used as a flux density calibrator. In this work we have adopted the Baars et al. (1977) value of 3.44 Jy for its flux density. An updated flux density value (3.37 Jy) for this source is given by Ott et al. (1994), so care must be taken when comparing our reported flux densities to programs that may use the updated value.

3C 48. The very compact quasar 3C 48 is well known and widely used as a flux density calibrator. It is known to vary slightly over long timescales, but is suitably stable for our purposes (e.g., Ott et al. 1994).

3C 161. The radio galaxy 3C 161 is also useful as a calibrator, but has been reported to vary by as much as 10% in 2.8 cm flux density (Andrew et al. 1978). We have had frequent anomalous pointing failures with this source, particularly in 2008 and 2009. Although we are uncertain of the mechanism for these failures, we have several hypotheses for the cause. As this is a rather southerly source (about -6° declination) it is typically observed at low elevations, so may be subject to atmospheric interference. The China Lake Naval Air Weapons Station is located south of the telescope and radio transmissions from their operations occasionally produce interference. However, other low-declination sources do not seem to exhibit this problem. Another possibility is contamination due to the low galactic latitude ($b = -8^\circ$) of this source. As described in section 3.2.1.1, we have manually removed instances of 3C 161 pointing failures and dropped any corrupted measurements that resulted.

3C 274. The bright radio galaxy 3C 274 (M87) normally exhibits only slow changes in flux density. We use this source as part of our procedure to remove residual systematic variations, described in section 3.2.2.2.

3C 295. We have used the bright, compact radio galaxy 3C 295 for measuring beam maps.

DR 21. The compact, galactic H II region DR 21 is useful as a bright, steady calibration source. Its low galactic latitude ($b = 1^\circ$) makes it susceptible to contamination from galactic emission at some parallactic angles.

4.2 Classifications and Redshifts

In the classification scheme we have adopted, blazars are a class of AGN that includes flat-spectrum radio quasars (FSRQs) and BL Lacertae objects (BL Lacs). In general, we have not attempted to verify the classifications or redshifts for the sources in this program, rather we have accepted the values in the publications from which we draw our sample. That is, for CGRaBS sources, we use the classifications redshifts from Healey et al. (2008) and for 1LAC sources, we use the classifications and redshifts from Abdo et al. (2010b). These publications agree on the values for most of the sources common to the two samples, but there are a

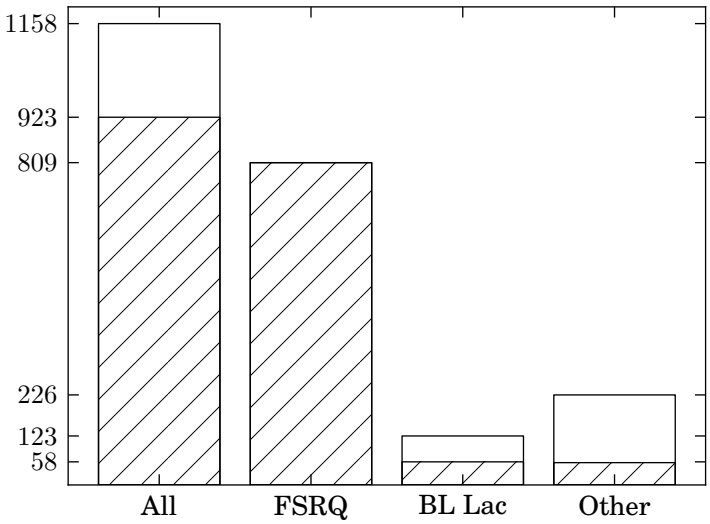

Figure 4.3. Optical classifications for the CGRaBS sample. Shaded regions indicate the number of sources in each class with measured redshifts. The "Other" category has 56 sources with known redshifts.

few differences. In most cases, the differences were due to additional observations after the publication of the CGRaBS paper, although in a few cases, redshift values changed substantially or optical classifications disagreed. We have accepted the values from the 1LAC paper in these cases. The values we have adopted are tabulated in appendix C.

In parallel with this radio monitoring program, we and our collaborators have been observing sources in our samples with optical telescopes to obtain spectroscopic redshifts and classifications. Some results from this program appeared in Abdo et al. (2010b) and were adopted from there. In a few cases, we have adopted redshifts based on unpublished results from this optical observing program. These results will be published in forthcoming papers separately describing the treatment of BL Lac and FSRQ sources (Shaw et al. 2012b, 2012a, *in preparation*) .

The core sample for our monitoring program consists of the 1158 CGRaBS sources north of declination -20°. As published, our subset of the CGRaBS sample contains 812 FSRQs, 111 BL Lacs, and 235 radio galaxies and objects without spectroscopic identification. In our analysis we use redshifts from the CGRaBS publication, which covered 93.9% of the sample (100% of FSRQs, 49% of BL Lacs). With the updates from our optical programs and from Abdo et al. (2010b), the 40 m CGRaBS sample now consists of 809 FSRQs, 123 BL Lacs, and 226 radio galaxies or unidentified objects. Among the FSRQ and BL Lacs with which we are concerned in this thesis, redshifts are available for 93.0% of the sample (100% of FSRQs; 47% of BL Lacs). The redshift completeness has fallen slightly due to the identification of 12 objects as BL Lacs with unknown redshift. Our 1LAC sample consists of the 454 objects north of declination -20°. Of these, 183 are classified as FSRQs and 223 as BL Lacs. Among the BL Lacs and FSRQs, redshifts are known for 68.9% of the sources (100% of FSRQs, 43% of BL Lacs).

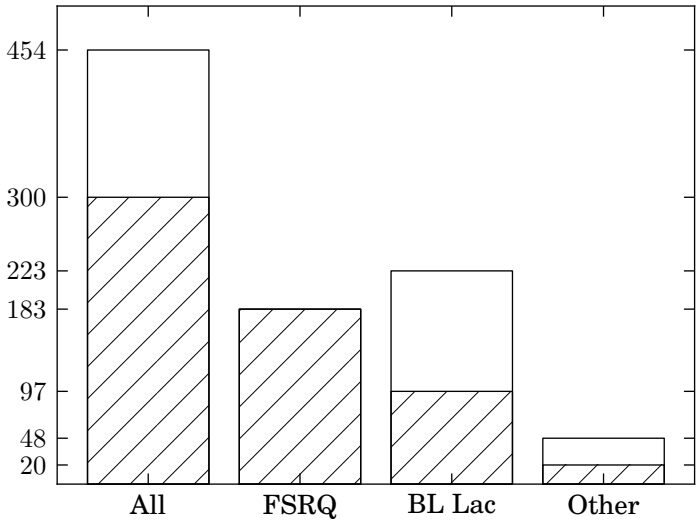

Figure 4.4. Optical classifications for the 1LAC sample. Shaded regions indicate the number of sources in each class with measured redshifts.

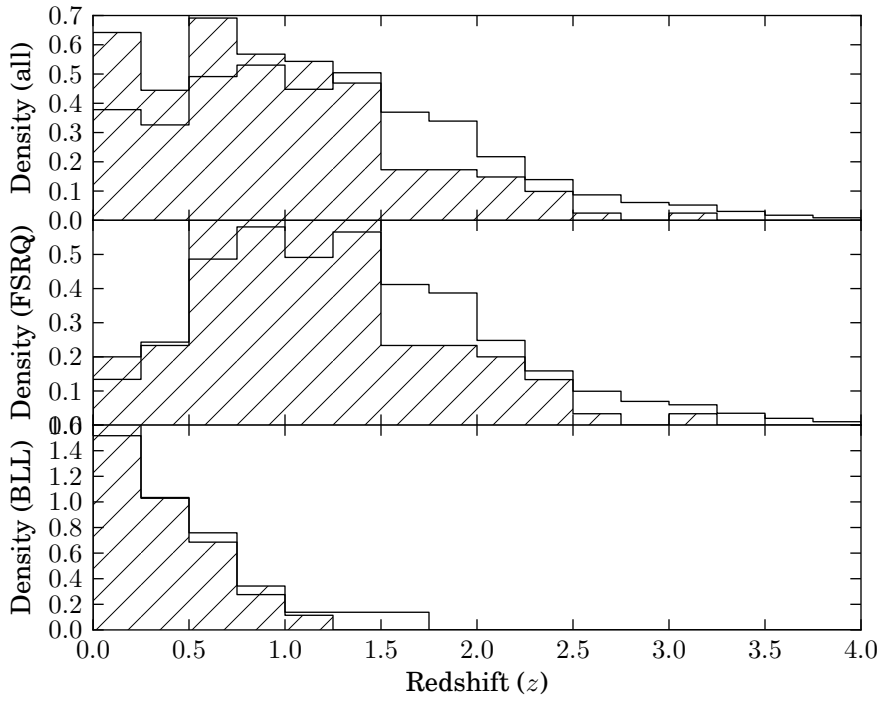

Figure 4.5. Histograms of redshifts for the CGRaBS sample (top), the FSRQ subset (middle), and the BL Lac subset (bottom). In each plot, the subset of sources that are also in the 1LAC sample is shown by the shaded region. Each histogram is normalized to integrate to unity.

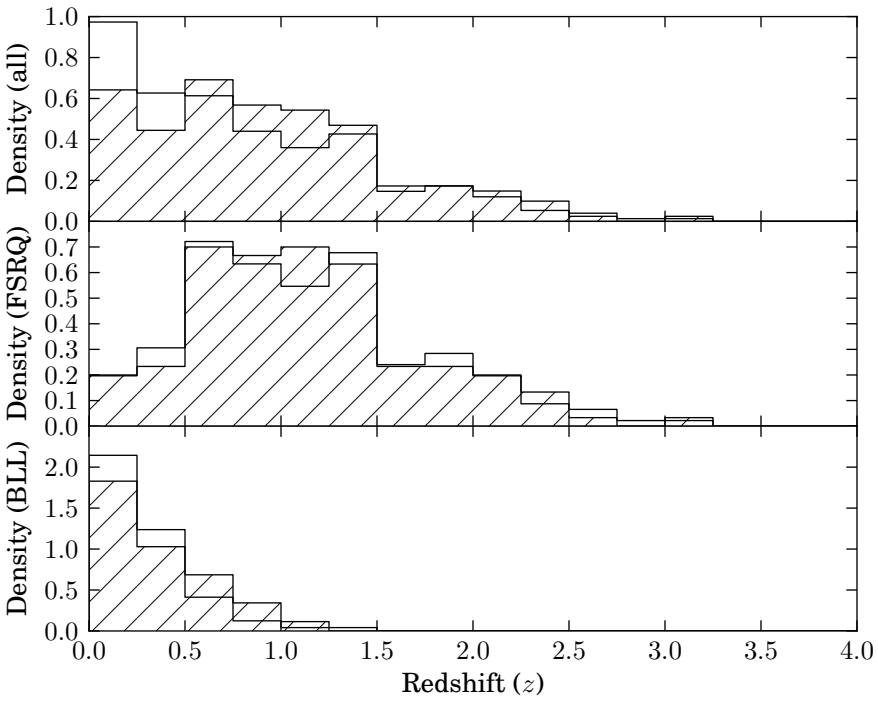

Figure 4.6. Histograms of redshifts for the 1LAC sample (top), the FSRQ subset (middle), and the BL Lac subset (bottom). In each plot, the subset of sources that are also in the CGRaBS sample is shown by the shaded region. Each histogram is normalized to integrate to unity.

How well did the CGRaBS program do at selecting gamma-ray–loud blazars? We know from the 1LAC results above that it did not include a majority of 1LAC sources, nor did a majority of the CGRaBS sources show up in 1LAC. Did it select a similar population? In figure 4.3 we show a chart of the source classifications for the CGRaBS sample. Sources classified as FSRQs dominated the sample with BL Lac objects outnumbered by nearly 7 to 1. A similar chart of classifications for the 1LAC sample is shown in figure 4.4. The 1LAC BL Lac population outnumbers the FSRQs in that sample. So clearly CGRaBS did not predict the ratio of optical classifications. As we will discuss in section 5.6, this can be attributed at least partially to the CGRaBS figure of merit being tuned using results from EGRET, which had a substantially different spectral sensitivity than the LAT. This makes the LAT much more efficient at detecting BL Lac objects, whereas EGRET was a better detector of FSRQs.

Next we consider the redshift distributions of the two samples. In figure 4.5, we show histograms of the known redshifts for the CGRaBS sample and its BL Lac and FSRQ subpopulations. The shaded regions represent the sources that are also in the 1LAC sample. A Kolmogorov-Smirnov (K-S) test[1] rejects the hypothesis that the CGRaBS and the 1LAC subset are drawn from the same redshift population for all sources ($p = 5.8\times10^{-5}$) and for the FSRQ subset ($p = 0.0031$). There appears to be an excess of CGRaBS at $z > 1.5$ compared to the 1LAC sample, and perhaps a deficit in the $z = 0.5$ to 1.5 range. The null hypothesis is *not* rejected for the BL Lac subset ($p = 0.80$).

The redshift distribution of the 1LAC sources in the 40 m program is shown in figure 4.6, for all sources as well as for the FSRQ and BL Lac subsets. The histograms for the 1LAC sources that are also in the CGRaBS sample are shown by the shaded regions (note that these shaded regions are the same as those in figure 4.5). A K-S test comparing the distributions marginally rejects the hypothesis that the overall 1LAC sample and the CGRaBS subset are drawn from the same distribution ($p = 0.043$). For the sources identified as FSRQs or as BL Lacs, the K-S test does not reject the null hypothesis, with $p = 1.0$ and $p = 0.52$, respectively.

What do these comparisons tell us? Clearly the BL Lac populations in CGRaBS and 1LAC are drawn from similar parent samples—in both cases, the K-S test accepted the null hypothesis. The situation is more complicated for FSRQs, with the CGRaBS subset of 1LAC FSRQs *inconsistent* with the total CGRaBS FSRQ population, while the 1LAC subset of CGRaBS FSRQs is *consistent* with the total 1LAC FSRQ population. This can be simply explained if the CGRaBS sample contains two populations of FSRQs: one drawn from the same population as the 1LAC sample, and another that is disjoint from it.

The K-S test soundly rejected the hypothesis that the CGRaBS total sample (BL Lac and FSRQ) matched its 1LAC subset. This is easy to understand: the overall sample cannot represent the same distribution if the FSRQ subsamples are so different. The different BL Lac and FSRQ fractions between the two samples reinforces the disagreement.

[1] All K-S tests performed for this work were computed using the `ks_2samp` routine from the SciPy package.

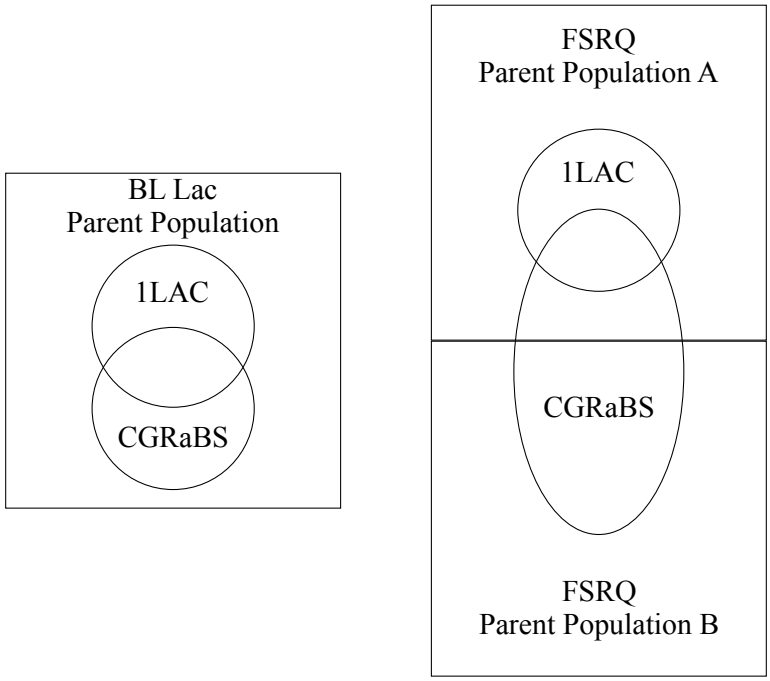

Figure 4.7. Venn diagram showing the relationship between the CGRaBS and 1LAC samples suggested by comparison of optical classification and redshift distributions.

For the 1LAC total sample versus its CGRaBS subset, the K-S test only marginally rejected the hypothesis of matching distributions. In this case, both the FSRQ and BL Lac subsets were compatible with coming from the same distribution. Thus, it is entirely the BL Lac/FSRQ fraction that causes the difference between the 1LAC total sample and the CGRaBS subset. In figure 4.7 we show a Venn diagram illustrating the apparent relationship between the 1LAC and CGRaBS samples and their various subpopulations.

4.3 Observation Scheduling

Because of the large number of targets and the continuous nature of the observing that constitutes this monitoring program, automated generation of observing schedules is a necessity. In this section we explain the requirements and constraints of this system and describe the solution we have developed. The original scheduling algorithm used from the inception of the program through early 2009 was written by Lawrence Weintraub, with substantial later developments by Walter Max-Moerbeck. The new algorithm used from 2009 until the present was developed and written by Walter Max-Moerbeck. In this section we first describe the newer algorithm, then briefly discuss the original approach.

The large number of sources being observed requires the development of strategies to optimize the use of the telescope and minimize the effect of known systematic errors. The principal systematic errors we try to minimize are gain variations, atmospheric optical depth variations, and pointing errors. To achieve this optimization while minimizing slew times and dead times between observations requires careful planning.

Schedules are arranged to ensure that sources are observed between zenith angles of 20° and 60° whenever possible. This is done for a number of reasons:

- the figure of the telescope was set for maximum gain in this elevation range;
- at zenith angles less than about 20° the telescope has to move rapidly to track an object and pointing accuracy can be compromised;
- at zenith angles greater than about 60° ground spillover increases significantly with decreasing elevation;
- it is desirable to minimize the variation in atmospheric optical depth on our sources so as to minimize this particular source of error; and
- we try to minimize telescope slew times by observing to the south and east in a limited elevation range.

In the scheme we have developed, the sky is divided into 192 cells, each with a diameter $\lesssim 20°$, using the HEALPix[2] mesh with $N_{side} = 4$ (Górski et al. 2005). Each source is assigned to a cell. From the sources in each cell, a pointing calibrator is selected using the following criteria, applied in order.

1. If there is a flux calibrator in the region, this source is selected.
2. If one or more sources in the region have a flux density larger than 500 mJy, the one which minimizes the average angular distance to all the sources in that region is selected.
3. The source with the largest flux density in the region is selected. For these flux density comparisons, the median flux density of the source during the previous year's observations is used.

Sources within the region are scheduled to minimize slew time, using a direct search to find the optimal order for regions with fewer than 9 sources and simulated annealing for regions with 9 or more sources. A second optimization step determines the order in which the regions are scheduled using a heuristic algorithm in which regions are observed within a fixed zenith angle range and regions to the south have priority. The total sample is observed in three days.

Prior to 16 March, 2009 (MJD 54906) a different scheduling system was used. This system used a genetic algorithm to find an ordering of pointing offset, calibration diode, and flux density observations that ensured the entire schedule could be observed with an adequate cadence. Pointing offsets and calibration diode measurements were scheduled approximately every 45 min. The scheduler ensured that observations near the sun and moon were avoided. Because the sun and moon move across the sky, schedules were updated regularly to avoid observations in the sun and moon regions.

The original algorithm suffered from several drawbacks that became apparent after studying the first year of data from the monitoring program. First, while the genetic algorithm tended to produce similar solutions

[2] `http://healpix.jpl.nasa.gov/`

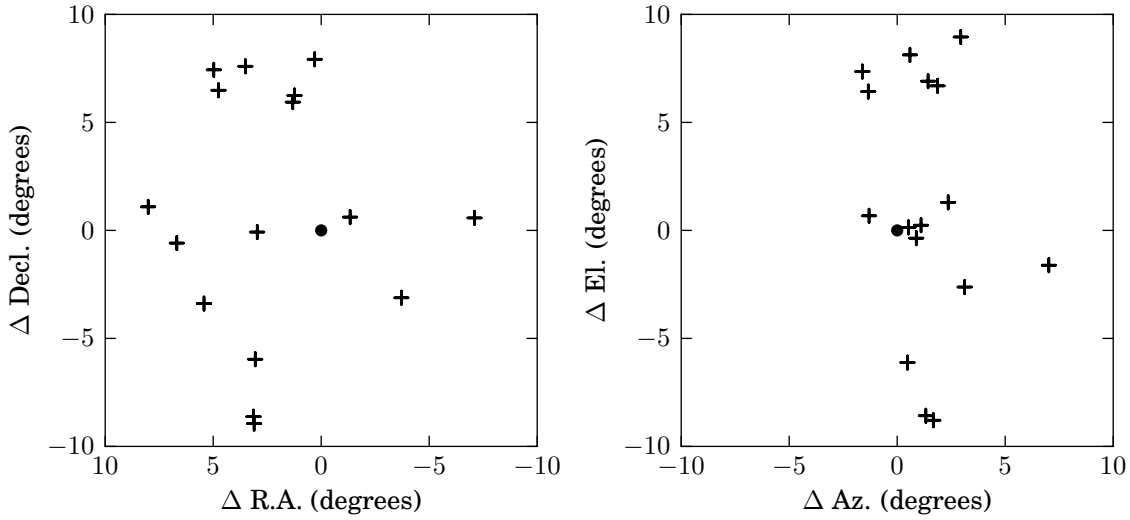

Figure 4.8. Example of an observation region, Region 83 in the 11 April 2011 schedule revision. *Left:* Source positions in equatorial coordinates offset from the pointing calibrator, indicated by the filled circle. *Right:* Region as it was observed on 15 April 2011, plotted in telescope azimuth and elevation coordinates. The solid circle indicates the location of the end of the POINT procedure, crosses indicate the location of the FLUX procedures, including the FLUX procedure on the pointing calibrator.

from one schedule iteration to the next, there was not an explicit attempt to observe sources at the same local sidereal time (LST) each time. As a result, sources sometimes varied substantially in parallactic angle, leading to increased exposure to contamination from confused sources. Second, and more critically, pointing offset measurements were not constrained to be near in azimuth and zenith angle to the subsequent flux density measurements. As a result, some sources were measured at more than 30° from the position of the pointing offset, which led to unreliable pointing. These shortcomings were addressed in the new system.

4.4 Sun and Moon Interference

The sun and moon are extremely bright 15 GHz emitters relative to the blazar sources we observe, both easily saturating (although not damaging) the 40 m receiver when the variable attenuator is set to its normal observing level. In this section we describe the method we used to determine the minimum solar elongation at which reliable observations can be made.

To detect solar contamination, we need to identify FLUX procedures that are artificially high. This task is greatly simplified if we examine one of the three combinations of the A, B, C, and D segments of the FLUX procedure that cancels the source field contribution. We use the source-nulled flux, $S_{\rm null}$, defined in equation (2.47). This signal measures the difference in brightness between the two reference fields of the FLUX procedure, which are unlikely to be exactly balanced when the sun or moon is present in a sidelobe. Thus, it is a good measure of contamination.

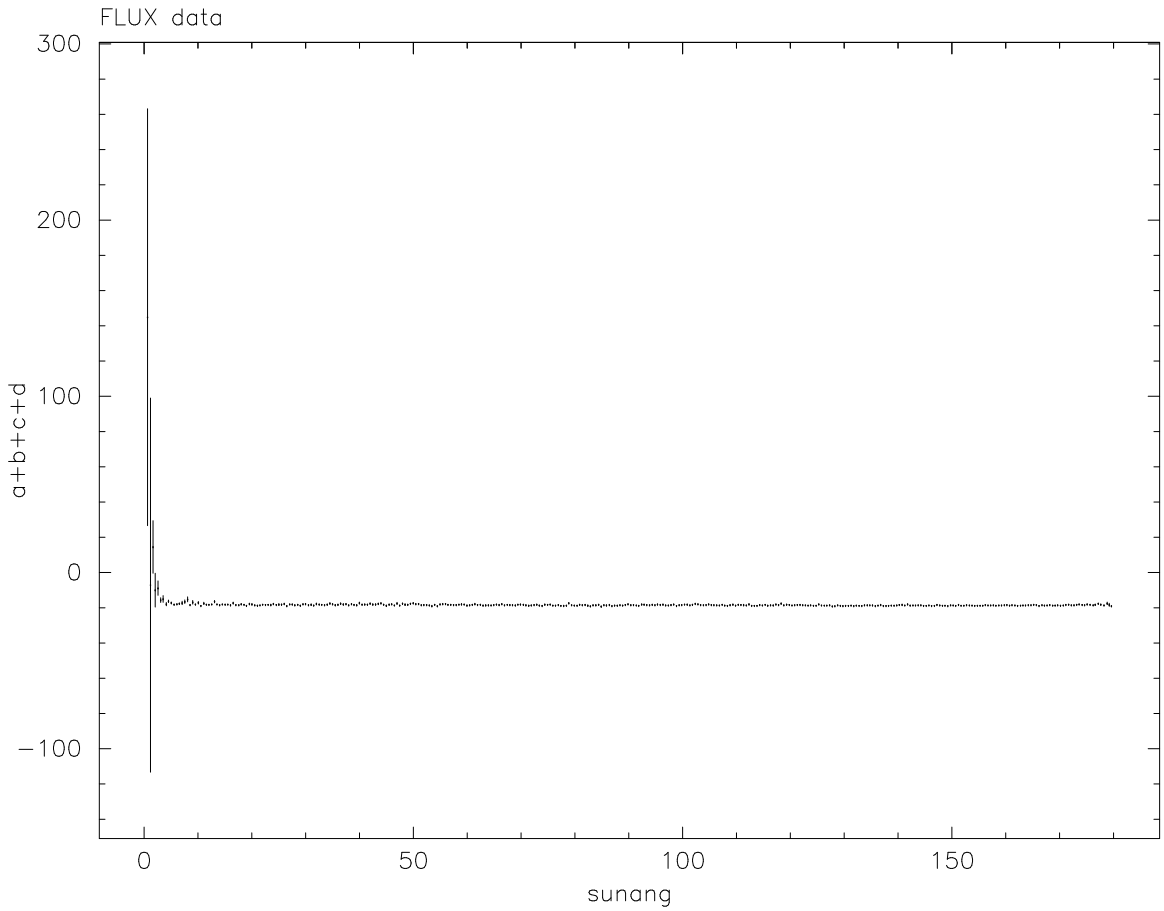

Figure 4.9. Source-nulled flux ($S_{\rm null}$) (arbitrary units) as a function of solar elongation (degrees). Plotted points and error bars are the means and standard deviations of 2009 data in $\sim$0.5° bins.

Figure 4.9 shows $S_{\rm null}$ as a function of solar elongation using partially calibrated data from 2009. The data were binned in approximately 0.5° bins and the mean in each bin is plotted, with the standard deviation of the data in the bin used for the error bar. Not surprisingly, severe contamination is found very near 0° solar elongation. Figure 4.10 shows a detail of this plot between 0° and 50°. A strong contaminating signal is apparent up to 10°, and there are several bins out to near 20° that suggest additional contamination may be present.

In figure 4.9, there is a hint of increased scatter near 180°elongation. This is certainly *not* due to actual solar interference—at an elongation of 180°, the sun must be below the horizon! This is because the zenith angle at which a very large solar elongation is possible is restricted. Figure 4.11 shows that small zenith angles are not sampled at high elongations. The increased scatter in these bins is therefore ascribed to increased atmospheric signal.

Solar contamination is also evident in the total power (i.e., the average of the Dicke-switched samples). We arbitrarily choose the A segment of the FLUX procedure and plot its total power signal as a function of solar elongation in figures 4.11 and 4.12.

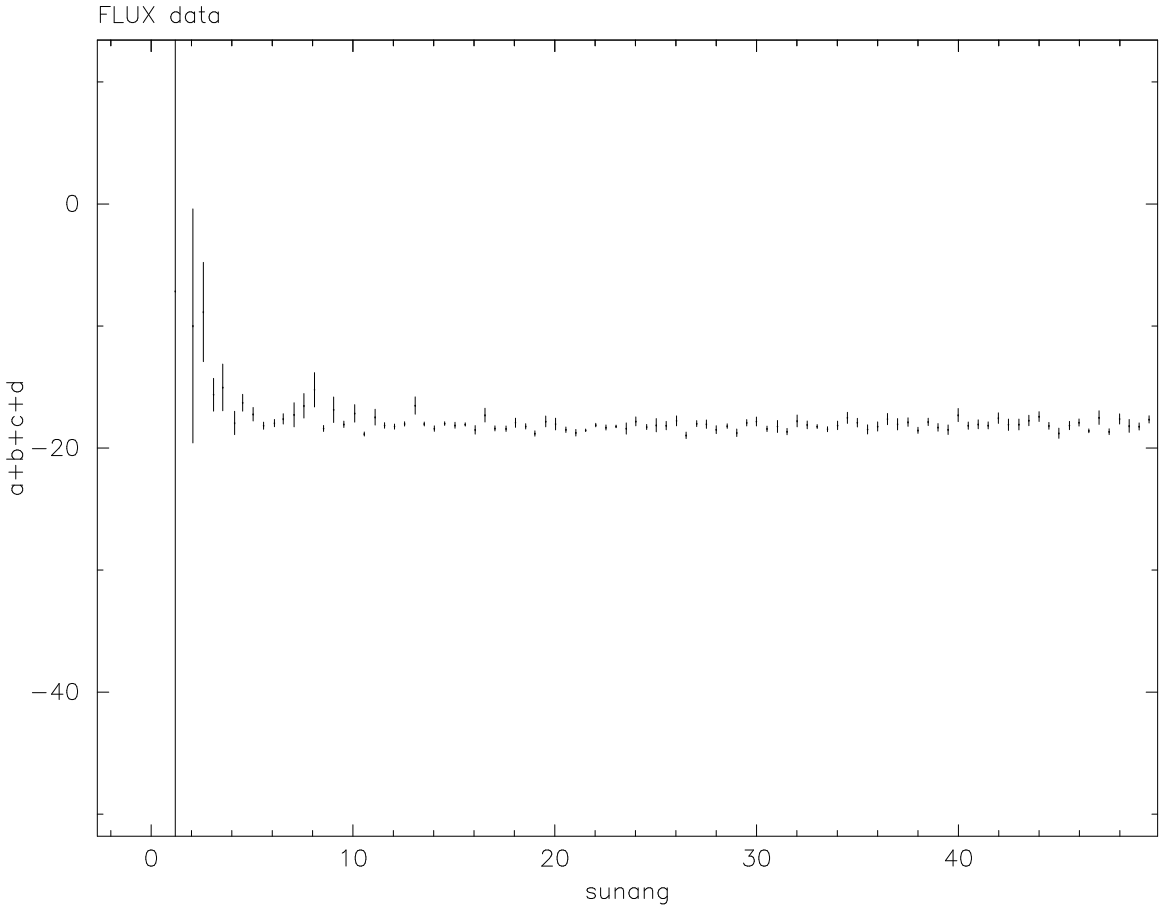

Figure 4.10. Detail of figure 4.9, showing features at small solar elongation.

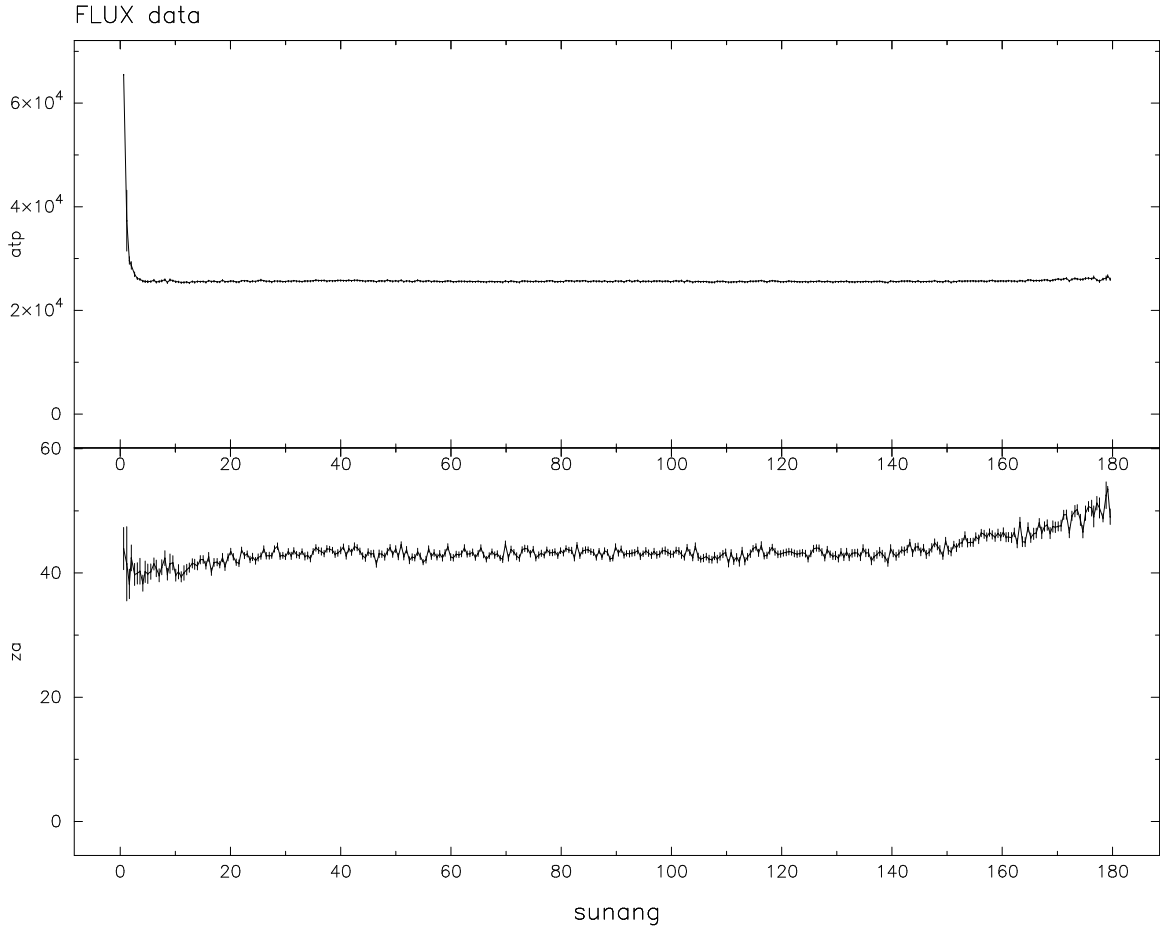

Figure 4.11. Top panel: Total power FLUX A segment data plotted against solar elongation. Data points (errors) are the means (standard deviations) within 0.5° bins. Bottom panel: Zenith angle of the observations in the upper panel, using the same binning. The absence of low-zenith angle observations at high solar elongations is a geometric effect and the increased atmospheric optical depth at higher zenith angles explains the high-elongation behavior of the upper plot.

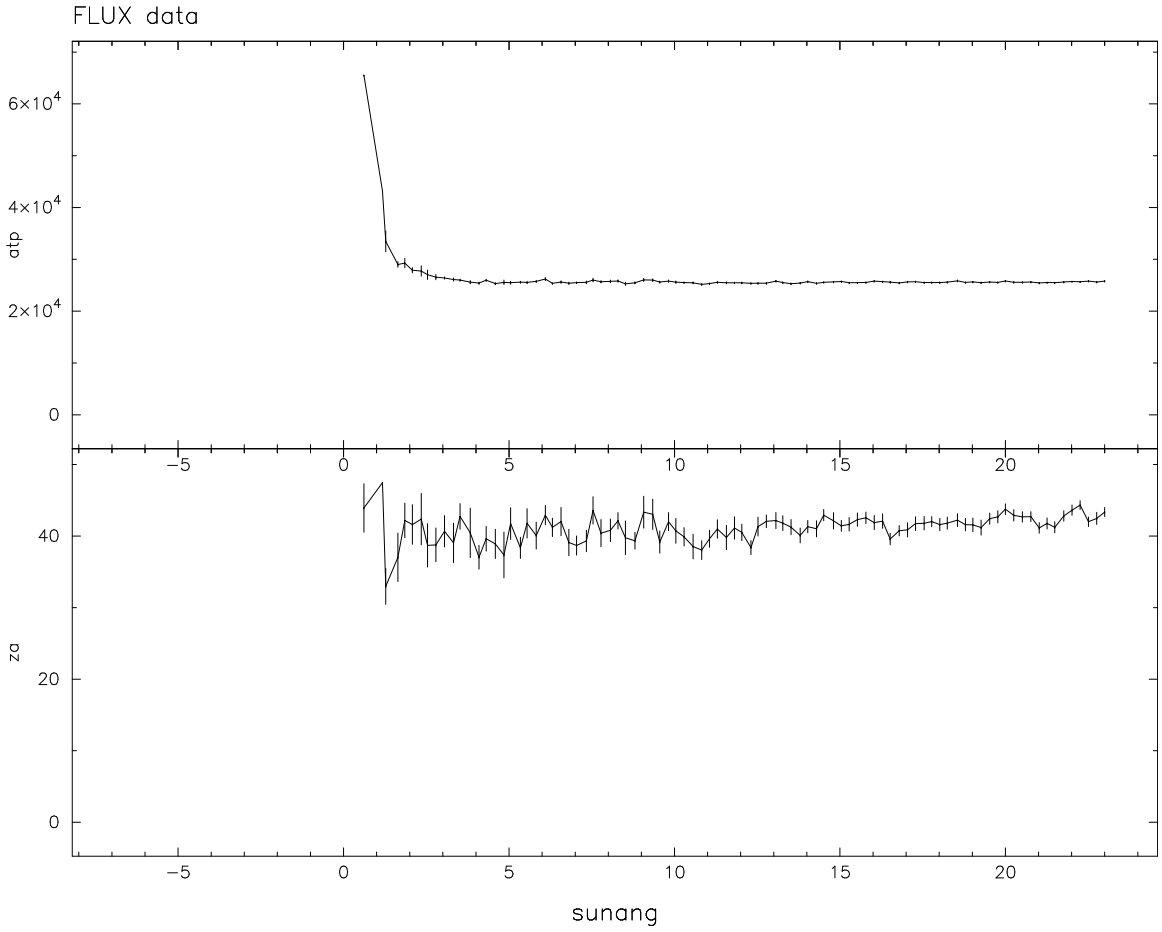

Figure 4.12. Detail of figure 4.11 with bins of $\sim 0.25°$.

Based on these results, we placed a threshold of 10° solar elongation below which FLUX procedures are discarded. Although there is evidence of contamination out to 20°, only a small amount of data in that range of elongation is actually affected so a disproportionate amount of perfectly sound data was discarded by increasing the threshold. Separate studies for the moon were not performed, but because its angular size is similar to that of the sun and because its brightness temperature is much lower, the effect should be similar but smaller in scale. We have therefore adopted the same threshold for lunar elongation.

4.5 Confusion

One of the chief reasons for the galactic latitude cut we have enforced on our program sources, $|b| > 10^\circ$, is to avoid confusion—contamination of our measurements from nearby sources. Even at high galactic latitudes, however, radio sources are numerous enough that inevitably some of our program sources will be affected by confusion. In this section, we use a model of the 15 GHz differential source count to estimate the number of confused program sources at various flux density limits. We will find that because our sources are relatively bright, the number of sources likely to be significantly affected by confusion is small enough that we need not be concerned about this effect when studying samples of our sources.

4.5.1 Basic Calculation

To estimate the number of confused sources, we begin by assuming a differential source count given by

$$n(S) \equiv \frac{dN}{dS} = n_0 \left(\frac{S}{\mathrm{Jy}} \right)^\beta , \tag{4.1}$$

where n_0 is a fiducial number of sources per Jy per sr at 1 Jy. Waldram et al. (2010) find a differential source count at 15.2 GHz with $n_0 = 51\ \mathrm{Jy}^{-1}\ \mathrm{sr}^{-1}$ and $\beta = -2.15$. They find no evidence for deviation down to a completeness limit of 5.5 mJy.

For a confusion limit of S_c, we need to consider only sources brighter than the limit. A source with flux density $S = S_c$ will only be detected at the center of the beam, while brighter sources will be detected farther from the beam center. The solid angle for contamination by a source of flux density S is

$$\Omega(S) = \iint_{g_A \geq S_c/S} d\Omega. \tag{4.2}$$

Here, g_A is the antenna gain, normalized so the total receiver gain is 1 at the center of the beam. Assuming a $157''$ FWHM Gaussian beam, this is $g_A(\theta, \phi) = \exp\left(-\theta^2/\theta_0^2\right)$ with $\theta_0 = 4.57 \times 10^{-4}$ rad. The region to integrate is bounded by $\theta_{max} = \theta_0 \left[\ln(S/S_c)\right]^{1/2}$, so

$$\Omega(S) = 2\pi \int_0^{\theta_{max}} \sin\theta \, d\theta = 2\pi \left[1 - \cos\left(\theta_0 \sqrt{\ln(S/S_c)} \right) \right] . \tag{4.3}$$

Assuming that source clustering is negligible, the number of contaminating sources detected as brighter than S_C in a single field is then given by

$$N(S_c) = \int_{S_c}^{\infty} n(S)\Omega(S)\, dS = \int_{S_c}^{\infty} 2\pi n_0 \left(\frac{S}{\mathrm{Jy}}\right)^{\beta} \left[1 - \cos\left(\theta_0 \sqrt{\ln(S/S_c)}\right)\right]\, dS. \tag{4.4}$$

The number of fields expected to be contaminated at various confusion limits is tabulated in table 4.3.

4.5.2 Contaminated FLUX Procedures

The preceding considered a single field corresponding to either the *ant* or the *ref* beam of the telescope at a single pointing. We actually wish to know how many FLUX procedures are contaminated. A FLUX procedure consists of three fields: a main field with t seconds of integration and two reference fields each with $t/2$ seconds of integration. Ignoring mismatch between the *ant* and *ref* beams, we can compute this from the above by treating the FLUX procedure as three independent fields and summing the expected number of contaminating sources in the three. Because the reference fields are integrated only half as long as the main field, the contamination limit for their contribution must be increased by $\sqrt{2}$. The resulting equation is then

$$N_{FLUX}(S_c) = N(S_c) + 2N(S_c\sqrt{2}). \tag{4.5}$$

The results for this are also presented in table 4.3.

To estimate the number of contaminated sources in our program, we note that distributing contaminating sources in the FLUX measurement fields is a Poisson process. The process is parameterized by a mean, $\bar{\nu} = N_{FLUX}(S_c)$. The probability that a particular source has one or more contaminating sources in its FLUX fields is then

$$p = \sum_{n=1}^{\infty} \frac{\bar{\nu}^n}{n!} e^{-\bar{\nu}} = 1 - e^{-\bar{\nu}}, \tag{4.6}$$

where the latter equality results from the normalization of the Poisson distribution. The expected number of program sources with one or more contaminating sources is then just $np = n(1 - \exp[-N_{FLUX}(S_c)])$, where n is the total number program sources. For the 40 m program, $n = 1413$. The results of this calculation are tabulated in table 4.3.

In these calculations we have neglected reference field rotations with parallactic angle, which will increase our exposure to confusion. In figure 4.13 we plot the complete reference field coverage during the 42 months of observing for four sources. The area covered is typically a few times the area of the reference beams for a single parallactic angle, and in any case cannot exceed that area by more than about a factor of 10 before the entire ring of reference fields has been covered. Thus our confusion estimates should be reliable to within a small factor. Because only a few percent of our sources are likely to be contaminated even at 10 mJy level

Table 4.3. Contamination estimates at various flux density limits, S_c

S_c (mJy)	$N(S_c)$[a]	$N_{FLUX}(S_c)$[b]	# Affected[c]
100	3.6×10^{-4}	8.4×10^{-4}	1
50	7.9×10^{-4}	1.9×10^{-3}	3
20	2.3×10^{-3}	5.3×10^{-3}	7
10	5.1×10^{-3}	1.2×10^{-2}	17
5	1.1×10^{-2}	2.6×10^{-2}	37
2	3.2×10^{-2}	7.5×10^{-2}	106
1	7.1×10^{-2}	1.7×10^{-1}	240

Note: Calculated from equations (4.4) and (4.5) with $n_0 = 51\ \mathrm{Jy}^{-1}\ \mathrm{sr}^{-1}$, $\beta = -2.15$, $\theta_0 = 4.57 \times 10^{-4}$ rad, and an upper integration limit of 1000 Jy.
[a] Expected number of contaminating sources per field.
[b] Expected number of contaminating sources per FLUX procedure.
[c] Contaminated FLUX procedures assuming 1413 sources.

($\sim 3\%$ of the median flux density of sources in our sample), we may safely ignore the effects of confusion in our statistical analyses.

4.6 Observation Results

In this section we describe the outcome of our observing efforts and the basic results of our monitoring program for the CGRaBS and 1LAC samples. More sophisticated variability analyses and results are described in chapter 5.

4.6.1 Observing Efficiency

Our target cadence was two flux density measurements per source per week, or about 365 measurements per CGRaBS source in the 42-month data set. For non-CGRaBS sources, the expected number of observations depends on when the source was added to the program. The number of successful observations for each source is listed in table C.1. Figures 4.14 and 4.15 show histograms of the per-source observing efficiencies relative to the nominal cadences. For non-CGRaBS 1LAC sources, we counted the expected number of observations from the date of the first successful observation for this calculation. The mean efficiency relative to the nominal cadence for CGRaBS sources was $202/365 = 55\%$ and for 1LAC sources was $147/234 = 54\%$. These efficiencies include all telescope outages, engineering time, and the effects of the data filters. The number of successful observations per week of the observing program is plotted in figure 4.16.

High winds are the biggest single cause of lost observing time in this program. As described in section 2.1.1.4, data collected when the wind exceeds 15 mph (6.7 $\mathrm{m\,s^{-1}}$) must be discarded. This is, of course, an unavoidable loss of observing time. Unpredictable hardware failures and power outages have also caused

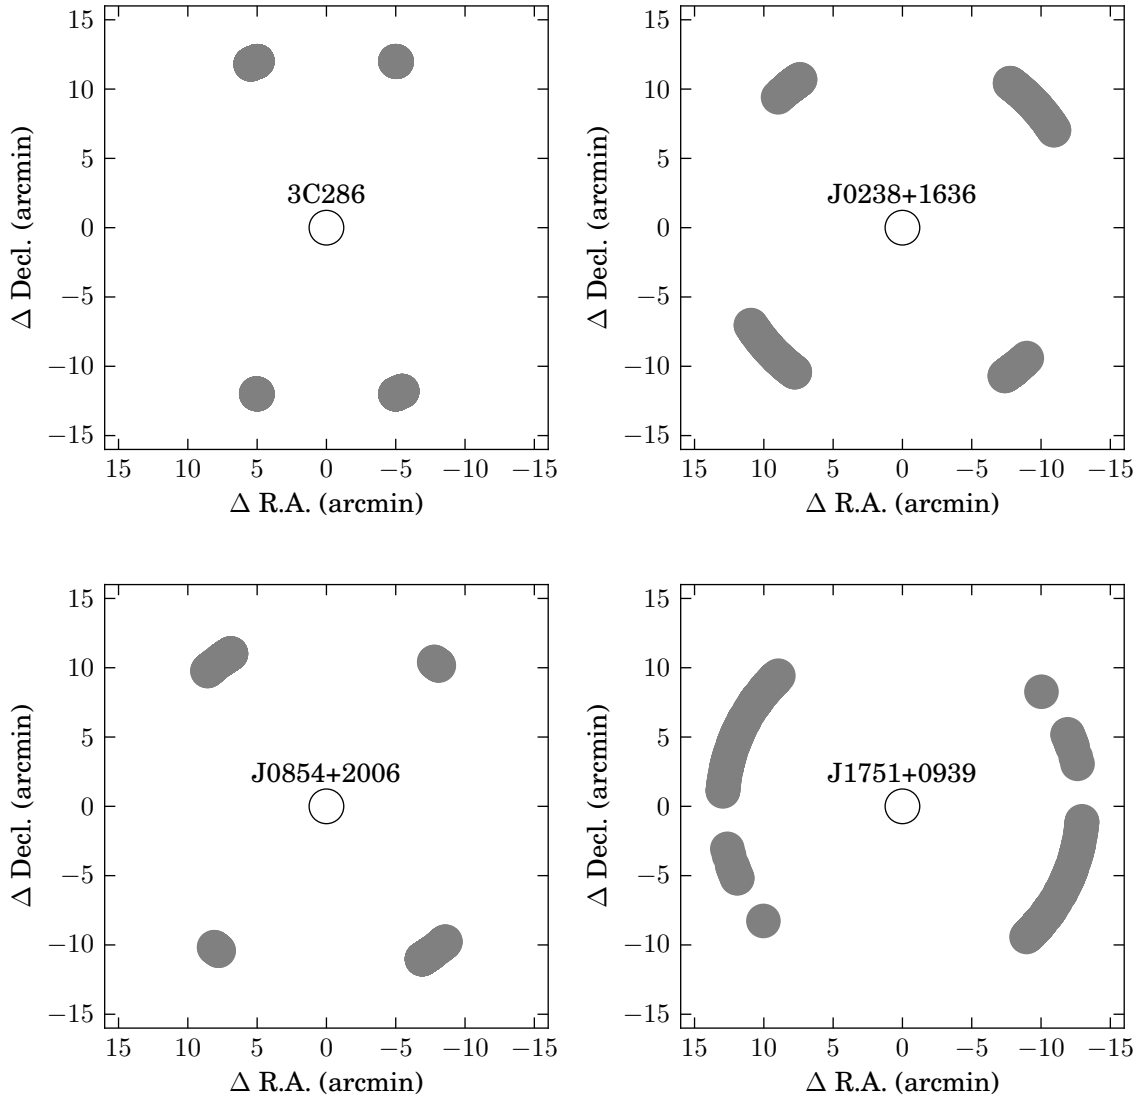

Figure 4.13. Reference field coverage during the full data set for four sources. The size (FWHM) of a single field is illustrated by the open circle at the source field location. J1751+0939 is a particularly extreme case, the other sources are more typical.

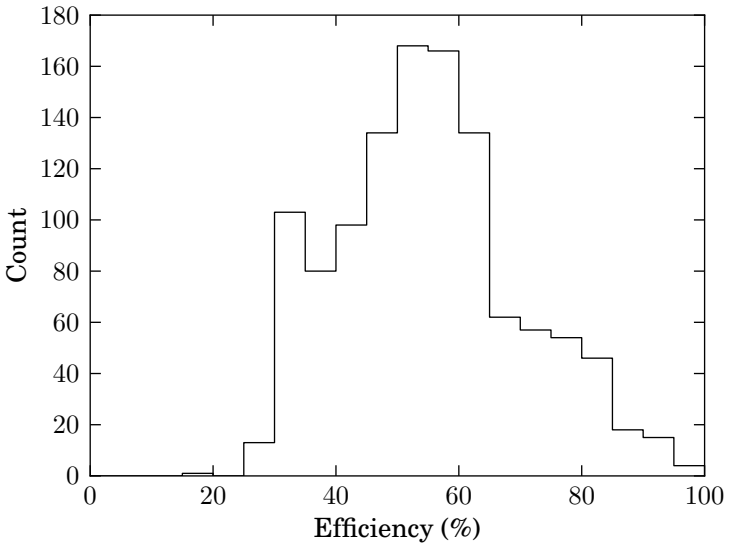

Figure 4.14. Histogram of per-source observing efficiency for CGRaBS sources. The expected number of observations per source was 365. Five sources exceeded 100% efficiency. The mean per-source efficiency was 55%.

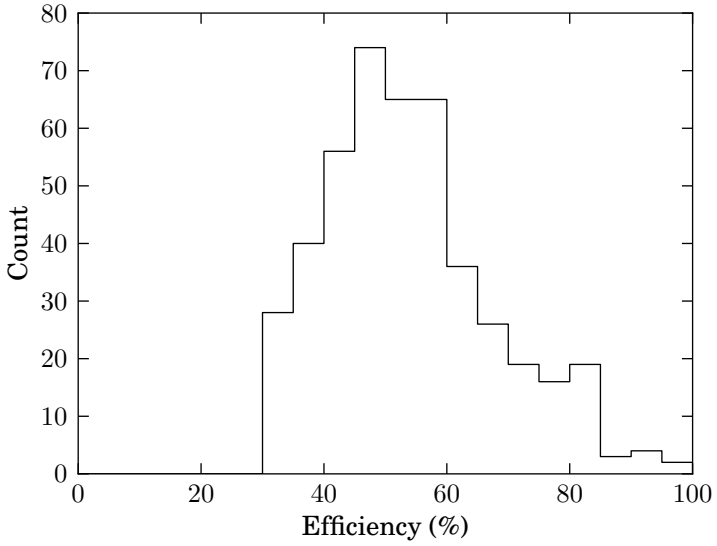

Figure 4.15. Histogram of per-source observing efficiency for 1LAC sources. The expected number of observations per source was 365 for CGRaBS sources. For non-CGRaBS sources, the expected number was computed using the time between the first successful observation and the end of the 42-month interval, assuming the nominal cadence of two observations per week. One source, CGRaBS J1321+2216, exceeded 100% efficiency. The mean per-source efficiency was 54%.

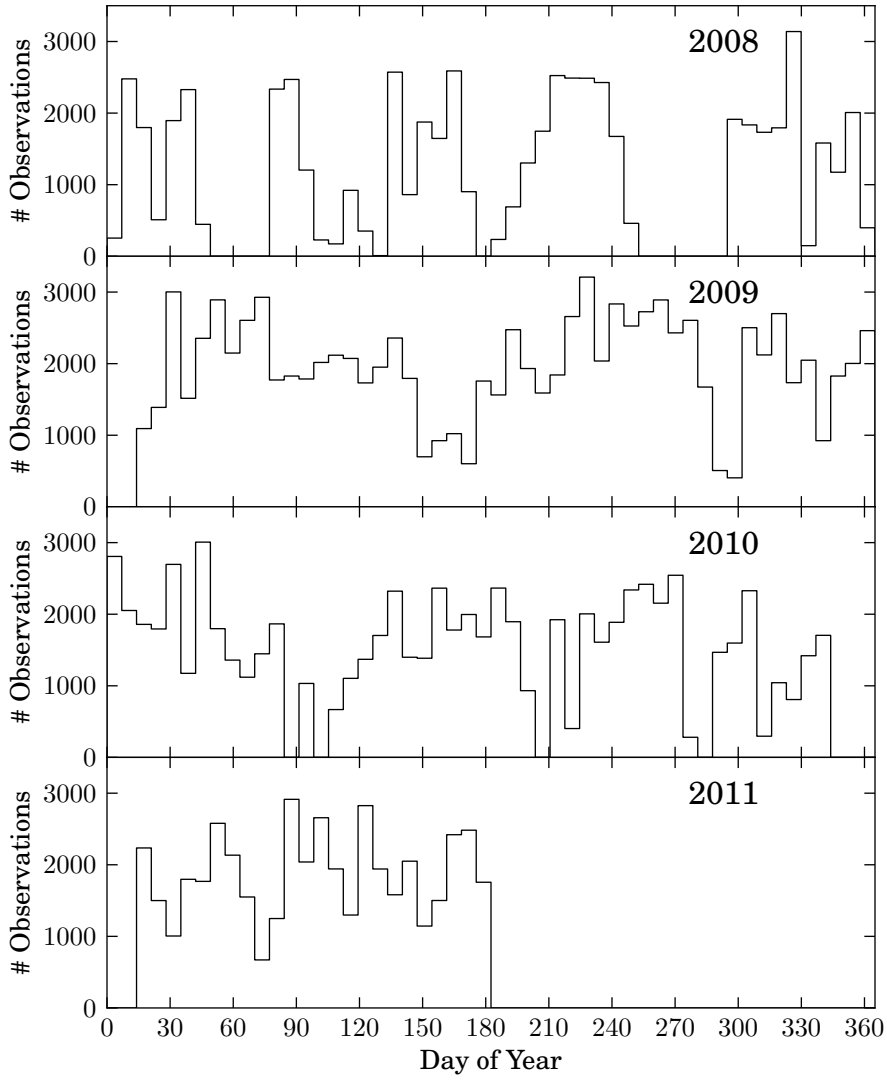

Figure 4.16. Weekly observation counts for each year of observations. Data plotted are the total numbers of flux density observations that survive to the end of the reduction pipeline.

occasional losses of observing time. Such events are inevitable during any long-term observing program, and one can only hope to minimize the impact through regular scheduled maintenance and planning.

In some cases, however, observing time has been lost due to preventable causes, and it is important to learn from these incidents to avoid them in the future. As we discussed in section 2.2.2.1, early in the program, the importance of measuring a pointing offset at a position near to the ensuing FLUX procedures was not recognized. This resulted in a complete loss of data for a few sources for several periods during 2008 and early 2009. The new scheduling algorithm prevents this from recurring. Another instance occurred in September 2008, when the programmable attenuator failed and no spare was available. Unfortunately, a long vendor lead time on a replacement component resulted in an outage lasting about 6 weeks. After this event, we have been careful to ensure that spares for critical system components are readily available at all times.

4.6.2 Flux Density Results

The distributions of per-source median flux densities for the CGRaBS and 1LAC samples are plotted in figures 4.17 and 4.18. These plots also show the distributions for each year of the 42-month data set, including only the 6 months of observations in the 2011 plot. Within both the CGRaBS and 1LAC samples, K-S tests comparing each individual year to the overall distribution and to each other year fail to reject the null hypotheses that all the flux density distributions are equivalent ($p > 0.88$ in all cases for the CGRaBS, $p > 0.99$ in all cases for 1LAC). A K-S test rejects the hypothesis that the CGRaBS and 1LAC flux distributions are drawn from the same distribution at the $p < 10^{-3}$ level. Comparisons of the flux density distributions between the CGRaBS and 1LAC samples are presented and discussed in section 5.6.1.

4.6.3 Future Prospects

This monitoring program is designed for long-term operation, and it is hoped that it will continue into the indefinite future. At the present twice-weekly cadence with almost 1600 sources, about one-half day in each three-day schedule cycle is available for scheduling more sources. Depending on the distribution in right ascension, an additional ~ 200 sources can probably be accommodated. Beyond that, it will be necessary to either eliminate some sources from the program (e.g., some of the CGRaBS sources not detected by *Fermi*, although this would complicate the study of differences between gamma-ray–loud and gamma-ray–quiet sources) or reduce the cadence for some or all sources. In any case, care must be taken to ensure that the sample retains a statistically meaningful definition.

The data reduction pipeline has been designed to facilitate frequent incremental data reduction and release, but this mode of operation has not yet seen much use. As regular data reductions and releases are made, changes and enhancements will be needed. For example, deciding when to begin new flux calibration or calibration spline epochs will be a challenge. In the present data set, these epochs were determined months after the change in calibration became evident. When reducing the data on a monthly or bimonthly schedule,

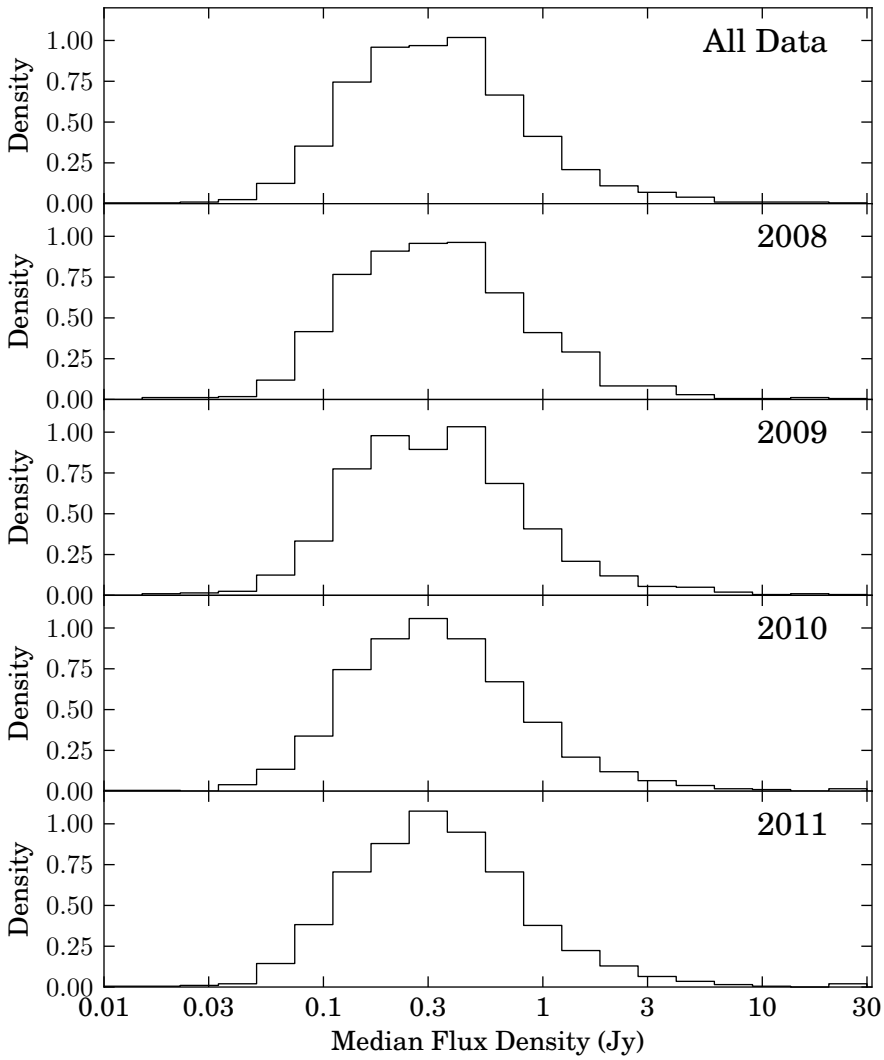

Figure 4.17. Histograms of the distributions of per-source median flux density in the CGRaBS sample for the entire 42-month data set and for individual years. Note that 2011 includes only 6 months of data. The visual similarity of the distributions is confirmed by K-S tests, which do not reject the hypotheses that any two distributions are equal ($p > 0.88$ in all permutations).

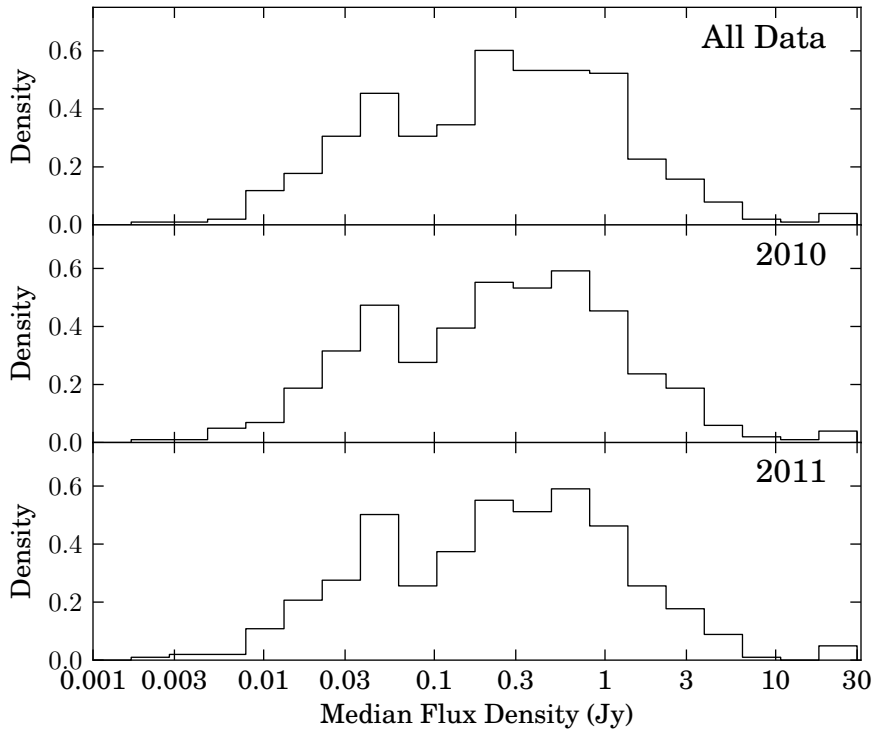

Figure 4.18. Histograms of the distributions of per-source median flux density in the 1LAC sample for the entire data set and for individual years. Note that 2011 includes only 6 months of data. K-S tests do not reject the hypotheses that these data are drawn from identical distributions ($p > 0.99$ in all permutations). The K-S test flatly rejects the hypothesis that the 1LAC sample flux densities are drawn from the same distribution as the CGRaBS flux densities ($p \ll 10^{-3}$).

it may be difficult to determine whether a new epoch is required based on only a few new data points per source. While these epochs can be revised at any time, if the data have been posted to the public website, adjusting already-released data values is undesirable (though perhaps necessary).

Finally, as we will discuss in section 5.1.2.2, even after the data filtering and editing described in section 3.2.1, unreliable data points survive into the final data set. There are certainly improvements to be made, either by better tuning the filter parameters or by devising new tests. Increased vigilance in monitoring observing conditions is also likely to help eliminate unreliable data. An increased frequency of data reductions will likely also help with this since events that impact the observing condition (e.g., thunder or snow storms) are easier to identify soon after they occur. Looking back beyond a few months to find causes of anomalous behavior can be extremely challenging. These, and surely other, enhancements to the observing problem should improve the quality of future monitoring data.

Chapter 5

Variability Analysis

With the light curve data from the monitoring program in hand, we now turn to analyzing the 15 GHz variability of each source. In this thesis, we will focus on the amplitude of that variability, treating the light curve as a population of samples drawn from a distribution defined by the processes responsible for the radio emission. For the most part, we will ignore the time coordinate in these studies. Other studies, both current and future, will examine the time dependence of these data more directly. Here, we will use a variability amplitude metric as a simple measure of source activity and aim to use this to identify a connection between gamma-ray emission and radio variability. Firmly and reliably establishing this connection is a prerequisite for performing more detailed correlation studies. The methods and the two-year CGRaBS results described in this chapter were first published in Richards et al. (2011). My colleague Dr. Vasiliki Pavlidou conceived of and developed computer codes to implement the likelihood analysis tools that we describe here. My contributions included early discussions of the standard variability methods that inspired this method, producing the light curve data for the tests, selecting the subpopulation samples for the comparisons, and the analysis and interpretation of the results of the tests.

In this chapter, we will first define the analytical tools we will use to characterize the variability amplitude of a light curve and to compare the variability of source populations. We will then present the results of a number of population studies, first demonstrating that our methods produce null results when given null inputs, then comparing the variability amplitudes of gamma-ray–loud sources with gamma-ray–quiet ones, of BL Lac objects with FSRQs, and of high-redshift with low-redshift FSRQs. We will then assess the impact of cosmological time dilation on this latter effect, and finally examine the differences between our CGRaBS and 1LAC source samples in detail.

We now have 18 months of data past the end of the set used for the Richards et al. (2011) results. These data, collected between January 2010 and June 2011, make a number of additional tests possible. First, we can verify that the properties of the various subpopulations we studied in the two-year results are stable as new data are added. Second, we now have enough data to characterize the variability of the *Fermi*-LAT–detected 1LAC sources, which were added to the OVRO sample in March 2010 (although some had been added earlier). Thus, we can now compare the 1LAC sample properties to the CGRaBS sample properties.

Finally, we have now sampled high-redshift ($z \sim 3$) sources for nearly a year of their rest frame time. In section 5.5, we will use these additional data to examine the effect of time dilation on the observed radio variability.

5.1 Analytical Tools

In this work, we are chiefly interested in characterizing and studying the *amplitude* of variability for a given blazar. This is a simple characterization of the behavior of a source that can readily be applied to the study of the collective properties of a large population. In this section, we begin by introducing the intrinsic modulation index, a variability amplitude metric that is particularly well suited to the data from our observing program. We will examine how this metric performs on our data set, then introduce a likelihood analysis framework for comparing the intrinsic modulation indices of subpopulations within our sample. We will then use these tools to explore the variability characteristics of the blazars in our program.

5.1.1 Intrinsic Modulation Index

Characterizing the variability amplitude of a source and assessing the confidence with which this can be measured are complex problems that have been addressed using a variety of measures and tests, such as the variability index (e.g., Aller et al. 1992); the fluctuation index (e.g., Aller et al. 2003); the modulation index (e.g., Kraus et al. 2003); the fractional variability amplitude (e.g., Edelson et al. 2002; Soldi et al. 2008); and χ^2 tests of a null hypothesis of nonvariability. Each of these tools provides different insights to the variability properties of sources and is sensitive to different uncertainties, biases, and systematic errors.

In this work, we will measure variability using the *intrinsic modulation index*, $\overline{m}$, which we introduced in Richards et al. (2011), where a full explanation of the likelihood analysis used to compute $\overline{m}$ is presented. Here we will give only a brief explanation of this variability measure and its properties. The intrinsic modulation index is based on the standard modulation index, defined as the standard deviation of the flux density measurements in units of the mean measured flux density, i.e.,

$$m_{\rm data} = \frac{\sqrt{\frac{1}{N}\sum_{i=1}^{N}\left(S_i - \frac{1}{N}\sum_{i=1}^{N} S_i\right)^2}}{\frac{1}{N}\sum_{i=1}^{N} S_i}\,. \tag{5.1}$$

The modulation index is reasonably well behaved: it is always nonnegative and is reasonably robust against outliers. However, it measures a convolution of intrinsic source variation and observational uncertainties—a large modulation index could be indicative of either a strongly variable source or a faint source with high uncertainties in individual flux density measurements. For this reason, the correct interpretation of the modulation index requires that measurement errors and the uncertainty in $m_{\rm data}$ due to the finite number of flux density measurements be properly estimated.

Our intrinsic modulation index is defined as

$$\overline{m} = \frac{\sigma_0}{S_0}\,, \tag{5.2}$$

and like the ordinary modulation index is defined by a standard deviation divided by the mean. In this case, however, σ_0 and S_0 represent *intrinsic* quantities—properties of the light curves before they are affected by observational noise, imperfect sampling, etc. Because we cannot directly measure these intrinsic quantities, we will use a likelihood analysis to estimate them from the data we collect. Observational uncertainties will affect the accuracy with which we can estimate these quantities, but we will quantify these uncertainties and propagate them into our later analysis as errors in our estimated values.

Evaluating the significance of a difference between measured values requires a good estimate of the uncertainty in those values. Thus, for the population comparisons we perform in this work, we require a rigorous estimate of the uncertainty in each intrinsic modulation index we calculate. Other methods for assessing the uncertainty in variability measures have been employed. One method that has been widely used is to evaluate each measure for a set of constant-flux-density calibrators which are known to have a flux density constant in time and which have been observed with the same instrument over the same periods of time. The value of the variability measure obtained for the calibrators is then used as a threshold value, so that any source with variability measure equal to or lower than that of the calibrators is considered consistent with being nonvariable. However, a variability measure value higher than that of the calibrators is a necessary but not sufficient condition for establishing variability. Calibrators are generally bright sources, with relative flux density measurement uncertainties lower than the majority of monitored sources; additionally, variability measures are affected by the sampling frequency, which is not necessarily the same for all monitored sources and the calibrators.

Alternatively, the *significance* of variability in a given source can be established through tests (such as a χ^2 test) evaluating the consistency of the obtained set of measurements with the hypothesis that the source was constant over the observation interval. However, such tests provide very little information on sources for which statistically significant variability cannot be established, as they cannot distinguish between *intrinsically nonvariable sources* and sources that could be *intrinsically variable* but inadequately observed for their variability to be revealed.

For our studies with the intrinsic modulation index, we will directly and rigorously propagate the measurement uncertainty for each flux density into an uncertainty on $\overline{m}$. In this way we will estimate both a measure of the intrinsic variability amplitude and our uncertainty in that estimate due to the measurement process.

5.1.1.1 Calculating the Intrinsic Modulation Index

As in Richards et al. (2011), we will assume that the "true" flux densities for each source are normally distributed with mean S_0, standard deviation σ_0, and intrinsic modulation index $\overline{m} = \sigma_0/S_0$. That is, we assume that the probability density for the "true" flux density S_t is

$$p(S_t, S_0, \sigma_0) = \frac{1}{\sigma_0\sqrt{2\pi}} \exp\left[-\frac{(S_t - S_0)^2}{2\sigma_0^2}\right]. \tag{5.3}$$

Similarly, we assume the observation process for the $j\mathrm{th}$ data point adds normally distributed error with mean S_t and standard deviation σ_j. Then, the likelihood for a single observation is given by

$$\ell_j = \int_{\mathrm{all\ S_t}} dS_t \frac{\exp\left[-\frac{(S_t-S_j)^2}{2\sigma_j^2}\right]}{\sigma_j\sqrt{2\pi}} \frac{\exp\left[-\frac{(S_t-S_0)^2}{2\sigma_0^2}\right]}{\sigma_0\sqrt{2\pi}}, \tag{5.4}$$

which after combining $j = 1, \ldots, N$ measurements and substituting $\overline{m}S_0 = \sigma_0$, gives

$$\begin{aligned}\mathcal{L}(S_0, \overline{m}) &= S_0 \left(\prod_{j=1}^{N} \frac{1}{\sqrt{2\pi(\overline{m}^2 S_0^2 + \sigma_j^2)}}\right) \times \\ &\quad \exp\left[-\frac{1}{2}\sum_{j=1}^{N} \frac{(S_j - S_0)^2}{\sigma_j^2 + \overline{m}^2 S_0^2}\right].\end{aligned} \tag{5.5}$$

By maximizing the likelihood given by equation (5.5), we find our estimates of S_0 and $\overline{m}$. In figure 5.1, we plot the most likely values and the 1σ, 2σ, and 3σ isolikelihood contours for the two-year data for CGRaBS source J1243–0218. These contours were computed to contain 68.26%, 95.45%, and 99.73% of the volume beneath the likelihood surface.

We finally obtain our estimate of $\overline{m}$ by marginalizing over all values of S_0, which yields a one-dimensional likelihood distribution, $\mathcal{L}(\overline{m})$, like that shown in figure 5.2. We compute the 1σ uncertainty by finding equal-likelihood points $\overline{m}_1$ and $\overline{m}_2$ on either side of the maximum likelihood value where

$$\frac{\int_{\overline{m}_1}^{\overline{m}_2} \mathcal{L}(\overline{m}) d\overline{m}}{\int_0^\infty \mathcal{L}(\overline{m}) d\overline{m}} = 0.6826. \tag{5.6}$$

When the maximum likelihood value $\overline{m}$ is less than 3σ from zero, we compute a 3σ upper limit $\overline{m}_3$ defined by

$$\frac{\int_0^{\overline{m}_3} \mathcal{L}(\overline{m}) d\overline{m}}{\int_0^\infty \mathcal{L}(\overline{m}) d\overline{m}} = 0.9973. \tag{5.7}$$

This method makes the assumption that the distribution of flux densities from a source are distributed normally. For many sources, this is a good description of the data. In figure 5.3, we plot the histogram of the two-year data set flux densities from J1243–0218 along with the maximum-likelihood Gaussian model we estimated from this analysis. In other sources, however, the distributions are clearly non-Gaussian, with

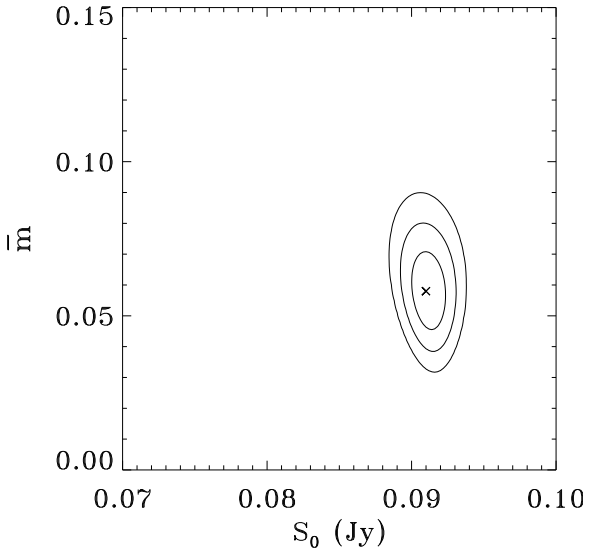

Figure 5.1. Likelihood parameter space, showing 1σ, 2σ, and 3σ contours of the joint likelihood $\mathcal{L}(S_0, \overline{m})$ for blazar J1243−0218.

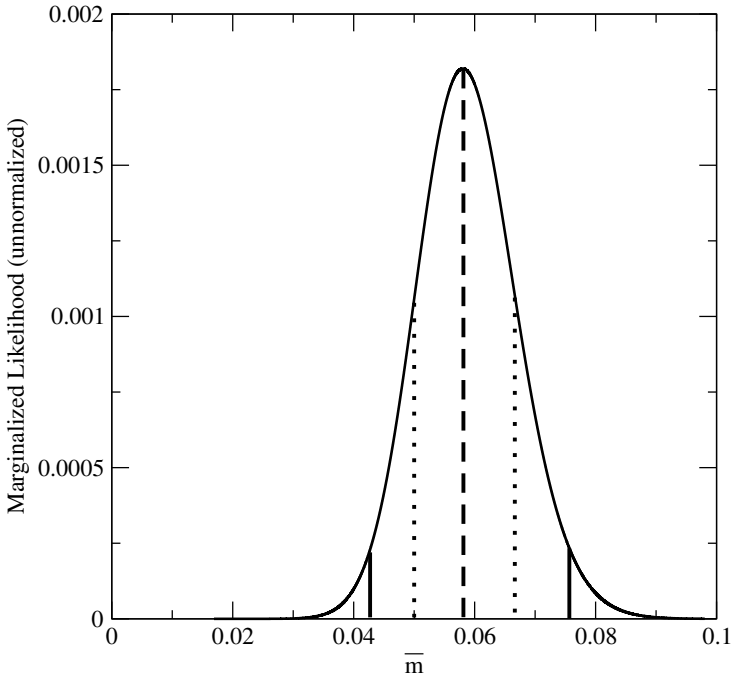

Figure 5.2. Marginalized likelihood $\mathcal{L}(m)$ for J1243−0218 (solid curve). Dashed vertical line: best-estimate $\overline{m}$; dotted vertical lines: 1σ $\overline{m}$ range; solid vertical lines: 2σ $\overline{m}$ range.

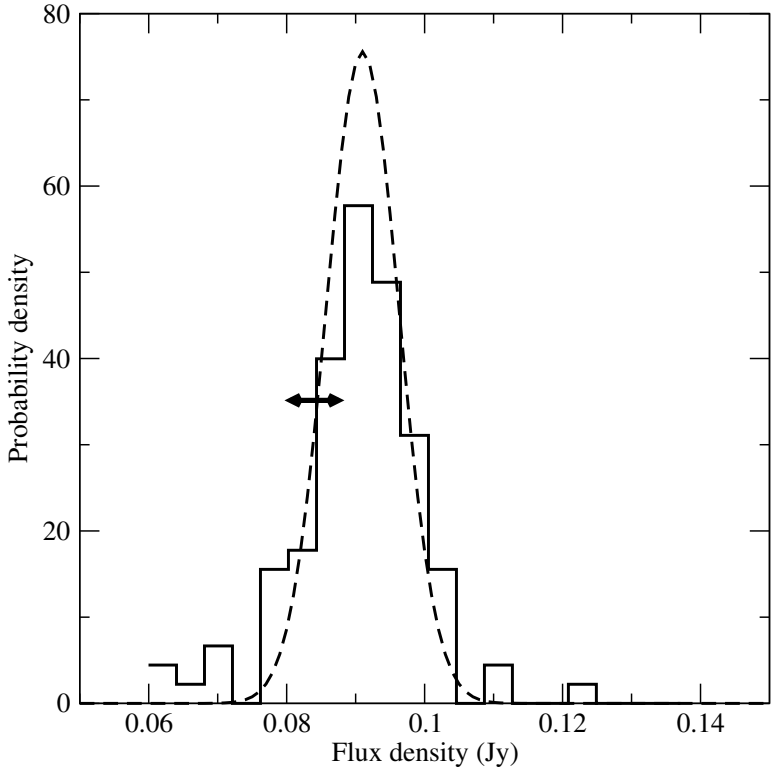

Figure 5.3. Maximum-likelihood Gaussian model for the flux density distribution (dashed line), plotted over the histogram of measured flux densities (solid line) for blazar J1243−0218. The arrow indicates the size of the typical measurement uncertainty.

possible evidence for, e.g., bimodality as shown in figure 5.4. However, we have verified that even when the true distribution is not well described as Gaussian, the modulation index and uncertainty is a reasonable description of the data. In future work, more sophisticated distributions of true flux density can be applied to this method.

5.1.2 Properties of the Intrinsic Modulation Index

We now evaluate the properties of the intrinsic modulation index using the two-year data set. Table C.1 in appendix C includes our measured values for $\overline{m}$, S_0, and their 1σ errors for the 42-month data set. Results from the two-year data are available in Richards et al. (2011). In figure 5.5 we plot the intrinsic modulation index $\overline{m}$ and associated 1σ uncertainty against the intrinsic, maximum-likelihood average flux density, S_0, for all our CGRaBS and calibrator sources. The error bar on S_0 corresponds to the 1σ uncertainty in mean flux density, calculated from the joint likelihood (equation (5.5)) marginalized over $\overline{m}$. CGRaBS sources are shown as black or magenta points or blue triangles for upper limits, while calibrators are shown as green points.

Variability could only be established at the 3σ confidence level or higher for 1139 out of 1158 CGRaBS blazars in our sample. For this study, we considered only sources for which at least 3 flux densities were mea-

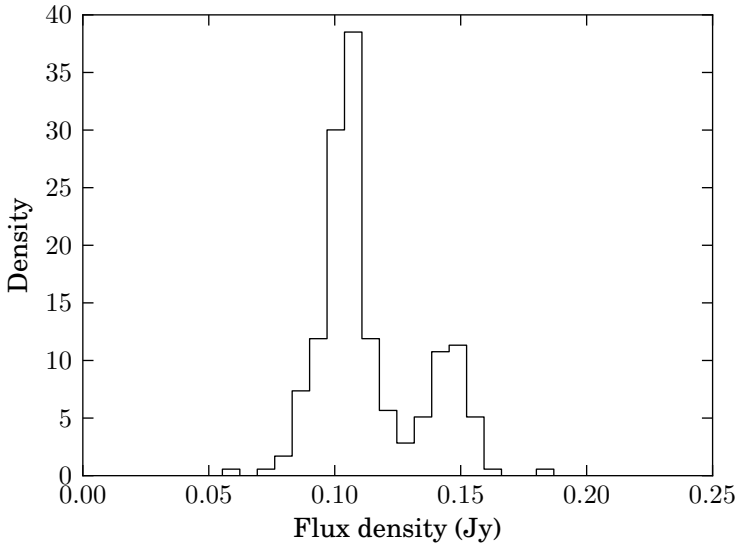

Figure 5.4. Histogram of flux density measurements for J0237+3022, demonstrating a non-Gaussian bimodal distribution.

sured, a positive mean flux density $\geq 2\sigma$ from zero was found, and at least 90% of the individual flux density measurements were $\geq 2\sigma$ from zero. These criteria excluded two sources (J1310+3233 and J1436–1846). For the other 17 sources we have calculated 3σ upper limits for $\overline{m}$. We plot these upper limits with blue triangles.

Calibration sources 3C 286, DR 21, and 3C 274 are shown in green. Although, as expected, these sources are the least variable of all sources in which variability can be established and a nonzero $\overline{m}$ can be measured, $\overline{m}$ for these sources is finite and measurable. This means that some residual long-term variability remains in our calibrators beyond what can be justified by statistical errors alone. This could conceivably result from true calibrator source variation, but more likely reflects incomplete removal of small-amplitude calibration trends. Because $\overline{m} < 1\%$ for these three sources, we quote a systematic uncertainty $\Delta\overline{m}_{\rm syst} = 0.01$ for the values of the intrinsic modulation index we produce through our analysis.

To ensure that our population studies are not affected by this residual systematic variability, in all analyses discussed in sections 5.2–5.5 only sources with $\overline{m} \geq 0.02$ will be used, so that we remain comfortably above this 1% systematic uncertainty limit. In addition, for sources with $S_0 < 60$ mJy, the number of sources for which variability can be established is of the same order as the number of sources (both CGRaBS and non-CGRaBS) for which we could only measure an upper limit, and these upper limits are very weak and nonconstraining. For this reason, we also exclude from our population studies any source with $S_0 < 60$ mJy. The part of the parameter space excluded due to these two criteria is shown in figure 5.5 as the yellow shaded area bounded by the solid black lines.

For $S_0 \geq 0.4$ Jy, no obvious correlation between flux density and modulation index is apparent, and no CGRaBS sources exist with upper limits above our cut of $\overline{m} = 2\%$. However, for sources with $S_0 < 0.4$ Jy, there is an absence of points in the lower-left corner of the allowed parameter space defined by the thick solid

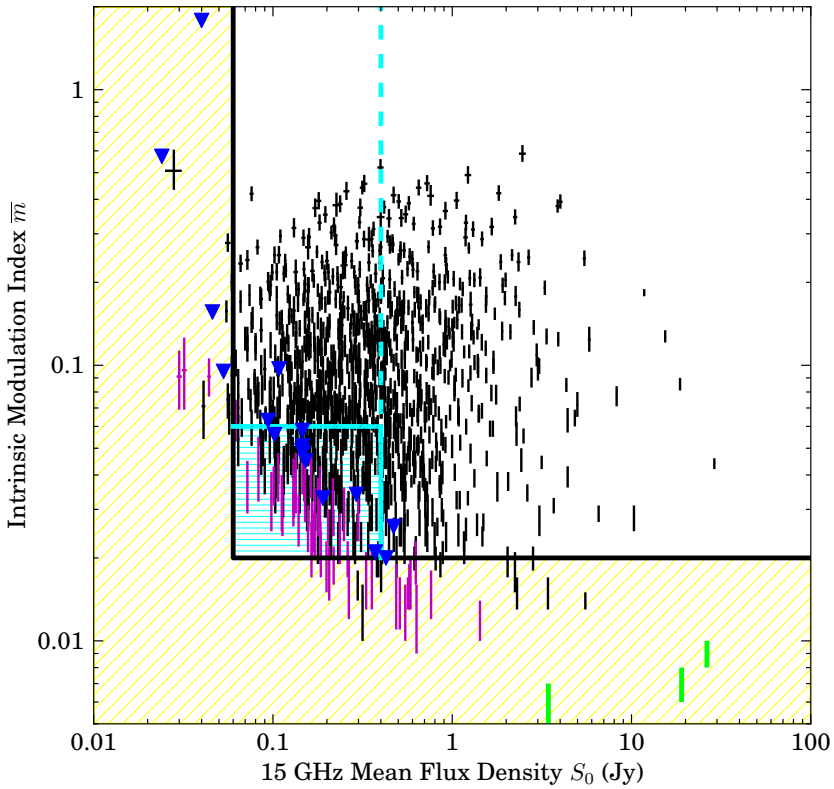

Figure 5.5. Two-year intrinsic modulation index $\overline{m}$ and associated 1σ uncertainty, plotted against intrinsic maximum-likelihood average flux density, S_0, for all sources in the program which have enough (more than 3) acceptable, nonnegative flux density measurements. Black points: CGRaBS sources found to be variable with 3σ confidence by χ^2 test; magenta points: CGRaBS sources found consistent with nonvariable by χ^2 test; green points: calibrators 3C 286, DR 21, and 3C 274; blue triangles: 3σ upper limits for CGRaBS sources for which variability could not be established at the $\geq 3\sigma$ confidence level. The error bar on S_0 corresponds to the 1σ uncertainty in mean flux density, calculated from the joint likelihood (equation (5.5)) marginalized over $\overline{m}$. Data, except for upper limits, outside the yellow and cyan shaded areas are used in the population studies of sections 5.2–5.5.

lines: for faint sources, we can only confidently establish variability if that variability is strong enough. The effect disappears for variability amplitudes greater than about 6%. In addition, there are only two CGRaBS sources with upper limits higher than 6% for sources brighter than 60 mJy (J0722+3722 and J0807+5117), $<0.5\%$ of the 452 sources measured in this region of parameter space. We conclude that we are able to measure variability at the level of 6% or higher for virtually all ($>99\%$ of) sources brighter than 60 mJy.

To ensure that our population studies are not affected by our decreased efficiency in measuring variability in sources with 60 mJy $\leq S_0 < 0.4$ Jy and $2\% \leq \overline{m} < 6\%$, we will also exclude this part of the $(S_0, \overline{m})$ parameter space from our analyses in sections 5.2–5.5. The part of the parameter space excluded due to these criteria is shown in figure 5.5 as the cyan shaded area.

For comparison, we also computed the χ^2-per degrees of freedom for each source and tested whether we could reject the hypothesis of a constant flux density at the 3σ level. Because of the long-term residual trend described in section 3.2.2.2, we added 1% of each flux density in quadrature to the reported uncertainty when computing χ^2. Of the 1139 CGRaBS sources for which we calculated the intrinsic modulation index, 51 (4.5%) are found to be nonvariable (i.e., we cannot reject the hypothesis of constant flux) with $>3\sigma$ confidence. These are plotted as magenta points in figure 5.5. All but one of these lie within the low-flux density and low-variability regions we have excluded from our population studies. The one such source not excluded, J2148+0211, is very near both the flux density and intrinsic modulation index cut lines. Of the 17 sources for which we report $\overline{m}$ upper limits, 15 are judged nonvariable by the χ^2 test. The two others are dim sources with a single outlier (J1613+4223) or very few measurements (J1954+6153), which led to large uncertainties in the estimate for $\overline{m}$ and resulted in an upper limit. Calibrator sources 3C 286, DR 21, and 3C 274 are found to be nonvariable while 3C 48 and 3C 161 (which were not used to fit the long-term calibration trend) are found to be variable by the χ^2 test, probably due to imperfect removal of the long-term calibration trend. Our estimates of $\overline{m}$ for both these calibrators are below our 2% intrinsic modulation cut level. We conclude that our analysis is generally consistent with the χ^2 test for identifying significant variability and that our data cuts for our population studies conservatively exclude the regions of parameter space where disagreements occur.

In figure 5.6 we plot the intrinsic modulation index $\overline{m}$ and associated 1σ uncertainty against the "raw" modulation index $m_{\rm data}$ of equation (5.1). The $\overline{m} = m_{\rm data}$ line is shown in blue. Green triangles are the 3σ upper limits of sources for which variability could not be established. Calibrators 3C 286, DR 21, and 3C 274 are plotted in red. Since apparent variability due to the finite accuracy with which individual flux densities can be measured has been corrected out of $\overline{m}$, the expectation is that deviations from the $\overline{m} = m_{\rm data}$ will be more pronounced for sources that are not intrinsically very variable (so that the scatter in the flux density measurements is appreciably affected, and even dominated, by measurement error). In addition, deviations are expected to be below the line, as $\overline{m}$ should be *smaller* than $m_{\rm data}$. Both these expectations are verified by figure 5.5. Note that upper limits need not satisfy this criterion, as the "true" value of the modulation index can take any value below the limit. Upper limits above the blue line are weak, indicating that the reason

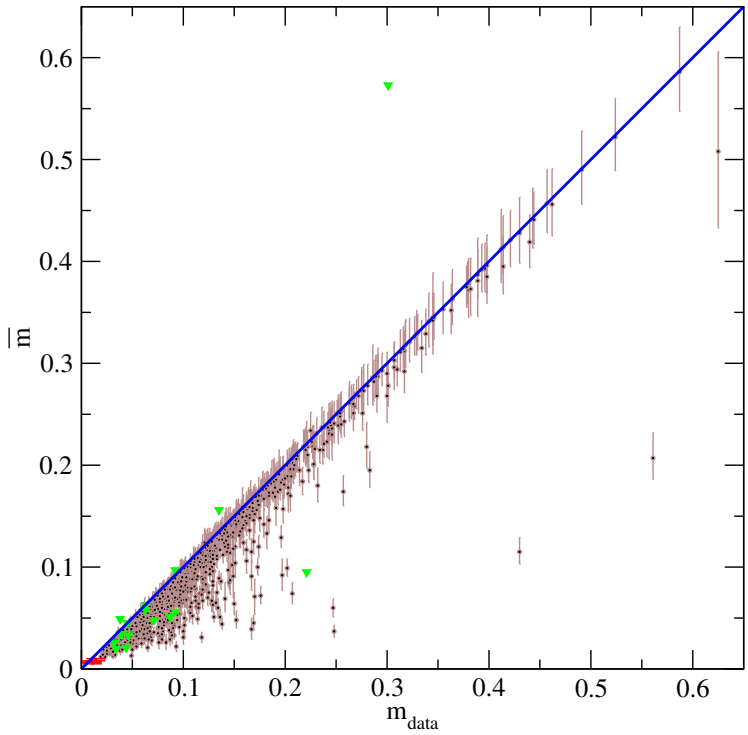

Figure 5.6. Two-year intrinsic modulation index $\overline{m}$ and associated 1σ uncertainty, plotted against the "raw" modulation index, $m_{\rm data}$, of equation (5.1) as black points with brown error bars. The $\overline{m} = m_{\rm data}$ line is shown in blue. Green triangles are the 3σ upper limits of sources for which variability could not be established. Calibrators 3C 286, 3C 274, and DR 21 are plotted in red.

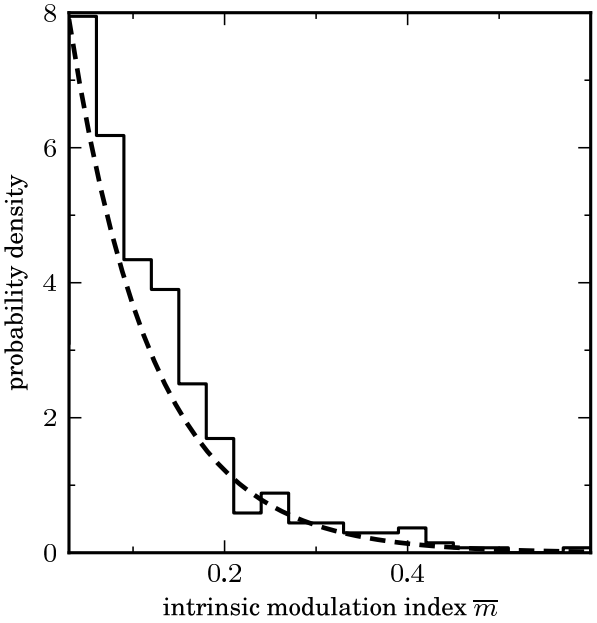

Figure 5.7. Histogram of two-year maximum-likelihood intrinsic modulation indices $\overline{m}$, for the 453 CGRaBS blazars with $S_0 > 400$ mJy, normalized as a probability density that integrates to unity. The dashed line represents an exponential distribution with $\langle \overline{m} \rangle = 0.091$.

variability could not be established is the poor sampling or quality of the data, and not necessarily a low intrinsic variation in the source flux density.

For the 453 CGRaBS objects which have $S_0 > 400$ mJy and for which variability can be established, we plot, in figure 5.7, a histogram of their intrinsic modulation indices $\overline{m}$ normalized so that the vertical axis has units of probability density. The dashed line represents an exponential distribution of mean $\langle \overline{m} \rangle = 0.091$ which, as we can see, is an excellent description of the data. Motivated by this plot, we will be using the monoparametric exponential family of distributions:

$$f(m)dm = \frac{1}{m_0} \exp\left[-\frac{m}{m_0}\right] dm, \tag{5.8}$$

with mean m_0 and variance m_0^2, to characterize various subsamples of our blazar sample.

5.1.2.1 Impact of Longer Time Series

We now compare the results of the intrinsic modulation indices computed from the 42-month data with those computed from the two-year data to look for systematic changes in apparent variability amplitude. The expectation based on the two-year data set was that additional data would tend to increase the variability amplitude on average, since many sources appeared to switch between periods of relatively steady quiescence and periods of active variability. This is because the addition of a a period of steady flux to a source with a

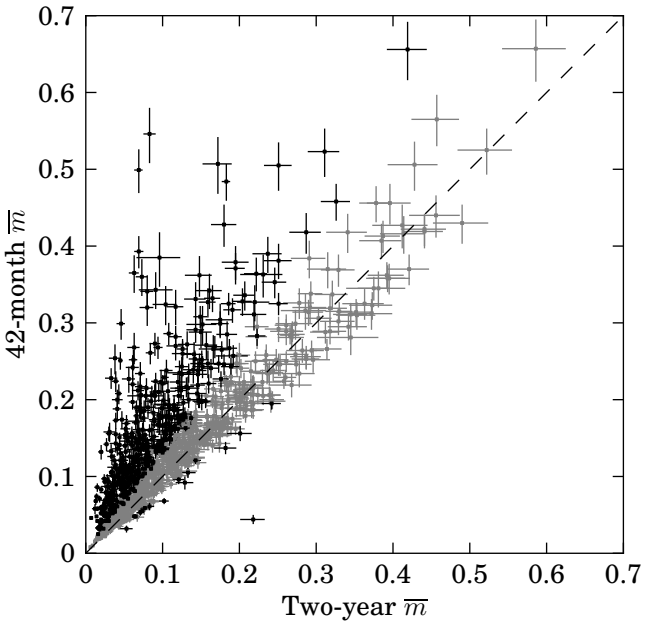

Figure 5.8. Scatter plot of 42-month versus two-year modulation indices for 1135 CGRaBS sources and five calibrators. The single outlier well below the dashed 1:1 reference line is J1154+1225. This source was affected by a single extreme high outlier in the two-year data set which was removed in the 42-month analysis. Sources for which the difference in intrinsic modulation index was less than 3σ are plotted in grey.

history of strong variability will reduce its intrinsic modulation only slightly. On the other hand, a source that has only been observed in a weakly variable state will see a large increase in its intrinsic modulation index if it begins to vary strongly.

The data confirm our expectations. In figure 5.8 we plot the 42-month $\overline{m}$ values against the two-year $\overline{m}$ values for the 1135 CGRaBS sources with measured $\overline{m}$ in both data sets, plus 3C 48, 3C 161, 3C 274, 3C 286, and DR 21. Figure 5.9 shows histograms of the two data sets. Clearly, most sources have 42-month intrinsic modulation indices that are either consistent with or greater than their two-year values. Of the 1140 sources compared, 513 changed by more than 3σ. These are plotted in black in figure 5.8. One significant exception is J1154+1225, a CGRaBS source. In the two-year light curve for this source, a single very large high outlier survived the data cuts. This outlier was eliminated in the 42-month data set, resulting in this large reduction in $\overline{m}$. The actual behavior of this source does not appear to have changed substantially.

This systematic increase in the variability index suggests that the two year interval was insufficiently long to capture the full range of behaviors in many CGRaBS sources. This is not surprising. Based on more than 25 years of monitoring at the University of Michigan Radio Observatory and the Metsähovi Radio Observatory in Finland, Hovatta et al. (2008) report typical flare durations of 2.5 years at 22 and 37 GHz. In Hovatta et al. (2007), typical flaring timescales of 4–6 years, sometimes with evidence for changes on 10 year or longer timescales, were reported. In gamma rays, Abdo et al. (2010b) found higher peak-to-mean

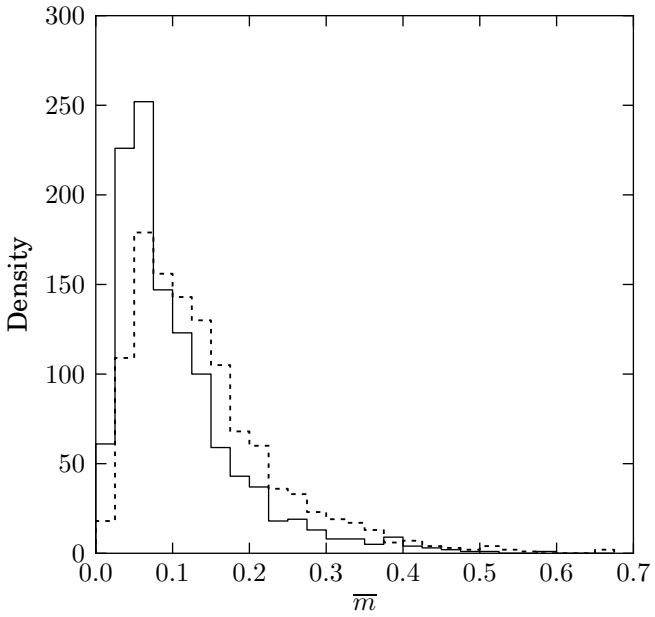

Figure 5.9. Histogram of the intrinsic modulation indices for the two-year (solid line) and 42-month (dashed line) data sets for 1135 CGRaBS sources and five calibrators.

flux ratios among *Fermi* sources compared to those reported by EGRET. It was speculated that this difference could result from the longer EGRET data collection period (4.5 years versus 11 months) permitting sources to visit more emission states. Thus, our results are consistent with these explicit timescale studies and suggest that even with the 42 month data set it is likely somewhat premature to expect our results to characterize the full behavior of the entire sample. However, unless there is, in fact, a connection between radio variability and gamma-ray emission, underestimates of typical variability amplitudes should affect different subpopulations equally. This may diminish our power to compare populations, but it should not lead us to conclude that false correlations exist.

Figure 5.10 shows a histogram of the change in intrinsic modulation index between the two-year and 42-month data sets for each source in the CGRaBS sample. The mean (median) change is 0.035 (0.021). In figure 5.11 we show the light curves of the four sources that showed the most significant change in intrinsic modulation index between the two data intervals. Not surprisingly, these sources all show steady flux density prior to 2010, then either gradually increase or decrease in brightness or exhibit a sharp change in behavior in the more recent data.

5.1.2.2 Impact of Data Outliers

The 40 m data set is not completely free of outlier points that are unlikely to represent intrinsic source variations since the filters described in chapter 3 do not detect all such points. Although most of these outliers are almost surely attributable to poor observing conditions, pointing offset measurement failures, and

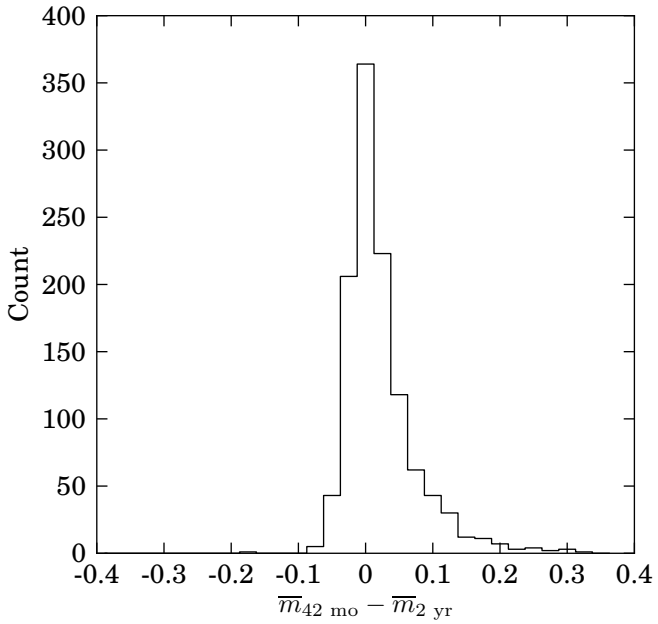

Figure 5.10. Change in intrinsic modulation indices between the two-year and 42-month data sets for 1135 CGRaBS sources and five calibrators.

other causes not related to actual behavior of the astronomical source, we do not delete a data point merely based on its apparently improbable flux density value. Unless we have an unbiased criterion by which we can eliminate the data point, we risk seriously biasing our results by rejecting source variability that does not meet our preconceptions of "reasonable." In many cases, detailed exploration of these data points has led to the discovery of such independent criteria—this is how we developed the set of data filters we employ—but at present, a fair number of these suspected unphysical outliers remain. Perhaps future work will improve the filtering and remove these, but for now we choose to study the impact of such faulty data on our analysis to ensure that our results are robust against their effect.

The most common extreme outliers we encounter are zero or near-zero (in some cases slightly negative) flux density values reported for bright sources. These probably result when the telescope obtains an incorrect pointing offset and measures blank or contaminated sky instead of the desired source. To measure the effect of such outliers, we computed the intrinsic modulation indices for each source using the 42-month data set with the addition of a single flux density value that was twice the average error above zero. Although in some cases light curves are affected by multiple outliers, in general the impact of the first such outlier is much greater than the addition of the second or third—once the probability for such an event is nonzero, the change in likelihood due to a few additional incidents is small. This was verified by comparing the modulation indices of the data set with one false outlier added to the same data with a second false outlier added. The mean change in modulation index was an increase by only 0.001, which is negligible compared to the uncertainties. We therefore only need to consider the impact of a single outlier on the intrinsic modulation index.

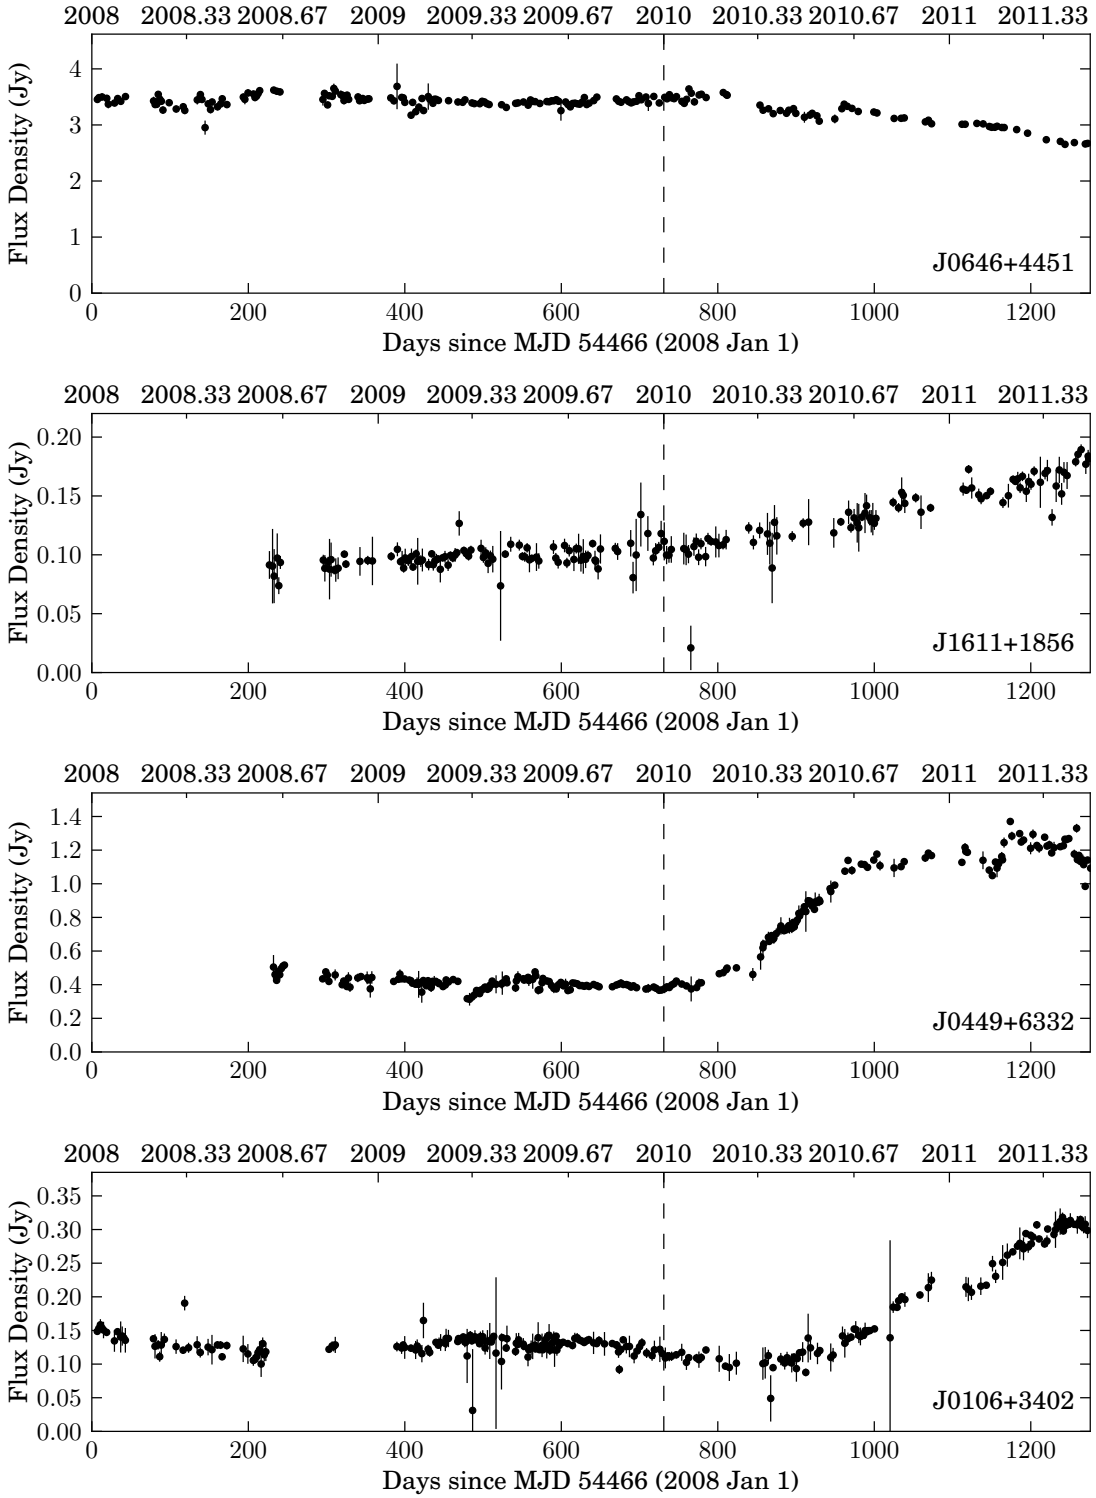

Figure 5.11. Light curves for the four sources with the most significant changes in intrinsic modulation index between the two-year and 42-month data sets. All four sources exhibit an increase in variability after 2010, the end of the period included in the two-year data set, indicated by the dashed line.

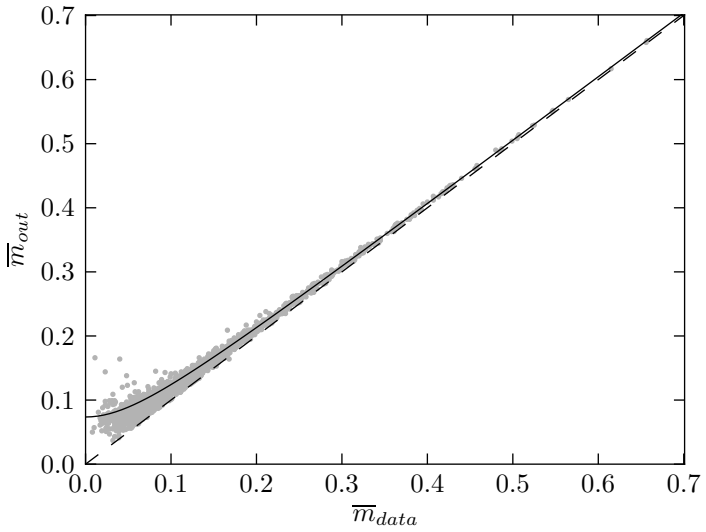

Figure 5.12. Grey points show the modulation index data computed with the addition of an extreme outlier data point plotted against the modulation index for the same source calculated from the actual data. The dashed line shows the ideal $y = x$ line. The solid line shows the effect of adding 0.066 in quadrature with the measured modulation index.

In figure 5.12, the grey points show the computed intrinsic modulation index using the data with the false outlier versus the actual intrinsic modulation index for the source. In figure 5.13, the histograms of these data sets are shown. The obvious visual difference is confirmed by a two-sample K-S test, which rejects the hypothesis of a common parent distribution at the $p \sim 10^{-20}$ level for arising by chance. The mean of the distribution for the real data is 0.143 and for the modified set is 0.160.

If the outliers induced a simple fixed increase in the modulation index for each source, to match the means this would be an increase by 0.018. Applying this increase to the real data, the two sample K-S test rejects the null hypothesis with $p = 6.9 \times 10^{-4}$. A fixed modulation increase in quadrature with the measured value would require an increase of 0.074 to match the means of the two data sets. The two-sample K-S test rejects this possibility with $p = 5.0 \times 10^{-7}$. Apparently, and not surprisingly, the impact of an extreme outlier depends on the properties of the light curve to which it is added.

First, we examine the change in intrinsic modulation index for a trend with the value measured for each source. Figure 5.14 shows the quadrature difference between the intrinsic modulation index with and without the false outlier point plotted against the actual measured value for each source. Although there is significant scatter, it appears that the binned mean of the change is well described by a constant addition of 0.066 in quadrature with the actual measured intrinsic modulation index. A K-S test using this constant quadrature increase is marginally compatible, with $p = 0.047$ to reject the hypothesis of a common distribution. The effect of adding this constant additional error is also shown in figure 5.12 by the solid line.

In figure 5.15 we show the same quadrature difference, now as a function of the maximum likelihood average source flux density. A linear fit to binned mean square of the quadrature difference as a function of

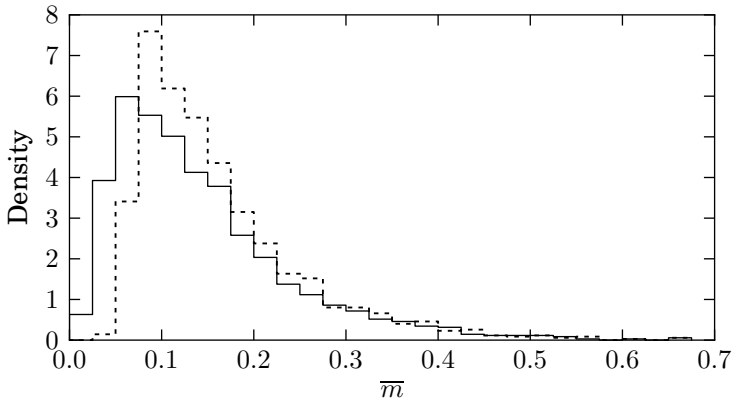

Figure 5.13. Histograms of intrinsic modulation indices for all sources using the 42-month data (solid line) and the 42-month data with an additional extreme outlier added to each light curve (dotted line). Includes 1396 sources in each, excluding those for which only upper limits were calculated. Both curves are normalized to integrate to unity.

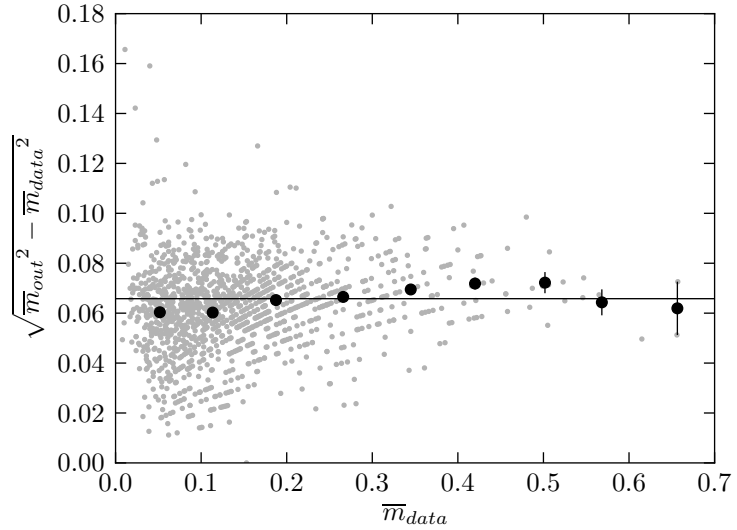

Figure 5.14. Grey points show the quadrature difference between the modulation indices calculated with the addition of an extreme outlier data point and the modulation indices computed from the actual data, plotted against the actual modulation index for the source. The black points show the binned mean difference in modulation index. The line is the constant mean quadrature difference, 0.066. The striping evident in the grey points is the result of rounding the modulation index values before computing the quadrature difference.

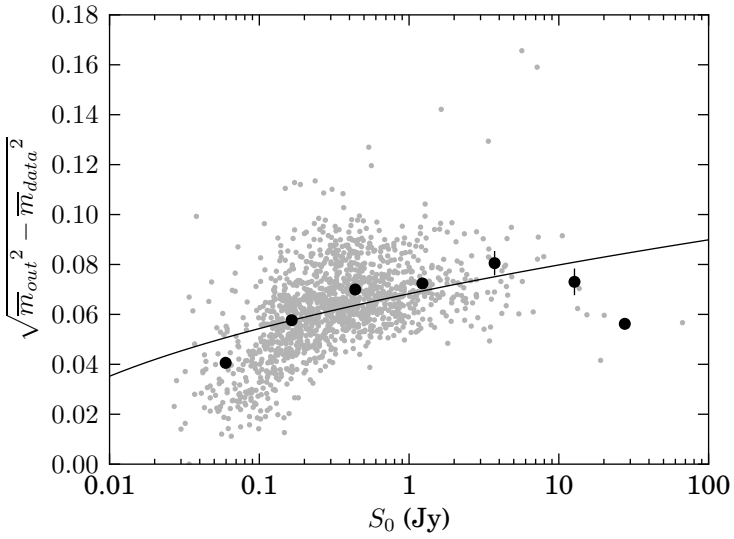

Figure 5.15. Grey points show the quadrature difference between the modulation indices calculated with the addition of an extreme outlier data point and the modulation indices computed from the actual data, plotted against the maximum likelihood average flux density for the source. The black points show the binned mean difference in modulation index using logarithmic bins in flux density. The line is a linear fit to the square of the binned y data as a function of the logarithm of the x data, excluding the last data point.

the logarithm of the flux density is shown by the solid line, which corresponds to

$$\Delta(\overline{m}^2) = (0.041)^2 \log_{10}\left(\frac{S_0}{1\text{ Jy}}\right) + (0.068)^2. \tag{5.9}$$

In this fit, we have excluded the last binned point because only a few data are in the bin. The lowest flux density bin lies significantly below the trend, suggesting that sources with mean flux densities below about 100 mJy were less affected by the addition of an outlier.

Although a trend with source flux density is apparent, the systematic quadrature change in intrinsic modulation index induced by the addition of an outlier varies only between about 0.04 and 0.08. Thus, when comparing two populations, unless the flux density distributions differ widely, this trend is unlikely to distort the results of our population studies.

It is difficult to define a robust criterion for identifying an outlier in a data set which includes sources that exhibit extreme actual variability. To estimate the number of such outliers, we choose a subset of the sources for which it is relatively easy to detect an outlier. Since the processes that result in a nonphysical outlier are connected to the instrument or observation conditions rather than astronomical effects, we can safely assume that the fraction of affected sources in such a subset is representative of the entire sample. We therefore examine bright sources with median flux density $\mathrm{median}(S) \geq S_{min}$, where low outliers are more prominent. To avoid extremely variable sources, we further restrict this study to sources with most measurements near the median. We require that at least a fraction R_{stable} of the flux densities for the source

Table 5.1. Fraction of sources determined to be affected by an extreme low outlier

Test Parameters				Test Results			
S_{min} (Jy)	$R_{central}$	R_{stable}	$R_{outlier}$	N_{total}	N_{stable}	$N_{outlier}$	Outlier Fraction (%)
1.0	0.5	0.5	0.3	163	162	8	4.9
1.0	0.5	0.5	0.2	163	162	7	4.3
0.5	0.5	0.5	0.3	408	407	34	8.4
0.5	0.8	0.5	0.3	408	407	34	8.4
0.5	0.5	0.8	0.3	408	387	30	7.8
0.5	0.5	0.8	0.2	408	387	22	5.7

Note: The estimates do not vary rapidly for small changes in test parameters. We adopt an approximate outlier fraction of 8% for our population, meaning that about 8% of sources are affected by at least one significant outlier.

be within a $\pm R_{central} \times \mathrm{median}(S)$ of the median. We then count the source as having an outlier if any flux density in its light curve is below $R_{outlier} \times \mathrm{median}(S)$. In table 5.1 we tabulate the results of this test for various values of the parameters. We conclude that no more than about 8% of sources are affected by extreme outliers.

It is extremely unlikely that the incidence of such outliers is tied to physical properties of the sources being observed, so we can reasonably assume that about 8% of the sources in any of the subsamples we select for population studies are affected by outliers. We conclude from these studies that the net effect of outliers in the data set is to add a false modulation of about 0.066 in quadrature with the intrinsic modulation index that would be measured in the absence of such outliers. This increase will affect about 8% of the sources in any physically selected sample. When comparing samples, these conclusions are valid as long as the flux density distributions of the two samples have similar dynamic ranges. If one sample is substantially different in flux density, particularly if it is clustered below about 100 mJy, then the inclusion of outliers in the data set may affect the samples differently.

More sources are likely to be affected by random outliers in the 42-month data set than the two-year data set, simply because each additional observation provides an opportunity for an outlier. Assuming the outlier incidence rate is constant, a rough estimate of the number of sources affected in the two-year data set is $8\% \times 24/42 = 5\%$. In other words, 3% fewer sources are likely to be affected in the two-year data set. For simplicity, if we assume the 0.066 average modulation index increase were added linearly, the additional 3% of affected sources would be expected to add about $0.03 \times 0.066 = 0.002$ to the mean modulation index of the population. Because the addition is in quadrature, this is an overestimate of the likely impact. The actual mean increase of 0.035 discussed in section 5.1.2.1, therefore, cannot be explained by the additional exposure to outliers and can be safely attributed to changes in the observed source behavior.

5.1.3 A Formalism for Population Studies

We now turn our attention to whether the intrinsic variability amplitude at 15 GHz, as quantified by $\overline{m}$, correlates with the physical properties of the sources in our sample. To this end, we will determine the *distribution* of intrinsic variability indices $\overline{m}$ for various subsets of our monitoring sample, and we will examine whether the various subsets are consistent with being drawn from the same distribution.

We will do so using again a likelihood analysis. We will assume that the distribution of $\overline{m}$ in any subset is an exponential distribution of the form given in equation (5.8). Since distributions of this family are uniquely described by the value of the mean, m_0, our aim is to determine m_0, or rather the probability distribution of possible m_0 values, in any specific subset.

The likelihood of a single observation of a modulation index $\overline{m}_i$ of Gaussian uncertainty σ_i drawn from an exponential distribution of mean m_0 is

$$\begin{aligned} \ell_i &= \int_{\overline{m}=0}^{\infty} d\overline{m} \frac{1}{m_0} \exp\left(-\frac{\overline{m}}{m_0}\right) \frac{1}{\sigma_i\sqrt{2\pi}} \exp\left[-\frac{(\overline{m}-\overline{m}_i)^2}{2\sigma_i^2}\right] \\ &= \frac{1}{m_0\sigma_i\sqrt{2\pi}} \exp\left[-\frac{\overline{m}_i}{m_0}\left(1-\frac{\sigma_i^2}{2m_0\overline{m}_i}\right)\right] \times \\ &\quad \int_{\overline{m}=0}^{\infty} d\overline{m} \exp\left[-\frac{[\overline{m}-(\overline{m}_i-\sigma_i^2/m_0)]^2}{2\sigma_i^2}\right], \end{aligned} \tag{5.10}$$

where, to obtain the second expression, we have completed the square in the exponent of the integrand. The last integral can be calculated analytically, yielding

$$\begin{aligned} \ell_i &= \frac{1}{2m_0} \exp\left[-\frac{\overline{m}_i}{m_0}\left(1-\frac{\sigma_i^2}{2m_0\overline{m}_i}\right)\right] \times \\ &\quad \left\{1+\mathrm{erf}\left[\frac{\overline{m}_i}{\sigma_i\sqrt{2}}\left(1-\frac{\sigma_i^2}{m_0\overline{m}_i}\right)\right]\right\}. \end{aligned} \tag{5.11}$$

If we want (as is the case for our data set) to implement data cuts that restrict the values of $\overline{m}_i$ to be larger than some limiting value m_l, the likelihood of a single observation of a modulation index $\overline{m}_i$ will be the expression above multiplied by a Heaviside step function, and renormalized so that the likelihood $\ell_{i,\mathrm{cuts}}$ to obtain any value of $\overline{m}_i$ *above* m_l is 1:

$$\ell_{i,\mathrm{cuts}}[m_l] = \frac{H(\overline{m}_i - m_l)\ell_i}{\int_{\overline{m}_i=m_l}^{\infty} d\overline{m}_i \ell_i}. \tag{5.12}$$

This renormalization enforces that there is no probability density for observed events "leaking" in the parameter space of rejected $\overline{m}_i$ values. In this way, it "informs" the likelihood that the reason why no objects of $\overline{m}_i < m_l$ are observed is not because such objects are not found in nature, but rather because we have excluded them "by hand."

The integral in the denominator is analytically calculable,

$$\int_{\overline{m}_i=m_l}^{\infty} d\overline{m}_i \ell_i = \frac{1}{2}\left\{\exp\left(\frac{\sigma_i^2}{2m_0^2} - \frac{m_l}{m_0}\right)\times \left[1+\mathrm{erf}\left(\frac{m_l}{\sigma_i\sqrt{2}} - \frac{\sigma_i}{m_0\sqrt{2}}\right)\right] +1-\mathrm{erf}\left(\frac{m_l}{\sigma_i\sqrt{2}}\right)\right\}. \quad (5.13)$$

The likelihood of N observations of this type is

$$\mathcal{L}(m_0) = \prod_{i=1}^{N} \ell_{i,\mathrm{cuts}}[m_l]\,. \quad (5.14)$$

If we wish to study two parts of the S_0 parameter space with different cuts (as in, for example, figure 5.5, where we have a cut of $m_l = 0.02$ for $S_0 > 0.4$ Jy, and a different cut of $m_u = 0.06$ for 0.06 Jy $\leq S_0 \leq 0.4$ Jy), we can implement this in a straightforward way, by considering each segment of the S_0 parameter space as a distinct "experiment," with its own data cut. If the first "experiment" involves N_l objects surviving the m_l cut, and the second "experiment" involves N_u objects surviving the m_u cut, then the overall likelihood will simply be

$$\mathcal{L}(m_0) = \prod_{i=1}^{N_l} \ell_{i,\mathrm{cuts}}[m_l] \prod_{i=1}^{N_u} \ell_{i,\mathrm{cuts}}[m_u]\,. \quad (5.15)$$

Maximizing equation (5.15) we obtain the maximum-likelihood value of m_0, $m_{0,\mathrm{maxL}}$. Statistical uncertainties on this value can also be obtained in a straightforward way, as equation (5.15), assuming a flat prior on m_0, gives the probability density of the mean intrinsic modulation index m_0 of the subset under study.

5.2 Null Tests

Here we begin to apply the formalism introduced in section 5.1.3 to examine whether the intrinsic modulation index $\overline{m}$ correlates with the properties of the sources in our sample. We will be testing whether the distributions of $\overline{m}$-values in subsets of our monitoring sample split according to some source property are consistent with each other. Before considering physically motivated population splits, we need to be certain that our formalism is correctly implemented. To verify that our analysis does not yield spurious results, we first discuss three null test cases where the likelihood analysis *should not* find a difference in the variability properties of the different subsets considered.

Because these null tests are tests of the method rather than tests of the source populations, we do not expect any significant change between the two-year and 42-month data sets. We therefore performed two of the three tests with only the two-year data set. As a sanity check we verified that the test described in section 5.2.2 below returned a null result with both two-year and the 42-month data sets.

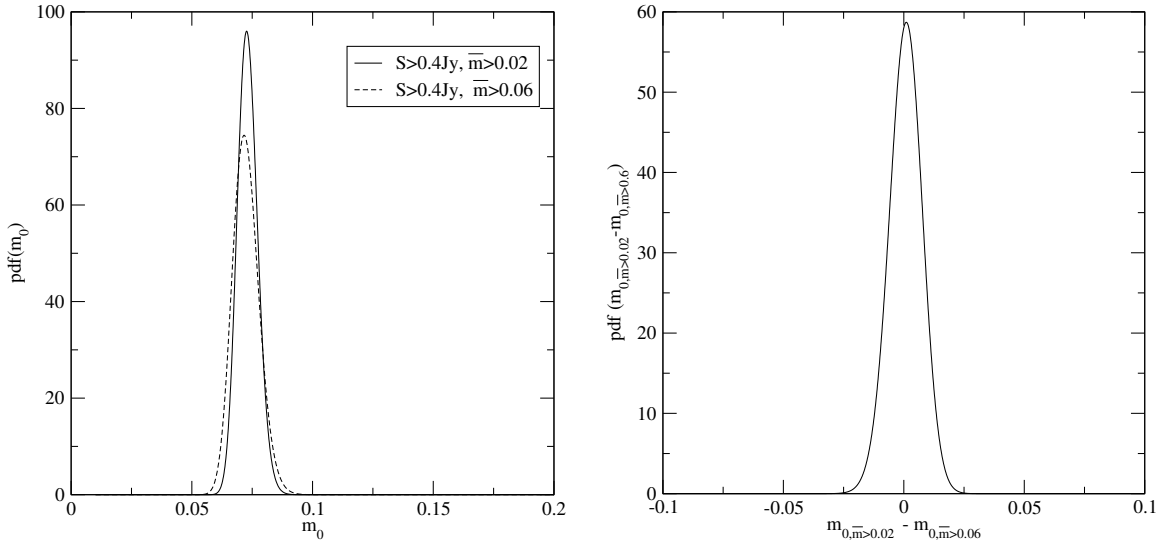

Figure 5.16. Verification that the data cuts described in section 5.1.2 are correctly implemented, using the two-year CGRaBS data. *Left:* Probability density of m_0 for the subset of bright CGRaBS blazars not found in 1LAC, for two values of the cutoff for data acceptance: $m_l = 0.02$ (solid line, maximum-likelihood value and 1σ error $m_0 = 0.073 \pm 0.004$), and $m_l = 0.06$ (dashed line, maximum-likelihood value and 1σ error $m_0 = 0.072^{+0.006}_{-0.005}$). The two distributions are consistent with a single value. *Right:* Probability density of the *difference* between the mean modulation index m_0 for the two sets. The difference (0.001 ± 0.007) is consistent with zero within 1σ.

This and the following sections include a number of histograms of modulation index for various subpopulations of the sample. It is important to remember that we have employed the data cuts described in section 5.1.2 prior to constructing these subpopulations. This distorts the appearance of the lowest bins in these histograms where we have excluded the parameter space region in which our sampling is incomplete. The likelihood analysis is properly "informed" that this parameter space region has been excluded, so this is merely an aesthetic effect.

5.2.1 Verifying Data Cuts

The first case tests whether the data cuts discussed in section 5.1.2 are implemented correctly in section 5.1.3. To this end, we calculate $\mathcal{L}(m_0)$ for the set of gamma-ray–quiet CGRaBS blazars (blazars not found in 1LAC) in our monitoring sample with $S > 0.4$ Jy, in two different ways: first, by applying an $\overline{m}$ cut at $m_l = 0.02$; second, by applying an $\overline{m}$ cut at $m_l = 0.06$ (a much more aggressive cut than necessary for the particular bright blazar population). The increased value of m_l in the second case should not affect the result other than by reducing the number of data points and thus resulting in a less constraining likelihood for m_0. This is indeed the case, as we see in the left panel of figure 5.16, where we plot the probability density of m_0 for the two subsets. That the two distributions are consistent with each other is explicitly demonstrated in the right panel of figure 5.16, where we plot the probability density of the *difference* between the means m_0 of the two

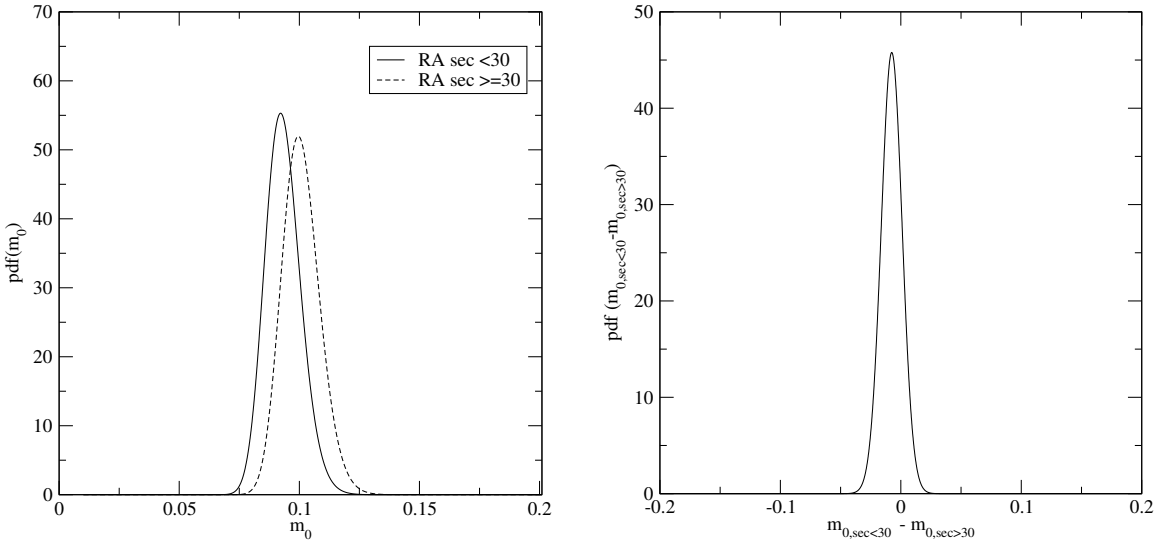

Figure 5.17. Null test verification using the two-year data. *Left:* Probability density of m_0 for bright CGRaBS blazars with seconds of right ascension < 30 s (solid line, maximum-likelihood value and 1σ error $m_0 = 0.088 \pm 0.006$) or ≥ 30 s (dashed line, maximum-likelihood value and 1σ error $m_0 = 0.096^{+0.007}_{-0.006}$). The two distributions are consistent with a single value. *Right:* Probability density of the *difference* between the mean modulation index m_0 for the two sets. The difference (-0.008 ± 0.009) is consistent with zero within 1σ.

subsets (which is formally equal to the cross-correlation of their individual distributions). The difference is consistent with zero within 1σ.

5.2.2 Physically Insignificant Population Split

The second case tests whether a split according to a source property *without* physical meaning and with the *same* value for the cutoff modulation index m_l will yield probability densities for the m_0 that are consistent with each other. For this reason, we split the population of bright ($S > 0.4$ Jy) blazars in the sample into two subsets in the following way: we divide the right ascension of each source by 1 min. If the remainder of this operation is <30 s, we include this source in the first subsample; if the remainder is ≥ 30 s we include the source in the second subsample.

We first applied this test to the two-year CGRaBS data set. The results are shown in figure 5.17. As expected, the probability distributions of m_0 for the two subsamples, shown in the left panel of figure 5.17, are consistent with each other. This is explicitly demonstrated in the right panel, which shows the probability density of the difference between the m_0 in the two subsamples. The difference is consistent with zero within 1σ.

As a sanity check, we repeated this null test with the 42-month data set including bright ($S > 0.4$ Jy) sources from both the CGRaBS and 1LAC samples. The resulting probability distributions are shown in figure 5.18. Again, the populations are consistent with each other to well within 1σ. The most likely values for the m_0 have increased for both bins relative to the two-year results. This is because we now include the

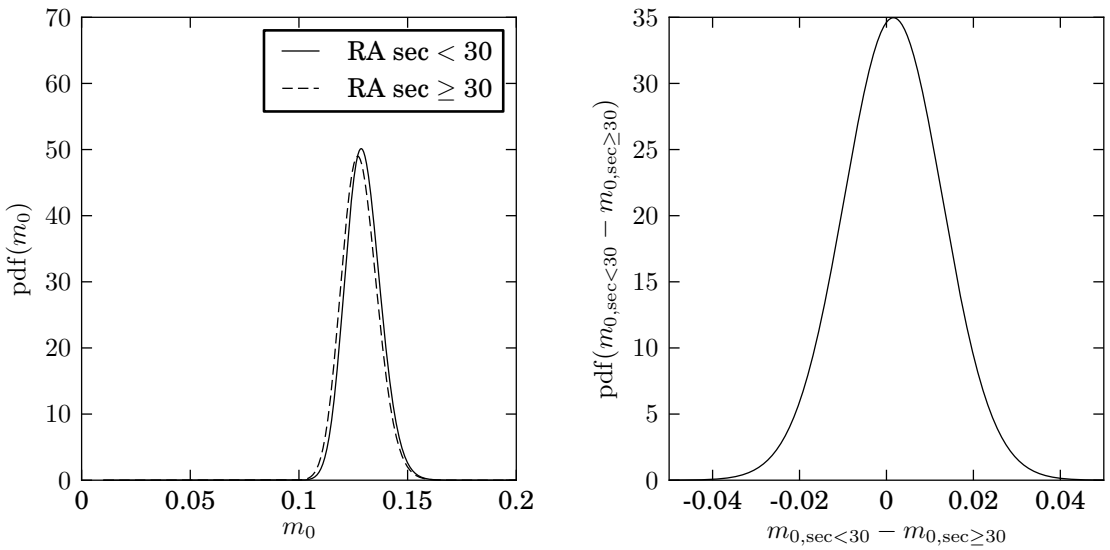

Figure 5.18. Null test verification using the 42-month data. *Left:* Probability density of m_0 for sources with seconds of right ascension <30 s (solid line, maximum likelihood value and 1σ error $0.127^{+0.009}_{-0.008}$) and ≥ 30 s (dashed line, maximum likelihood value and 1σ error 0.129 ± 0.008). The two distributions *are* consistent with a single value. *Right:* Probability density of the *difference* between the mean modulation index m_0 for the two sets. The peak of the distribution ($0.002^{+0.011}_{-0.012}$) is much less than 1σ from zero, as expected.

more variable 1LAC sample (see section 5.3) and because we have increased the time series length for the CGRaBS sample which, in section 5.1.2.1, we showed leads to a systematic increase in measured variability amplitude. In figure 5.19 we show the histograms of the intrinsic modulation indices in the two populations for the 42-month data set.

5.2.3 Galactic Latitude Split

In the final test case, we use the two-year CGRaBS data set to examine whether a split in galactic latitude yields consistent probability densities for the two subsamples. Again, we expect consistent distributions because this division does not reflect an intrinsic physical property of the sources. For this test, we restrict the sample to bright ($S > 0.4$ Jy) FSRQs and use the cutoff modulation index $m_l = 0.02$. We split between low- and high-galactic latitude at $|b| = 39°$. This produces similarly sized subsamples (181 and 168 for low- and high-latitude, respectively). The left panel in figure 5.20 shows the probability distributions for m_0 for these two subsamples, which, as anticipated, are consistent with each other. The right panel of figure 5.20 shows the probability density for the difference between m_0 for the two subsamples, which is consistent with zero to within 1σ.

This test is not a pure null test because there is a potentially significant observational difference between the sources in the two populations. Sources at lower galactic latitudes are more likely to be affected by scintillation due to the galactic interstellar medium (e.g., Rickett et al. 2006). However, due to the galactic

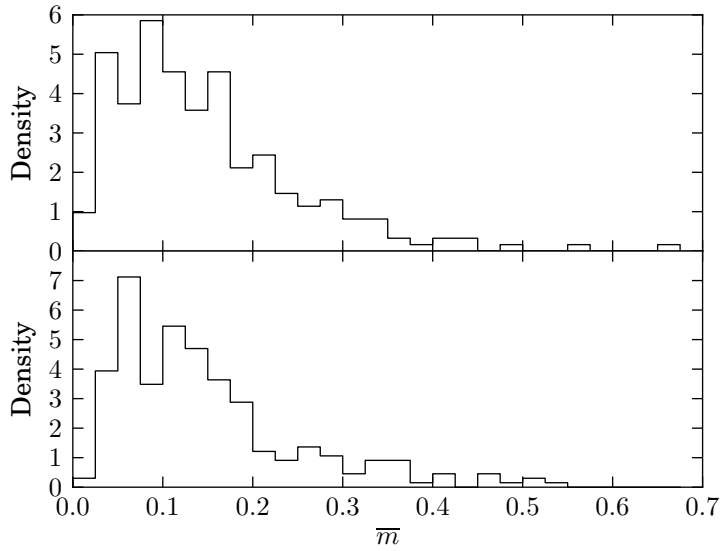

Figure 5.19. Histograms of 42-month intrinsic modulation index values for bright ($S_0 > 0.4$ Jy) sources with seconds of right ascension ≥ 30 s (top, 246 sources) and < 30 s (bottom, 264 sources). Each histogram is normalized to integrate to unity.

latitude cut $|b| \geq 10^\circ$ that defines our samples and the relatively high observation frequency, this was not expected to be a significant effect. This expectation was confirmed by our results.

5.3 Gamma-Ray Loud versus Quiet Populations

We now turn from null tests to studying subsets defined according to physical properties of the sources. The first criterion we apply is whether the source has been detected by *Fermi*-LAT at a significance level high enough to warrant inclusion in the 1LAC catalog. For sources with $S_0 < 0.4$ Jy we apply a cut $\overline{m} > m_u = 0.06$ and for sources with $S_0 \geq 0.4$ Jy a cut $\overline{m} > m_l = 0.02$. The results for the two-year data set are shown in figure 5.21. The set of sources that are included in 1LAC is depicted by a solid line, while the set of sources that are not in 1LAC is depicted by a dashed line. The two are not consistent with each other at a confidence level of 6σ (right panel of figure 5.21), with a maximum-likelihood difference of 5.7 percentage points, with gamma-ray–loud blazars exhibiting, on average, a higher variability amplitude *by almost a factor of 2* versus gamma-ray–quiet blazars.

This significant difference persists in the 42-month data set. Here again, we have considered a CGRaBS source to be gamma-ray loud if it appeared in the 1LAC sample. The likelihood distributions for the population mean intrinsic modulation indices are shown in the left panel of figure 5.22. The right panel shows the probability density for the difference in m_0 between the two populations. The values for both subpopulations are somewhat larger than before, but the most likely difference has increased and continues to be significant

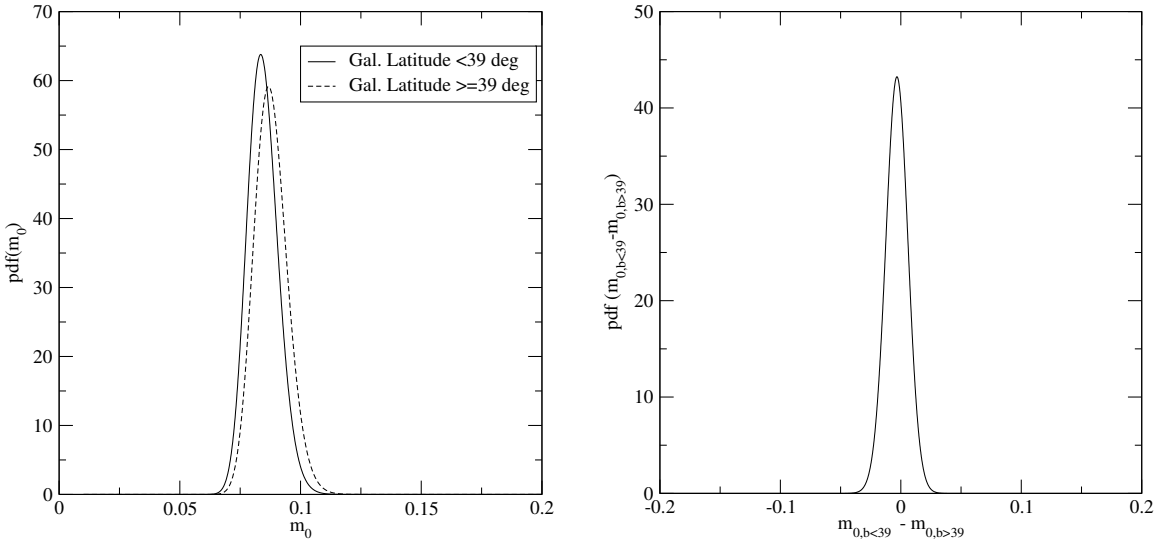

Figure 5.20. Comparison of bright ($S > 0.4$ Jy) FSRQs at high and low galactic latitude using the two-year data set. *Left:* Probability density of m_0 for the subset of bright ($S > 0.4$ Jy) CGRaBS FSRQs with $\overline{m} > 0.02$ at low galactic latitude ($|b| < 39°$, solid line, maximum-likelihood value and 1σ error $m_0 = 0.084^{+0.007}_{-0.006}$) or high galactic latitude ($|b| \geq 39°$, dashed line, maximum-likelihood value and 1σ error $m_0 = 0.087^{+0.007}_{-0.006}$). The two distributions are consistent with a single value. *Right:* Probability density of the *difference* between the mean modulation index m_0 for the two sets. The difference (-0.003 ± 0.009) is consistent with zero within 1σ.

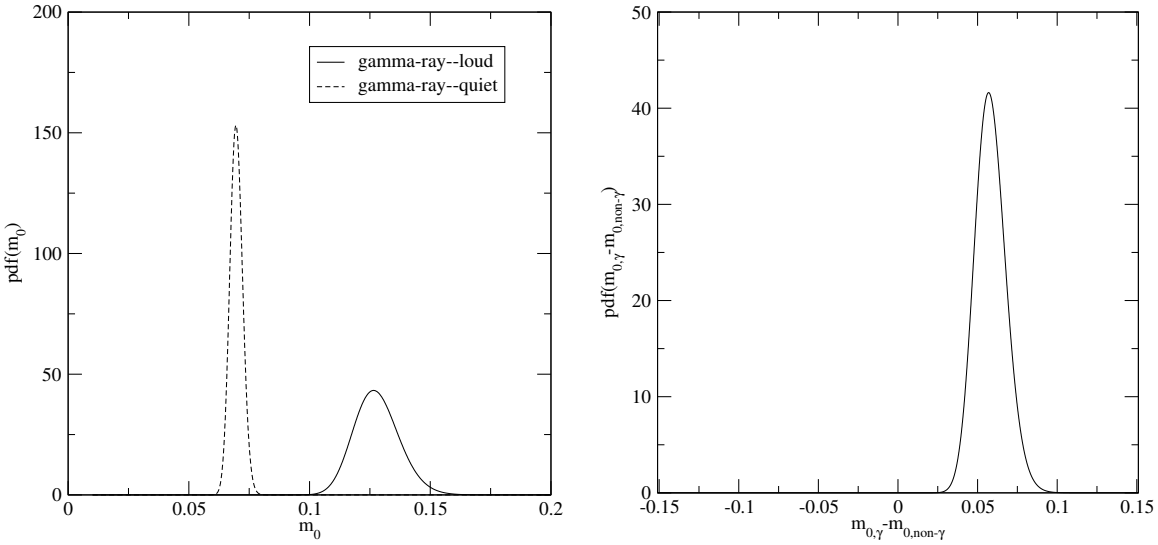

Figure 5.21. Comparison of gamma-ray–loud and gamma-ray–quiet CGRaBS populations using the two-year data set. *Left:* Probability density of m_0 for CGRaBS blazars in our monitoring sample that are (solid line, maximum-likelihood value and 1σ error $m_0 = 0.127^{+0.010}_{-0.009}$) and are not (dashed line, maximum-likelihood value and 1σ error $m_0 = 0.070 \pm 0.003$) included in 1LAC. The two distributions are *not* consistent with a single value. *Right:* Probability density of the *difference* between the mean modulation index m_0 for the two sets. The peak of the distribution ($0.057^{+0.010}_{-0.009}$) is 6σ away from zero.

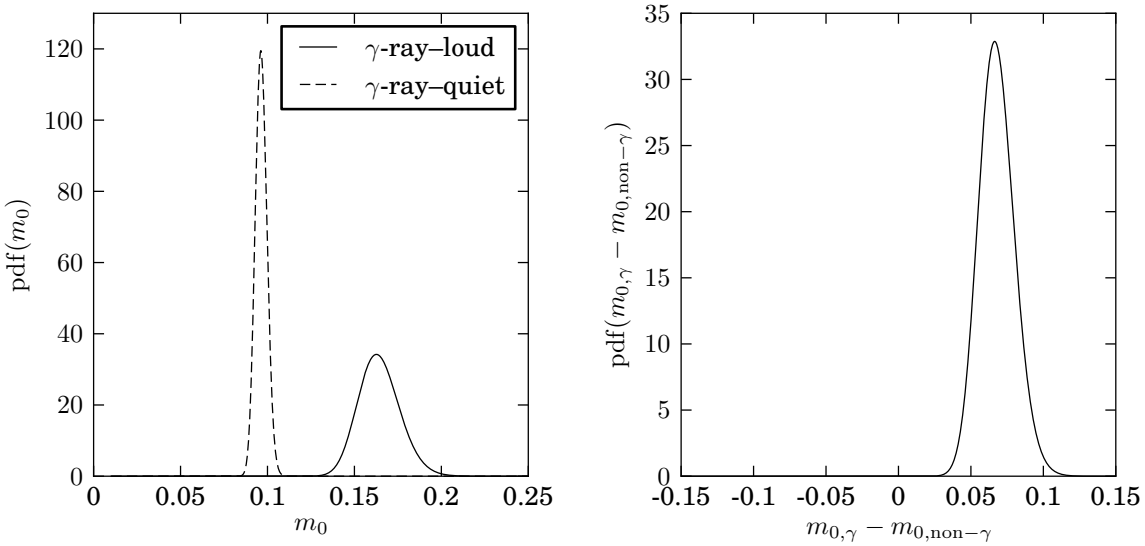

Figure 5.22. Comparison of gamma-ray–loud and gamma-ray–quiet CGRaBS populations using the 42-month data set. *Left:* Probability density of m_0 for the CGRaBS blazars in our monitoring sample that are (solid line, maximum likelihood value and 1σ error $0.163^{+0.012}_{-0.011}$) and are not (dashed line, maximum likelihood value and 1σ error 0.096 ± 0.003) included in 1LAC. The two distributions are *not* consistent with a single value. *Right:* Probability density of the *difference* between the mean modulation index m_0 for the two sets. The peak of the distribution ($0.066^{+0.013}_{-0.012}$) is more than 6σ away from zero.

at the 6σ level. In figure 5.23, we plot the histograms of intrinsic modulation indices for the gamma-ray–loud and gamma-ray–quiet CGRaBS subpopulations using the 42-month data.

Clearly the gamma-ray–loud subset of the CGRaBS sample is much more variable, on average, than the gamma-ray–quiet subset. While this is likely an important clue about the physical conditions necessary for the production of observable gamma rays in a blazar, before further discussing this result we will examine the variability difference between other subsets of our sample. Both the CGRaBS and 1LAC samples contain a variety of source types and properties, so this may shed further light on the significant differences between gamma-ray–loud and gamma-ray–quiet sources.

5.4 BL Lac Object versus FSRQ Populations

We next examine the variability amplitude as a function of optical spectral classification. In this section we consider the BL Lac and FSRQ subsets of our samples and examine whether they differ in terms of 15 GHz variability. First, we examine the CGRaBS sample using the two-year data set. The probability densities for the mean m_0 of the two subsets are shown in the left panel of figure 5.24. The results for BL Lacs (FSRQs) are plotted as a solid (dashed) line. The two curves are not consistent with each other—the BL Lacs appear to have, on average, higher variability amplitude than the FSRQs. We verify this finding by plotting, in the right panel, the probability density of the difference between the m_0 of BL Lacs and FSRQs. The most likely difference is 3.2 percentage points, and it is more than 3σ away from zero. Note that the difference between

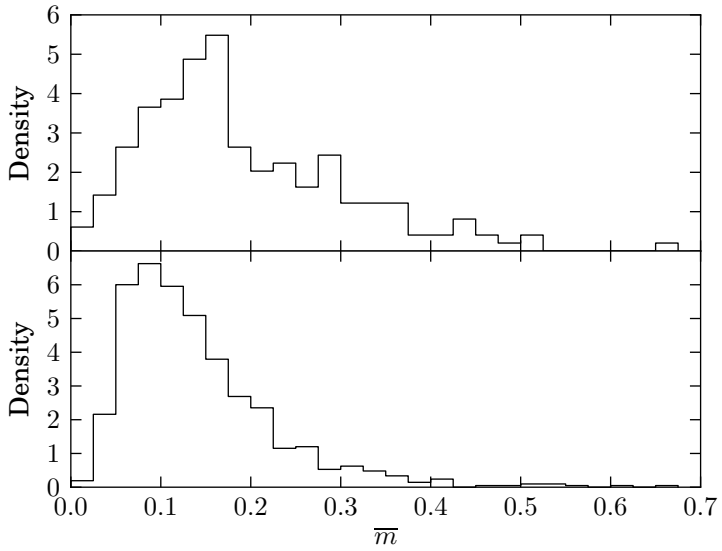

Figure 5.23. Histograms of 42-month intrinsic modulation index values for CGRaBS sources detected (top, 197 sources) and not detected (bottom, 833 sources) in gamma rays by the LAT in the 1LAC catalog. Each histogram is normalized to integrate to unity.

BL Lacs and FSRQs is less significant than that between gamma-ray–loud and gamma-ray–quiet blazars. This is both because the most likely difference in m_0 values between the BL Lac and FSRQ subsets is smaller and because the BL Lac sample is smaller than the gamma-ray–loud blazar sample: only 94 BL Lacs satisfy the data cuts we impose, versus 191 gamma-ray–loud blazars of all types. As a result, the constraints on the intrinsic distribution of modulation indices (i.e., on m_0) are stronger in the latter case.

Turning now to the 42-month data set, we first examine the CGRaBS population. The likelihood distributions for the population mean intrinsic modulation indices are shown in figure 5.25 and figure 5.26 shows the histogram of the intrinsic modulation indices for sources identified as BL Lacs or as FSRQs within the CGRaBS sample. The significant distinction in variability between CGRaBS BL Lac and FSRQ sources remains, and continues to appear at the $>3\sigma$ confidence level.

In figure 5.27, we show the 42-month results for the 1LAC sample, and in figure 5.28 we show the modulation index histograms. Using this sample, the most likely difference in variability amplitude between the two samples is now found to be less than 2σ significant. Remarkably, the difference has not only faded in significance, but the sign of the most likely difference has switched. In the 1LAC sample, the BL Lac subpopulation is found to be *less* variable than the FSRQ subpopulation, whereas BL Lacs were found to be *more* variable in the CGRaBS sample. We will discuss this further in section 5.6 below.

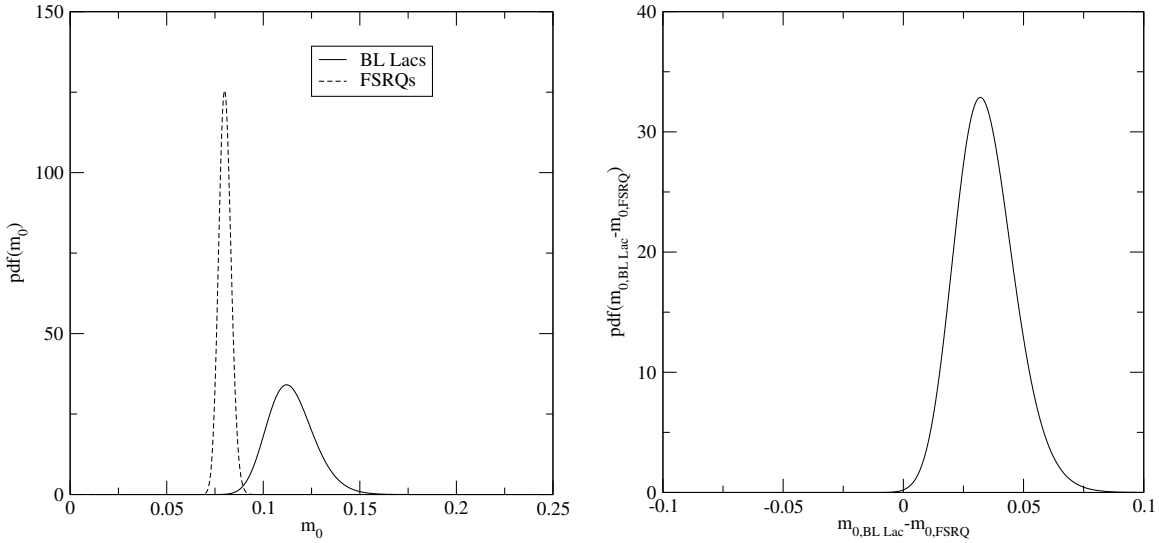

Figure 5.24. Comparison of CGRaBS BL Lac and FSRQ populations using two-year data. *Left:* Probability density of m_0 for BL Lac (solid line, maximum-likelihood value and 1σ error $m_0 = 0.112^{+0.013}_{-0.011}$) and FSRQ (dashed line, maximum-likelihood value and 1σ error $m_0 = 0.080 \pm 0.003$) CGRaBS blazars in our monitoring sample. The two distributions are *not* consistent with a single value. *Right:* Probability density of the *difference* between the mean modulation index m_0 for the two sets. The peak of the distribution ($0.032^{+0.013}_{-0.011}$) is more than 3σ away from zero.

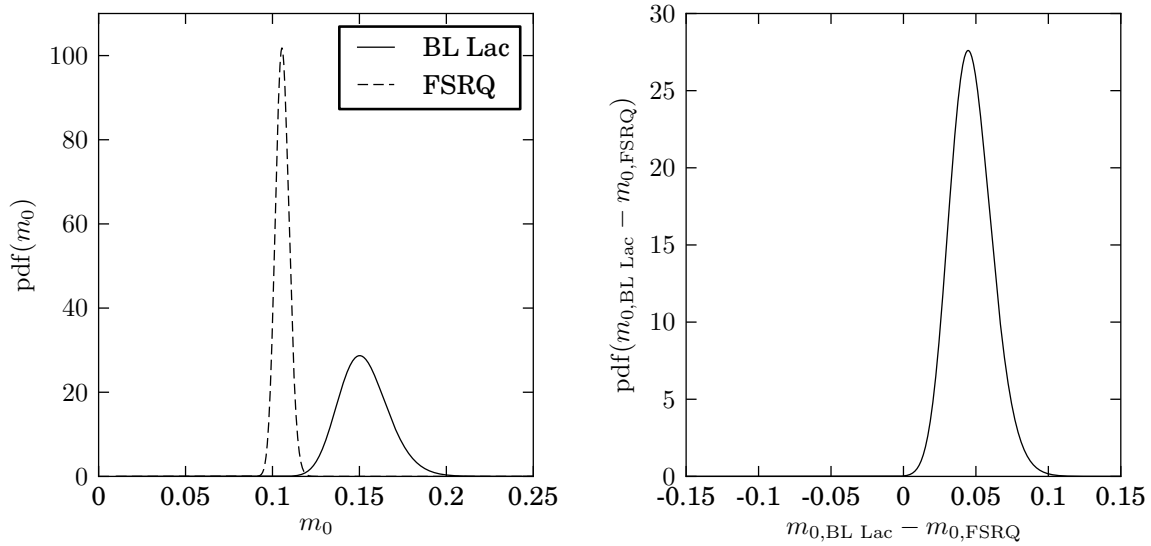

Figure 5.25. Comparison of CGRaBS BL Lac and FSRQ populations using 42-month data. *Left:* Probability density of m_0 for the CGRaBS blazars in our monitoring sample that are identified as BL Lac (solid line, maximum likelihood value and 1σ error $0.150^{+0.015}_{-0.013}$) and as FSRQ (dashed line, maximum likelihood value and 1σ error 0.105 ± 0.004). The two distributions are *not* consistent with a single value. *Right:* Probability density of the *difference* between the mean modulation index m_0 for the two sets. The peak of the distribution ($0.045^{+0.015}_{-0.014}$) is about 3.5σ away from zero.

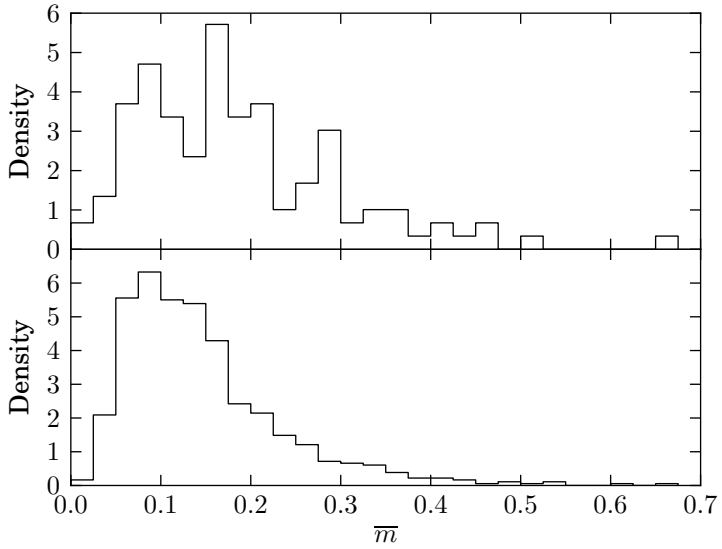

Figure 5.26. Histograms of 42-month intrinsic modulation index values for CGRaBS sources classified as BL Lac (top, 119 sources) and as FSRQ (bottom, 727 sources). Each histogram is normalized to integrate to unity.

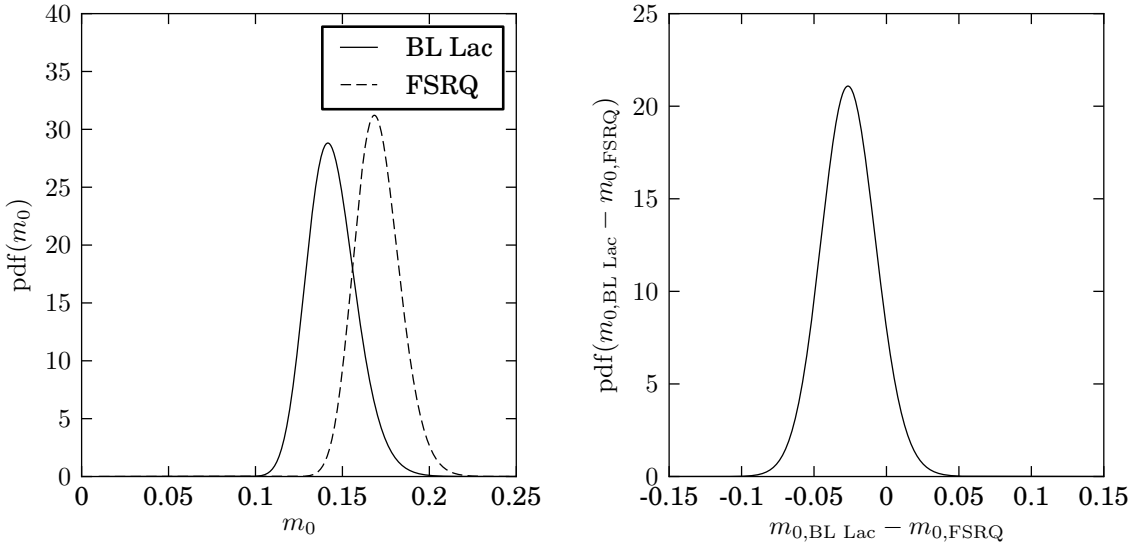

Figure 5.27. Comparison of 1LAC BL Lac and FSRQ populations using 42-month data. *Left:* Probability density of m_0 for the 1LAC blazars in our monitoring sample that are identified as BL Lac (solid line, maximum likelihood value and 1σ error $0.142^{+0.015}_{-0.013}$) and as FSRQ (dashed line, maximum likelihood value and 1σ error $0.168^{+0.014}_{-0.012}$). The two distributions *are* consistent with a single value. *Right:* Probability density of the *difference* between the mean modulation index m_0 for the two sets. The peak of the distribution (0.027 ± 0.019) is less than 2σ away from zero.

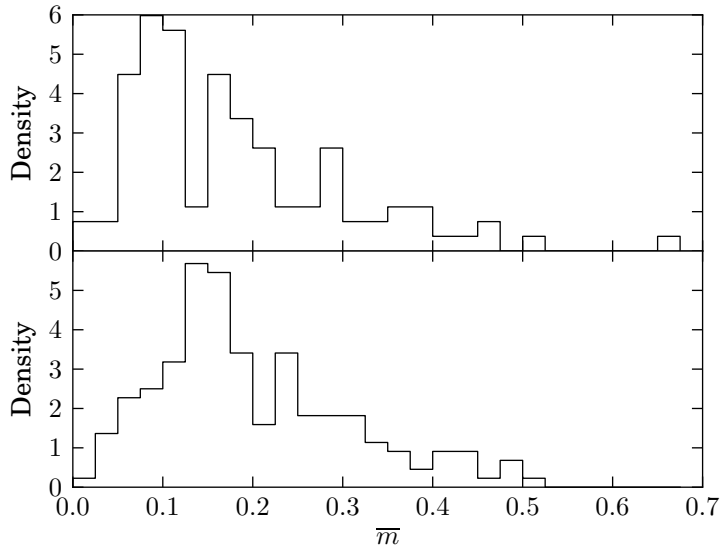

Figure 5.28. Histograms of 42-month intrinsic modulation index values for 1LAC sources classified as BL Lac (top, 107 sources) and as FSRQ (bottom, 176 sources). Each histogram is normalized to integrate to unity.

5.5 Redshift Trend

Finally, we examine the dependence of variability amplitude on redshift. In the left panel of figure 5.29 we plot the mean $\overline{m}$ (as calculated by a simple average rather than the likelihood analysis) in redshift bins of $\Delta z = 0.5$ for bright ($S \geq 400$ mJy) CGRaBS FSRQs with known redshifts in our monitoring sample, using the two-year data set. We exclude BL Lacs from this analysis so as not to bias the result, as BL Lacs with known redshifts are located at low z, and we have also already shown that they have a higher mean $\overline{m}$ compared to FSRQs within the CGRaBS sample. Although the errors are large, there is a hint of a trend toward decreasing variability amplitude with increasing redshift. We further test the significance of this result by splitting sources in our monitored sample in high- and low-redshift subsets with the dividing redshift at $z = 1$ (dashed line in figure 5.29). In the two subsets we also include faint ($S < 400$ mJy) sources, with the usual cut at $m_u = 0.06$. The probability density for the mean m_0 of each subset is shown in the left panel of figure 5.30, where the solid curve corresponds to low-redshift blazars and the dashed curve to high-redshift FSRQs. We find that low-redshift FSRQs have *higher*, on average, intrinsic modulation indices. The result is shown to be statistically significant in the right panel of figure 5.30, where we plot the probability density of the difference between m_0 in each subset. The most likely difference is found to be about 2.4 percentage points, and more than 3σ away from zero.

In figures 5.31 and 5.32, we plot the likelihood distributions and m_0 difference probability distributions for high- and low-redshift subpopulations of the CGRaBS and 1LAC samples, now computed using the 42-month data set. The histograms of the modulation indices in each sample are shown in figures 5.33 and 5.34. Although in both cases we continue to find that the low-redshift sources are characterized by greater average

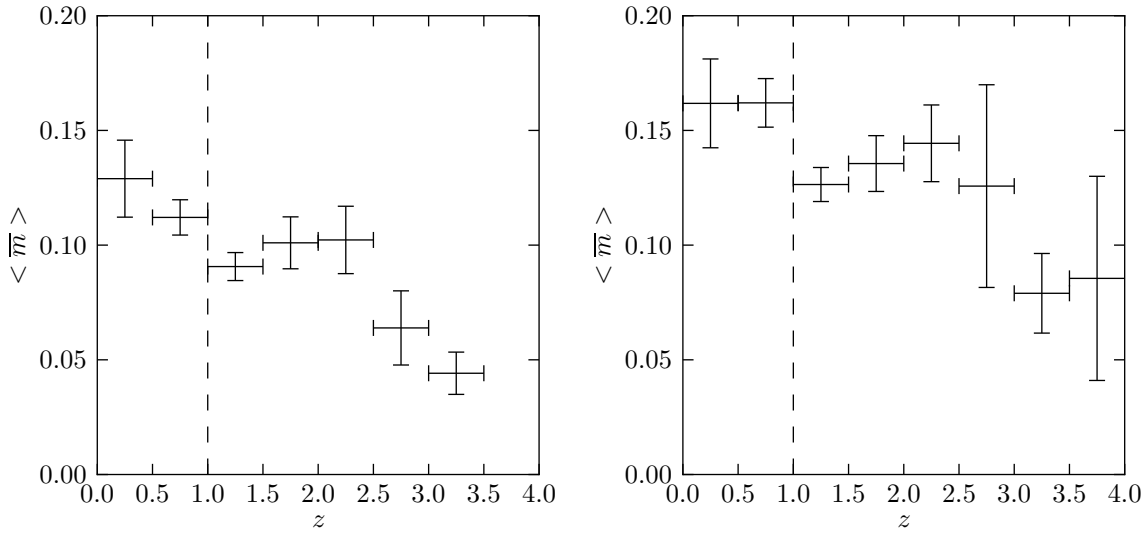

Figure 5.29. Mean $\overline{m}$ in redshift bins of 0.5 for bright ($S > 400$ mJy) FSRQs in our CGRaBS monitoring sample using the two-year (left) and 42-month (right) data. Horizontal error bars indicate the bin width, vertical error bars indicate the uncertainty in the mean computed from the scatter in the data in that bin. The dashed line indicates the $z = 1$ split between high and low redshift sources used for the population comparison.

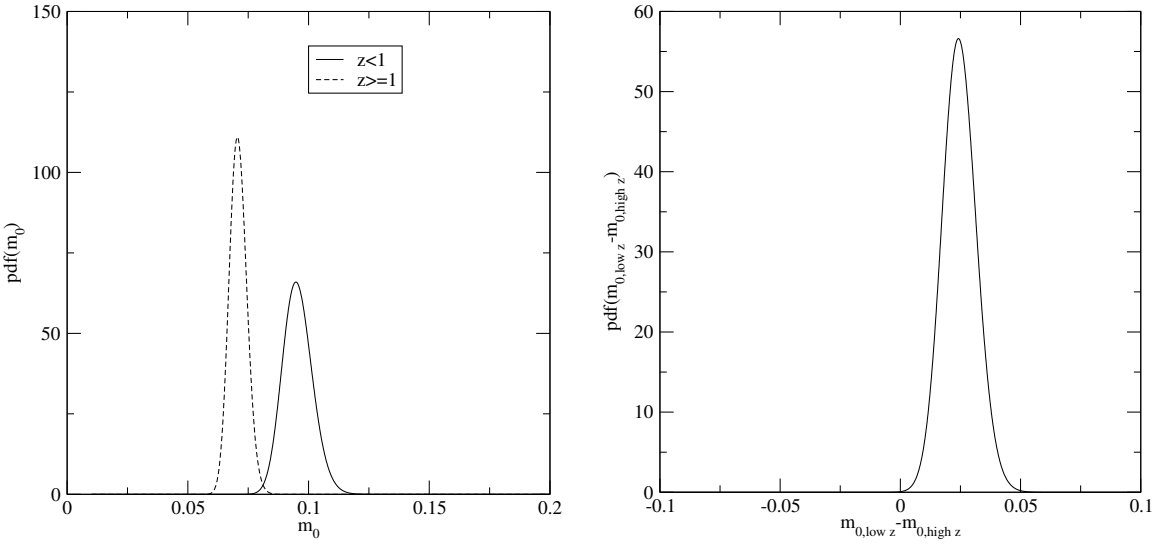

Figure 5.30. Comparison of high- and low-redshift CGRaBS FSRQs using the two-year data set. *Left:* Probability density of m_0 for FSRQs in our monitoring sample with $z < 1.0$ (solid line, maximum-likelihood value and 1σ error $m_0 = 0.095 \pm 0.006$) and $z \geq 1.0$ (dashed line, maximum-likelihood value and 1σ error $m_0 = 0.071 \pm 0.004$). The two distributions are *not* consistent with a single value. *Right:* Probability density of the *difference* between the mean modulation index m_0 for the two sets. The peak of the distribution (0.024 ± 0.007) is more than 3σ away from zero.

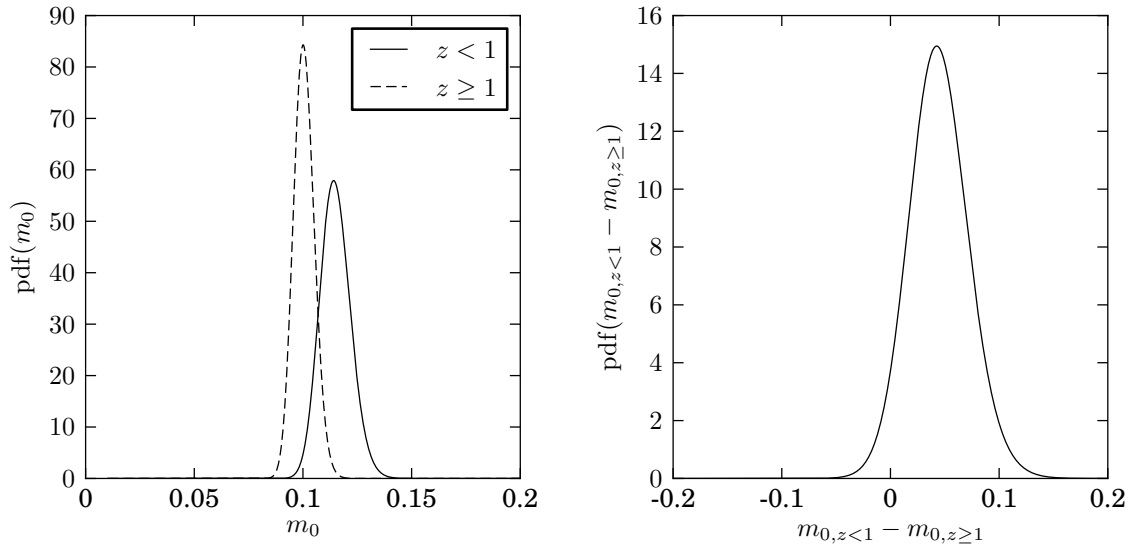

Figure 5.31. Comparison of high- and low-redshift CGRaBS FSRQ populations using 42-month data. *Left:* Probability density of m_0 for the CGRaBS FSRQs with known redshift in our monitoring sample with $z < 1$ (solid line, maximum likelihood value and 1σ error 0.114 ± 0.007) and $z \geq 1$ (dashed line, maximum likelihood value and 1σ error 0.100 ± 0.005). The two distributions *are* consistent with a single value. *Right:* Probability density of the *difference* between the mean modulation index m_0 for the two sets. The peak of the distribution ($0.014^{+0.009}_{-0.008}$) is less than 2σ away from zero.

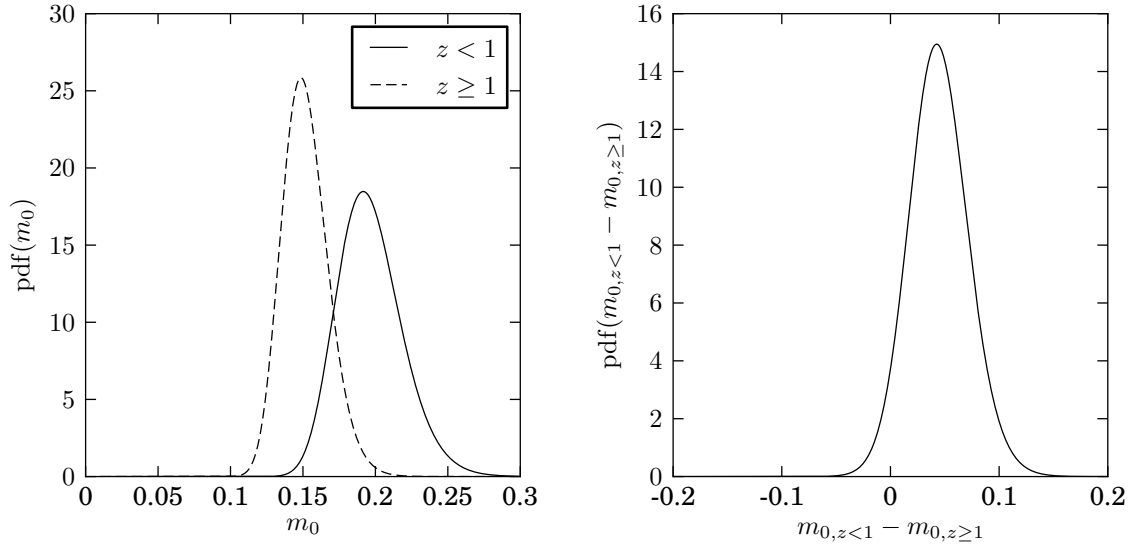

Figure 5.32. Comparison of high- and low-redshift 1LAC FSRQ populations using 42-month data. *Left:* Probability density of m_0 for the 1LAC FSRQs with known redshift in our monitoring sample with $z < 1$ (solid line, maximum likelihood value and 1σ error $0.192^{+0.023}_{-0.020}$) and $z \geq 1$ (dashed line, maximum likelihood value and 1σ error $0.149^{+0.017}_{-0.014}$). The two distributions *are* consistent with a single value. *Right:* Probability density of the *difference* between the mean modulation index m_0 for the two sets. The peak of the distribution ($0.043^{+0.028}_{-0.026}$) is less than 2σ away from zero.

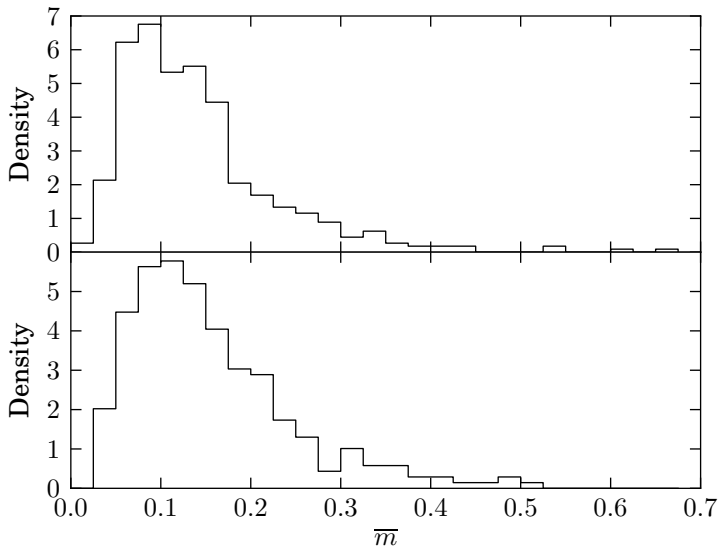

Figure 5.33. Histograms of 42-month intrinsic modulation index values for CGRaBS sources at $z \geq 1$ (top, 450 sources) and $z < 1$ (bottom, 277 sources). Each histogram is normalized to integrate to unity.

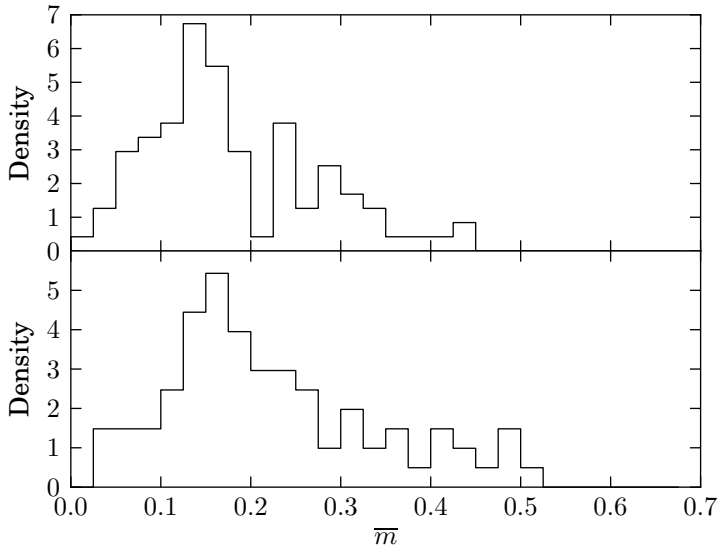

Figure 5.34. Histograms of 42-month intrinsic modulation index values for 1LAC sources at $z \geq 1$ (top, 95 sources) and $z < 1$ (bottom, 81 sources). Each histogram is normalized to integrate to unity.

intrinsic modulation index, with 42 months of data this difference is no longer even 2σ significant for either the CGRaBS or the 1LAC samples. In the right panel of figure 5.29, we plot the binned average modulation index versus redshift using the 42-month CGRaBS data set. Although the mean value in each bin still shows a hint of a trend, the scatter has also increased in the higher-z bins, diluting the appearance of a trend. That the additional data has reduced, not enhanced, the significance of this result suggests that it was merely a chance effect. Nonetheless, because evidence for cosmological evolution of blazar variability would be an important finding, we will examine this more closely.

Understanding the implications of the presence or absence of a trend of variability amplitude with redshift is complicated by competing observational and selection effects that can push the trend in either direction. To claim or rule out cosmological source evolution will require quantitative measurement or modeling of these effects. For example, at higher redshift, the 15 GHz observation frequency corresponds to a higher rest frame emission frequency. Because variability amplitude generally increases with increasing cm-wave radio frequency Stevens et al. (e.g., 1994), this would lead to an increase in variability amplitude with increasing z. The effects of Doppler beaming of the blazar jet emission also gives rise to selection effects which are somewhat more difficult to quantify (Lister & Marscher 1997). In this section, we will examine in detail the effect of one effect—cosmological time dilation—which leads to an underestimate of variability amplitude at higher z.

5.5.1 Cosmological Time Dilation

The redshifts assigned to the sources in our sample are determined spectroscopically—that is, by identifying the observer-frame wavelengths, $\lambda_{\rm obs}$, of photons that are emitted or absorbed in the rest frame of the source. By identifying patterns of emission or absorption lines that correspond to atomic or molecular line spectra, the rest-frame wavelengths, $\lambda_{\rm rest}$, of those lines can be determined. Then the redshift z of the source is obtained as

$$z = \frac{\lambda_{\rm obs}}{\lambda_{\rm rest}} - 1 = \frac{\nu_{\rm rest}}{\nu_{\rm obs}} - 1. \tag{5.16}$$

This redshift is normally dominated by the expansion of space, so this redshift is ascribed to the Doppler effect due to the cosmological motion of the source away from the observer.

Because this effect is due to the expansion of space, it is expected that the observed spectral redshift will be accompanied by a cosmological time dilation effect. This has been confirmed by, e.g., Blondin et al. (2008), who compared the evolution of Type Ia supernovae at various redshifts. Thus, a rest-frame time interval $\Delta t_{\rm rest}$ for a source at redshift z will correspond to an observed interval

$$\Delta t_{\rm obs} = (1+z)\,\Delta t_{\rm rest}. \tag{5.17}$$

Inverting equation (5.17), it is clear that, for equal observer time intervals, a source that is at a lower redshift has been observed for a longer rest frame interval. The observer is thus more likely to observe a full cycle[1] of variability behavior from a source at lower z since more rest-frame time has passed during which variability-causing events, whatever their physical nature, can occur. As a result, we are likely to sample more flares and observe more periods of variability in a source for which we have observed a long rest-frame time interval.

We need not rely on our intuition to conclude that time dilation will tend to reduce variability amplitude as redshift increases—we have data to prove it. In section 5.1.2.1, we demonstrated that in almost all cases, the intrinsic modulation index for a source increased or remained level between the two-year and the 42-month data sets. That is, by increasing the observed time period for a given source, a larger variability amplitude typically results.

5.5.2 Compensating for Time Dilation

Virtually all monitoring programs, including ours, observe their sources either over a time interval that is the same for each source—the total length of operation of the program—or that may be shorter for sources added or dropped during operation, but not selected with regard to the redshift of the source. To compensate for this effect, we simply discard some data from low-z sources in order to compare the intrinsic modulation index found in equal rest-frame time intervals.

For sources at near-zero redshift, the time dilation effect will not be significant so the rest-frame and observer-frame intervals will be nearly equal. Our sample includes redshifts as high as 5.47 (J0906+6930), however, for which the dilation factor will be $1 + z = 6.47$, which is clearly significant. Our 42 months of data correspond to only 6.5 months in the rest frame at this redshift. Because the level of short-timescale variability in a given blazar seems to vary with time, sometimes apparently switching from a stable state to a variable one, we want to keep our measurement intervals long enough to give each source a fair chance to enter a variable state. We therefore use only sources with a redshift $z \leq 3.0$, for which we have sampled a rest-frame interval of 10.5 months, or 315 d.

5.5.3 Selecting Time Intervals

A subset of the data for each source to use for the equal-Δt_{rest} comparison was selected as follows. First, from the redshift z, the time dilation factor $(1 + z)$ determines the observer-frame time period to sample, $\Delta t_{\mathrm{obs}} = (1 + z) \times 315$ d. Next, we must choose a segment of the data for the source with this length. To

[1]The term "cycle" is used loosely here: although it is clear from our light curve data that many blazars exhibit long periods of low variability followed by flares or periods of rapid variation, there is no particular reason to believe this is cyclical in a periodic sense. Searches for periodicity have not found convincing evidence, see, e.g., Hovatta et al. (2007).

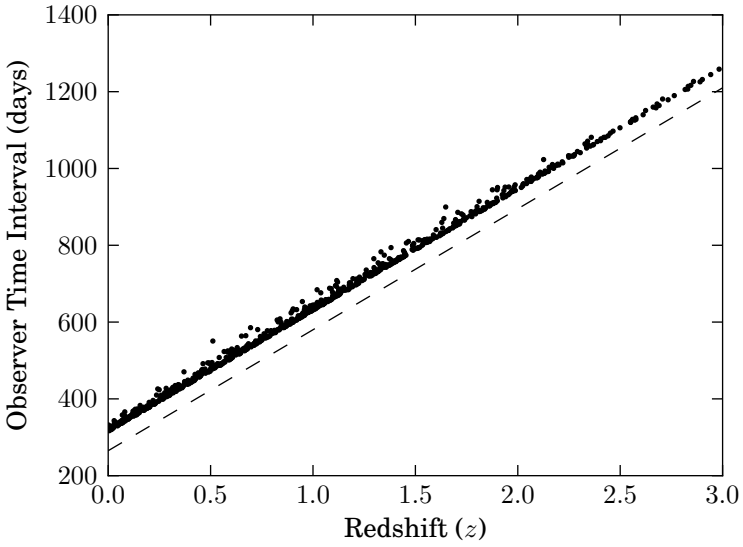

Figure 5.35. Observer time interval for each source plotted versus redshift. The dashed line shows the expected slope, $dt/dz = 315$ d, for reference.

do this, we choose a random start date within the data set for the source, at least $\Delta t_{\rm obs}$ days before the last measurement.

We choose a random start date for each source to avoid correlating the redshift of a source with the observing conditions. If, say, we simply chose an interval of the desired length that ended on the last measurement for each source, low-redshift sources would all include data observed in the winter and spring of 2011. If the conditions were slightly different during that time than the average over the 42 months of the program, this would lead to a bias. By choosing the interval at random, any such changing conditions are averaged as much as possible over the population of low-z sources.

After selecting a start date, we must ensure that interruptions in our program have not drastically reduced the number of measurements in the interval. First, we extend the end date of the selection to ensure that an outage $\Delta t_{\rm obs}$ days after the start date does not reduce the actual observed interval. Figure 5.35 shows the actual observer time interval selected for each source. We also require that the selected interval contains at least an average of one data point per 11 d (in the observer frame). If fewer data than this have been chosen, we select another random start date and repeat this process until an adequate number of data survive. Figure 5.36 shows the number of data points selected per source as a function of redshift.

From our 1LAC sample, 4 sources were dropped from this analysis because they were observed for less than 315 d in their rest frame. As shown in figure 5.37, these did not substantially alter the redshift distribution. A K-S test comparing the full data set with that of the surviving data does not reject the hypothesis that the two are drawn from the same distribution, with $1 - p = 2 \times 10^{-15}$ (i.e., $p \sim 1$).

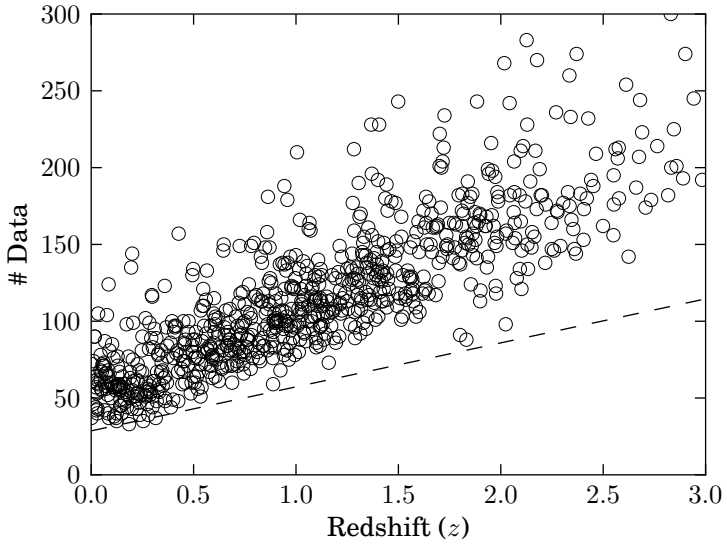

Figure 5.36. Number of data points in the equal-$\Delta t_{\rm rest}$ data sample for each source. The dashed line shows the minimum accepted number of data at each z.

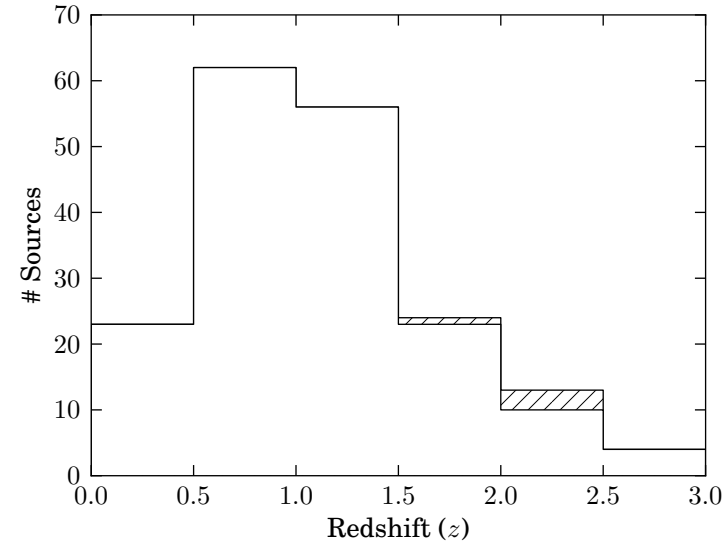

Figure 5.37. Histogram of redshifts for the 1LAC FSRQs with known $z < 3.0$. Hatched regions show the effect of dropping 4 sources that were sampled less than 315 d in their rest frames. A K-S test comparing these distributions does not reject the null hypothesis that the distributions are equal.

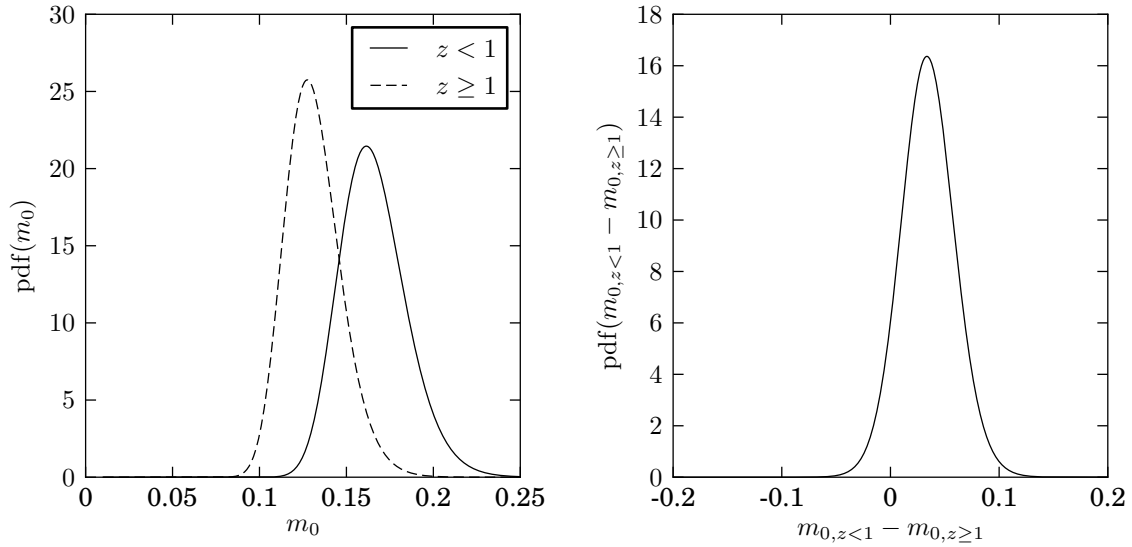

Figure 5.38. Comparison of high- and low-redshift 1LAC FSRQ populations using equal rest-frame time periods from the 42-month data. *Left:* Probability density of m_0 for the 1LAC FSRQs with known redshift in monitoring sample with $z < 1$ (solid line, maximum likelihood value and 1σ error $0.162^{+0.020}_{-0.017}$) and $z \geq 1$ (dashed line, maximum likelihood value and 1σ error $0.128^{+0.017}_{-0.014}$). The two distributions *are* consistent with a single value. *Right:* Probability density of the *difference* between the mean modulation index m_0 for the two sets. The peak of the distribution ($0.034^{+0.024}_{-0.025}$) is less than 2σ away from zero.

5.5.4 Equal Rest-Frame Time Interval Results

Using the shortened data sets to provide equal rest-frame time intervals for all sources in the 1LAC population with $z < 3$, we remove the effect of time dilation from the observed modulation index. In figure 5.38, we plot the likelihood distribution for the m_0 parameters and for the difference in m_0 between the 1LAC FSRQs at $z < 1$ from those at $z \geq 1$. The low-redshift source population is still found to have a larger average modulation index, and this is still not a statistically significant result.

As expected, the average intrinsic modulation index in both the high- and low-redshift bins were *smaller* when the shorter equal rest-frame time intervals were used, although in both cases the difference was not a statistically significant change beyond the 1σ level. For the high-redshift population, the difference was

$$\Delta m_{z\geq1} = 0.128 - 0.149 = -0.021 \pm 0.021, \tag{5.18}$$

where we have computed the uncertainty in this difference approximately, using the average of the upper and lower error bars for the quantities being subtracted. For the low-redshift population, the difference was

$$\Delta m_{z<1} = 0.162 - 0.192 = -0.030 \pm 0.029, \tag{5.19}$$

again computing uncertainties using the average of the upper and lower error bars for each value.

We conclude that the effect of cosmological time dilation on our earlier result was not significant, but our results are consistent with there being a small effect on our data. The change was smaller for high-z sources, as expected because nearly the entire time series was used, whereas time series for low-z sources were trimmed more substantially.

5.6 CGRaBS versus 1LAC

It is clear that the population of sources predicted to be gamma-ray loud in the CGRaBS paper differed substantially from the sources actually detected by the *Fermi*-LAT instrument during the first year of operation. In section 4.2, we demonstrated BL Lac objects made up a much larger fraction of the 1LAC sample than of the CGRaBS sample. This can be attributed to instrumental differences between EGRET and the LAT. This became evident from the reported EGRET "GeV excess" (Hunter et al. 1997). This excess diffuse gamma-ray emission was not confirmed by *Fermi* (Abdo et al. 2009b) and is now believed to be a systematic artifact due to uncertain instrumental response above about 1 GeV (Stecker et al. 2008). As a result, high-energy results from the EGRET instrument are suspect, limiting its ability to detect spectrally hard sources. The LAT, which is especially efficient at detecting sources via 1 GeV photons, is thus more efficient than EGRET at detecting BL Lacs because they typically exhibit hard gamma-ray spectra (Abdo et al. 2010a, 2010c).

We found in section 5.3 that gamma-ray–loud CGRaBS sources that were part of the 1LAC sample are significantly—more than 6σ—more variable than are gamma-ray–quiet sources. Then in section 5.4 we showed that in the CGRaBS sample, BL Lac objects are significantly more variable than FSRQs, whereas in 1LAC, the situation is reversed (and no longer statistically significant). What can we conclude from these observations?

We know that the CGRaBS and 1LAC samples are not drawn from the same parent distribution. Is this true of subpopulations within those two samples? First, let us examine whether the BL Lac objects are consistent with coming from a common parent sample. In the left panel of figure 5.39 we compare the likelihood distributions for m_0 for the BL Lac objects in the CGRaBS and 1LAC samples. These samples are consistent with a common value for m_0 to within 1σ. Since the BL Lac populations appear to exhibit the same average variability amplitude, the FSRQ variability amplitudes must differ to explain the overall disagreement between the CGRaBS and 1LAC samples.

This is just what we find. In the right panel of figure 5.40, we compare the FSRQ variability amplitudes between the CGRaBS and 1LAC samples. The 1LAC FSRQ population is much more variable with $m_0 = 0.168^{+0.014}_{-0.012}$ whereas for the CGRaBS FSRQs $m_0 = 0.105 \pm 0.004$. This most likely difference of $0.063^{+0.014}_{-0.014}$ is significant at nearly the 6σ level. The CGRaBS FSRQs are, on average, much less variable at 15 GHz than the gamma-ray–loud 1LAC FSRQs.

This enhanced radio variability amplitude among gamma-ray–loud FSRQs together with the equivalence of the variability in the BL Lac population would be expected if the gamma-ray–loud FSRQs are more

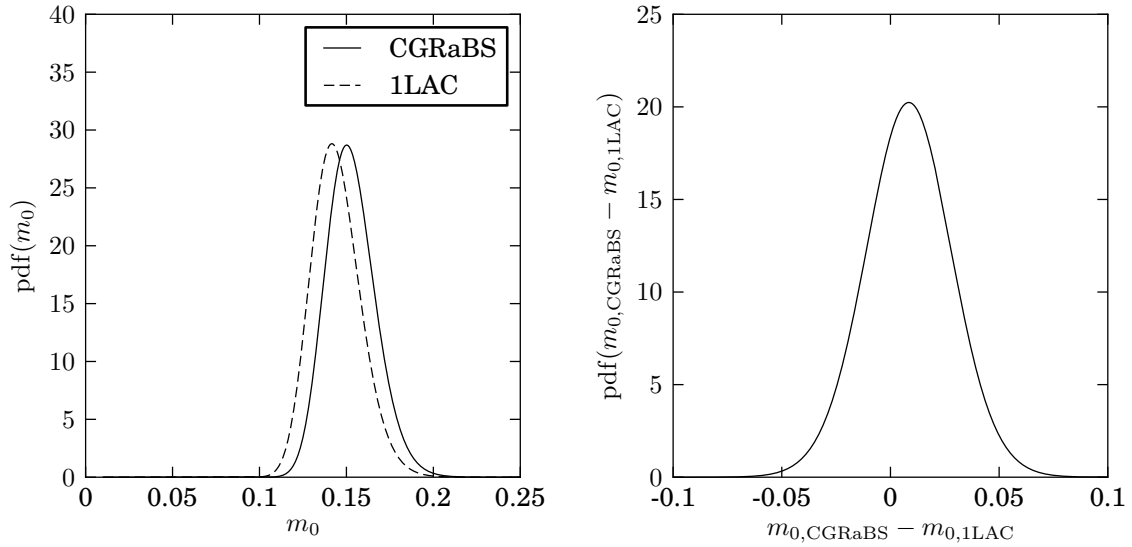

Figure 5.39. Comparison of CGRaBS and 1LAC BL Lac populations using the 42-month data. *Left:* Probability density of m_0 for the CGRaBS BL Lacs (solid line, maximum likelihood value and 1σ error $0.150^{+0.015}_{-0.013}$) and 1LAC BL Lacs (dashed line, maximum likelihood value and 1σ error $0.142^{+0.015}_{-0.013}$). The two distributions *are* consistent with a single value. *Right:* Probability density of the *difference* between the mean modulation index m_0 for the two sets. The peak of the distribution (0.008 ± 0.020) is less than 1σ away from zero.

strongly beamed than their gamma-ray–quiet counterparts. This could be due to different inverse Compton seed photon sources between the two classes. The Doppler beaming pattern for the Compton component of EC models is narrower about the jet axis than that for the synchrotron component, while both components are affected by the same beaming pattern in SSC sources (Dermer 1995). The SEDs of FSRQ sources frequently require EC models, while BL Lac objects with high-frequency SED peaks (HBLs) are usually consistent with SSC models (e.g., Böttcher 2007; Abdo et al. 2010c). While low spectrally peaked BL Lacs (LBLs) more often resemble FSRQs in this regard, *Fermi* was much more efficient at detecting the HBLs, which may be reflected in the statistics. Thus, it seems that for FSRQ sources, radio spectral properties alone are insufficient to predict gamma-ray emission. Additional data sensitive to the beaming angle, such as variability statistics, are also needed to discriminate between the gamma-ray–loud and gamma-ray–quiet FSRQs.

5.6.1 Flux Density Comparisons

We can also compare the brightness distributions of the CGRaBS and 1LAC samples. Figure 5.41, we compare the measured average flux density (S_0) distributions for subsets of CGRaBS sources in 1LAC and not in 1LAC (left panel), and for subsets of 1LAC sources in CGRaBS and not in CGRaBS (right panel). The overlap of the samples clearly preferentially selects the brighter sources in each. That is, the brightest CGRaBS predictions are more likely to be in 1LAC and the brightest 1LAC detections are more likely to be in CGRaBS.

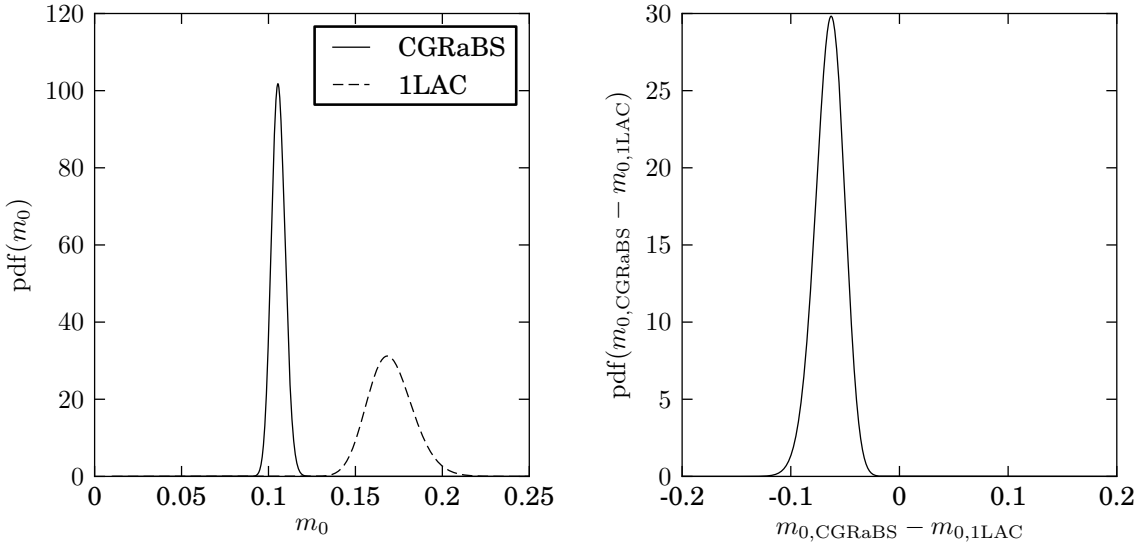

Figure 5.40. Comparison of CGRaBS and 1LAC FSRQ populations using the 42-month data. *Left:* Probability density of m_0 for the CGRaBS FSRQs (solid line, maximum likelihood value and 1σ error 0.105 ± 0.004) and 1LAC FSRQs (dashed line, maximum likelihood value and 1σ error $0.168^{+0.014}_{-0.012}$). The two distributions are *not* consistent with a single value. *Right:* Probability density of the *difference* between the mean modulation index m_0 for the two sets. The peak of the distribution ($0.063^{+0.014}_{-0.013}$) is almost 6σ away from zero.

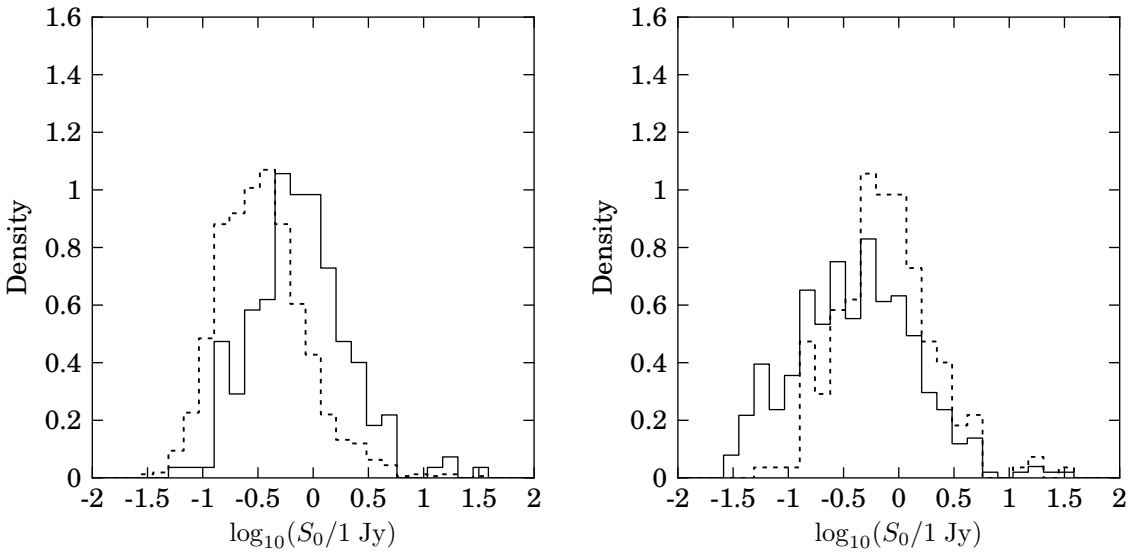

Figure 5.41. Histograms of S_0, normalized to integrate to unity. *Left:* Distribution of CGRaBS sources in 1LAC (solid) and not in 1LAC (dashed). A K-S test rejects the null hypothesis that these samples are drawn from the same distribution with $p < 5 \times 10^{-21}$. *Right:* Distribution of 1LAC sources in CGRaBS (dashed) and not in CGRaBS (solid). A K-S test rejects the null hypothesis with $p < 10^{-7}$. The dashed line in the right panel is the same as the solid line in the left panel.

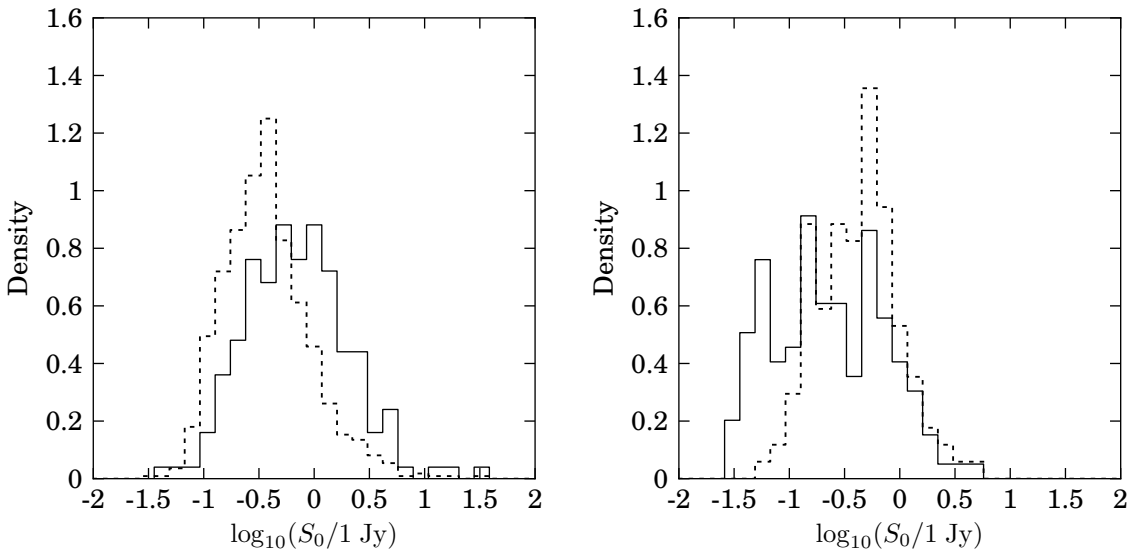

Figure 5.42. B
L Lac populations.] Histograms of S_0 for FSRQs (left) and BL Lacs (right) in the 1LAC (solid) and CGRaBS (dashed) samples. Each histogram is normalized to integrate to unity. A K-S test gives $p < 7 \times 10^{-12}$ (FSRQ) and $p < 4 \times 10^{-5}$ (BL Lac) for these data to come from the same distribution.

In figure 5.42, we again plot histograms of the distributions of the average flux density, S_0, now separating the FSRQ and BL Lac subpopulations of the two samples. The mean S_0 values for the FSRQs in the 1LAC and CGRaBS samples are (1.35 ± 0.20) Jy and (0.66 ± 0.05) Jy, respectively, and a K-S test gives $p < 7 \times 10^{-12}$ that the two data sets are drawn from the same distribution. Note that, as can be seen in figure 5.44, the 1LAC FSRQs in CGRaBS (mean $S_0 = (1.70 \pm 0.29)$ Jy) are brighter, on average, than the overall 1LAC FSRQ population. Thus, as we saw for the samples overall, brighter CGRaBS FSRQs were more likely to be detected by *Fermi* and brighter *Fermi*-detected FSRQs were more likely to have been predicted by the CGRaBS figure of merit (FoM) as gamma-ray emitters. Among BL Lacs, the mean S_0 for the 1LAC and CGRaBS samples are (0.46 ± 0.05) Jy and (0.62 ± 0.06) Jy, respectively, and a K-S test gives $p < 4 \times 10^{-5}$ that the two data sets come from the same distribution. Additionally, a number of sources were excluded from this histogram because the radio flux density measurements gave a $< 2\sigma$ detection. Most of these sources were 1LAC BL Lacs, so inclusion of these would further separate the 1LAC and CGRaBS BL Lac mean flux densities.

Among FSRQs, the result that gamma-ray–bright sources are on average brighter in the radio is consistent with the general trend found in radio-gamma connection studies, as discussed in section 1.3. If the increased variability we detect in the gamma-ray–loud FSRQs is a result of preferentially enhanced beaming in these sources, the higher Doppler factor in these sources could also give rise to this increased radio flux density. Some of this difference is probably attributable to the lower average redshift of 1LAC FSRQs, rather than intrinsic differences between sources. The reverse trend among BL Lacs—LAT-detected BL Lacs are somewhat dimmer on average than the CGRaBS BL Lacs—is most likely a result of the CGRaBS being a poor

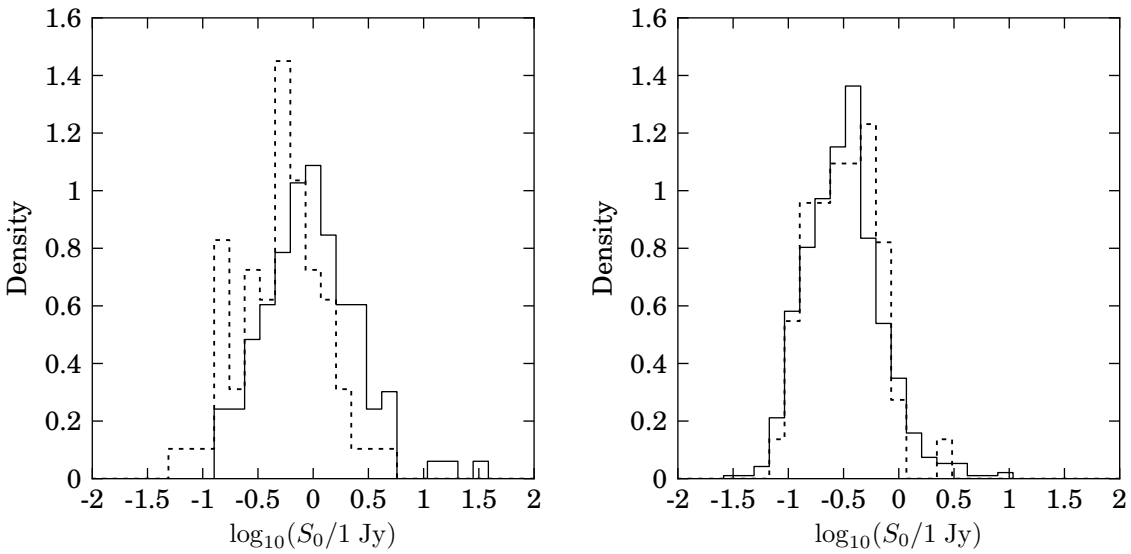

Figure 5.43. Histograms of S_0 for CGRaBS sources in 1LAC (left) and not in 1LAC (right), separately showing the distributions of FSRQ (solid) and BL Lac (dashed) samples. Each histogram is normalized to integrate to unity. A K-S test gives $p = 0.001$ (in 1LAC) and $p = 0.991$ (not in 1LAC) for these data to come from the same distribution.

prediction of the BL Lac sources likely to be detected by the LAT, reflecting the influence of the EGRET sensitivity bias on the CGRaBS selection. Thus, the LAT may simply be detecting a sample of BL Lac sources dimmer in gamma rays, and thus likely to be dimmer in radio as well.

This explanation is consistent with the average flux density distributions shown in figure 5.43, which show that among CGRaBS sources, the BL Lacs and FSRQs detected in gamma rays are quite different ($p = 0.001$ to come from the same distribution), while the gamma-ray–quiet sources in both classes are indistinguishable ($p = 0.991$ to come from the same distribution). Since the selection criteria for the CGRaBS predominantly chose FSRQs, it seems that many of the BL Lac sources in the were selected because their properties resembled those of FSRQs, not necessarily those of gamma-ray–loud BL Lacs.

Our data do not explain the population of 1LAC FSRQs that were not included in CGRaBS. It is possible this results from the use of archival radio data, which would reflect the emission state of the source at some past time rather than during the *Fermi* observation era. We do not find evidence that this is caused by an overall radio brightening: in figure 5.44 we see 1LAC detections that were well above the CGRaBS minimum radio flux density. There is no evidence that 1LAC sources are preferentially brighter in the concurrent 15 GHz data than non-1LAC sources. However, as we are comparing two different radio frequencies, it is possible that spectral index variation is masking an effect. This possibility is enhanced because the CGRaBS FoM depended strongly on the radio spectral index of the source.

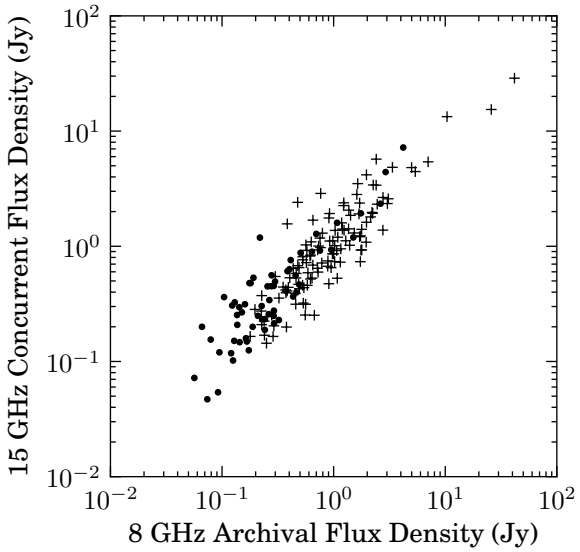

Figure 5.44. OVRO 40 m 15 GHz average flux densities (S_0) versus the archival 8 GHz flux densities tabulated in Ackermann et al. (2011) for 1LAC FSRQs. FSRQs in CGRaBS are plotted as crosses, those not in CGRaBS are plotted as points.

5.6.2 High Redshift FSRQ Populations

In addition to the differences in source classification, the redshift distributions of the 1LAC sources and CGRaBS differed. Within the BL Lac population, both 1LAC and CGRaBS were consistent with the same redshift distribution. Among FSRQs, however, the distributions were clearly not consistent. Examining the histograms in figures 4.5 and 4.6, it appears that the CGRaBS distribution is quite similar below $z \sim 1.5$ and the CGRaBS are clearly more numerous at high redshifts $z \gtrsim 1.5$.

To examine this, we compare the properties of the CGRaBS FSRQ populations in 1LAC and not in 1LAC at redshifts $z > 1.5$. Histograms of the intrinsic modulation indices are shown in figure 5.45. Although there are only 26 samples in the overlap between CGRaBS and 1LAC, the distribution extends to very high intrinsic modulation indices. The likelihood distributions for the population parameters are shown in figure 5.46. Despite the small sample size, the high-redshift CGRaBS FSRQs in 1LAC are significantly more variable than those not in 1LAC, with more than 3.5σ significance. The maximum-likelihood difference ($0.096^{+0.045}_{-0.034}$) is larger than was found between gamma-ray–loud and gamma-ray–quiet CGRaBS including all redshifts ($0.066^{+0.013}_{-0.012}$), but is consistent within 1σ of being the same.

Thus, it would seem that the CGRaBS selection criteria simply were not well tuned to picking out the systematically higher-variability sources that are bright gamma-ray emitters. By comparing *Fermi* 1LAC detections against the Australia Telescope 20 GHz survey catalog (AT20G), Mahony et al. (2010) conclude that the AGN in 1LAC are predominantly characterized by flat radio spectral indices and that there is no missing steep-spectrum population among the gamma-ray–loud population. This confirms that the most important CGRaBS selection criterion—a flat radio spectral index—is an effective method for selecting gamma-ray–

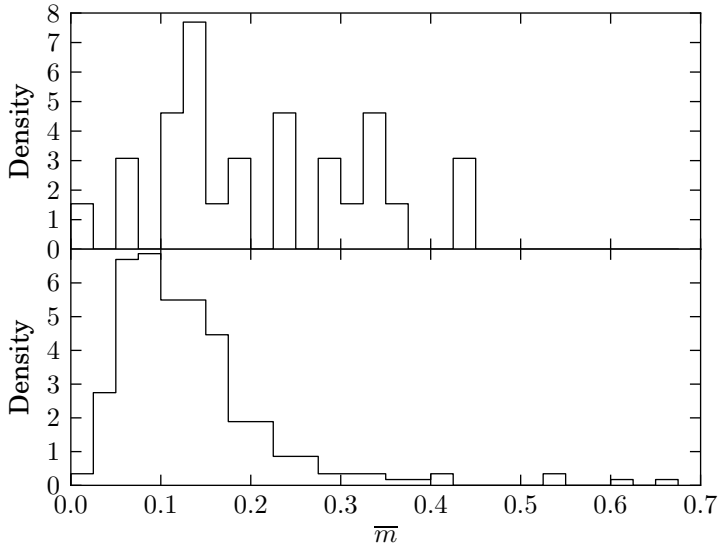

Figure 5.45. Histograms of 42-month intrinsic modulation index values for CGRaBS FSRQs at $z > 1.5$ in 1LAC (top, 26 sources) and those not in 1LAC (bottom, 233 sources). Each histogram is normalized to integrate to unity.

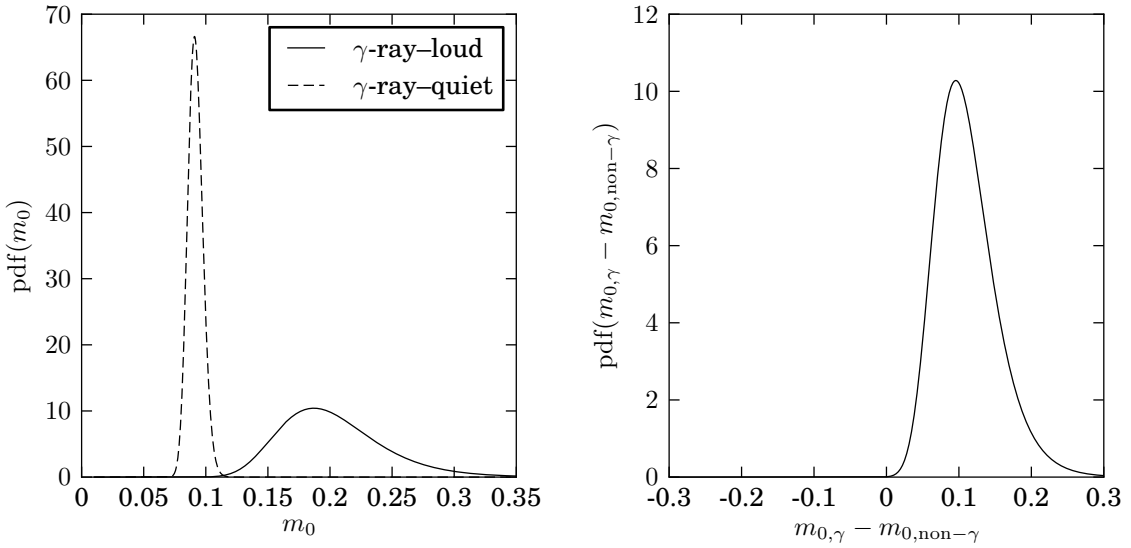

Figure 5.46. Comparison between high-redshift ($z \geq 1.5$) CGRaBS FSRQ populations in and not in the 1LAC sample, using the 42-month data. *Left:* Probability density of m_0 for the $z \geq 1.5$ CGRaBS FSRQs in the 1LAC sample (solid line, maximum likelihood value and 1σ error $0.187^{+0.044}_{-0.034}$) and those not in 1LAC (dashed line, maximum likelihood value and 1σ error 0.091 ± 0.006). The two distributions are *not* consistent with a single value. *Right:* Probability density of the *difference* between the mean modulation index m_0 for the two sets. The peak of the distribution ($0.096^{+0.045}_{-0.034}$) is more than 3.5σ away from zero.

loud AGN. That there is no evidence for a missing population is consistent with our conclusion that the CGRaBS sample is drawn from a superset of the 1LAC parent population. Although it is unclear precisely what mechanism connects between 15 GHz radio variability and gamma-ray emission, it seems evident that the presence of variability in radio is a good predictor of gamma-ray emission from a blazar.

Appendix A

User's Guide to Arcreduce

Arcreduce is a Python module that provides methods for reducing data from the 40 m data archives. Its interface is provided primarily through two classes: `CalManager`, a high-level data reduction interface, and `ArchiveReader`, a lower-level engine for directly accessing the archives to decode procedure data. In addition, a number of methods and classes are included in the module to support these interfaces. Arcreduce is built atop the readarc module provided by Martin C. Shepherd. The readarc module is a Python wrapper around a C library that provides raw access to the data in the archive. Arcreduce was written in order to simplify access to the archive and to provide a more "Pythonic" set of interfaces.

In addition to readarc, Arcreduce depends on several other Python modules. These include the public libraries NumPy,[1] SciPy,[2] and PyEphem.[3] The py40m module, a library of useful routines developed alongside the 40 m reduction pipeline, is also required.

A.1 High-Level Data Reduction: `CalManager`

`CalManager` is the main class intended for high-level analysis and reduction of data recorded with the MCS control system. It is built upon `ArchiveReader`, and uses that class to read data from the archive and convert it into reduced procedure result data. `CalManager` provides tools to simplify the operations needed to examine and calibrate those data.

A.1.1 `CalManager` Concepts

The work flow within `CalManager` is modeled after a small subset of CMBPROG (Leitch 1998). The results of all executions of a particular procedure, are stored in an instance of the `Procedure` class. Within the procedure are a set of vector-like `Member` objects, each of which represents a time series of values from the results of that procedure.

[1] `http://numpy.scipy.org/`
[2] `http://scipy.org/`
[3] `http://rhodesmill.org/pyephem/`

Table A.1. List of `CalManager` procedures

Attribute	Radiometry Procedure	Members
`flux`	FLUX	a, b, c, d, atp, btp, ctp, dtp, swd, swp
`cal`	CAL	diode, a, b, c, d, atp, btp, ctp, dtp, swd, swp
`point`	POINT	snr, hpbw, failed
—	common to all	source, flux, flags, time, mjd, azo, zao, az, za, taf, tlr, focus, pa, samples

Note: Data for each procedure are contained in the members `Procedure` accessed as an attribute of the `CalManager` object. Each procedure type contains the common members listed, as well as members unique to its function. Table A.2 defines the various members.

The reader should be aware that the word "procedure" is frequently used in this section to refer either to a specific execution of a radiometry procedure, to the `Procedure` class, or to an instance of that class. It should be clear from context which meaning is intended, but careful reading is warranted.

A.1.1.1 `Procedure` Class

When data are loaded from the archive, the radiometry procedures are processed and decoded into a stream of results. The `Procedure` class represents the time-ordered series of results of one of these procedures. Table A.1 lists the procedures that are available. These procedures are accessible as attributes of the `CalManager` instance, so results from, say, FLUX procedures are stored in `cm.flux`.

The `CalManager` class defines a number of methods needed to manage and calibrate the data stored in the `Procedure` objects. Generally, methods defined in the `CalManager` class affect the calibration process as a whole or operate on more than one `Procedure` object. Methods that affect a single `Procedure` or `Member` are normally defined within those classes. For historical reasons, a number of relevant methods are defined at the top level of the `arcreduce` module.

A.1.1.2 `Member` Class

The `Member` class represents a member of a procedure—a series of time-ordered data. All members within a procedure share a common set of time stamps. The time stamps themselves are represented by the `time` member of the procedure. The `Member` class supports data masking to allow the parent `Procedure` to restrict operations to a subset of the values in the time series. This is used, for example, when working with FLUX results for a single source or CAL results for only the NOISE diode. Table A.2 lists and briefly describes the members that belong to the various procedures.

Vector arithmetic operations are supported for `Member` objects that contain numerical values. To avoid wastefully copying data, unary operations (e.g., `x+=y`) operate in place on the data within the `Member`. When a copy is desired, binary operations (e.g., `z=x+y`) allocate a new `Member` to store the result. This behavior is illustrated in the examples in section A.1.2.3.

Table A.2. Descriptions of the `CalManager` procedure members

Member	Used by	Description
`source`	all	Name of source being observed
`time`	all	UTC time stamp
`mjd`	all	MJD time stamp (derived from `time`)
`flags`	all	Reduction flags
`az`	all	Telescope azimuth
`za`	all	Telescope zenith angle
`taf`	all	Aft-forward tilt meter reading
`tlr`	all	Left-right tilt meter reading
`focus`	all	Telescope z-axis focus setting
`pa`	all	Parallactic angle of source
`samples`	all	Number of samples in procedure (definition varies by procedure)
`diode`	`cal`	Name of diode measured ("CAL" or "NOISE")
`a`, `b`, `c`, `d`	`flux`, `cal`	Measured switched power in each segment of the procedure
`atp`, `btp`, `ctp`, `dtp`	`flux`, `cal`	Measured total power in each segment of the procedure
`swd`, `swp`	`flux`, `cal`	"Switched difference" and "switched power" diagnostic signals, as described in section 2.2.2.2.
`hpbw`	`point`	Half-power beam width used for pointing measurement
`snr`	`point`	Signal-to-noise ratio of the pointing measurement
`failed`	`point`	True if a pointing measurement failed or gave an unreliable result
`azo`	all	Azimuthal offset from the pointing model (contains the result of a pointing measurement)
`zao`	all	Zenith angle offset from the pointing model (contains the result of a pointing measurement)

When relevant, a `Member` represents both a value and its uncertainty. Arithmetic operations propagate uncertainties between members assuming they are random and uncorrelated. Care must be taken if this assumption is not true. Constant values are assumed to have no uncertainty.

If access to the data values or uncertainties is needed, the `get_val()` or `get_err()` methods should be used. These return a NumPy `ndarray` with the requested values, after applying any active masks. Data stored in the member can be directly updated using the `set_val()` or `set_err()` method.

A.1.1.3 Masking and Flagging

When working with data, it is frequently useful to restrict the working data set to one or a few sources of interest. For the CAL procedure, normally only one diode (CAL or NOISE) is of interest at a time. To support this, `CalManager` provides a mechanism for masking procedures based on source or diode names. Initially upon load, all data are active. If a mask is subsequently applied, only those procedures that satisfy the mask are operated on until the mask is changed. The source mask is specified by passing a source name

or list of source names to the `set_active_sources()` method. The diode mask is specified via the `set_active_diodes()` method.

During reduction, data that are unreliable must be identified and discarded. `CalManager` allows data points to be flagged (stored as a bit field in the `flags` member) to identify the reason the data point was discarded. Flagged data are not used during later operations nor retrieved via the `get_val()` or `get_err()` methods. The `remove_flag()` method can be used to remove all instances of a particular flag from a `Procedure`.

A.1.1.4 Plotting and Advanced Processing

Python supports powerful numerical processing (e.g., the NumPy and SciPy packages) and plotting (e.g., Matplotlib). Routines of this nature are not implemented in `CalManager`. Instead, only basic, common operations are directly implemented for the `Procedure` and `Member` classes. When necessary, the data contained within a `Member` can be easily extracted to a NumPy `ndarray` and plotted or processed using a suitable external routine.

A.1.2 `CalManager` Tutorial

This simple tutorial demonstrates the essential functions of `CalManager` and its related classes. Before starting, the `arcreduce` module must be loaded and a `CalManager` object must be instantiated.

```
>>> import arcreduce as ar
>>> cm = ar.CalManager()
```

A.1.2.1 Loading Data

To load procedure data from the archive, the `load_data` method of the `CalManager` object is used. The start and end dates to be loaded can be specified as calendar dates or MJD days.

```
>>> cm.load_data('2011-05-01 00:00:00', '2011-05-01 12:00:00')
```

This example loads and decodes the procedures for 12 hours of data, beginning at midnight UTC on May 1, 2011.

A.1.2.2 Examining a `Procedure`

We can find the names of the members within a `Procedure` object as follows.

```
>>> print cm.flux.members
['source', 'flags', 'time', 'mjd', 'flux', 'az', 'za', 'taf', 'tlr',
 'focus', 'samples', 'azo', 'zao', 'pa', 'a', 'b', 'c', 'd',
 'atp', 'btp', 'ctp', 'dtp']
```

Each member of the procedure is represented by a `Member` object, which can be accessed as an attribute of the `Procedure` object. `Member` objects provide a convenient pretty-printing interface, illustrated here.

```
>>> print len(cm.flux)
212
>>> print cm.flux.flux
[   485.656 +/- 0.363,
   148.228 +/- 0.524,
   45.524 +/- 0.394,
   ...]
```

We see that 212 FLUX procedures were decoded. When the `flux` member of `cm.flux` is printed, the first few values are shown, along with their errors. Not all members have errors defined—for those without errors, these are shown as zero.

A.1.2.3 Working with a `Member`

`Member` objects that contain numerical data support vector arithmetic operations. In-place operations are performed using the unary Python arithmetic operators.

```
>>> cm.flux.flux *= 2
>>> print cm.flux.flux
[   971.313 +/- 0.726,
   296.456 +/- 1.047,
   91.048 +/- 0.787,
   ...]
```

Note that the errors were propagated assuming that the constant value in the multiplication was exact.

If two `Member` objects have the same length, they can be combined arithmetically. Within a single `Procedure`, all `Member` objects are guaranteed to be the same length, so this is always possible.

```
>>> cm.flux.flux += cm.flux.flux
>>> print cm.flux.flux
[   1942.625 +/- 1.026,
   592.912 +/- 1.481,
   182.097 +/- 1.113,
   ...]
```

The errors of the two inputs (here both `cm.flux.flux`) are combined to produce the error in the output. Note that the errors are propagated assuming they are uncorrelated (which is incorrect in this case).

If a binary Python operator is used, a new `Member` is created to contain the result. This new `Member` can (and normally would) be assigned to a `Procedure` as we do here.

```
>>> cm.flux.flux_copy = cm.flux.flux / 4.0
>>> print cm.flux.flux_copy
[   485.656 +/- 0.257,
   148.228 +/- 0.370,
   45.524 +/- 0.278,
   ...]
```

We now have a new `Member` whose values are equal to the initial values we loaded. The errors are smaller than they were at the start because they were assumed to be uncorrelated in the previous example. The `cm.flux.flux` object is unchanged because we created a copy.

A.1.2.4 Accessing Data from a `Member`

While basic arithmetic operations on `Member` data are supported directly, it is frequently useful to use ordinary Python methods to work with the data. This is possible using the `get_val()` and `get_err()` methods of the `Member`.

```
>>> val = cm.flux.flux.get_val()
>>> err = cm.flux.flux.get_err()
>>> val[0:3]
array([ 1942.62528571,   592.91225   ,   182.09658333])
>>> err[0:3]
array([ 1.02642661,  1.48117483,  1.11319623])
```

Python data may be inserted into a `Member` using the `set_val()` and `set_err()` methods. Here, we'll correct the uncertainty that was incorrectly propagated when we added `cm.flux.flux` to itself.

```
>>> import numpy as np
>>> cm.flux.flux.set_err(np.sqrt(2)*err)
>>> cm.flux.flux.err[0:3]
array([ 1.45158643, 2.09469753, 1.57429721])
```

When masks are applied to a procedure, such as when a specific source or diode is selected, these get and set methods act only on the actively selected elements of the procedure.

A.1.2.5 Selecting a Source or Diode

To restrict a procedure to a particular source, we call the `CalManager` object's `set_active_sources()` method. A similar method, `set_active_diodes()`, allows CAL procedures to be restricted to either the CAL or the NOISE diode.

```
>>> sources=cm.flux.source.get_val()
>>> print sources[0:3]
['j0750+1231' 'j0811+0146' 'j0805-0111']
>>> cm.set_active_sources('j0750+1231')
>>> print cm.flux.flux.get_val()
[ 1942.62528571]
```

To select multiple sources or diodes, simply pass these as a list of source names. To reenable all available sources or diodes, call this procedure with an empty list.

```
>>> cm.set_active_sources([])
```

A.1.2.6 Applying a Calibration Factor

Calibration by dividing by a filtered member or by a polynomial function of a member are supported. For example, we can normalize the flux density by the interpolated value of the CAL diode to remove the effect of gain fluctuations. As an example, here we simply calibrate the CAL diode flux member by itself. The deviation from 1.0 results from the averaging of nearby values.

```
>>> cm.set_active_diodes('CAL')
>>> print cm.cal.flux.get_val()[0:3]
[ 471.95144286  472.43159286  471.22714286]
>>> cm.apply_flux_cal(cm.cal, cm.cal.flux, 1.0)
>>> print cm.cal.flux.get_val()[0:3]
[ 0.99949928  1.00051614  0.99920581]
```

A.1.2.7 Flagging Data

Unreliable data points can be flagged to remove them from further processing and to indicate the nature of the problem with the data point. Table A.3 lists the supported flags. Several processing routines will implicitly flag data points when problems are encountered. Additionally, explicit flagging based on data or uncertainty values can be performed. In this example, we flag CAL values between 0 and 1, leaving only values greater than 1 unflagged.

```
>>> print cm.cal.flux.get_val()[0:3]
```

Table A.3. Flag values supported by `CalManager`

Name	Value	Description
FLAG_OUTLIER	1	Outlying data point flagged by `arcreduce.flag_outliers()`.
FLAG_SWEEP	2	Outlying data point flagged by `arcreduce.flag_sweep_outliers()`.
FLAG_VALUE_CLIP	4	Value was explicitly flagged by the user via `arcreduce.clip_values()`.
FLAG_ERROR_CLIP	8	Uncertainty was explicitly clipped by the user via `arcreduce.clip_errors()`.
FLAG_FAILED_POINT	16	Procedure was preceded by a failed (or no) point.
FLAG_BAD_CAL	32	Problem interpolating calibration diode to the data time stamp.

```
[ 0.99949928  1.00051614  0.99920581]
>>> ar.clip_values(cm.cal.flux, 0, 1.0, inclusive=True)
>>> print cm.cal.flux.get_val()[0:3]
[ 1.00051614  1.00238618  1.00077693]
```

A.1.2.8 Example Reduction Script

Listings A.1 and A.2 are excerpts of the Python reduction pipeline script and demonstrate actual use of the `CalManager` class.

A.1.3 Module Reference

- `class arcreduce.`**`CalManager`**: Interface class to manage high-level reduction and calibration.

  - **`load_data`**(*`start_date, end_date`*)

    Loads data from the archive.

  - **`apply_flux_cal`**(*`target_proc, source_mem, cal_value, inverse=False`*)

    Scales appropriate flux density-like members of *`target_proc`* by dividing by the value of *`source_mem`* at each sample. This is normally used to calibrate flux densities by dividing by an interpolated CAL member. *Note: this method implicitly applies a 7200 s boxcar interpolation to the source member to reconcile its time base with that of the target procedure.*

  - **`apply_polynomial_cal`**(*`target_mem, source_mem, poly, inverse=False`*)

    Divide flux density-like members of *`target_proc`* by a polynomial evaluated at the value of *`source_mem`* at each sample. This is normally used to apply the gain-versus-elevation curve using the `flux.za` member as *`source_mem`*.

  - **`set_active_diodes`**(*`diodes`*)

```
def reduce_data(start_date, stop_date):
    """Reduce data between given dates."""
    cm = arcreduce.CalManager()
    cm.load_data(start_date, stop_date)

    # clip low-SNR POINTs
    arcreduce.clip_values(cm.point.snr, -np.inf, 2, inclusive=True)

    # flag FLUXes with bad pointings
    arcreduce.flag_failed_point(cm.flux, cm.point, 4800)

    # clip total powers to plausible levels
    arcreduce.clip_values(cm.flux.atp, 10000, 50000, inclusive=False)
    arcreduce.clip_values(cm.flux.btp, 10000, 50000, inclusive=False)
    arcreduce.clip_values(cm.flux.ctp, 10000, 50000, inclusive=False)
    arcreduce.clip_values(cm.flux.dtp, 10000, 50000, inclusive=False)

    # clean the cal procedure
    cm.set_active_diodes(['CAL']) # use small cal diode
    clean_cal(cm.cal)

    # apply the CAL diode
    cm.apply_flux_cal(cm.flux, cm.cal.flux, 8.33/2.0)

    # All done.
    return cm
```

Listing A.1. The CalManager class is the foundation for calibration of data in the reduction pipeline.

Limits all CAL procedures to those using the specified diodes (CAL or NOISE). To activate all diodes, call with *diodes*=None. By default, all diodes are active.

- **set_active_sources**()

  Limits all procedures to those with the specified source name. To activate all sources, call with *sources*=None. By default, all sources are active.

- **boxcar_interpolate**(*source_time, source_y, target_time, dt_max, n_min, dy_max*)

  Creates a new Member containing the $\pm$*dt_max* seconds boxcar interpolation of the *source_y* member with sample times from *source_time*. The new samples are evaluated at times from *target_time*. If any such interpolation has fewer than *n_min* input samples or if

  $$|\max(y) - \min(y)| \,/\, |\min(y)| \geq \texttt{dy\_max}$$

  within that bin, that sample is flagged with the BAD_CAL flag.

- class arcreduce.**Procedure**: Container for the results from a particular radiometer procedure. These results are stored in Member objects that are accessible as attributes of the Procedure.

```
def clean_cal(cal):
    """Apply the cal cleaning process to the procedure."""
    # first flag data that's equal to zero (with a little wiggle-room
    # since they're float values)
    arcreduce.clip_values(cal.a, -1e-6, 1e-6, inclusive=True)
    arcreduce.clip_values(cal.b, -1e-6, 1e-6, inclusive=True)
    arcreduce.clip_values(cal.c, -1e-6, 1e-6, inclusive=True)
    arcreduce.clip_values(cal.d, -1e-6, 1e-6, inclusive=True)

    # Now clip values less than 5 which are completely unbelievable.
    arcreduce.clip_values(cal.flux, -np.inf, 5, inclusive=True)

    # Clip values with implausible measured uncertainties.  The range
    # 0 to 6 is acceptable; flag OUTSIDE that range.
    arcreduce.clip_errors(cal.flux, 0, 6, inclusive=False)

    # Apply iterative outlier filters.  We apply to swp twice because
    # the swd filter sometimes triggers a few new swp outliers.
    arcreduce.flag_outliers(cal.swp, 4)
    arcreduce.flag_outliers(cal.swd, 4)
    arcreduce.flag_outliers(cal.swp, 4)

    # Now sweep along in day-long buffers and flag outliers
    arcreduce.flag_sweep_outliers(cal.mjd, cal.flux, 1.0, 3.5)
```

Listing A.2. Helper function used by the reduction routine in Listing A.1.

`Member` objects can be added to a `Procedure` using simple Python assignment, e.g., `proc.new_mem=Member(...)`.

  - **`remove_flag`**(*flag*)

    Remove the specified flag from all data points in each `Member` of the `Procedure`.

- `class arcreduce.`**`Member`**: Representation of one parameter or result of the execution of a procedure.

  - **`get_err`**()

    Return a NumPy `ndarray` containing the uncertainties of the active elements of the `Member`.

  - **`set_err`**(*err*)

    Set the uncertainties of the active elements of the `Member` to the values given.

  - **`get_val`**()

    Return a NumPy `ndarray` containing the data values of the active elements of the `Member`.

  - **`set_val`**(*val*)

    Set the data values of the active elements of the `Member` to the values given.

- `arcreduce.`**`clip_errors`**(*mem, emin, emax, inclusive=False*)

Apply the ERROR_CLIP flag to data when the uncertainty in member *mem* is outside (inside if *inclusive* is True) the specified range.

- arcreduce.**clip_values**(*mem, xmin, xmax, inclusive=False*)

  Apply the VALUE_CLIP flag to data when the value in member *mem* is outside (inside if *inclusive* is True) the specified range.

- arcreduce.**flag_failed_point**(*target_proc, point_proc, dt_max*)

  Apply the FAILED_POINT flag to data in *target_proc* if the procedure in *point_proc* that immediately precedes it failed, or if there is no procedure within *dt_max* seconds prior.

- arcreduce.**flag_outliers**(*mem, nsigma*)

  Iteratively apply the OUTLIER flag to data that lie more than *nsigma* standard deviations from the mean of the *mem* member until no further data are flagged.

- arcreduce.**flag_sweep_outliers**(*xmem, ymem, dx, nsigma*)

  Apply the SWEEP flag to data in a sliding window spanning $\pm dx/2$ along *xmem* when the value in member *ymem* exceeds *nsigma* standard deviations from the mean in the window.

A.2 Low-Level Data Processing: `ArchiveReader`

The ArchiveReader class is the low-level Python interface to the data archive. It relies on the readarc module to extract data from the binary archive files recorded by the MCS control system. Details about the architecture and use of this class are described in section 3.1.2.2. In this section, we provide a detailed module reference for the classes and methods relevant to the ArchiveReader class.

A.2.1 `ArchiveReader` Class

- class arcreduce.**ArchiveReader**: Pythonic object-oriented wrapper around the readarc library. Instances of this class manage a data reduction session.

  - **add_frame_handler**(*self, fh*)

    Add a decoder that should be notified when data from a new frame are read. The object must implement the decoder interface and should normally be a subclass of the GenericDecoder class.

  - **add_output_handler**(*self, oh*)

    Add an output handler to be notified when a completed procedure result is available from any of the attached decoders. The object must implement the object handler interface and should normally be a subclass of the OutputHandler class.

– **add_register**(*self, reg_name, reg_type, index=0, length=None, our_name=None*)

Add a register to the set that will be extracted from each frame.

* *reg_name*: name of the register in the archive
* *type*: one of the RT_* types indicating the type of the register
* *index*: index of the register element to read (default: 0)
* *length*: max length of the array to be returned, only used for arrays where it must be set
* *our_name*: name by which to refer to the register (see below)

If a register is added with an existing name, an `ArchiveError` will be raised if the parameters for the new register do not exactly match the existing entry. Otherwise the second addition of the register will have no effect.

If a local name is needed (e.g., to provide a shorthand name for one index in a multi-index register), specify it with `our_name`. By default, registers are known by their official names. For example, to extract `mount.tracker.horiz_actual[1]` and refer to it by the name `mount.tracker.horiz_actual_el`, one would call

```
add_register('mount.tracker.horiz_actual', RT_FLOAT, index=1,
    our_name='mount.tracker.horiz_actual_el')
```

– **dispatch_frame**(*self*)

Call all registered frame handlers to notify of new frame.

Might be able to allow handlers to notify this class which frame labels are of interest, but would have to add a protocol to let them detect the end of a sequence of frames. For now easier to just let them implement that as needed.

Handler should return a `ProcedureData` object when a complete procedure has been processed. It will then be output using the current output method. If a procedure is not complete, return `None`.

– **handle_completed_procedure**(*self, p*)

Manage output of a completed procedure by calling output handlers.

– **handle_frame**(*self*)

Process the next frame. Returns `True` until all frames are consumed. After each call, the label of the frame just processed will be available from the `self.last_frame_label` member. (This will be `None` before the first or after the last frame).

– **read_register**(*self, our_name*)

Read a register, properly managing the data type.

  - **update_registers**(*self*)

    Read next frame and all the configured registers in an unspecified order. Returns `True` on success, `False` if no more frames were available.

A.2.2 Decoders

- class arcreduce.**GenericDecoder**: Generic base class for decoder implementations.

  A decoder is an object that can be called by the `ArchiveReader` to process frames for a particular type of procedure (or, perhaps, for other reasons). The `GenericDecoder` base class provides generally useful features, such as a dynamically sized buffer for storing samples.

  Buffers are stored as NumPy `ndarray` objects with an array size that is generally larger than the valid data stored. The `buffer_index` member stores the index of the next element to be written, which equals the number of valid entries in the array. Buffers added through `add_buffer()` will be managed (via `reset_buffers()` and `grow_buffers()`). These can be accessed as member variables (i.e., `self.x`). There is some risk of namespace collisions, but this is unlikely.

  Procedure decoders should be subclasses of `GenericDecoder`. To implement the decoder interface, they must provide `handle_frame(`*`self, timestamp, frame_label, registers`*`)` to process data from each frame and a `install_registers(`*`self, ar`*`)` method to notify their parent `ArchiveReader` instance of the registers from which they require data. These methods are not explicitly listed in the references below for subclasses of this class.

  - **add_buffer**(*self, buffer_name, dtype, register_name=None*)

    Add a buffer to be managed.

  - **finalize_buffers**(*self*)

    Clip the buffers to their actual length so that they can be accessed without considering the `buffer_index` member. Do not load any additional data after doing this.

  - **grow_buffers**(*self*)

    Increase the size of the managed buffers by a factor of 2.

  - **reset_buffers**(*self*)

    Resets the data buffers and associated data to prepare for a fresh decoding attempt.

  - **update_buffers**(*self, registers, data_len*)

    Pull data out of the register associated with each buffer. Pulls `data_len` samples out of each register. Grows buffers as needed.

- class arcreduce.**MillisecondSampleDumpDecoder**(GenericDecoder): Decoder for dumping millisecond samples. Implements the decoder interface.

- class arcreduce.**Point2dDecoder**(GenericDecoder): Class to decode POINT2D procedures. The FWHM beam width must be specified when instantiating this object. Implements the decoder interface.

  - **decode**(*self*)

    Decode the results and compute flux / bg levels and some statistics.

    We do not fit the beam ourselves, we just take the results from the on-line procedure. We compute the signal-to-noise ratio (SNR) in a method roughly analogous to the old control system's SNR method, which was the minimum of the peak or half-power SNRs. Here, we take the minimum among any points measured at half-power or greater.

    Raises PointError if there is a problem decoding the data.

  - **find_result_frame_index**(*self*)

    Find the frame that indicates the completion of the procedure.

  - **find_valid_data**(*self, start, end, acq_delay=6*)

    Find valid data for integrating.

    Looks only between start and end indexes. Currently just identifies the last run of source-acquired data according to tracker_state. Returns indexes of valid data.

    Note: indexes are relative to the same zero-point as start/end, i.e., the absolute origin of the buffers.

    Waits acq_delay frames (seconds) after acquisition before it considers the source really acquired.

    Raises:

    * AcquisitionLostError: lost acquisition during the segment.
    * NoDataError: no valid data available.

  - **finish_proc**(*self, timestamp*)

    Called when the end of a procedure has been found. Processes the data and fills in a PointProcedureData object with the result.

  - **fit_gauss_amp**(*self, x, y, y_err, fwhm, amp0, bg0*)

    Helper function. Fits amplitude and background level of Gaussian.

    * fwhm: FWHM of the Gaussian
    * amp0: initial guess for amplitude
    * bg0: initial guess for background level

  - **identify_segments**(*self*)

    Identify pointing segments in the data.

This only identifies the start and end index of the period during which the telescope was intended to track a particular offset. It does not check that the telescope was acquired, that is done elsewhere.

- **integrate_segment**(*self, start, end*)

  Integrate a single segment.

- **integrate_segments**(*self, segs*)

  Integrate valid data within all the segments.

- class arcreduce.**PointDecoder**(GenericDecoder): Class to decode POINT procedures. Implements the decoder interface.

  - **find_segments**(*self*)

    Locate the start and and indexes of each of the POINT segments.

    Since the details of when tracker acquisitions occur is a bit fragile, we look at the commanded tracker offsets as a more reliable indicator of what is going on. We count the final acquisition within each segment as the one intended for integration.

  - **finish_proc**(*self, timestamp*)

    Called when the end of a procedure has been found. Processes the data and fills in a PointProcedureData object with the result.

- class arcreduce.**RegisterDumpDecoder**(GenericDecoder): Decoder class for dumping samples. Only works for one-sample-per-frame registers (will read only the first element from each frame for array registers). Implements the decoder interface.

- class arcreduce.**SimpleAbcdDecoder**(GenericDecoder): Class for decoding ABCD procedures. This includes FLUX and CAL procedures, and eventually perhaps others. Implements the decoder interface.

  To create a decoder, subclass this and call its __init__() routine. Pass a four-element list with the names of the labels that correspond to the four states of the procedure. For a FLUX, these are 'flux:A' through 'flux:D'.

  - **finish_proc**(*self*)

    An entire procedure has been received, process it.

  - **flux_a_handle_frame**(*self, timestamp, frame_label, registers*)

    Handle 'flux:A' state and keep accepting 'flux:A' frames until 'flux:B' comes along.

  - **flux_b_handle_frame**(*self, timestamp, frame_label, registers*)

    Handle 'flux:B' state and keep accepting 'flux:B' frames until 'flux:C' comes along.

- **flux_c_handle_frame**(*self, timestamp, frame_label, registers*)

  Handle 'flux:C' state and keep accepting 'flux:C' frames until 'flux:D' comes along.

- **flux_d_handle_frame**(*self, timestamp, frame_label, registers*)

  Handle 'flux:D' state and keep accepting 'flux:D' frames, then wrap up.

- **handle_frame**(*self, timestamp, frame_label, registers*)

  An entire procedure has been received, process it.

- **identify_segment**(*self, segment_num, i_start, i_end*)

  Pick out a valid segment and return the valid indexes for the segment. Override this to change the behavior of identifying segments.

- **init_abcd**(*self*)

  Set or reset data to prepare for a new procedure.

- **store_samples**(*self, registers*)

  Store samples to the appropriate location in `self.switched_samples`.

  Updates auxiliary variables as necessary. Data to copy should be passed in the `new_samples` argument. Extra samples will be ignored.

- **waiting_handle_frame**(*self, timestamp, frame_label, registers*)

  Method to implement the "waiting" state: goes from anything to 'flux:A'.

- class arcreduce.**SimpleCalDecoder**(SimpleAbcdDecoder): Class for decoding CAL procedures. This uses only the mean value from each frame rather than the low-level millisecond samples. Implements the decoder interface.

  - **finish_proc**(*self*)

    Finish the procedure. Overrides the parent class. Uses the `SimpleAbcdDecoder` to do the real work, but need to do a few CAL-specific checks also.

  - **identify_segment**(*self, segment_num, i_start, i_end*)

    Identify segment indices for a CAL procedure.

- class arcreduce.**SimpleFluxDecoder**(SimpleAbcdDecoder): Class for decoding FLUX procedures. This uses only the mean value from each frame rather than the low-level samples. Implements the decoder interface.

  - **finish_proc**(*self*)

    Finish the procedure. Overrides the parent class. Uses the `SimpleAbcdDecoder` to do the real work, but need to do a few FLUX-specific checks also.

A.2.3 Output Handlers

- class arcreduce.**OutputHandler**: Generic output handler class, just defines the interface. Subclasses must provide a handle_procedure(*self, p*) to process ProcedureData objects containing processed results.

- class arcreduce.**CallbackOutputHandler**(OutputHandler): Output class that passes the output procedure to a callback. The callback should accept two parameters. First is the procedure just completed, second is the argument passed to the __init__() routine. Implements the output handler interface. This is the output handler used for actual data reduction in the pipeline.

- class arcreduce.**FileOutputHandler**: Output class that dumps to a file handle. Default output file is stdout.

  - **handle_cal_procedure**(*self, p*)

    Output details of a CAL procedure.

  - **handle_flux_procedure**(*self, p*)

    Output details of a FLUX procedure.

  - **handle_point_procedure**(*self, p*)

    Output details of a POINT procedure.

A.2.4 Procedure Data Structures

- class arcreduce.**ProcedureData**: Generic class to contain procedure output data.

- class arcreduce.**AbcdProcedureData**(ProcedureData): Class to contain output data relevant to ABCD procedures. Not normally used directly.

- class arcreduce.**CalProcedureData**(AbcdProcedureData: CAL procedure data output class.

- class arcreduce.**FluxProcedureData**(AbcdProcedureData): FLUX procedure data output class.

- class arcreduce.**PointProcedureData**: Point procedure data output class.

A.2.5 Exceptions

- class arcreduce.**ArchiveReduceError**(Exception): Superclass for all the exceptions defined in this module.

- `class arcreduce.`**`ArchiveError`**`(ArchiveReduceError)`: Generic exception for errors during archive data processing.
- `class arcreduce.`**`AcquisitionLostError`**: Irrecoverably lost acquisition of a source during a procedure.
- `class arcreduce.`**`FluxError`**: Error encountered while processing a procedure with a `SimpleAbcdDecoder` (FLUX or CAL).
- `class arcreduce.`**`NoDataError`**: Encountered a procedure segment with no valid data.
- `class arcreduce.`**`OutputHandlerError`**: Error encountered during output handling.
- `class arcreduce.`**`PointError`**: Error occurred during handling of a POINT procedure.

Appendix B

Detailed Database Specification

This appendix contains diagrams for the database tables that were not shown in section 3.3 and the contents of the domain tables.

B.1 Database Table Diagrams

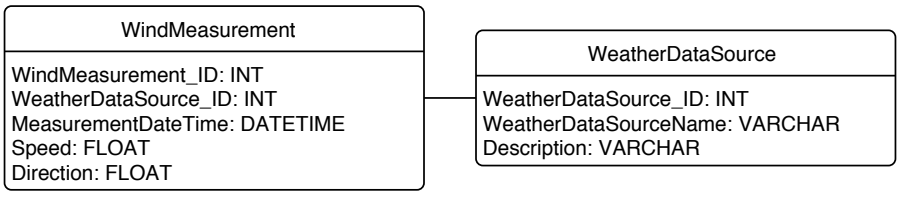

Figure B.1. Diagram illustrating the database tables used for storing wind speed data.

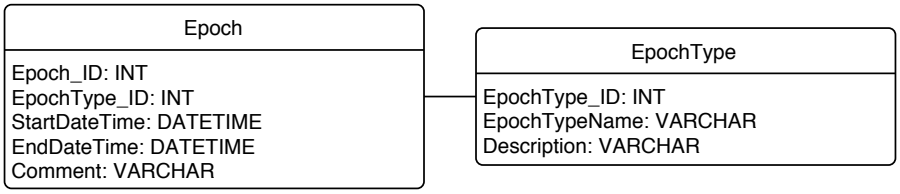

Figure B.2. Diagram illustrating the database tables used for storing epoch definitions.

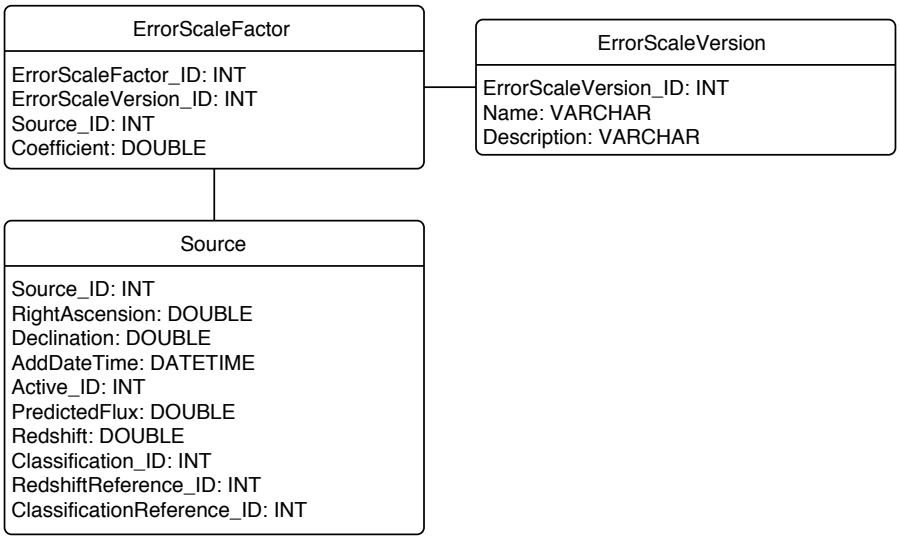

Figure B.3. Diagram illustrating the database tables used for storing error scale factors.

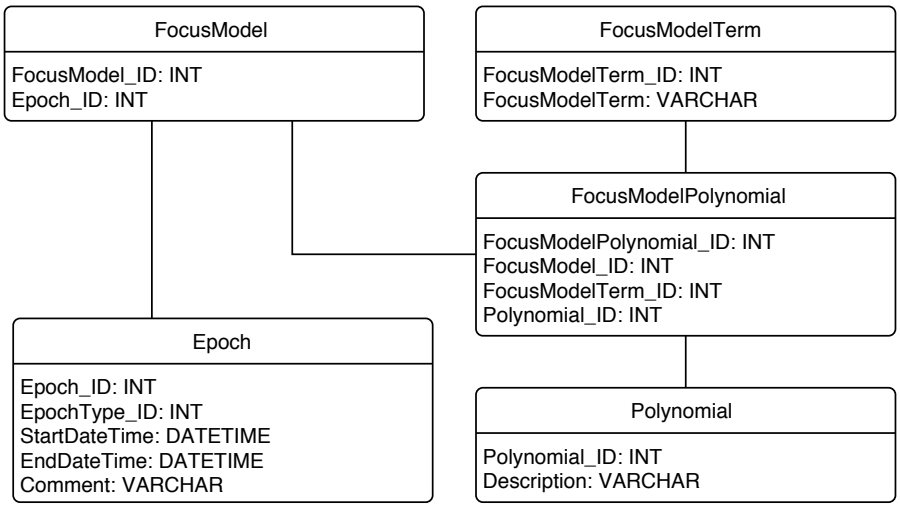

Figure B.4. Diagram illustrating the database tables used for storing focus model data.

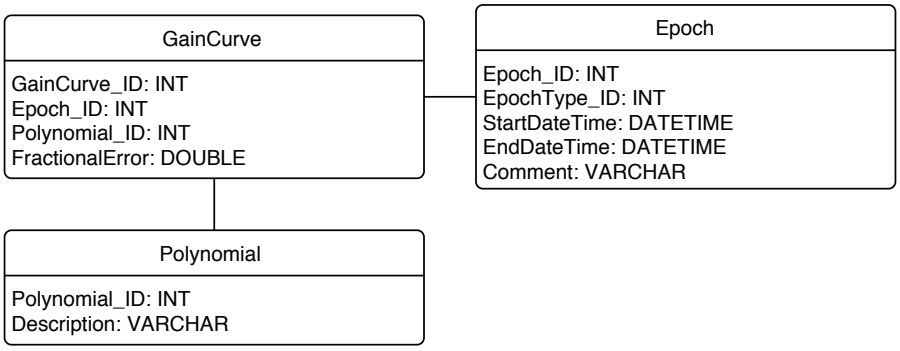

Figure B.5. Diagram illustrating the database tables used for storing gain curves.

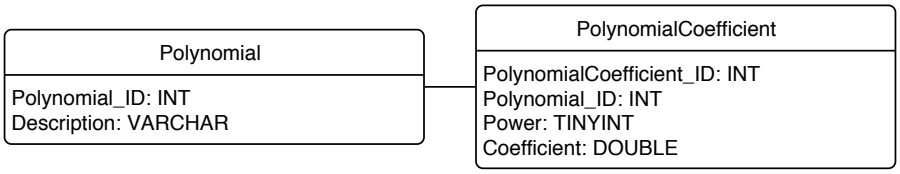

Figure B.6. Diagram illustrating the database tables used for storing polynomial definitions.

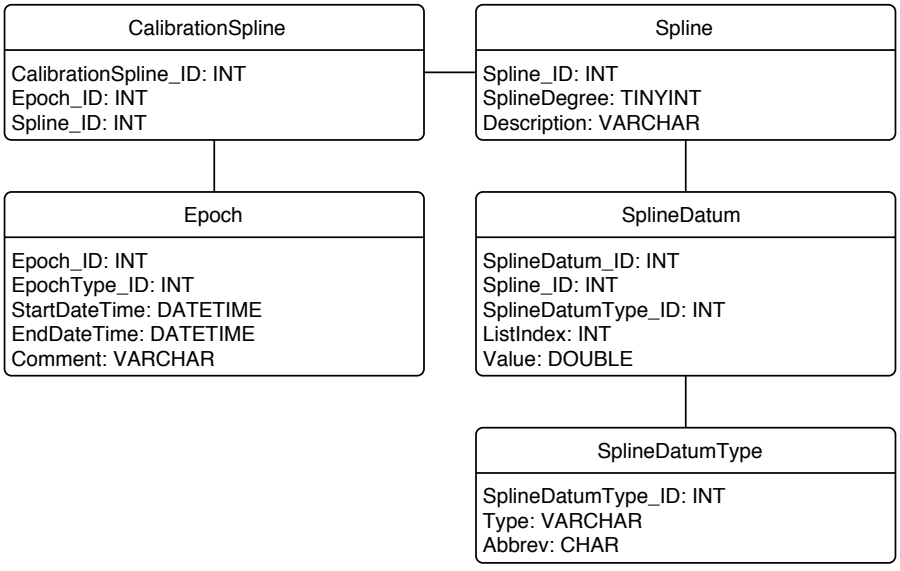

Figure B.7. Diagram illustrating the database tables used for storing calibration spline definitions.

Table B.1. Contents of the `Active` table

Active_ID	Code	Description
1	Active	Source is currently being observed.
2	Inactive	Source is not currently being observed.

Note: This table defines the options for describing whether a source is being actively observed.

Table B.2. Contents of the `Catalog` table

Catalog_ID	Code	Description	Reference_ID
1	OVRO	Name used in OVRO 40m schedules	NULL
2	OVRO_Obsolete	Obsolete name used in old OVRO 40m schedules	NULL
3	1FGL	First Fermi-LAT Catalog	3
4	0FGL	Fermi/LAT Bright Source List	4
5	CGRaBS	Candidate Gamma-Ray Blazar Survey	5
6	CRATES	Combined Radio All-Sky Targeted Eight GHz Survey	6
7	3C	Third Cambridge catalog	7
8	CRTS	Catalina Real-Time Transient Survey	11
9	1LAC	The First Catalog of AGN Detected by the Fermi LAT	9
10	1LAC_CLEAN	Clean subset of First Catalog of AGN Detected by the Fermi LAT	9

Note: This table defines the catalogs to which source names can belong.

B.2 Domain Table Contents

Table B.3. Contents of the `Classification` table

Classification_ID	Code	Description
1	AGN	Active Galactic Nucleus
2	BLL	BL Lac
3	FSRQ	Flat-Spectrum Radio Quasar
4	GAL	Galaxy
5	NLRG	Narrow-Lined Radio Galaxy
6	PN	Planetary Nebula
7		Unclassified

Note: This table defines the recognized source classifications.

Table B.4. Contents of the `Reference` table

Reference_ID	Code	URL	Comment
3	1FGL	2010ApJS..188..405A	First Fermi-LAT catalog
4	0FGL	2009ApJS..183...46A	Fermi/LAT Bright Source List
5	CGRaBS	2008ApJS..175...97H	Candidate Gamma-Ray Blazar Survey
6	CRATES	2007ApJS..171...61H	Combined Radio All-Sky Targeted Eight GHz Survey
7	3C	1962MmRAS..68..163B	Third Cambridge catalog
9	1LAC	2010ApJ...715..429A	The First Catalog of AGN Detected by the Fermi LAT
11	CRTS_WEB	http://crts.caltech.edu/	Catalina Real-Time Transient Survey web page

Note: This table stores the references used for source catalogs, redshifts, and classifications.

Table B.5. Contents of the `FluxType` table

FluxType_ID	FluxType
1	FL
2	SP
3	TP

Note: This table defines the types that can be associated with entries in the `FluxResults` table.

Table B.6. Contents of the `ReductionParameterName` table

ReductionParameter-Name_ID	ReductionParameterName	ParameterDataType_ID	PythonParameterName
3	DB_IMPORT_VERSION	4	NULL
4	DETECTION_THRESH	2	no_detection
5	EPOCH LIST ID	4	NULL
16	POINT_TEST_MAX_DELAY	2	max_delay
24	FLUX_SWD_THRESH	2	thresh
25	FLUX_SWD_SWD0	2	swd0
26	FLUX_SWD_SIG0	2	sig0
27	FLUX_SWD_KAPPA_S	2	kappa_s
28	FLUX_SWD_KAPPA_T	2	kappa_t
29	FLUX_SD_THRESH	2	thresh
30	FLUX_SD_KAPPA	2	kappa
31	DATEFLAG_SET_NAME	4	setname
32	PDIST_THRESH	2	max_pdist
33	MOON_ANGLE_THRESH	2	min_angle
34	SUN_ANGLE_THRESH	2	min_angle
35	WIND_THRESH	2	max_wind
36	ZA_LIMIT_MIN	2	za_min
37	ZA_LIMIT_MAX	2	za_max
38	FOCUS_MISS_MAX	2	focus_miss_max

Note: Several obsolete entries are not shown.

Table B.7. Contents of the `ReductionType` table

Reduction-Type_ID	ReductionType	Description
1	DB_IMPORT	Loaded log file into CMBPROG, applied low-level cuts/calibration, dumped, and inserted into database.
2	INTERNAL_CALIBRATION	Performed high-level calibration and corrections.

Table B.8. Contents of the `ParameterDataType` table

ParameterDataType_ID	ParameterDataTypeName
1	bool
2	float
3	int
4	string

Table B.9. Contents of the `DataSource` table

DataType_ID	DataSourceName	DataSourceDescription
1	vax	Data from "old" VAX-based control system, reduced via CMBPROG.
2	mcs_simple	Data from MCS control system, reduced using 1-second sample processing.
3	db	Data from the database.

Table B.10. Contents of the `FlagType` table

FlagType_ID	FlagTypeName	Comment	PythonFlagTypeName
4	DATE_FLAG	Flag fluxes by date/time intervals. Parameter is name of FlaggedDateSet	dateflag
5	FLUX_SD	Flux SD outlier filter. Parameter is "thresh,epsilon" where thresh is the SD threshold and epsilon is the FLUX->SD coefficient.	flux_sd
6	FLUX_SWD	Flag FLUXes with SWD outside a threshold. Arguments is string with values "thresh,swd0,sigma0,eps_s,eps_t"	flux_swd
11	NO_DETECTION	FLUX below detection threshold. Param is number of measured SD required for detection.	no_detection
12	POINTED	Indicates a POINT was performed on-source before the FLUX.	pointed
13	POINTING_DISTANCE	Limit on distance between FLUX and POINT, in degrees.	pdist
14	POINTING_MODEL_ONLY	FLUX was observed without a correction from a preceding POINT.	pointing_model_only
15	POINT_TEST	FLUX was measured as part of a test of the pointing model. Parameter is max time in minutes after the FLUX where the POINT can occur.	pointtest
18	SUN_ANGLE	Cut measurements too near the sun.	sun_angle
19	MOON_ANGLE	Cut measurements too near the moon.	moon_angle
20	WIND	Cut windy measurements.	wind
21	ZA_LIMIT	Cut measurements outside ZA range..	za_limit
22	FOCUS_MISS	Cut measurements with bad focus.	focus_miss
23	NEGATIVE_POINT_FLUX	flux depends on a POINT that had negative flux	negative_point_flux

Note: Several obsolete entries are not shown.

Table B.11. Contents of the `FlaggedDateType` table

FlaggedDate-Type_ID	FlaggedDate-TypeName	Description
1	OUT_OF_PROGRAM	Periods where the monitoring program was not operating. All data should be dropped entirely.
2	INTERFERENCE	Periods where severe atmosphere, RFI, or other interference render the data entirely untrustworthy. All data should be dropped entirely.
3	WIND	Periods of heavy or gusty wind.
4	POINTING_3C161	3C 161 frequently has major pointing failures. These flag instances of this problem.
5	OTHER	Miscellaneous problems.
6	BAD_POINTING	Bad measurement on a pointing calibrator, flag period should include entire region.

Table B.12. Contents of the `EpochType` table

Epoch-Type_ID	EpochType-Name	Description
1	ABS_CAL	Indicates a shift in the flux density calibration, e.g., due to a change in the ratio of the CAL diode to astronomical inputs.
2	FOCUS	Indicates a change in the focus curve.
3	POINTING	Indicates a change in the pointing performance.
4	GAIN_CURVE	Indicates a change in the gain versus elevation behavior.
5	CAL_SPLINE	Interval of applicability for a particular systematic calibration correction spline.

Table B.13. Contents of the `FocusModelTerm` table

FocusModel-Term_ID	FocusModelTerm
1	ZA
2	SUN_ANG
3	SUN_ZA
4	ERROR_POLY

Note: The entries define the names for the polynomial components of the focus model. These are the zenith angle-dependent term, the solar elongation and zenith angle terms, and the polynomial to estimate the gain error due to missed focus.

Table B.14. Contents of the `SplineDatumType` table

SplineDatum-Type_ID	Type	Abbrev
1	Knot	t
2	Coefficient	c

Table B.15. Contents of the `WeatherDataSource` table

WeatherData-Source_ID	WeatherData-SourceName	Description
1	Legacy	"Old" weather station in use from before inception of program until April, 2010.
2	Capricorn	Columbia Weather Systems Capricorn 2000EX installed Nov 2009.
3	TCS	Data from the telescope control system.

Appendix C

Source List

In this section, we present table C.1, which lists the monitoring program sources. Coordinates observed for the CGRaBS sources were those given in Healey et al. (2008). For non-CGRaBS sources, the radio positions of the associated sources in Abdo et al. (2010b) were observed. The "CGR?" column contains "Y" for sources in the CGRaBS sample. For these sources, the name in the "Source" column matches Healey et al. (2008). Sources with a 1FGL name are part of the 1LAC clean sample. The "Class" column gives the optical classification we have adopted. Only BLL (BL Lac object) and FSRQ are included in the analyses in this thesis. Explanations of the other classifications are in Healey et al. (2008) or Abdo et al. (2010b). The "$N_{\rm obs}$" column gives the number of observations *after* flagging and editing. The S_0 and $\overline{m}$ columns give values computed from the 42-month data. The values corresponding to the two-year data for the CGRaBS sample are available in the online supplements to Richards et al. (2011). Where S_0 is given as "no det.," this means that either the mean flux was less than twice the average per-flux error, more than 10% of the individual flux density measurements were less than twice the error, or the mean flux was <0. All uncertainties are 1σ values. Where $\overline{m}$ upper limits are given, these are 3σ upper limits.

Table C.1. List of sources in the monitoring program

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
0317+185	···	J0319.7+1847	BLL	0.19	121	no det.	···
0502+675	···	J0507.9+6738	BLL	0.416	194	no det.	···
0836+710	···	J0842.2+7054	FSRQ	2.218	104	1.932 ± 0.029	$0.153^{+0.012}_{-0.010}$
1406–076	···	J1408.9–0751	FSRQ	1.494	114	0.933 ± 0.007	$0.074^{+0.006}_{-0.005}$
1ES0737+746	···	J0745.2+7438	BLL	0.315	41	no det.	···
1ES0806+524	···	J0809.5+5219	BLL	0.138	263	0.113 ± 0.001	$0.103^{+0.006}_{-0.005}$
1ES1218+304	···	J1221.3+3008	BLL	0.182	221	0.044 ± 0.001	$0.167^{+0.013}_{-0.012}$
1ES1421+582	···	J1422.2+5757	BLL	···	61	no det.	···
1ES2321+419	···	J2323.5+4211	BLL	···	87	no det.	···

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
2230+114	···	J2232.5+1144	FSRQ	1.037	154	4.407 ± 0.085	$0.238^{+0.015}_{-0.014}$
2E1415+2557	···	J1417.8+2541	BLL	0.237	51	no det.	···
3C274	···	J1230.8+1223	AGN	0.004	306	26.354 ± 0.021	0.010 ± 0.001
3C380	···	J1829.8+4845	AGN	0.692	58	3.384 ± 0.023	0.048 ± 0.005
3C407	···	J2008.6–0419	AGN	0.589	130	0.354 ± 0.003	$0.083^{+0.006}_{-0.005}$
3C66A	···	J0222.6+4302	BLL	···	294	0.931 ± 0.006	$0.099^{+0.005}_{-0.004}$
4C–06.46	···	J1624.7–0642	···	···	77	0.262 ± 0.002	$0.036^{+0.006}_{-0.005}$
B20200+30	···	J0203.5+3044	···	···	89	0.298 ± 0.005	$0.135^{+0.013}_{-0.012}$
B20242+23	···	J0245.4+2413	FSRQ	2.243	73	0.233 ± 0.005	$0.152^{+0.017}_{-0.015}$
B20321+33	···	J0325.0+3403	AGN	0.061	256	0.427 ± 0.010	$0.366^{+0.019}_{-0.018}$
B20437+27B	···	J0440.6+2748	BLL	···	48	no det.	···
B30707+476	···	J0711.4+4731	BLL	···	57	0.538 ± 0.012	$0.166^{+0.018}_{-0.016}$
B30757+441	···	J0800.5+4407	···	···	60	0.140 ± 0.003	$0.125^{+0.017}_{-0.014}$
B30908+416B	···	J0912.3+4127	FSRQ	2.563	48	0.149 ± 0.004	$0.138^{+0.021}_{-0.017}$
B31518+423	···	J1519.7+4216	FSRQ	0.484	65	0.072 ± 0.002	$0.139^{+0.021}_{-0.018}$
BBJ0018+2947	···	J0018.6+2945	BLL	···	81	no det.	···
BBJ0045+2127	···	J0045.3+2127	BLL	···	124	0.052 ± 0.002	$0.296^{+0.025}_{-0.023}$
BBJ0115+2519	···	J0115.5+2519	BLL	···	147	no det.	···
BBJ0136+3905	···	J0136.5+3905	BLL	···	143	no det.	···
BBJ0154+4433	···	J0155.0+4433	BLL	···	148	0.041 ± 0.001	0.099 ± 0.014
BBJ0159+1047	···	J0159.5+1047	BLL	0.195	101	no det.	···
BBJ0208+3523	···	J0208.6+3522	BLL	0.318	96	no det.	···
BBJ0316+0904	···	J0316.1+0904	BLL	···	108	0.075 ± 0.001	$0.078^{+0.014}_{-0.013}$
BBJ0321+2326	···	J0322.1+2336	BLL	···	107	0.048 ± 0.001	$0.099^{+0.015}_{-0.014}$
BBJ0326+0225	···	J0326.2+0222	BLL	0.147	103	no det.	···
BBJ0448–1632	···	J0448.5–1633	BLL	···	57	no det.	···
BBJ0710+5908	···	J0710.6+5911	BLL	0.125	93	0.061 ± 0.001	$0.078^{+0.014}_{-0.013}$
BBJ0723+5841	···	J0722.3+5837	BLL	···	56	no det.	···
BBJ0730+3307	···	J0730.0+3305	BLL	0.112	50	no det.	···
BBJ0816+5739	···	J0816.7+5739	BLL	···	90	0.055 ± 0.001	$0.093^{+0.016}_{-0.015}$
BBJ0816–1311	···	J0816.4–1311	BLL	···	88	no det.	···
BBJ0842+0252	···	J0842.2+0251	BLL	0.425	53	no det.	···
BBJ0844+5312	···	J0844.0+5314	BLL	···	95	no det.	···

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
BBJ0847+1133	···	J0847.2+1134	BLL	0.198	51	no det.	···
BBJ0940+6148	···	J0941.2+6149	BLL	0.211	46	no det.	···
BBJ0952+3936	···	J0952.2+3926	BLL	···	49	no det.	···
BBJ1006+3454	···	J1007.0+3454	BLL	···	52	no det.	···
BBJ1022–0113	···	J1022.8–0115	BLL	···	53	no det.	···
BBJ1026–1748	···	J1027.1–1747	BLL	0.114	52	no det.	···
BBJ1031+5053	···	J1031.0+5051	BLL	0.361	90	no det.	···
BBJ1051+0103	···	J1051.9+0106	BLL	0.265	45	no det.	···
BBJ1053+4929	···	J1053.6+4927	BLL	0.14	97	0.043 ± 0.001	$0.127^{+0.019}_{-0.018}$
BBJ1107+1502	···	J1107.8+1502	BLL	0.259	58	no det.	···
BBJ1110+7133	···	J1109.9+7134	BLL	···	50	no det.	···
BBJ1117+5355	···	J1118.0+5354	BLL	···	47	no det.	···
BBJ1120+4212	···	J1121.0+4209	BLL	0.124	99	no det.	···
BBJ1125–0742	···	J1126.0–0741	BLL	0.279	47	no det.	···
BBJ1136+2550	···	J1136.9+2551	BLL	0.156	61	no det.	···
BBJ1136+6737	···	J1136.2+6739	BLL	0.136	44	0.034 ± 0.001	$0.153^{+0.044}_{-0.040}$
BBJ1150+4154	···	J1150.5+4152	BLL	···	49	no det.	···
BBJ1154–0010	···	J1154.0–0008	BLL	0.254	42	no det.	···
BBJ1204+1145	···	J1204.4+1139	BLL	0.296	48	no det.	···
BBJ1215+5002	···	J1214.9+5004	BLL	···	43	no det.	···
BBJ1226+0638	···	J1226.8+0638	BLL	···	43	no det.	···
BBJ1233–0144	···	J1233.6–0146	BLL	···	52	no det.	···
BBJ1253+6242	···	J1254.0+6236	BLL	···	63	no det.	···
BBJ1309+4305	···	J1309.5+4304	BLL	···	99	0.052 ± 0.001	$0.125^{+0.020}_{-0.018}$
BBJ1340+4410	···	J1340.6+4406	BLL	0.546	59	no det.	···
BBJ1341+3959	···	J1341.3+3951	BLL	0.172	67	no det.	···
BBJ1351+1114	···	J1351.5+1115	BLL	···	95	no det.	···
BBJ1426+3404	···	J1426.0+3403	BLL	···	112	no det.	···
BBJ1436+5639	···	J1437.0+5640	BLL	···	58	no det.	···
BBJ1442+1200	···	J1442.8+1158	BLL	0.163	107	0.035 ± 0.001	$0.168^{+0.017}_{-0.000}$
BBJ1503–1541	···	J1503.5–1544	BLL	···	52	no det.	···
BBJ1517+6525	···	J1517.8+6530	BLL	0.702	116	no det.	···
BBJ1725+5851	···	J1725.5+5854	BLL	···	92	0.057 ± 0.001	$0.114^{+0.015}_{-0.014}$

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
BBJ1829+5402	···	J1829.8+5404	BLL	···	59	no det.	···
BBJ1838+4802	···	J1838.6+4756	BLL	···	113	0.028 ± 0.001	$0.279^{+0.033}_{-0.029}$
BBJ2150–1410	···	J2150.3–1410	BLL	0.229	77	no det.	···
BBJ2247+0000	···	J2247.3+0000	BLL	···	97	0.495 ± 0.005	$0.089^{+0.008}_{-0.007}$
BBJ2334+1408	···	J2334.7+1429	BLL	···	128	0.044 ± 0.002	$0.361^{+0.033}_{-0.029}$
BBJ2338+2124	···	J2339.0+2123	BLL	0.291	130	no det.	···
BL0647+250	···	J0650.7+2503	BLL	···	153	0.060 ± 0.001	$0.165^{+0.013}_{-0.012}$
BLLacertae	···	J2202.8+4216	BLL	0.069	152	4.334 ± 0.059	$0.165^{+0.010}_{-0.009}$
BQJ1258–1800	···	J1258.4–1802	FSRQ	1.956	84	0.366 ± 0.013	$0.322^{+0.031}_{-0.026}$
BQJ1514+4450	···	J1514.7+4447	FSRQ	0.57	105	no det.	···
BUJ0645+6024	···	J0645.5+6033	AGN	0.832	57	0.068 ± 0.001	$0.071^{+0.020}_{-0.018}$
BUJ1532+3016	···	J1531.8+3018	BLL	0.065	55	0.038 ± 0.002	$0.300^{+0.043}_{-0.036}$
BUJ2313+1444	···	J2314.1+1444	···	···	78	no det.	···
C0058+3311	···	J0058.0+3314	BLL	1.371	180	0.187 ± 0.003	$0.191^{+0.012}_{-0.011}$
C0100+0745	···	J0100.2+0747	···	···	123	0.061 ± 0.001	$0.102^{+0.015}_{-0.014}$
C0144+2705	···	J0144.6+2703	BLL	···	185	0.278 ± 0.004	$0.181^{+0.011}_{-0.010}$
C0424+0036	···	J0424.8+0036	BLL	···	129	0.649 ± 0.017	$0.283^{+0.021}_{-0.018}$
C0509+1011	···	J0509.2+1015	FSRQ	0.621	125	0.450 ± 0.010	$0.232^{+0.017}_{-0.015}$
C0608–1520	···	J0608.0–1521	···	···	121	0.207 ± 0.003	$0.150^{+0.012}_{-0.011}$
C0719+3307	···	J0719.3+3306	FSRQ	0.779	120	0.479 ± 0.008	$0.179^{+0.013}_{-0.012}$
C0835+0937	···	J0835.4+0936	BLL	···	111	0.054 ± 0.001	$0.128^{+0.019}_{-0.017}$
C0850+4854	···	J0849.9+4852	···	···	97	0.120 ± 0.003	$0.191^{+0.019}_{-0.016}$
C0909–0231	···	J0909.6–0229	FSRQ	0.957	109	0.229 ± 0.001	$0.037^{+0.005}_{-0.004}$
C0953–0840	···	J0953.0–0838	BLL	···	105	0.047 ± 0.001	$0.162^{+0.022}_{-0.019}$
C0957+5522	···	J0957.7+5523	FSRQ	0.896	101	1.193 ± 0.004	0.028 ± 0.003
C1012+2439	···	J1012.7+2440	FSRQ	1.805	118	0.047 ± 0.001	$0.112^{+0.017}_{-0.016}$
C1037+5711	···	J1037.7+5711	BLL	···	108	0.143 ± 0.006	$0.376^{+0.034}_{-0.030}$
C1047+7238	···	J1048.5+7239	···	···	116	0.069 ± 0.000	< 0.078
C1112+3446	···	J1112.8+3444	FSRQ	1.949	121	0.314 ± 0.008	$0.265^{+0.020}_{-0.018}$
C1117+2014	···	J1117.1+2013	BLL	0.138	127	0.056 ± 0.001	$0.135^{+0.015}_{-0.014}$
C1224+2122	···	J1224.7+2121	FSRQ	0.435	117	1.599 ± 0.043	$0.288^{+0.022}_{-0.020}$
C1228+4858	···	J1228.2+4855	FSRQ	1.722	98	$0.341^{+0.006}_{-0.005}$	$0.150^{+0.013}_{-0.011}$
C1239+0443	···	J1239.5+0443	FSRQ	1.761	99	0.453 ± 0.011	$0.242^{+0.021}_{-0.018}$

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
C1243+3627	···	J1243.1+3627	BLL	···	143	0.077 ± 0.001	$0.074^{+0.009}_{-0.008}$
C1303+2433	···	J1303.0+2433	BLL	···	111	0.206 ± 0.008	$0.381^{+0.033}_{-0.029}$
C1345+4452	···	J1345.4+4453	FSRQ	2.534	132	0.449 ± 0.012	$0.307^{+0.023}_{-0.020}$
C1440+0610	···	J1440.9+0613	BLL	···	133	no det.	···
C1454+5124	···	J1454.6+5125	BLL	···	115	0.102 ± 0.001	$0.124^{+0.011}_{-0.010}$
C1607+1551	···	J1607.1+1552	AGN	0.496	141	0.474 ± 0.003	$0.079^{+0.006}_{-0.005}$
C1724+4004	···	J1724.0+4002	AGN	1.049	152	0.656 ± 0.003	0.040 ± 0.003
C1743+1935	···	J1744.2+1934	BLL	0.083	128	0.202 ± 0.001	$0.047^{+0.006}_{-0.005}$
C2025–0735	···	J2025.6–0735	FSRQ	1.388	157	0.917 ± 0.013	$0.177^{+0.011}_{-0.010}$
C2116+3339	···	J2116.1+3338	BLL	···	172	no det.	···
C2121+1901	···	J2120.9+1901	FSRQ	2.18	178	0.249 ± 0.001	0.042 ± 0.004
C2225–0457	···	J2225.8–0457	FSRQ	1.404	145	7.200 ± 0.102	$0.168^{+0.011}_{-0.010}$
C2311+3425	···	J2311.0+3425	FSRQ	1.817	175	0.897 ± 0.013	$0.193^{+0.012}_{-0.010}$
C2325+3957	···	J2325.2+3957	BLL	···	179	0.151 ± 0.003	$0.249^{+0.016}_{-0.015}$
CLJ0128+4439	···	J0128.6+4439	FSRQ	0.228	74	no det.	···
CLJ0212+2244	···	J0213.2+2244	BLL	0.459	113	0.046 ± 0.002	$0.300^{+0.046}_{-0.039}$
CLJ0713+1935	···	J0714.0+1935	FSRQ	0.534	95	0.249 ± 0.013	$0.480^{+0.048}_{-0.040}$
CLJ1054+2210	···	J1054.5+2212	BLL	···	117	no det.	···
CLJ1333+5057	···	J1333.2+5056	···	···	41	no det.	···
CLJ1423+3737	···	J1422.7+3743	BLL	···	107	no det.	···
CLJ1424+3615	···	J1425.0+3614	BLL	···	64	0.043 ± 0.001	$0.102^{+0.024}_{-0.022}$
CLJ1442+4348	···	J1442.1+4348	BLL	···	116	no det.	···
CLJ1450+5201	···	J1451.0+5204	BLL	···	58	no det.	···
CLJ1503+4759	···	J1503.3+4759	BLL	···	55	0.149 ± 0.005	$0.204^{+0.025}_{-0.021}$
CLJ1537+8154	···	J1536.6+8200	···	···	54	0.088 ± 0.002	$0.107^{+0.020}_{-0.017}$
CR0154+0823	···	J0154.1+0823	BLL	···	124	0.109 ± 0.004	$0.390^{+0.034}_{-0.030}$
CR1016+0513	···	J1016.1+0514	FSRQ	1.713	109	0.494 ± 0.014	$0.280^{+0.023}_{-0.020}$
CR1059–1134	···	J1059.3–1132	BLL	···	105	0.323 ± 0.011	$0.327^{+0.028}_{-0.024}$
CR1208+5441	···	J1209.3+5444	FSRQ	1.344	96	0.257 ± 0.006	$0.198^{+0.018}_{-0.016}$
CR1354–1041	···	J1354.9–1041	FSRQ	0.332	109	0.632 ± 0.011	$0.179^{+0.014}_{-0.012}$
CR1427+2347	···	J1426.9+2347	BLL	···	120	0.273 ± 0.002	$0.062^{+0.006}_{-0.005}$
CR1542+6129	···	J1542.9+6129	BLL	···	142	0.137 ± 0.002	$0.122^{+0.009}_{-0.008}$
CR1553+1256	···	J1553.4+1255	FSRQ	1.308	143	0.759 ± 0.008	$0.121^{+0.008}_{-0.007}$

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
CR1903+5540	···	J1903.0+5539	BLL	···	127	0.152 ± 0.003	$0.222^{+0.018}_{-0.016}$
CR2243+2021	···	J2244.0+2021	BLL	···	158	0.075 ± 0.001	$0.122^{+0.012}_{-0.011}$
CRJ0001–0746	···	J0000.9–0745	BLL	···	64	0.172 ± 0.001	$0.049^{+0.007}_{-0.006}$
CRJ0305–0607	···	J0305.0–0601	BLL	···	61	no det.	···
CRJ0505+0415	···	J0505.2+0420	BLL	···	59	no det.	···
CRJ0505–0419	···	J0505.8–0416	FSRQ	1.481	61	0.155 ± 0.004	$0.156^{+0.021}_{-0.018}$
CRJ1012+0630	···	J1012.2+0634	BLL	0.727	53	0.212 ± 0.002	$0.032^{+0.008}_{-0.007}$
CRJ1142+1547	···	J1141.8+1549	···	···	61	0.141 ± 0.002	$0.063^{+0.011}_{-0.010}$
CRJ1151+5859	···	J1151.6+5857	BLL	···	44	0.069 ± 0.002	$0.178^{+0.030}_{-0.025}$
CRJ1151–1347	···	J1151.4–1345	···	···	42	no det.	···
CRJ1204–0710	···	J1204.3–0714	BLL	0.185	42	0.140 ± 0.002	$0.044^{+0.014}_{-0.013}$
CRJ1220+7105	···	J1221.5+7106	FSRQ	0.451	67	0.200 ± 0.005	$0.174^{+0.020}_{-0.017}$
CRJ1305+7854	···	J1306.0+7852	···	···	57	0.174 ± 0.002	0.041 ± 0.009
CRJ1321+8316	···	J1321.3+8310	···	···	55	0.269 ± 0.004	$0.093^{+0.012}_{-0.010}$
CRJ1730+3714	···	J1730.8+3716	BLL	···	62	0.057 ± 0.000	< 0.143
CRJ1803+0341	···	J1804.1+0336	FSRQ	1.42	71	0.362 ± 0.009	$0.188^{+0.020}_{-0.017}$
CRJ1925–1018	···	J1925.1–1018	BLL	···	74	no det.	···
CRJ2001+7040	···	J2001.9+7040	···	···	72	0.051 ± 0.001	$0.093^{+0.022}_{-0.019}$
CRJ2108–0250	···	J2108.5–0249	···	···	78	0.095 ± 0.001	$0.050^{+0.010}_{-0.009}$
CRJ2322+3436	···	J2322.6+3435	BLL	0.098	92	no det.	···
FRBAJ1219–03	···	J1219.8–0309	BLL	0.299	49	no det.	···
FRBAJ1338+11	···	J1338.9+1153	BLL	···	48	no det.	···
FRBAJ2304+37	···	J2304.3+3709	BLL	···	92	no det.	···
FRBAJ2340+80	···	J2341.6+8015	BLL	0.274	65	0.042 ± 0.001	$0.166^{+0.030}_{-0.026}$
H1426+428	···	J1428.7+4239	BLL	0.129	179	0.027 ± 0.001	$0.267^{+0.026}_{-0.024}$
J0001+1914	Y	···	FSRQ	3.1	231	0.266 ± 0.003	$0.134^{+0.008}_{-0.007}$
J0001–1551	Y	···	FSRQ	2.044	228	0.223 ± 0.002	0.110 ± 0.006
J0003+2129	Y	···	AGN	0.45	259	0.082 ± 0.001	$0.108^{+0.008}_{-0.007}$
J0004+2019	Y	···	BLL	0.677	248	0.357 ± 0.004	0.162 ± 0.008
J0004+4615	Y	···	FSRQ	1.81	230	0.172 ± 0.004	$0.358^{+0.021}_{-0.019}$
J0004–1148	Y	···	BLL	···	173	0.688 ± 0.008	0.157 ± 0.009
J0005+0524	Y	···	FSRQ	1.9	208	0.114 ± 0.001	0.034 ± 0.006
J0005+3820	Y	J0005.7+3815	FSRQ	0.229	237	0.528 ± 0.006	0.179 ± 0.009

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
J0005–1648	Y	···	···	···	148	0.165 ± 0.001	0.040 ± 0.005
J0006+2422	Y	···	FSRQ	1.684	193	0.248 ± 0.002	0.076 ± 0.005
J0006–0623	Y	···	BLL	0.347	146	2.225 ± 0.012	0.061 ± 0.004
J0009+0628	···	J0008.9+0635	BLL	···	115	0.222 ± 0.001	$0.048^{+0.005}_{-0.004}$
J0010+1058	Y	···	AGN	0.089	238	0.743 ± 0.027	$0.565^{+0.035}_{-0.032}$
J0010+1724	Y	···	FSRQ	1.601	239	0.487 ± 0.002	$0.066^{+0.004}_{-0.003}$
J0010+2047	Y	···	FSRQ	0.6	243	0.090 ± 0.001	$0.223^{+0.014}_{-0.012}$
J0011+0057	Y	J0011.1+0050	FSRQ	1.492	135	0.252 ± 0.005	$0.239^{+0.018}_{-0.016}$
J0012+3353	Y	···	FSRQ	1.682	210	0.142 ± 0.002	0.150 ± 0.009
J0013+1910	Y	···	BLL	···	248	0.111 ± 0.002	0.215 ± 0.012
J0013–0423	Y	···	FSRQ	1.075	153	0.280 ± 0.008	$0.360^{+0.025}_{-0.023}$
J0013–1513	Y	···	FSRQ	1.838	227	0.154 ± 0.002	$0.154^{+0.009}_{-0.008}$
J0015–1812	Y	···	FSRQ	0.743	123	0.367 ± 0.003	0.096 ± 0.007
J0016–0015	Y	···	FSRQ	1.574	134	0.431 ± 0.004	$0.107^{+0.008}_{-0.007}$
J0017+8135	Y	···	FSRQ	3.387	221	0.916 ± 0.003	0.047 ± 0.003
J0017–0512	Y	J0017.4–0510	FSRQ	0.227	171	0.274 ± 0.005	$0.211^{+0.014}_{-0.013}$
J0019+2021	Y	J0019.3+2017	BLL	···	251	0.614 ± 0.004	0.110 ± 0.005
J0019+2602	Y	···	FSRQ	0.284	200	0.391 ± 0.005	$0.160^{+0.009}_{-0.008}$
J0019+7327	Y	···	FSRQ	1.781	242	1.238 ± 0.017	$0.209^{+0.011}_{-0.010}$
J0022+0608	Y	J0022.5+0607	BLL	···	144	0.409 ± 0.009	$0.270^{+0.019}_{-0.017}$
J0022+4525	Y	···	FSRQ	1.897	227	0.223 ± 0.001	0.072 ± 0.004
J0023+4456	Y	J0023.0+4453	FSRQ	1.062	237	0.169 ± 0.002	$0.151^{+0.008}_{-0.007}$
J0024+2439	Y	···	FSRQ	1.444	188	0.173 ± 0.003	$0.217^{+0.013}_{-0.012}$
J0027+2241	Y	···	FSRQ	1.108	240	0.145 ± 0.001	$0.082^{+0.006}_{-0.005}$
J0028+2408	Y	···	BLL	0.373	188	0.081 ± 0.001	$0.097^{+0.009}_{-0.008}$
J0029+0554	Y	···	FSRQ	1.317	174	0.620 ± 0.006	$0.123^{+0.008}_{-0.007}$
J0035–1305	Y	···	···	···	204	0.054 ± 0.001	$0.234^{+0.018}_{-0.017}$
J0037+0808	Y	···	···	···	175	0.190 ± 0.001	0.072 ± 0.005
J0037+1109	Y	···	···	···	185	0.164 ± 0.002	0.121 ± 0.008
J0037+2141	Y	···	BLL	0.598	239	0.125 ± 0.001	0.090 ± 0.006
J0038+1856	Y	···	FSRQ	1.193	237	0.132 ± 0.003	$0.282^{+0.016}_{-0.014}$
J0038+4137	Y	···	FSRQ	1.353	240	0.535 ± 0.002	0.050 ± 0.003
J0038–0329	Y	···	FSRQ	1.858	164	0.119 ± 0.001	0.090 ± 0.007

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
J0039+1411	Y	···	BLL	1.738	257	0.315 ± 0.002	0.069 ± 0.004
J0040–0146	Y	···	FSRQ	1.176	165	0.562 ± 0.003	0.059 ± 0.004
J0042+1009	Y	···	FSRQ	1.657	173	0.107 ± 0.001	0.057 ± 0.007
J0042+2320	Y	J0041.9+2318	FSRQ	1.426	236	0.473 ± 0.002	0.046 ± 0.003
J0046+3900	Y	···	FSRQ	0.958	288	0.113 ± 0.002	$0.221^{+0.012}_{-0.011}$
J0047+2435	Y	···	FSRQ	0.62	181	0.294 ± 0.002	0.063 ± 0.004
J0048+3157	Y	···	AGN	0.014	274	0.847 ± 0.009	0.175 ± 0.008
J0049+0237	Y	J0050.2+0235	BLL	···	177	0.231 ± 0.004	$0.206^{+0.013}_{-0.012}$
J0049+5128	Y	···	FSRQ	0.87	202	0.344 ± 0.003	0.114 ± 0.006
J0050–0452	Y	J0050.0–0446	FSRQ	0.92	134	0.321 ± 0.006	$0.211^{+0.015}_{-0.014}$
J0050–0929	Y	J0050.6–0928	BLL	···	217	1.026 ± 0.037	$0.523^{+0.033}_{-0.030}$
J0051–0650	Y	J0051.1–0649	FSRQ	1.975	163	1.178 ± 0.010	$0.106^{+0.007}_{-0.006}$
J0052+4402	Y	···	FSRQ	2.624	282	0.188 ± 0.001	$0.061^{+0.004}_{-0.003}$
J0056+1625	Y	···	BLL	0.206	262	0.305 ± 0.006	$0.286^{+0.015}_{-0.014}$
J0057+2218	Y	···	BLL	···	188	0.108 ± 0.001	$0.147^{+0.010}_{-0.009}$
J0057+3021	Y	···	AGN	0.016	211	0.857 ± 0.004	0.059 ± 0.003
J0058+0620	Y	···	FSRQ	0.592	159	0.195 ± 0.004	$0.235^{+0.016}_{-0.015}$
J0100+3345	Y	···	FSRQ	2.144	254	0.126 ± 0.001	0.061 ± 0.005
J0102+4214	···	J0102.2+4223	FSRQ	0.874	151	$0.151^{+0.006}_{-0.005}$	$0.421^{+0.033}_{-0.029}$
J0105+4819	Y	···	···	···	297	0.568 ± 0.002	0.065 ± 0.003
J0106+1300	Y	···	···	···	242	0.191 ± 0.001	0.076 ± 0.005
J0106+2539	Y	···	NLRG	0.199	225	0.071 ± 0.000	0.050 ± 0.008
J0106+3402	Y	···	BLL	0.579	259	0.155 ± 0.004	$0.393^{+0.022}_{-0.020}$
J0106–0315	Y	···	···	···	131	0.298 ± 0.003	$0.097^{+0.008}_{-0.007}$
J0107+2611	Y	···	FSRQ	0.522	211	0.032 ± 0.001	$0.343^{+0.025}_{-0.023}$
J0108+0135	Y	J0108.6+0135	FSRQ	2.107	129	3.406 ± 0.046	$0.150^{+0.010}_{-0.009}$
J0109+1816	···	J0109.0+1816	BLL	0.145	118	0.064 ± 0.001	$0.123^{+0.016}_{-0.014}$
J0110–0415	Y	···	···	···	124	0.085 ± 0.001	$0.105^{+0.013}_{-0.012}$
J0110–0741	Y	···	FSRQ	1.776	206	0.430 ± 0.009	$0.299^{+0.017}_{-0.016}$
J0111+3906	Y	···	NLRG	0.669	283	0.515 ± 0.001	0.025 ± 0.002
J0111–1317	Y	···	FSRQ	2.42	208	0.279 ± 0.004	$0.205^{+0.012}_{-0.011}$
J0112+2244	Y	J0112.0+2247	BLL	0.265	276	0.523 ± 0.013	$0.413^{+0.022}_{-0.020}$
J0112+3208	···	J0112.9+3207	FSRQ	0.603	152	0.608 ± 0.005	$0.103^{+0.007}_{-0.006}$

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
J0112+3522	Y	···	FSRQ	0.45	288	1.050 ± 0.012	$0.197^{+0.009}_{-0.008}$
J0113+0222	Y	···	BLL	0.047	134	0.576 ± 0.004	0.068 ± 0.005
J0113+4948	Y	J0113.8+4945	FSRQ	0.395	293	0.763 ± 0.006	0.137 ± 0.006
J0115–0127	Y	···	FSRQ	1.365	129	0.505 ± 0.004	$0.089^{+0.007}_{-0.006}$
J0116+2422	Y	···	···	···	317	0.256 ± 0.001	0.084 ± 0.005
J0116–1136	Y	J0115.5–1132	FSRQ	0.672	205	1.020 ± 0.008	0.115 ± 0.006
J0117+1418	Y	···	FSRQ	0.839	253	0.368 ± 0.007	$0.288^{+0.015}_{-0.014}$
J0121+0422	Y	···	FSRQ	0.637	131	0.945 ± 0.005	$0.062^{+0.005}_{-0.004}$
J0121+1127	Y	···	FSRQ	2.465	250	0.106 ± 0.001	$0.054^{+0.006}_{-0.005}$
J0121+1149	Y	···	FSRQ	0.57	243	2.822 ± 0.058	$0.317^{+0.017}_{-0.015}$
J0122+2502	Y	···	FSRQ	2.025	321	0.556 ± 0.002	0.042 ± 0.002
J0123+2615	Y	···	FSRQ	0.849	323	0.237 ± 0.002	$0.116^{+0.006}_{-0.005}$
J0124+2805	Y	···	FSRQ	0.71	277	0.074 ± 0.001	$0.076^{+0.007}_{-0.006}$
J0125–0005	Y	···	FSRQ	1.077	127	1.016 ± 0.009	$0.093^{+0.007}_{-0.006}$
J0126+2559	Y	···	FSRQ	2.37	333	0.549 ± 0.003	$0.081^{+0.004}_{-0.003}$
J0127–0821	Y	···	BLL	0.362	202	0.169 ± 0.002	$0.150^{+0.009}_{-0.008}$
J0128+4901	Y	···	AGN	0.067	292	0.667 ± 0.017	$0.422^{+0.022}_{-0.020}$
J0130+0842	Y	···	FSRQ	0.725	257	0.187 ± 0.002	0.125 ± 0.007
J0131+3834	Y	···	FSRQ	1.277	289	0.304 ± 0.002	0.077 ± 0.004
J0132+4325	Y	···	···	···	289	0.360 ± 0.007	$0.308^{+0.015}_{-0.014}$
J0132–1654	Y	J0132.6–1655	FSRQ	1.02	177	1.760 ± 0.043	$0.319^{+0.020}_{-0.018}$
J0136+4751	Y	J0137.0+4751	FSRQ	0.859	292	3.498 ± 0.044	$0.212^{+0.010}_{-0.009}$
J0137+3122	Y	···	FSRQ	1.716	285	0.218 ± 0.001	0.051 ± 0.003
J0140–1532	Y	···	FSRQ	0.819	180	0.248 ± 0.002	$0.108^{+0.007}_{-0.006}$
J0141–0202	Y	···	FSRQ	1.281	213	0.276 ± 0.004	$0.206^{+0.012}_{-0.011}$
J0141–0928	Y	J0141.7–0929	BLL	0.733	210	0.785 ± 0.009	$0.162^{+0.009}_{-0.008}$
J0143+4129	Y	···	FSRQ	0.825	287	0.191 ± 0.001	0.119 ± 0.006
J0148+3854	Y	···	FSRQ	1.442	290	0.321 ± 0.001	0.037 ± 0.002
J0148+4215	Y	···	FSRQ	3.242	280	0.130 ± 0.001	0.102 ± 0.006
J0149+0555	Y	···	FSRQ	2.345	198	0.810 ± 0.005	$0.078^{+0.005}_{-0.004}$
J0149+1857	Y	···	FSRQ	0.584	262	0.394 ± 0.003	0.103 ± 0.005
J0151+2744	Y	···	FSRQ	1.26	333	0.515 ± 0.001	0.028 ± 0.002
J0151–1732	Y	···	···	···	200	0.346 ± 0.002	$0.051^{+0.004}_{-0.003}$

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
J0152+2207	Y	···	FSRQ	1.32	269	0.898 ± 0.005	$0.096^{+0.005}_{-0.004}$
J0154+4743	Y	···	FSRQ	1.026	284	0.576 ± 0.009	$0.251^{+0.012}_{-0.011}$
J0155+2230	Y	···	FSRQ	1.456	265	0.119 ± 0.001	0.112 ± 0.007
J0200–1356	Y	···	···	···	196	0.067 ± 0.001	$0.129^{+0.012}_{-0.011}$
J0200–1542	Y	···	···	···	207	0.135 ± 0.001	0.084 ± 0.006
J0202+3943	Y	···	FSRQ	0.78	286	0.193 ± 0.002	0.160 ± 0.008
J0202+4205	Y	···	BLL	···	280	0.428 ± 0.005	0.198 ± 0.009
J0202–0559	Y	···	FSRQ	0.189	182	0.113 ± 0.001	$0.115^{+0.011}_{-0.010}$
J0202–1948	Y	···	FSRQ	0.493	215	0.102 ± 0.001	$0.055^{+0.007}_{-0.006}$
J0203+1134	Y	···	FSRQ	3.61	264	0.507 ± 0.004	0.130 ± 0.006
J0203+7232	Y	J0203.5+7234	BLL	···	225	0.451 ± 0.004	$0.123^{+0.007}_{-0.006}$
J0204+1514	Y	J0204.5+1516	AGN	0.405	258	1.081 ± 0.007	0.107 ± 0.005
J0204+4005	Y	···	NLRG	0.072	281	0.163 ± 0.002	$0.143^{+0.008}_{-0.007}$
J0204–1701	Y	J0205.0–1702	FSRQ	1.74	212	1.411 ± 0.007	$0.065^{+0.004}_{-0.003}$
J0205+3212	Y	J0205.3+3217	FSRQ	1.466	272	2.875 ± 0.027	$0.151^{+0.007}_{-0.006}$
J0205+3932	Y	···	FSRQ	0.454	284	0.075 ± 0.001	$0.117^{+0.009}_{-0.008}$
J0206–1150	Y	···	FSRQ	1.663	211	0.263 ± 0.006	$0.337^{+0.020}_{-0.018}$
J0209+1352	Y	···	FSRQ	0.631	259	0.404 ± 0.005	0.177 ± 0.009
J0209+7229	Y	···	FSRQ	0.895	226	0.396 ± 0.003	$0.117^{+0.007}_{-0.006}$
J0210–1444	Y	···	···	···	217	0.123 ± 0.002	$0.174^{+0.010}_{-0.009}$
J0211+1051	Y	J0211.2+1049	BLL	···	266	0.801 ± 0.023	$0.458^{+0.025}_{-0.023}$
J0211–1558	Y	···	FSRQ	0.177	197	0.208 ± 0.002	$0.136^{+0.009}_{-0.008}$
J0213+1820	Y	···	FSRQ	1.818	255	0.295 ± 0.001	0.059 ± 0.004
J0213+8717	Y	···	···	···	172	0.062 ± 0.001	0.085 ± 0.010
J0215–0222	Y	···	FSRQ	1.178	210	0.614 ± 0.002	0.036 ± 0.003
J0217+0144	Y	J0217.9+0144	FSRQ	1.715	208	1.597 ± 0.015	0.132 ± 0.007
J0217+0837	Y	J0217.2+0834	BLL	0.085	239	0.512 ± 0.006	$0.189^{+0.010}_{-0.009}$
J0217+7349	Y	···	FSRQ	2.367	222	3.544 ± 0.016	0.064 ± 0.003
J0219+0120	Y	···	FSRQ	1.623	192	0.700 ± 0.003	$0.054^{+0.004}_{-0.003}$
J0219+4727	Y	···	···	···	280	0.154 ± 0.001	0.076 ± 0.005
J0219–1842	Y	···	BLL	···	185	0.433 ± 0.007	$0.210^{+0.013}_{-0.011}$
J0220–1305	Y	···	FSRQ	1.445	207	0.330 ± 0.002	0.082 ± 0.005
J0221+3556	Y	J0221.0+3555	FSRQ	0.944	279	1.378 ± 0.006	$0.074^{+0.004}_{-0.003}$

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
J0222–1615	Y	J0222.1–1618	FSRQ	0.698	150	0.253 ± 0.003	$0.133^{+0.010}_{-0.009}$
J0224+0659	Y	···	FSRQ	0.511	157	0.854 ± 0.021	$0.311^{+0.021}_{-0.018}$
J0225+1846	Y	···	FSRQ	2.69	245	0.405 ± 0.014	$0.525^{+0.032}_{-0.028}$
J0226–1843	Y	···	···	···	140	0.176 ± 0.003	$0.198^{+0.014}_{-0.013}$
J0230+4032	Y	J0230.8+4031	FSRQ	1.019	278	0.458 ± 0.004	$0.137^{+0.007}_{-0.006}$
J0231+1322	Y	···	FSRQ	2.065	257	1.670 ± 0.012	$0.116^{+0.006}_{-0.005}$
J0237+0526	Y	···	FSRQ	0.562	172	0.137 ± 0.003	$0.247^{+0.017}_{-0.015}$
J0237+2848	Y	J0237.9+2848	FSRQ	1.213	248	2.654 ± 0.024	$0.142^{+0.007}_{-0.006}$
J0237+3022	Y	···	···	···	255	0.114 ± 0.001	0.162 ± 0.009
J0238+1636	Y	J0238.6+1637	BLL	0.94	247	2.117 ± 0.089	$0.657^{+0.043}_{-0.038}$
J0239+0416	Y	···	FSRQ	0.978	173	0.399 ± 0.003	$0.097^{+0.007}_{-0.006}$
J0239–0234	Y	···	FSRQ	1.116	164	1.038 ± 0.016	$0.190^{+0.012}_{-0.011}$
J0240+1848	Y	···	FSRQ	1.297	210	0.343 ± 0.003	$0.120^{+0.007}_{-0.006}$
J0240+4216	Y	···	FSRQ	1.701	284	0.095 ± 0.001	$0.070^{+0.006}_{-0.005}$
J0241–0815	Y	···	AGN	0.005	213	1.119 ± 0.010	$0.126^{+0.007}_{-0.006}$
J0242+1101	Y	···	FSRQ	2.68	262	0.899 ± 0.003	0.041 ± 0.002
J0242+1742	Y	···	···	···	218	0.283 ± 0.003	$0.136^{+0.008}_{-0.007}$
J0242+2653	Y	···	FSRQ	1.851	252	0.122 ± 0.002	$0.204^{+0.011}_{-0.010}$
J0243+7120	···	J0243.5+7116	BLL	···	106	0.193 ± 0.001	$0.042^{+0.010}_{-0.008}$
J0243–0550	Y	···	FSRQ	1.801	150	0.363 ± 0.005	$0.164^{+0.011}_{-0.010}$
J0245–1107	Y	···	GALAXY	0.099	202	0.158 ± 0.001	0.042 ± 0.005
J0246–1236	Y	···	FSRQ	2.201	203	0.296 ± 0.003	$0.114^{+0.007}_{-0.006}$
J0249+0619	Y	···	FSRQ	1.881	179	0.518 ± 0.001	0.022 ± 0.002
J0251+3734	Y	···	FSRQ	1.818	252	0.159 ± 0.001	$0.080^{+0.006}_{-0.005}$
J0251+4315	Y	···	FSRQ	1.31	272	0.704 ± 0.002	0.042 ± 0.002
J0251+7226	Y	···	···	···	225	0.147 ± 0.005	$0.505^{+0.033}_{-0.030}$
J0254+2343	Y	···	FSRQ	1.988	237	0.178 ± 0.003	$0.235^{+0.013}_{-0.012}$
J0254+3931	Y	···	FSRQ	0.291	276	0.385 ± 0.005	$0.203^{+0.010}_{-0.009}$
J0256+1542	Y	···	···	···	290	0.121 ± 0.002	$0.206^{+0.011}_{-0.010}$
J0257+1847	Y	···	FSRQ	0.427	279	0.318 ± 0.008	$0.419^{+0.022}_{-0.020}$
J0257+7843	Y	···	···	···	228	0.211 ± 0.002	$0.103^{+0.007}_{-0.006}$
J0257–1212	Y	J0257.8–1204	FSRQ	1.391	208	0.224 ± 0.003	$0.148^{+0.009}_{-0.008}$
J0258+0541	Y	···	FSRQ	1.381	173	0.248 ± 0.004	$0.178^{+0.012}_{-0.011}$

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
J0259+0747	···	J0259.5+0743	FSRQ	0.893	106	0.557 ± 0.009	$0.160^{+0.013}_{-0.011}$
J0259–0018	Y	···	···	···	159	0.129 ± 0.003	$0.215^{+0.018}_{-0.017}$
J0305+0523	Y	···	···	···	172	0.128 ± 0.001	0.065 ± 0.007
J0309+1029	Y	···	FSRQ	0.863	299	0.825 ± 0.007	0.142 ± 0.006
J0309–0559	Y	···	FSRQ	0.745	150	0.140 ± 0.003	$0.205^{+0.016}_{-0.014}$
J0310+3814	Y	J0310.6+3812	FSRQ	0.816	281	0.531 ± 0.012	$0.358^{+0.018}_{-0.017}$
J0312+0133	Y	J0312.6+0131	FSRQ	0.664	194	0.390 ± 0.005	$0.164^{+0.010}_{-0.009}$
J0313+4120	Y	···	AGN	0.136	288	0.894 ± 0.006	0.120 ± 0.005
J0315–1031	···	J0315.9–1033	FSRQ	1.565	101	0.159 ± 0.002	$0.113^{+0.012}_{-0.011}$
J0315–1656	Y	···	···	···	162	0.237 ± 0.002	$0.112^{+0.008}_{-0.007}$
J0318–0029	Y	···	···	···	168	0.072 ± 0.001	$0.198^{+0.015}_{-0.014}$
J0319+1901	Y	···	FSRQ	0.296	281	0.252 ± 0.001	0.037 ± 0.003
J0319+4130	Y	J0319.7+4130	AGN	0.018	289	20.084 ± 0.143	0.120 ± 0.005
J0319+6949	Y	···	···	···	275	0.151 ± 0.001	0.072 ± 0.005
J0319–1613	Y	···	FSRQ	1.618	148	0.296 ± 0.005	$0.205^{+0.014}_{-0.013}$
J0322+3948	Y	···	···	···	282	0.081 ± 0.001	$0.081^{+0.007}_{-0.006}$
J0323+0145	Y	···	···	···	197	0.246 ± 0.001	0.050 ± 0.004
J0325+2224	Y	J0325.9+2219	FSRQ	2.066	279	0.811 ± 0.014	$0.289^{+0.014}_{-0.013}$
J0327+0044	Y	···	FSRQ	1.357	203	0.329 ± 0.002	0.053 ± 0.004
J0329+3510	Y	···	···	···	272	0.773 ± 0.009	0.198 ± 0.009
J0334+0800	Y	···	FSRQ	1.982	194	0.285 ± 0.002	$0.105^{+0.007}_{-0.006}$
J0336–1302	Y	···	FSRQ	1.303	131	0.453 ± 0.007	$0.159^{+0.012}_{-0.010}$
J0338+3106	Y	···	FSRQ	1.662	248	0.074 ± 0.001	$0.251^{+0.016}_{-0.014}$
J0339–0146	Y	···	FSRQ	0.852	212	2.242 ± 0.024	0.156 ± 0.008
J0341+3352	Y	···	FSRQ	0.725	270	0.083 ± 0.001	0.098 ± 0.008
J0343+3622	Y	···	FSRQ	1.484	245	0.303 ± 0.002	$0.103^{+0.006}_{-0.005}$
J0345+1453	Y	···	FSRQ	1.557	178	0.477 ± 0.005	0.121 ± 0.007
J0348–1610	Y	···	BLL	···	191	0.781 ± 0.004	$0.055^{+0.004}_{-0.003}$
J0351–1153	Y	···	FSRQ	1.52	144	0.429 ± 0.003	$0.073^{+0.006}_{-0.005}$
J0354+8009	···	J0354.6+8009	···	···	106	0.325 ± 0.003	$0.086^{+0.009}_{-0.008}$
J0357+0542	Y	···	FSRQ	2.164	197	0.101 ± 0.002	$0.198^{+0.014}_{-0.013}$
J0357+2319	Y	···	···	···	270	0.482 ± 0.010	$0.338^{+0.017}_{-0.016}$
J0359+2758	Y	···	···	···	286	0.119 ± 0.002	$0.250^{+0.014}_{-0.013}$

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
J0359+3220	Y	···	FSRQ	1.332	219	0.634 ± 0.003	0.052 ± 0.003
J0400+0550	Y	···	FSRQ	0.761	193	0.226 ± 0.003	$0.189^{+0.012}_{-0.011}$
J0401+0413	Y	···	GALAXY	0.306	205	0.534 ± 0.010	$0.250^{+0.014}_{-0.013}$
J0401+2110	Y	···	FSRQ	0.834	232	0.223 ± 0.004	$0.237^{+0.014}_{-0.013}$
J0401–1606	Y	···	FSRQ	0.031	178	0.386 ± 0.003	$0.105^{+0.007}_{-0.006}$
J0403+2600	Y	···	FSRQ	2.109	278	1.770 ± 0.015	0.141 ± 0.006
J0405–0739	Y	···	FSRQ	2.817	194	0.086 ± 0.001	$0.061^{+0.010}_{-0.009}$
J0406+0637	Y	···	FSRQ	0.666	200	0.388 ± 0.004	$0.128^{+0.008}_{-0.007}$
J0407+0742	Y	···	FSRQ	1.133	194	0.493 ± 0.011	$0.295^{+0.018}_{-0.016}$
J0408–0122	Y	···	···	···	199	0.103 ± 0.001	$0.136^{+0.010}_{-0.009}$
J0409+1217	Y	···	BLL	1.02	175	0.262 ± 0.003	$0.109^{+0.008}_{-0.007}$
J0409–1238	Y	···	FSRQ	1.563	182	0.426 ± 0.006	0.168 ± 0.010
J0412+0010	Y	···	FSRQ	1.13	196	0.173 ± 0.002	$0.124^{+0.008}_{-0.007}$
J0412+0438	Y	···	FSRQ	1.081	208	0.119 ± 0.001	0.070 ± 0.006
J0412+1856	Y	···	···	···	214	0.166 ± 0.001	0.028 ± 0.004
J0414+3418	Y	···	···	···	231	0.885 ± 0.004	$0.066^{+0.004}_{-0.003}$
J0416+0105	···	J0416.8+0107	BLL	0.287	100	0.056 ± 0.000	< 0.092
J0416–1851	Y	J0416.5–1851	FSRQ	1.536	120	0.919 ± 0.011	0.126 ± 0.009
J0422+0219	Y	J0422.1+0211	FSRQ	2.277	208	0.860 ± 0.004	$0.053^{+0.004}_{-0.003}$
J0422–0643	···	J0422.0–0647	FSRQ	0.242	96	0.125 ± 0.007	$0.487^{+0.050}_{-0.042}$
J0423–0120	Y	J0423.2–0118	FSRQ	0.915	191	5.705 ± 0.099	$0.238^{+0.014}_{-0.012}$
J0424+0805	Y	···	FSRQ	3.086	212	0.278 ± 0.003	$0.156^{+0.010}_{-0.009}$
J0426+2327	Y	···	NLRG	0.55	234	0.278 ± 0.001	0.046 ± 0.003
J0426+2350	Y	···	···	···	191	0.108 ± 0.001	$0.087^{+0.007}_{-0.006}$
J0428+1732	Y	···	FSRQ	3.317	239	0.134 ± 0.002	$0.176^{+0.011}_{-0.010}$
J0428+3259	Y	···	FSRQ	0.476	223	0.340 ± 0.001	0.047 ± 0.003
J0432–1614	Y	···	···	···	123	0.289 ± 0.001	0.035 ± 0.004
J0433+0521	Y	···	AGN	0.033	201	2.809 ± 0.029	$0.145^{+0.008}_{-0.007}$
J0433+2905	Y	J0433.5+2905	BLL	···	206	0.276 ± 0.006	$0.295^{+0.018}_{-0.016}$
J0434–1442	Y	···	FSRQ	1.899	135	0.272 ± 0.004	$0.155^{+0.011}_{-0.010}$
J0435+2923	Y	···	FSRQ	2.336	204	0.103 ± 0.001	0.079 ± 0.007
J0437–1844	Y	···	FSRQ	2.702	129	0.423 ± 0.002	0.046 ± 0.004
J0438+3004	Y	···	FSRQ	1.454	211	0.445 ± 0.002	0.066 ± 0.004

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
J0438–1251	Y	J0438.8–1250	FSRQ	1.286	134	0.443 ± 0.004	$0.099^{+0.008}_{-0.007}$
J0439+0520	Y	···	FSRQ	0.202	217	0.295 ± 0.002	0.091 ± 0.005
J0439+3045	Y	···	···	···	211	0.252 ± 0.002	0.070 ± 0.005
J0442–0017	Y	···	FSRQ	0.844	181	1.364 ± 0.011	0.100 ± 0.006
J0449+1121	Y	···	FSRQ	1.207	226	0.995 ± 0.017	$0.252^{+0.013}_{-0.012}$
J0449+6332	Y	···	FSRQ	0.781	242	0.651 ± 0.021	$0.499^{+0.030}_{-0.027}$
J0449–1814	Y	···	···	···	127	0.143 ± 0.001	$0.066^{+0.007}_{-0.006}$
J0452+1236	Y	···	FSRQ	1.177	226	0.146 ± 0.001	$0.120^{+0.009}_{-0.008}$
J0455+0655	Y	···	···	···	208	0.073 ± 0.001	$0.193^{+0.015}_{-0.013}$
J0456+0400	Y	···	FSRQ	1.349	197	0.405 ± 0.002	$0.048^{+0.004}_{-0.003}$
J0457+0645	Y	J0457.9+0649	FSRQ	0.405	229	0.642 ± 0.007	$0.166^{+0.009}_{-0.008}$
J0457–1819	Y	···	···	···	212	0.095 ± 0.001	$0.161^{+0.011}_{-0.010}$
J0501+1356	Y	···	···	···	212	0.278 ± 0.002	0.092 ± 0.006
J0501+7128	Y	···	···	···	279	0.155 ± 0.001	0.077 ± 0.005
J0501–0159	Y	J0501.0–0200	FSRQ	2.291	181	1.085 ± 0.012	$0.148^{+0.009}_{-0.008}$
J0502+1338	Y	···	BLL	···	213	0.364 ± 0.007	$0.258^{+0.015}_{-0.014}$
J0503+6600	Y	···	FSRQ	1.696	228	0.190 ± 0.002	$0.155^{+0.009}_{-0.008}$
J0505+0459	Y	···	FSRQ	0.954	192	0.611 ± 0.014	$0.310^{+0.019}_{-0.017}$
J0508+8432	Y	···	BLL	1.34	185	0.210 ± 0.003	$0.165^{+0.010}_{-0.009}$
J0509+0541	Y	J0509.3+0540	BLL	···	192	0.538 ± 0.007	$0.170^{+0.010}_{-0.009}$
J0510+1800	Y	···	AGN	0.416	200	0.924 ± 0.018	$0.269^{+0.015}_{-0.014}$
J0511+1357	Y	···	FSRQ	1.696	210	0.415 ± 0.004	$0.140^{+0.008}_{-0.007}$
J0517+0648	Y	···	FSRQ	0.84	184	0.232 ± 0.001	0.048 ± 0.004
J0521–1737	Y	···	···	···	216	0.276 ± 0.001	0.058 ± 0.004
J0522–0725	Y	···	FSRQ	0.164	126	no det.	···
J0527+0106	Y	···	···	···	185	0.055 ± 0.001	0.087 ± 0.012
J0527+0331	Y	···	BLL	···	185	0.561 ± 0.011	$0.252^{+0.015}_{-0.014}$
J0529–0519	Y	···	FSRQ	0.685	133	0.163 ± 0.002	$0.135^{+0.012}_{-0.011}$
J0530+1331	Y	J0531.0+1331	FSRQ	2.07	203	2.320 ± 0.058	$0.353^{+0.021}_{-0.019}$
J0532+0732	Y	J0532.9+0733	FSRQ	1.254	186	1.381 ± 0.007	0.067 ± 0.004
J0538+5107	Y	···	···	···	153	0.167 ± 0.002	$0.106^{+0.009}_{-0.008}$
J0541+5312	Y	···	FSRQ	1.275	152	0.460 ± 0.002	0.054 ± 0.004
J0541–0541	Y	···	FSRQ	0.839	127	0.895 ± 0.014	$0.168^{+0.012}_{-0.011}$

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
J0542–0913	Y	···	···	···	126	0.437 ± 0.013	$0.320^{+0.024}_{-0.022}$
J0551–1909	Y	···	FSRQ	0.73	209	0.103 ± 0.001	0.075 ± 0.007
J0552+0313	Y	···	···	···	191	0.457 ± 0.004	$0.110^{+0.007}_{-0.006}$
J0554+6857	Y	···	FSRQ	1.373	225	0.203 ± 0.003	$0.186^{+0.011}_{-0.010}$
J0558–1317	Y	···	FSRQ	1.725	207	0.521 ± 0.005	$0.117^{+0.007}_{-0.006}$
J0559+5804	Y	···	FSRQ	0.904	145	0.328 ± 0.003	0.101 ± 0.007
J0559–1817	Y	···	···	···	198	0.112 ± 0.001	0.132 ± 0.011
J0606–0724	Y	···	FSRQ	1.227	197	0.315 ± 0.002	$0.077^{+0.006}_{-0.005}$
J0607+4739	Y	J0607.2+4739	BLL	···	208	0.309 ± 0.003	$0.105^{+0.007}_{-0.006}$
J0607+6720	Y	···	FSRQ	1.97	158	0.640 ± 0.002	0.039 ± 0.003
J0607–0834	Y	J0608.2–0837	FSRQ	0.872	196	1.616 ± 0.039	$0.336^{+0.020}_{-0.018}$
J0609–1542	Y	···	FSRQ	0.324	120	4.508 ± 0.077	$0.184^{+0.013}_{-0.012}$
J0610–1847	Y	···	···	···	125	0.338 ± 0.006	$0.189^{+0.015}_{-0.013}$
J0612+4122	Y	J0612.7+4120	BLL	···	125	0.311 ± 0.009	$0.321^{+0.024}_{-0.022}$
J0616–1041	Y	···	···	···	190	0.182 ± 0.004	$0.320^{+0.021}_{-0.019}$
J0617+5701	···	J0616.9+5701	BLL	···	93	0.266 ± 0.003	$0.088^{+0.008}_{-0.007}$
J0617+7816	Y	···	FSRQ	1.43	199	0.213 ± 0.002	$0.084^{+0.006}_{-0.005}$
J0617–1715	···	J0617.7–1718	BLL	0.32	99	0.238 ± 0.002	0.048 ± 0.006
J0618+4620	Y	···	FSRQ	0.607	209	0.221 ± 0.002	0.133 ± 0.008
J0619–1140	Y	···	FSRQ	0.97	182	0.331 ± 0.006	$0.235^{+0.015}_{-0.013}$
J0623+3830	Y	···	NLRG	0.421	125	0.168 ± 0.001	$0.030^{+0.005}_{-0.004}$
J0624+3856	Y	···	FSRQ	3.469	129	0.421 ± 0.007	$0.169^{+0.012}_{-0.011}$
J0625+4440	Y	J0625.4+4440	BLL	···	209	0.175 ± 0.004	$0.302^{+0.018}_{-0.016}$
J0626+8202	Y	···	FSRQ	0.71	167	0.593 ± 0.007	$0.138^{+0.009}_{-0.008}$
J0629–1959	Y	J0629.6–2000	BLL	···	119	1.109 ± 0.012	$0.117^{+0.009}_{-0.008}$
J0630–1323	Y	···	FSRQ	1.021	200	0.398 ± 0.004	$0.124^{+0.008}_{-0.007}$
J0631–1410	Y	···	FSRQ	1.017	211	0.564 ± 0.003	0.063 ± 0.004
J0632+3200	Y	···	FSRQ	1.832	97	0.133 ± 0.001	$0.089^{+0.009}_{-0.008}$
J0637+3322	Y	···	···	···	129	0.168 ± 0.003	$0.144^{+0.012}_{-0.011}$
J0638+5933	Y	···	···	···	124	0.791 ± 0.002	0.026 ± 0.002
J0639+7324	Y	J0639.9+7325	FSRQ	1.854	126	1.004 ± 0.021	$0.234^{+0.017}_{-0.015}$
J0642+3509	Y	···	NLRG	0.269	127	0.178 ± 0.003	$0.142^{+0.011}_{-0.010}$
J0642+6758	Y	···	FSRQ	3.177	127	0.198 ± 0.001	0.036 ± 0.004

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
J0642+8811	Y	···	···	···	172	0.172 ± 0.004	$0.312^{+0.021}_{-0.019}$
J0644+3914	Y	···	FSRQ	1.266	128	0.399 ± 0.002	$0.057^{+0.005}_{-0.004}$
J0646+4451	Y	···	FSRQ	3.408	198	$3.350^{+0.015}_{-0.014}$	0.058 ± 0.003
J0650+5616	Y	···	FSRQ	1.428	204	0.185 ± 0.002	0.170 ± 0.010
J0650+6001	Y	···	FSRQ	0.455	124	0.729 ± 0.002	0.026 ± 0.002
J0653+3705	Y	···	FSRQ	1.982	131	0.720 ± 0.005	0.067 ± 0.005
J0654+4514	Y	J0654.3+4514	FSRQ	0.933	205	0.440 ± 0.013	$0.416^{+0.026}_{-0.023}$
J0654+5042	Y	J0654.4+5042	···	···	201	0.224 ± 0.003	$0.153^{+0.010}_{-0.009}$
J0655+4100	Y	···	NLRG	0.022	129	0.383 ± 0.002	0.051 ± 0.004
J0657+2423	Y	···	FSRQ	1.926	110	0.255 ± 0.002	$0.065^{+0.007}_{-0.006}$
J0702+2644	Y	···	···	···	108	0.458 ± 0.005	$0.106^{+0.009}_{-0.008}$
J0702+8549	Y	···	FSRQ	1.059	176	0.184 ± 0.001	0.089 ± 0.006
J0712+5033	Y	J0712.7+5033	BLL	···	209	0.355 ± 0.005	$0.214^{+0.012}_{-0.011}$
J0717+4538	Y	···	FSRQ	0.94	205	0.706 ± 0.008	0.149 ± 0.008
J0720+4737	Y	···	···	···	205	0.410 ± 0.002	0.065 ± 0.004
J0721+7120	Y	J0721.9+7120	BLL	0.31	230	1.980 ± 0.049	$0.370^{+0.021}_{-0.019}$
J0722+3722	Y	···	FSRQ	1.629	125	0.094 ± 0.001	0.039 ± 0.008
J0725+1425	Y	J0725.3+1431	FSRQ	1.038	176	0.728 ± 0.009	$0.154^{+0.009}_{-0.008}$
J0726+0636	Y	···	FSRQ	1.95	127	0.119 ± 0.002	$0.146^{+0.014}_{-0.013}$
J0726+2153	Y	···	FSRQ	1.857	116	0.431 ± 0.003	$0.068^{+0.006}_{-0.005}$
J0726+7911	Y	···	···	···	116	0.591 ± 0.002	$0.020^{+0.003}_{-0.002}$
J0728+2153	Y	···	···	···	116	0.351 ± 0.003	$0.092^{+0.008}_{-0.007}$
J0728+5701	Y	···	FSRQ	0.424	202	0.406 ± 0.004	0.131 ± 0.007
J0730+4049	Y	···	FSRQ	2.501	135	0.102 ± 0.001	0.053 ± 0.008
J0731+2451	Y	···	FSRQ	1.089	111	0.139 ± 0.001	$0.044^{+0.006}_{-0.005}$
J0732+2548	Y	···	FSRQ	1.443	128	0.351 ± 0.002	$0.057^{+0.005}_{-0.004}$
J0733+0456	Y	···	FSRQ	3.01	123	0.323 ± 0.005	$0.148^{+0.011}_{-0.010}$
J0733+5022	Y	···	FSRQ	0.72	214	0.640 ± 0.003	0.068 ± 0.004
J0735+4750	Y	···	FSRQ	0.782	203	0.378 ± 0.002	0.065 ± 0.004
J0736+2604	Y	···	FSRQ	0.997	129	0.159 ± 0.001	$0.053^{+0.006}_{-0.005}$
J0738+1742	Y	J0738.2+1741	BLL	0.424	118	0.674 ± 0.004	$0.053^{+0.005}_{-0.004}$
J0739+0137	Y	J0739.1+0138	FSRQ	0.191	128	1.210 ± 0.019	$0.170^{+0.012}_{-0.011}$
J0740+2852	Y	···	FSRQ	0.711	137	0.156 ± 0.002	$0.092^{+0.009}_{-0.008}$

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
J0741+3112	Y	···	FSRQ	0.631	151	1.224 ± 0.005	0.042 ± 0.003
J0742+4900	Y	···	FSRQ	2.305	210	0.383 ± 0.003	0.109 ± 0.006
J0742+5444	···	J0742.2+5443	FSRQ	0.723	113	0.296 ± 0.013	$0.458^{+0.041}_{-0.035}$
J0743+1714	Y	···	BLL	···	105	0.176 ± 0.004	$0.191^{+0.016}_{-0.015}$
J0745+1011	Y	···	FSRQ	2.624	149	1.424 ± 0.003	$0.015^{+0.002}_{-0.001}$
J0745+3313	Y	···	FSRQ	0.61	153	0.125 ± 0.001	$0.116^{+0.009}_{-0.008}$
J0745–0044	Y	···	FSRQ	0.994	198	1.340 ± 0.005	0.050 ± 0.003
J0746+2549	···	J0746.6+2548	FSRQ	2.979	98	0.276 ± 0.007	$0.242^{+0.021}_{-0.018}$
J0746+2734	Y	···	···	···	169	0.543 ± 0.006	$0.130^{+0.008}_{-0.007}$
J0747+7639	Y	···	···	···	195	0.546 ± 0.002	0.036 ± 0.003
J0748+2400	Y	···	FSRQ	0.409	129	1.291 ± 0.015	$0.128^{+0.009}_{-0.008}$
J0749+7420	Y	···	FSRQ	1.629	195	0.348 ± 0.002	0.063 ± 0.004
J0750+1021	Y	···	FSRQ	1.117	169	0.212 ± 0.004	$0.252^{+0.017}_{-0.015}$
J0750+1231	Y	J0750.6+1235	FSRQ	0.889	158	4.170 ± 0.030	$0.089^{+0.006}_{-0.005}$
J0750+1823	Y	···	FSRQ	1.16	126	0.646 ± 0.009	$0.154^{+0.011}_{-0.010}$
J0750+4814	Y	···	FSRQ	1.956	218	0.609 ± 0.005	$0.109^{+0.006}_{-0.005}$
J0751+3313	Y	···	FSRQ	1.932	151	0.461 ± 0.004	$0.107^{+0.007}_{-0.006}$
J0752+3730	Y	···	AGN	0.44	162	0.227 ± 0.002	$0.122^{+0.009}_{-0.008}$
J0753+5352	Y	J0752.8+5353	BLL	0.2	207	0.734 ± 0.011	$0.219^{+0.012}_{-0.011}$
J0754+3033	Y	···	FSRQ	0.796	151	0.123 ± 0.001	0.065 ± 0.006
J0756+6347	Y	···	···	···	210	0.204 ± 0.001	0.036 ± 0.004
J0757+0956	Y	J0757.2+0956	BLL	0.266	159	1.420 ± 0.021	$0.186^{+0.012}_{-0.011}$
J0802+1809	Y	···	FSRQ	1.586	105	0.406 ± 0.002	$0.035^{+0.004}_{-0.003}$
J0805+6144	Y	J0806.2+6148	FSRQ	3.033	216	0.597 ± 0.005	0.111 ± 0.006
J0805–0111	Y	···	FSRQ	1.388	194	0.421 ± 0.007	$0.221^{+0.013}_{-0.012}$
J0806+4504	Y	···	FSRQ	2.102	164	0.464 ± 0.003	$0.069^{+0.005}_{-0.004}$
J0807+5117	Y	···	FSRQ	1.136	117	0.104 ± 0.001	0.065 ± 0.009
J0807–0541	···	J0807.0–0544	···	···	92	0.321 ± 0.001	0.023 ± 0.004
J0808+4052	Y	···	FSRQ	1.418	162	0.727 ± 0.006	0.103 ± 0.006
J0808+4950	Y	···	FSRQ	1.436	159	0.486 ± 0.007	$0.179^{+0.012}_{-0.010}$
J0808+7315	Y	···	FSRQ	0.496	193	0.291 ± 0.002	0.097 ± 0.006
J0808–0751	Y	J0808.2–0750	FSRQ	1.837	202	1.403 ± 0.043	$0.430^{+0.027}_{-0.024}$
J0809+3455	···	J0809.4+3455	BLL	0.082	108	0.097 ± 0.001	0.051 ± 0.008

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
J0809+5341	Y	···	FSRQ	2.133	186	0.212 ± 0.003	$0.187^{+0.012}_{-0.011}$
J0810+4134	Y	···	FSRQ	0.507	160	0.241 ± 0.003	$0.168^{+0.011}_{-0.010}$
J0811+0146	Y	J0811.2+0148	BLL	1.148	183	0.927 ± 0.020	$0.289^{+0.017}_{-0.016}$
J0811+4533	Y	···	FSRQ	1.017	160	0.181 ± 0.002	$0.161^{+0.010}_{-0.009}$
J0813+2542	Y	···	FSRQ	2.024	113	0.436 ± 0.007	$0.171^{+0.013}_{-0.012}$
J0814+6431	Y	J0815.0+6434	BLL	···	207	0.152 ± 0.002	0.125 ± 0.008
J0815+3635	Y	···	FSRQ	1.028	166	0.663 ± 0.004	0.069 ± 0.004
J0817+3227	Y	···	···	···	164	0.158 ± 0.001	0.038 ± 0.004
J0817–0933	Y	J0818.0–0938	BLL	···	201	0.248 ± 0.001	0.062 ± 0.005
J0818+4222	Y	J0818.2+4222	BLL	···	159	1.623 ± 0.021	$0.161^{+0.010}_{-0.009}$
J0818+4754	Y	···	···	···	158	0.067 ± 0.001	$0.080^{+0.012}_{-0.011}$
J0819+3226	Y	···	FSRQ	0.651	129	0.244 ± 0.005	$0.223^{+0.017}_{-0.016}$
J0823+2928	Y	···	FSRQ	2.368	110	0.396 ± 0.001	0.031 ± 0.003
J0824+3916	Y	···	FSRQ	1.216	164	1.265 ± 0.011	$0.105^{+0.007}_{-0.006}$
J0824+5552	Y	J0825.0+5555	FSRQ	1.417	188	0.735 ± 0.005	$0.096^{+0.006}_{-0.005}$
J0824–1527	Y	···	FSRQ	1.289	112	0.322 ± 0.005	$0.142^{+0.012}_{-0.010}$
J0824–1827	Y	···	···	···	107	0.158 ± 0.001	$0.062^{+0.008}_{-0.007}$
J0825+0309	Y	J0825.9+0309	BLL	0.506	204	1.069 ± 0.028	$0.369^{+0.022}_{-0.020}$
J0825+1332	Y	···	FSRQ	1.143	112	0.291 ± 0.007	$0.254^{+0.021}_{-0.018}$
J0825+6157	Y	···	FSRQ	0.542	168	0.481 ± 0.005	$0.122^{+0.008}_{-0.007}$
J0827+3525	Y	···	FSRQ	2.249	162	0.443 ± 0.001	0.027 ± 0.002
J0830+2410	Y	J0830.5+2407	FSRQ	0.94	119	1.300 ± 0.032	$0.266^{+0.020}_{-0.018}$
J0831+0429	Y	J0831.6+0429	BLL	0.174	207	0.594 ± 0.007	$0.151^{+0.009}_{-0.008}$
J0831+0847	Y	···	FSRQ	0.941	276	0.074 ± 0.001	0.098 ± 0.008
J0833+0350	Y	···	FSRQ	0.903	197	0.276 ± 0.002	0.090 ± 0.006
J0833+4224	Y	J0834.4+4221	FSRQ	0.249	159	0.314 ± 0.004	$0.169^{+0.011}_{-0.010}$
J0834+6019	Y	···	FSRQ	0.72	129	0.218 ± 0.002	$0.067^{+0.006}_{-0.005}$
J0835+6835	Y	···	FSRQ	1.414	128	0.137 ± 0.004	$0.327^{+0.025}_{-0.022}$
J0836+0052	Y	···	FSRQ	1.826	202	0.124 ± 0.001	0.060 ± 0.006
J0836+2728	Y	···	FSRQ	0.765	119	0.358 ± 0.002	0.062 ± 0.005
J0837+2454	Y	···	FSRQ	1.122	115	0.398 ± 0.009	$0.247^{+0.019}_{-0.017}$
J0837+5825	Y	···	FSRQ	2.101	158	1.001 ± 0.013	$0.164^{+0.010}_{-0.009}$
J0839+0104	Y	J0839.5+0059	FSRQ	1.123	189	0.520 ± 0.006	$0.141^{+0.008}_{-0.007}$

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
J0839+0319	Y	···	FSRQ	1.57	198	0.311 ± 0.001	$0.051^{+0.004}_{-0.003}$
J0842+1835	Y	···	FSRQ	1.272	101	0.469 ± 0.005	$0.099^{+0.008}_{-0.007}$
J0847–0703	Y	···	BLL	···	196	0.607 ± 0.009	$0.200^{+0.012}_{-0.011}$
J0849+5108	Y	···	FSRQ	0.584	181	0.285 ± 0.006	$0.264^{+0.016}_{-0.015}$
J0850–1213	Y	J0850.0–1213	FSRQ	0.566	185	0.716 ± 0.019	$0.364^{+0.023}_{-0.021}$
J0851+0845	Y	···	···	···	256	0.088 ± 0.001	$0.116^{+0.009}_{-0.008}$
J0853–0150	Y	···	FSRQ	1.498	191	no det.	···
J0854+2006	Y	J0854.8+2006	BLL	0.306	270	4.552 ± 0.127	$0.456^{+0.025}_{-0.022}$
J0854+5757	Y	···	FSRQ	1.318	172	0.471 ± 0.002	$0.053^{+0.004}_{-0.003}$
J0856+1739	Y	···	FSRQ	0.516	184	0.093 ± 0.001	$0.139^{+0.011}_{-0.010}$
J0856+2111	Y	···	FSRQ	2.098	201	0.298 ± 0.001	0.025 ± 0.002
J0856–1105	Y	J0856.6–1105	···	···	161	0.495 ± 0.006	$0.139^{+0.009}_{-0.008}$
J0900+4108	Y	···	FSRQ	1.629	126	0.233 ± 0.002	$0.093^{+0.008}_{-0.007}$
J0900–1242	Y	···	···	···	181	0.098 ± 0.001	0.122 ± 0.011
J0901+0448	Y	···	FSRQ	1.863	197	0.370 ± 0.006	$0.206^{+0.012}_{-0.011}$
J0902+4310	Y	···	FSRQ	2.41	118	0.255 ± 0.007	$0.304^{+0.024}_{-0.021}$
J0902+5402	Y	···	FSRQ	1.682	181	0.128 ± 0.001	$0.120^{+0.008}_{-0.007}$
J0903+6757	Y	···	FSRQ	1.499	118	0.400 ± 0.004	$0.096^{+0.008}_{-0.007}$
J0903–1721	Y	···	FSRQ	0.872	108	0.185 ± 0.002	$0.089^{+0.010}_{-0.009}$
J0905+1358	···	J0905.5+1356	···	···	92	0.061 ± 0.001	$0.086^{+0.017}_{-0.016}$
J0905+2849	Y	···	FSRQ	1.219	118	0.279 ± 0.002	0.055 ± 0.005
J0906+6930	Y	···	FSRQ	5.47	117	0.156 ± 0.003	$0.145^{+0.014}_{-0.012}$
J0908+1609	Y	···	···	···	240	0.199 ± 0.002	$0.120^{+0.008}_{-0.007}$
J0909+0121	Y	J0909.0+0126	FSRQ	1.024	199	1.373 ± 0.026	$0.268^{+0.015}_{-0.014}$
J0909+0200	Y	···	BLL	···	173	0.109 ± 0.002	$0.156^{+0.012}_{-0.011}$
J0910+2248	Y	J0911.0+2247	FSRQ	2.661	204	0.165 ± 0.002	0.159 ± 0.009
J0910+3329	···	J0910.7+3332	BLL	0.354	105	0.120 ± 0.002	$0.151^{+0.015}_{-0.013}$
J0914+0245	Y	···	···	···	171	0.731 ± 0.009	0.154 ± 0.009
J0915+2933	···	J0915.7+2931	BLL	···	105	0.109 ± 0.001	$0.090^{+0.010}_{-0.009}$
J0917–1345	Y	···	···	···	108	0.187 ± 0.001	$0.058^{+0.007}_{-0.006}$
J0919+3324	Y	···	···	···	118	0.168 ± 0.002	$0.126^{+0.011}_{-0.009}$
J0920+4441	Y	J0920.9+4441	FSRQ	2.19	121	2.058 ± 0.005	0.023 ± 0.002
J0921+6215	Y	J0919.6+6216	FSRQ	1.446	189	1.305 ± 0.007	0.068 ± 0.004

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
J0922–0529	Y	···	NLRG	0.974	193	0.303 ± 0.003	0.119 ± 0.007
J0923+2815	···	J0924.2+2812	FSRQ	0.744	108	1.189 ± 0.007	0.052 ± 0.004
J0923+3849	Y	···	···	···	124	0.311 ± 0.003	$0.113^{+0.009}_{-0.008}$
J0923+4125	Y	J0923.2+4121	AGN	0.028	126	0.359 ± 0.004	$0.119^{+0.009}_{-0.008}$
J0925+1658	Y	···	FSRQ	1.545	265	0.095 ± 0.002	$0.379^{+0.023}_{-0.021}$
J0926+4029	Y	···	FSRQ	1.879	122	0.146 ± 0.002	$0.094^{+0.009}_{-0.008}$
J0927+3902	Y	···	FSRQ	0.695	117	10.543 ± 0.030	0.029 ± 0.002
J0928+4446	Y	···	FSRQ	1.904	119	0.228 ± 0.002	$0.108^{+0.009}_{-0.008}$
J0929+5013	Y	···	FSRQ	0.37	200	0.414 ± 0.004	$0.121^{+0.007}_{-0.006}$
J0929+8612	Y	···	BLL	···	181	0.151 ± 0.002	$0.179^{+0.011}_{-0.010}$
J0930+7420	Y	···	···	···	193	0.248 ± 0.001	0.060 ± 0.004
J0932+5306	Y	···	FSRQ	0.597	191	0.231 ± 0.004	$0.208^{+0.012}_{-0.011}$
J0933–0819	Y	···	FSRQ	0.903	191	0.405 ± 0.004	$0.110^{+0.007}_{-0.006}$
J0933–1139	Y	···	···	···	189	0.320 ± 0.002	$0.064^{+0.005}_{-0.004}$
J0934+3926	Y	J0934.5+3929	BLL	···	124	0.131 ± 0.003	$0.194^{+0.017}_{-0.015}$
J0935–1939	Y	···	···	···	109	0.083 ± 0.003	$0.298^{+0.027}_{-0.024}$
J0936–0535	Y	···	···	···	189	0.163 ± 0.001	0.080 ± 0.007
J0937+5008	Y	J0937.7+5005	FSRQ	0.276	111	0.418 ± 0.011	$0.281^{+0.023}_{-0.020}$
J0938–0708	Y	···	FSRQ	1.281	192	0.326 ± 0.004	$0.162^{+0.010}_{-0.009}$
J0939+4141	Y	···	FSRQ	1.224	118	0.141 ± 0.002	$0.112^{+0.011}_{-0.010}$
J0940+2603	Y	···	FSRQ	0.498	262	0.341 ± 0.001	0.056 ± 0.003
J0941+2728	Y	J0941.2+2722	FSRQ	1.306	271	0.209 ± 0.002	0.117 ± 0.006
J0941–1335	Y	···	FSRQ	0.551	194	0.304 ± 0.004	0.153 ± 0.009
J0942+6403	Y	···	FSRQ	0.952	215	0.063 ± 0.001	$0.258^{+0.017}_{-0.016}$
J0942–0759	Y	···	···	···	189	0.185 ± 0.003	$0.196^{+0.013}_{-0.012}$
J0943+1702	Y	···	FSRQ	1.598	258	0.216 ± 0.001	0.036 ± 0.003
J0943+3614	Y	···	AGN	0.022	170	0.194 ± 0.001	0.029 ± 0.004
J0945+4636	Y	···	FSRQ	0.639	120	0.270 ± 0.002	$0.069^{+0.006}_{-0.005}$
J0945+5757	···	J0945.6+5754	BLL	0.229	95	0.066 ± 0.001	$0.106^{+0.017}_{-0.015}$
J0946+1017	···	J0946.6+1012	FSRQ	1.007	102	0.214 ± 0.002	0.070 ± 0.008
J0948+0022	Y	J0949.0+0021	FSRQ	0.585	189	0.373 ± 0.012	$0.440^{+0.029}_{-0.026}$
J0948+4039	Y	···	FSRQ	1.249	174	1.569 ± 0.013	0.104 ± 0.006
J0952+3512	Y	···	FSRQ	1.876	171	0.340 ± 0.002	0.053 ± 0.004

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
J0952+5048	Y	···	FSRQ	1.091	191	0.086 ± 0.001	$0.143^{+0.012}_{-0.011}$
J0953+3225	Y	···	FSRQ	1.574	172	0.201 ± 0.002	0.083 ± 0.006
J0954+2639	Y	···	···	···	260	0.167 ± 0.001	$0.104^{+0.006}_{-0.005}$
J0956+2515	Y	J0956.9+2513	FSRQ	0.712	268	0.931 ± 0.011	0.192 ± 0.009
J0957–0156	Y	···	FSRQ	0.86	191	0.094 ± 0.001	$0.077^{+0.008}_{-0.007}$
J0957–1350	Y	···	FSRQ	1.323	192	0.428 ± 0.003	$0.085^{+0.006}_{-0.005}$
J0958+4725	Y	···	FSRQ	1.882	163	0.840 ± 0.010	0.154 ± 0.009
J0958+5039	Y	···	FSRQ	1.154	162	0.219 ± 0.002	$0.099^{+0.007}_{-0.006}$
J0958+6533	Y	J1000.1+6539	BLL	0.367	219	1.427 ± 0.025	$0.253^{+0.014}_{-0.013}$
J1001+1015	Y	···	FSRQ	1.533	192	0.197 ± 0.001	0.037 ± 0.004
J1001+2911	Y	J1000.9+2915	BLL	···	169	0.469 ± 0.009	$0.246^{+0.015}_{-0.014}$
J1001+3424	Y	···	FSRQ	0.948	174	0.268 ± 0.002	0.099 ± 0.006
J1002+1216	Y	···	FSRQ	0.861	200	0.206 ± 0.001	0.074 ± 0.005
J1007+1356	Y	···	FSRQ	2.707	185	0.625 ± 0.003	0.052 ± 0.003
J1007–0207	Y	···	FSRQ	1.214	196	0.504 ± 0.005	$0.137^{+0.008}_{-0.007}$
J1008+0621	Y	J1007.9+0619	BLL	···	145	0.472 ± 0.008	$0.188^{+0.013}_{-0.012}$
J1010+3330	Y	···	FSRQ	2.068	172	0.242 ± 0.003	0.128 ± 0.008
J1010+8250	Y	···	FSRQ	0.322	196	0.404 ± 0.001	0.041 ± 0.003
J1010–0200	···	J1011.0–0156	FSRQ	0.887	90	0.466 ± 0.005	$0.100^{+0.009}_{-0.008}$
J1012+2312	Y	···	FSRQ	0.749	249	0.343 ± 0.003	$0.149^{+0.008}_{-0.007}$
J1013+2449	Y	···	FSRQ	1.636	202	0.806 ± 0.002	0.020 ± 0.002
J1013+3445	Y	···	FSRQ	1.414	176	0.475 ± 0.003	0.084 ± 0.005
J1014+2301	Y	···	FSRQ	0.565	238	0.727 ± 0.004	0.085 ± 0.004
J1015+1227	Y	···	BLL	···	197	0.403 ± 0.004	$0.146^{+0.009}_{-0.008}$
J1015+4926	Y	J1015.1+4927	BLL	0.2	164	0.278 ± 0.001	0.040 ± 0.004
J1015+6728	Y	···	···	···	120	0.108 ± 0.001	$0.059^{+0.008}_{-0.007}$
J1016+2037	Y	···	FSRQ	3.11	244	0.361 ± 0.001	0.048 ± 0.003
J1017+6116	Y	···	FSRQ	2.805	119	0.391 ± 0.002	0.059 ± 0.005
J1018+0530	Y	···	FSRQ	1.938	189	0.382 ± 0.005	$0.172^{+0.011}_{-0.010}$
J1018+3542	Y	···	FSRQ	1.226	179	0.411 ± 0.002	0.048 ± 0.003
J1019+6320	Y	···	BLL	···	121	0.206 ± 0.003	$0.164^{+0.013}_{-0.012}$
J1021+3437	Y	···	FSRQ	1.4	168	0.289 ± 0.001	0.046 ± 0.003
J1022+4239	Y	···	FSRQ	0.991	165	0.300 ± 0.002	$0.062^{+0.005}_{-0.004}$

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
J1023+2856	Y	···	FSRQ	0.671	201	0.144 ± 0.001	$0.116^{+0.008}_{-0.007}$
J1023+3948	Y	···	FSRQ	1.254	168	0.660 ± 0.004	$0.078^{+0.005}_{-0.004}$
J1024+1912	Y	···	FSRQ	0.828	238	0.398 ± 0.002	$0.080^{+0.005}_{-0.004}$
J1025+1253	Y	···	FSRQ	0.663	206	0.703 ± 0.007	$0.145^{+0.008}_{-0.007}$
J1025–0509	Y	···	···	···	189	0.202 ± 0.003	$0.183^{+0.011}_{-0.010}$
J1028+0255	Y	···	FSRQ	0.715	190	0.225 ± 0.002	$0.078^{+0.006}_{-0.005}$
J1029–1852	Y	···	FSRQ	1.784	111	0.332 ± 0.003	$0.082^{+0.008}_{-0.007}$
J1033+0711	Y	···	FSRQ	1.535	182	0.320 ± 0.004	$0.173^{+0.011}_{-0.010}$
J1033+3935	Y	···	FSRQ	1.095	163	0.209 ± 0.001	0.027 ± 0.003
J1033+4116	Y	J1033.2+4116	FSRQ	1.117	162	1.565 ± 0.048	$0.384^{+0.026}_{-0.023}$
J1033+6051	Y	···	FSRQ	1.401	122	0.400 ± 0.012	$0.316^{+0.025}_{-0.022}$
J1036+1440	Y	···	FSRQ	1.373	213	0.134 ± 0.001	$0.135^{+0.008}_{-0.007}$
J1036+2203	Y	···	FSRQ	0.595	239	0.293 ± 0.003	0.171 ± 0.009
J1036–0605	Y	···	···	···	188	0.156 ± 0.001	0.061 ± 0.005
J1038+0512	Y	···	FSRQ	0.473	179	1.428 ± 0.019	0.178 ± 0.010
J1039–1541	Y	···	FSRQ	0.525	111	0.390 ± 0.010	$0.268^{+0.022}_{-0.019}$
J1041+0610	Y	···	FSRQ	1.265	189	1.344 ± 0.005	0.051 ± 0.003
J1041+5233	Y	···	FSRQ	0.678	147	0.427 ± 0.003	$0.080^{+0.006}_{-0.005}$
J1043+2408	Y	J1043.1+2404	BLL	0.56	246	0.989 ± 0.011	0.169 ± 0.008
J1044+2959	Y	···	FSRQ	2.983	194	0.085 ± 0.001	$0.109^{+0.009}_{-0.008}$
J1044+5322	Y	···	FSRQ	1.901	131	0.331 ± 0.005	$0.176^{+0.013}_{-0.011}$
J1044+8054	Y	J1048.7+8054	FSRQ	1.26	195	0.736 ± 0.009	$0.162^{+0.009}_{-0.008}$
J1045+0624	Y	···	FSRQ	1.509	186	0.306 ± 0.001	0.032 ± 0.003
J1046+5354	Y	···	FSRQ	1.71	125	0.166 ± 0.003	$0.205^{+0.016}_{-0.014}$
J1047+2635	Y	···	FSRQ	0.99	231	0.112 ± 0.001	0.112 ± 0.008
J1047–1308	Y	···	FSRQ	1.288	183	0.230 ± 0.002	0.081 ± 0.006
J1048+7143	Y	J1048.8+7145	FSRQ	1.15	208	1.116 ± 0.013	$0.168^{+0.009}_{-0.008}$
J1048–1909	Y	···	FSRQ	0.595	117	1.133 ± 0.013	$0.119^{+0.009}_{-0.008}$
J1051+2027	Y	···	FSRQ	1.036	239	0.180 ± 0.003	$0.235^{+0.013}_{-0.012}$
J1051+2119	Y	···	FSRQ	1.3	233	0.404 ± 0.005	0.174 ± 0.009
J1054–0713	Y	···	FSRQ	1.618	181	0.137 ± 0.002	$0.224^{+0.014}_{-0.013}$
J1056+7011	Y	···	FSRQ	2.492	121	0.831 ± 0.025	$0.331^{+0.025}_{-0.022}$
J1058+0133	Y	J1058.4+0134	FSRQ	0.888	169	4.848 ± 0.025	0.063 ± 0.004

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
J1058+1951	Y	···	FSRQ	1.11	260	0.587 ± 0.003	$0.069^{+0.004}_{-0.003}$
J1058+5628	Y	J1058.6+5628	BLL	0.143	111	0.172 ± 0.001	$0.056^{+0.007}_{-0.006}$
J1058+8114	Y	···	FSRQ	0.706	191	0.916 ± 0.014	$0.205^{+0.012}_{-0.011}$
J1059+2057	Y	···	AGN	0.4	264	0.287 ± 0.003	0.188 ± 0.009
J1102+2757	Y	···	FSRQ	1.861	229	0.306 ± 0.003	0.137 ± 0.007
J1103+3014	Y	···	FSRQ	0.387	204	0.308 ± 0.004	$0.157^{+0.009}_{-0.008}$
J1104+0730	···	J1104.4+0734	BLL	···	68	0.130 ± 0.002	$0.107^{+0.014}_{-0.012}$
J1104+3812	Y	J1104.4+3812	BLL	0.03	319	0.455 ± 0.002	0.086 ± 0.004
J1106+2812	···	J1106.5+2809	···	···	105	0.241 ± 0.003	$0.119^{+0.011}_{-0.010}$
J1108+0811	Y	···	FSRQ	1.123	175	0.132 ± 0.001	0.081 ± 0.007
J1108+4330	Y	···	FSRQ	1.226	209	0.389 ± 0.002	0.067 ± 0.004
J1110+4403	Y	···	FSRQ	1.693	213	0.176 ± 0.002	$0.148^{+0.009}_{-0.008}$
J1110–1835	···	J1110.2–1839	···	···	90	no det.	···
J1113+1442	Y	···	FSRQ	0.866	256	0.283 ± 0.004	$0.195^{+0.010}_{-0.009}$
J1114–0816	Y	···	FSRQ	2.078	171	0.317 ± 0.004	$0.143^{+0.009}_{-0.008}$
J1116+0829	Y	···	AGN	0.486	172	0.312 ± 0.004	$0.176^{+0.011}_{-0.010}$
J1118+1234	Y	···	FSRQ	2.129	301	1.159 ± 0.007	$0.105^{+0.005}_{-0.004}$
J1119+0410	Y	···	FSRQ	0.736	170	0.179 ± 0.001	$0.091^{+0.007}_{-0.006}$
J1119+1656	Y	···	FSRQ	0.949	264	0.140 ± 0.001	0.064 ± 0.005
J1120–1420	Y	···	FSRQ	1.114	196	0.440 ± 0.003	$0.095^{+0.006}_{-0.005}$
J1121–0553	···	J1121.5–0554	FSRQ	1.297	75	0.409 ± 0.006	$0.105^{+0.012}_{-0.010}$
J1121–0711	Y	···	BLL	···	158	0.292 ± 0.002	0.077 ± 0.006
J1121–1722	Y	···	FSRQ	0.986	137	0.108 ± 0.001	$0.097^{+0.009}_{-0.008}$
J1122+1805	Y	···	FSRQ	1.04	266	0.506 ± 0.002	$0.044^{+0.003}_{-0.002}$
J1124+2336	Y	J1123.9+2339	BLL	···	256	0.453 ± 0.009	$0.297^{+0.015}_{-0.014}$
J1125+0001	Y	···	FSRQ	1.696	173	0.163 ± 0.001	0.068 ± 0.006
J1125+2610	Y	···	FSRQ	2.341	298	0.912 ± 0.009	0.163 ± 0.007
J1126+4516	Y	···	FSRQ	1.811	208	0.194 ± 0.001	0.055 ± 0.004
J1127+0555	Y	···	FSRQ	2.217	177	0.427 ± 0.014	$0.418^{+0.028}_{-0.025}$
J1127+3620	Y	···	FSRQ	0.884	212	0.095 ± 0.002	$0.218^{+0.015}_{-0.013}$
J1127+5650	Y	···	FSRQ	2.89	205	0.260 ± 0.003	$0.134^{+0.008}_{-0.007}$
J1127–1857	Y	J1126.8–1854	FSRQ	1.048	162	1.855 ± 0.022	$0.148^{+0.009}_{-0.008}$
J1128+1039	Y	···	···	···	187	0.240 ± 0.001	$0.058^{+0.005}_{-0.004}$

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
J1128+5925	Y	···	FSRQ	1.795	205	0.324 ± 0.003	0.140 ± 0.008
J1129–0240	Y	···	FSRQ	2.093	179	0.170 ± 0.001	0.087 ± 0.006
J1130+3815	Y	···	FSRQ	1.733	223	1.271 ± 0.009	0.103 ± 0.005
J1131–0500	Y	···	AGN	0.266	157	0.761 ± 0.002	0.021 ± 0.002
J1133+0040	Y	···	FSRQ	1.633	178	0.374 ± 0.003	0.104 ± 0.007
J1135–0428	Y	···	FSRQ	0.273	171	0.430 ± 0.011	$0.326^{+0.021}_{-0.019}$
J1136+3407	Y	···	FSRQ	1.332	221	0.113 ± 0.001	$0.148^{+0.010}_{-0.009}$
J1136+7009	Y	J1136.6+7009	BLL	0.045	205	0.173 ± 0.001	0.071 ± 0.005
J1136–0330	Y	···	FSRQ	1.648	171	0.415 ± 0.002	0.038 ± 0.004
J1141+6410	Y	···	···	···	207	0.168 ± 0.002	0.164 ± 0.010
J1143+6633	Y	···	FSRQ	2.328	210	0.086 ± 0.001	0.090 ± 0.009
J1145+0455	Y	···	FSRQ	1.342	171	0.242 ± 0.002	0.091 ± 0.006
J1146+3958	Y	J1146.8+4004	FSRQ	1.089	210	0.921 ± 0.009	0.135 ± 0.007
J1146+5356	Y	···	FSRQ	2.201	211	0.440 ± 0.004	$0.120^{+0.007}_{-0.006}$
J1146+5848	Y	···	FSRQ	1.982	205	0.393 ± 0.007	$0.261^{+0.015}_{-0.013}$
J1147+2635	Y	···	FSRQ	0.867	301	0.384 ± 0.003	0.140 ± 0.006
J1147+3501	Y	···	AGN	0.063	225	0.211 ± 0.002	0.104 ± 0.006
J1147–0724	Y	J1147.7–0722	FSRQ	1.342	176	0.671 ± 0.005	$0.088^{+0.006}_{-0.005}$
J1148+5254	Y	···	FSRQ	1.633	208	0.360 ± 0.002	0.086 ± 0.005
J1148+5924	Y	···	AGN	0.011	205	0.428 ± 0.002	0.060 ± 0.004
J1148–0404	Y	···	FSRQ	0.341	169	0.266 ± 0.002	$0.085^{+0.006}_{-0.005}$
J1150+2417	Y	J1150.2+2419	BLL	0.2	309	0.724 ± 0.001	0.024 ± 0.002
J1150+4332	Y	···	FSRQ	3.037	212	0.126 ± 0.001	$0.087^{+0.007}_{-0.006}$
J1150–0640	Y	···	···	···	177	0.422 ± 0.005	0.139 ± 0.008
J1152+4939	Y	···	FSRQ	1.093	208	0.715 ± 0.005	0.089 ± 0.005
J1152–0519	Y	···	FSRQ	1.983	175	0.132 ± 0.002	$0.166^{+0.012}_{-0.011}$
J1152–0841	Y	J1152.2–0836	FSRQ	2.367	172	1.028 ± 0.026	$0.325^{+0.021}_{-0.019}$
J1153+8058	Y	···	FSRQ	1.25	194	1.066 ± 0.006	$0.080^{+0.005}_{-0.004}$
J1154+1225	Y	···	FSRQ	0.081	185	0.102 ± 0.001	0.044 ± 0.006
J1154+5934	Y	···	FSRQ	0.871	198	0.162 ± 0.001	0.071 ± 0.005
J1157+5527	Y	···	AGN	0.004	210	0.184 ± 0.003	$0.223^{+0.013}_{-0.012}$
J1158+2450	Y	···	FSRQ	0.202	310	0.376 ± 0.002	0.104 ± 0.005
J1158+4825	Y	···	FSRQ	2.028	209	0.153 ± 0.001	0.118 ± 0.007

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
J1159+2914	Y	J1159.4+2914	FSRQ	0.729	230	2.383 ± 0.049	$0.311^{+0.017}_{-0.015}$
J1201+1431	Y	···	···	···	217	0.305 ± 0.002	0.063 ± 0.004
J1202–0528	Y	···	···	···	174	0.556 ± 0.011	$0.246^{+0.015}_{-0.014}$
J1203+4803	Y	···	FSRQ	0.817	241	0.552 ± 0.004	0.109 ± 0.005
J1203+6031	···	J1202.9+6032	AGN	0.065	79	0.168 ± 0.003	$0.151^{+0.015}_{-0.014}$
J1207+1211	Y	···	FSRQ	0.89	208	0.362 ± 0.006	$0.218^{+0.012}_{-0.011}$
J1207+2754	Y	···	FSRQ	2.177	310	0.274 ± 0.003	0.198 ± 0.009
J1209+1810	···	J1209.7+1806	FSRQ	0.845	97	0.253 ± 0.004	$0.144^{+0.013}_{-0.011}$
J1209+2547	Y	···	FSRQ	1.436	271	0.219 ± 0.001	$0.052^{+0.004}_{-0.003}$
J1209+4119	Y	J1209.4+4119	BLL	···	236	0.156 ± 0.003	$0.241^{+0.013}_{-0.012}$
J1210–1218	Y	···	···	···	178	no det.	···
J1214+0829	Y	···	FSRQ	2.359	180	0.162 ± 0.002	$0.117^{+0.008}_{-0.007}$
J1215+1654	Y	···	FSRQ	1.132	238	0.392 ± 0.005	$0.181^{+0.010}_{-0.009}$
J1215–1731	Y	···	···	···	108	1.799 ± 0.007	0.034 ± 0.003
J1217+3007	Y	J1217.7+3007	BLL	0.13	209	0.385 ± 0.003	$0.094^{+0.006}_{-0.005}$
J1217+5835	Y	···	FSRQ	2.552	194	0.312 ± 0.002	0.077 ± 0.005
J1219+4829	Y	···	FSRQ	1.076	241	0.801 ± 0.012	$0.224^{+0.012}_{-0.011}$
J1219+6344	Y	···	···	···	227	0.170 ± 0.000	0.024 ± 0.003
J1219+6600	Y	···	FSRQ	1.266	227	0.188 ± 0.001	0.080 ± 0.005
J1220+3431	Y	J1220.2+3432	BLL	···	216	0.138 ± 0.001	$0.099^{+0.008}_{-0.007}$
J1220+3808	Y	···	···	···	243	0.222 ± 0.002	0.119 ± 0.006
J1221+2813	Y	J1221.5+2814	BLL	0.102	300	0.419 ± 0.002	$0.096^{+0.005}_{-0.004}$
J1221+4411	Y	···	FSRQ	1.344	240	0.368 ± 0.001	0.045 ± 0.003
J1222+0413	Y	J1222.5+0415	FSRQ	0.965	154	1.013 ± 0.013	$0.160^{+0.010}_{-0.009}$
J1223+8040	···	J1224.8+8044	BLL	···	100	0.518 ± 0.004	0.062 ± 0.005
J1226–1328	Y	J1226.7–1332	BLL	0.456	195	0.099 ± 0.003	$0.428^{+0.029}_{-0.026}$
J1227+4932	Y	···	FSRQ	1.348	196	0.142 ± 0.001	$0.107^{+0.008}_{-0.007}$
J1228+3128	Y	···	FSRQ	2.195	210	0.187 ± 0.001	0.053 ± 0.004
J1228+3706	Y	···	FSRQ	1.515	236	0.455 ± 0.004	0.124 ± 0.006
J1229+0203	Y	J1229.1+0203	FSRQ	0.158	303	28.787 ± 0.081	0.048 ± 0.002
J1230+2518	Y	J1230.4+2520	BLL	0.135	185	0.294 ± 0.007	$0.327^{+0.020}_{-0.018}$
J1231+0418	Y	···	FSRQ	1.03	188	0.300 ± 0.003	$0.119^{+0.007}_{-0.006}$
J1231+2847	···	J1231.6+2850	BLL	0.236	116	0.113 ± 0.001	$0.085^{+0.009}_{-0.008}$

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
J1232+4821	Y	···	FSRQ	1.588	237	0.223 ± 0.002	$0.097^{+0.006}_{-0.005}$
J1235+3621	Y	···	FSRQ	1.598	229	0.212 ± 0.002	0.106 ± 0.006
J1238+0723	Y	···	FSRQ	1.172	179	0.327 ± 0.004	$0.175^{+0.011}_{-0.010}$
J1239+0730	Y	···	FSRQ	0.4	181	0.826 ± 0.005	0.084 ± 0.005
J1239–1023	Y	···	FSRQ	0.75	191	0.885 ± 0.003	$0.039^{+0.003}_{-0.002}$
J1242+3750	Y	···	FSRQ	1.316	232	0.320 ± 0.001	0.037 ± 0.003
J1243+7442	Y	···	FSRQ	0.782	122	0.191 ± 0.002	0.086 ± 0.008
J1243–0218	Y	···	···	···	162	0.101 ± 0.002	$0.242^{+0.018}_{-0.016}$
J1245–1617	Y	···	···	···	118	0.348 ± 0.002	$0.053^{+0.006}_{-0.005}$
J1247+1022	Y	···	FSRQ	1.37	259	0.073 ± 0.000	0.075 ± 0.007
J1248+5820	Y	J1248.2+5820	BLL	···	170	0.147 ± 0.002	$0.166^{+0.011}_{-0.010}$
J1248–0632	Y	···	FSRQ	0.762	165	0.502 ± 0.004	$0.103^{+0.007}_{-0.006}$
J1251–1717	Y	···	FSRQ	0.606	117	0.187 ± 0.002	$0.075^{+0.007}_{-0.006}$
J1253+0326	···	J1253.7+0326	BLL	0.065	84	0.055 ± 0.001	$0.087^{+0.017}_{-0.016}$
J1254+1141	Y	···	FSRQ	0.87	277	0.536 ± 0.006	0.179 ± 0.008
J1254–1317	Y	···	···	···	180	0.173 ± 0.005	$0.370^{+0.025}_{-0.022}$
J1255+1817	Y	···	FSRQ	1.367	290	0.381 ± 0.003	$0.117^{+0.006}_{-0.005}$
J1256–0547	Y	J1256.2–0547	FSRQ	0.536	278	$15.399^{+0.163}_{-0.162}$	$0.174^{+0.008}_{-0.007}$
J1256–1146	···	J1256.5–1148	···	···	88	0.101 ± 0.001	$0.051^{+0.009}_{-0.008}$
J1257+3229	Y	J1258.3+3227	FSRQ	0.806	304	0.679 ± 0.008	$0.192^{+0.009}_{-0.008}$
J1300+1206	Y	···	···	···	290	0.151 ± 0.002	$0.230^{+0.011}_{-0.010}$
J1300+2830	Y	···	FSRQ	0.645	295	0.362 ± 0.002	0.083 ± 0.004
J1300+5029	Y	···	FSRQ	1.561	186	0.293 ± 0.001	0.039 ± 0.003
J1301+4634	Y	···	GALAXY	0.205	223	0.129 ± 0.001	$0.038^{+0.005}_{-0.004}$
J1302+5748	Y	···	FSRQ	1.088	190	0.464 ± 0.009	$0.259^{+0.015}_{-0.014}$
J1303–1051	Y	···	BLL	···	183	0.146 ± 0.001	0.032 ± 0.005
J1305–1033	Y	···	FSRQ	0.286	176	0.465 ± 0.010	$0.289^{+0.018}_{-0.016}$
J1306+5529	Y	···	FSRQ	1.601	194	0.220 ± 0.003	$0.159^{+0.009}_{-0.008}$
J1306–1718	Y	···	···	···	113	0.246 ± 0.002	$0.068^{+0.007}_{-0.006}$
J1308+3546	Y	J1308.5+3550	FSRQ	1.055	289	0.459 ± 0.003	$0.120^{+0.006}_{-0.005}$
J1309+1154	Y	J1309.2+1156	BLL	···	297	0.559 ± 0.002	0.052 ± 0.003
J1309+5557	Y	···	FSRQ	1.629	185	0.265 ± 0.003	0.136 ± 0.008
J1310+0044	Y	···	FSRQ	1.603	108	0.192 ± 0.003	$0.158^{+0.013}_{-0.012}$

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
J1310+3220	Y	J1310.6+3222	FSRQ	0.997	297	2.349 ± 0.032	$0.234^{+0.011}_{-0.010}$
J1310+3233	Y	···	FSRQ	1.65	289	0.320 ± 0.012	$0.615^{+0.037}_{-0.033}$
J1310+4653	Y	···	FSRQ	0.972	226	0.127 ± 0.001	0.117 ± 0.007
J1311+5513	Y	···	FSRQ	0.924	181	0.189 ± 0.002	$0.115^{+0.008}_{-0.007}$
J1312+4828	Y	J1312.4+4827	FSRQ	0.501	179	0.144 ± 0.003	$0.238^{+0.016}_{-0.014}$
J1312–0424	Y	···	FSRQ	0.825	120	0.245 ± 0.004	$0.146^{+0.012}_{-0.011}$
J1313+1027	Y	···	FSRQ	2.901	288	0.089 ± 0.001	0.089 ± 0.006
J1314+2348	···	J1314.7+2346	BLL	···	191	0.153 ± 0.001	0.070 ± 0.006
J1314+5306	Y	···	···	···	181	0.172 ± 0.001	0.065 ± 0.006
J1317+3425	Y	J1317.8+3425	FSRQ	1.05	295	0.437 ± 0.004	0.136 ± 0.006
J1317–1345	Y	···	···	···	120	0.362 ± 0.003	$0.089^{+0.007}_{-0.006}$
J1318–0607	Y	···	FSRQ	2.734	193	0.293 ± 0.001	0.027 ± 0.003
J1319–1217	Y	···	FSRQ	0.86	116	0.301 ± 0.004	$0.143^{+0.012}_{-0.011}$
J1319–1239	Y	···	AGN	0.008	117	0.178 ± 0.001	$0.054^{+0.009}_{-0.008}$
J1321+2216	Y	J1321.1+2214	FSRQ	0.943	379	0.355 ± 0.006	$0.315^{+0.013}_{-0.012}$
J1322+3912	Y	···	FSRQ	2.985	227	0.086 ± 0.001	0.080 ± 0.007
J1322–0937	Y	J1322.7–0943	FSRQ	1.864	192	0.589 ± 0.011	$0.246^{+0.014}_{-0.013}$
J1323+7942	Y	···	FSRQ	1.97	118	0.482 ± 0.006	$0.133^{+0.010}_{-0.009}$
J1324+4743	Y	···	FSRQ	2.26	182	0.166 ± 0.001	$0.078^{+0.007}_{-0.006}$
J1324–1049	Y	···	FSRQ	0.872	199	0.594 ± 0.007	$0.168^{+0.010}_{-0.009}$
J1325–0804	Y	···	FSRQ	2.354	190	0.226 ± 0.001	$0.062^{+0.005}_{-0.004}$
J1326–0500	Y	···	FSRQ	1.882	189	0.123 ± 0.002	$0.246^{+0.015}_{-0.014}$
J1327+1223	Y	···	FSRQ	0.95	294	0.361 ± 0.004	$0.175^{+0.008}_{-0.007}$
J1327+2210	Y	J1326.6+2213	FSRQ	1.4	354	1.114 ± 0.009	0.151 ± 0.006
J1327+5008	Y	···	FSRQ	1.012	182	0.134 ± 0.002	$0.142^{+0.010}_{-0.009}$
J1327–1336	Y	···	FSRQ	1.33	119	0.127 ± 0.002	$0.172^{+0.015}_{-0.013}$
J1329+3154	Y	···	BLL	···	270	0.675 ± 0.002	0.036 ± 0.002
J1330+5202	Y	J1331.0+5202	AGN	0.688	181	0.152 ± 0.001	$0.068^{+0.007}_{-0.006}$
J1332+4722	Y	J1332.9+4728	FSRQ	0.669	181	0.314 ± 0.006	$0.256^{+0.016}_{-0.014}$
J1332–0509	Y	J1331.9–0506	FSRQ	2.15	188	1.092 ± 0.024	$0.292^{+0.017}_{-0.016}$
J1332–1256	···	J1332.6–1255	FSRQ	1.498	94	0.267 ± 0.012	$0.407^{+0.039}_{-0.033}$
J1333+1649	Y	···	FSRQ	2.097	290	0.377 ± 0.001	0.024 ± 0.002
J1333+2725	Y	···	FSRQ	2.126	384	0.532 ± 0.005	$0.195^{+0.008}_{-0.007}$

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
J1333–1950	Y	···	···	···	108	0.285 ± 0.004	$0.125^{+0.011}_{-0.010}$
J1334–1150	Y	···	FSRQ	1.402	118	0.336 ± 0.007	$0.227^{+0.018}_{-0.016}$
J1335+4542	Y	···	FSRQ	2.452	216	0.631 ± 0.003	0.068 ± 0.004
J1335+5844	Y	···	···	···	180	0.543 ± 0.003	0.058 ± 0.004
J1336–0829	Y	···	GALAXY	0.023	186	0.539 ± 0.002	0.041 ± 0.003
J1337+5501	Y	···	FSRQ	1.099	178	0.383 ± 0.004	$0.147^{+0.009}_{-0.008}$
J1337+6532	Y	···	FSRQ	0.946	139	0.180 ± 0.002	$0.100^{+0.009}_{-0.008}$
J1337–1257	Y	J1337.7–1255	FSRQ	0.539	110	4.821 ± 0.100	$0.215^{+0.016}_{-0.015}$
J1341+2816	Y	···	FSRQ	1.275	292	0.246 ± 0.001	$0.064^{+0.004}_{-0.003}$
J1342+2709	Y	···	FSRQ	1.185	377	0.370 ± 0.003	0.149 ± 0.006
J1343+6602	Y	···	FSRQ	0.766	141	0.412 ± 0.004	0.106 ± 0.007
J1344+6606	Y	···	FSRQ	1.351	140	0.400 ± 0.004	$0.100^{+0.007}_{-0.006}$
J1344–1723	Y	J1344.2–1723	FSRQ	2.49	113	0.547 ± 0.012	$0.235^{+0.018}_{-0.016}$
J1345+0706	···	J1346.0+0703	···	···	96	0.172 ± 0.007	$0.395^{+0.037}_{-0.032}$
J1347+1835	Y	···	FSRQ	2.169	282	0.219 ± 0.001	0.043 ± 0.003
J1349+5341	Y	···	FSRQ	0.979	184	0.765 ± 0.007	$0.128^{+0.008}_{-0.007}$
J1349–1110	Y	···	FSRQ	0.141	115	0.222 ± 0.003	$0.124^{+0.011}_{-0.010}$
J1349–1132	Y	···	FSRQ	0.341	116	0.254 ± 0.009	$0.363^{+0.030}_{-0.026}$
J1350+0940	Y	···	AGN	0.132	206	0.130 ± 0.001	$0.045^{+0.005}_{-0.004}$
J1350+3034	Y	J1351.0+3035	FSRQ	0.714	216	0.252 ± 0.005	$0.257^{+0.015}_{-0.014}$
J1350+6428	Y	···	···	···	138	0.235 ± 0.001	0.033 ± 0.004
J1350–1634	Y	···	FSRQ	0.086	114	0.169 ± 0.002	$0.101^{+0.010}_{-0.009}$
J1351+0830	Y	···	FSRQ	1.429	206	0.292 ± 0.001	$0.045^{+0.004}_{-0.003}$
J1351+5542	Y	···	FSRQ	0.389	185	0.081 ± 0.001	$0.169^{+0.014}_{-0.013}$
J1353+1435	···	J1353.3+1434	BLL	···	100	0.224 ± 0.001	$0.047^{+0.006}_{-0.005}$
J1353+7532	Y	···	FSRQ	1.619	116	0.378 ± 0.005	$0.148^{+0.011}_{-0.010}$
J1357+7643	Y	J1358.1+7646	FSRQ	1.585	158	0.644 ± 0.022	$0.427^{+0.030}_{-0.027}$
J1357–1527	Y	···	FSRQ	1.89	121	0.561 ± 0.007	$0.132^{+0.010}_{-0.009}$
J1359+4011	Y	···	FSRQ	0.407	178	0.244 ± 0.002	$0.107^{+0.007}_{-0.006}$
J1359+5544	···	J1359.1+5539	FSRQ	1.014	86	0.147 ± 0.004	$0.191^{+0.020}_{-0.018}$
J1400–1858	Y	···	FSRQ	0.114	113	0.208 ± 0.007	$0.324^{+0.027}_{-0.024}$
J1401+5835	Y	···	FSRQ	1.924	213	0.103 ± 0.001	$0.113^{+0.008}_{-0.007}$
J1405+0415	Y	···	FSRQ	3.215	208	0.653 ± 0.003	0.052 ± 0.003

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
J1405–1440	Y	···	FSRQ	1.096	118	0.156 ± 0.002	$0.087^{+0.008}_{-0.007}$
J1406+3433	Y	···	FSRQ	2.56	132	0.181 ± 0.001	$0.032^{+0.005}_{-0.004}$
J1407+2827	Y	···	AGN	0.077	211	0.576 ± 0.003	$0.059^{+0.004}_{-0.003}$
J1408+6854	Y	···	FSRQ	1.272	141	0.058 ± 0.001	$0.188^{+0.016}_{-0.015}$
J1410+0731	Y	···	FSRQ	0.901	206	0.198 ± 0.004	$0.314^{+0.019}_{-0.017}$
J1412+1334	Y	···	···	···	213	0.217 ± 0.001	0.084 ± 0.005
J1415+0832	Y	···	FSRQ	0.327	200	0.235 ± 0.003	$0.177^{+0.011}_{-0.010}$
J1415+1320	Y	···	BLL	0.247	185	0.822 ± 0.013	$0.201^{+0.012}_{-0.011}$
J1416–1705	Y	···	···	···	113	0.186 ± 0.003	$0.132^{+0.012}_{-0.010}$
J1419+2706	Y	···	FSRQ	0.536	120	0.454 ± 0.007	$0.161^{+0.012}_{-0.011}$
J1419+3821	Y	···	FSRQ	1.831	182	0.567 ± 0.018	$0.418^{+0.027}_{-0.024}$
J1419+5423	Y	···	BLL	0.152	210	0.900 ± 0.014	$0.220^{+0.012}_{-0.011}$
J1420+1703	Y	···	FSRQ	1.854	169	0.319 ± 0.005	$0.217^{+0.014}_{-0.012}$
J1421+4645	Y	···	FSRQ	1.668	220	0.200 ± 0.001	$0.090^{+0.006}_{-0.005}$
J1421–1931	Y	···	···	···	114	0.238 ± 0.005	$0.193^{+0.016}_{-0.014}$
J1422+3223	Y	···	FSRQ	0.685	128	0.301 ± 0.004	$0.154^{+0.011}_{-0.010}$
J1423+4802	Y	···	FSRQ	2.22	213	0.241 ± 0.003	$0.166^{+0.010}_{-0.009}$
J1423+5055	Y	···	FSRQ	0.276	201	0.223 ± 0.001	$0.062^{+0.005}_{-0.004}$
J1424+2256	Y	···	FSRQ	3.62	123	0.337 ± 0.001	$0.035^{+0.004}_{-0.003}$
J1425+1424	Y	···	FSRQ	0.78	181	0.545 ± 0.004	0.104 ± 0.006
J1426+3625	Y	···	FSRQ	1.091	181	0.262 ± 0.004	$0.220^{+0.013}_{-0.012}$
J1428+2724	Y	···	AGN	0.014	126	0.037 ± 0.001	$0.385^{+0.038}_{-0.033}$
J1430+1043	Y	···	FSRQ	1.71	303	0.620 ± 0.006	$0.176^{+0.008}_{-0.007}$
J1430+3649	Y	···	BLL	0.566	179	0.255 ± 0.005	$0.243^{+0.015}_{-0.014}$
J1431+3952	Y	···	FSRQ	1.215	189	0.606 ± 0.005	$0.101^{+0.006}_{-0.005}$
J1433–1548	Y	···	FSRQ	1.573	117	0.609 ± 0.004	$0.058^{+0.007}_{-0.006}$
J1434+1952	Y	···	FSRQ	1.382	126	0.281 ± 0.002	$0.070^{+0.006}_{-0.005}$
J1434+4203	Y	J1433.9+4204	FSRQ	1.24	179	0.165 ± 0.003	$0.238^{+0.015}_{-0.014}$
J1435+3012	Y	···	FSRQ	1.568	130	0.125 ± 0.002	$0.154^{+0.012}_{-0.011}$
J1436+2321	Y	J1436.9+2314	FSRQ	1.545	122	0.686 ± 0.012	$0.190^{+0.014}_{-0.012}$
J1436+6336	Y	···	FSRQ	2.066	230	1.500 ± 0.003	0.032 ± 0.002
J1436–1846	Y	···	···	···	114	0.045 ± 0.001	$0.155^{+0.021}_{-0.019}$
J1437+0405	Y	···	FSRQ	2.025	201	0.147 ± 0.001	$0.048^{+0.005}_{-0.004}$

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
J1437+3119	Y	···	FSRQ	1.357	135	0.143 ± 0.001	0.040 ± 0.007
J1438+3710	Y	···	FSRQ	2.401	136	0.410 ± 0.003	$0.079^{+0.006}_{-0.005}$
J1439+2114	Y	···	···	···	122	0.188 ± 0.002	$0.105^{+0.008}_{-0.007}$
J1439+4958	Y	···	GALAXY	0.174	197	0.334 ± 0.012	$0.506^{+0.034}_{-0.030}$
J1439–1531	Y	···	BLL	···	118	0.322 ± 0.003	$0.083^{+0.008}_{-0.007}$
J1442+3234	Y	···	FSRQ	2.12	130	0.241 ± 0.002	0.062 ± 0.005
J1443+2501	···	J1443.8+2457	FSRQ	0.939	107	0.232 ± 0.008	$0.363^{+0.031}_{-0.027}$
J1445–1614	Y	···	FSRQ	1.195	109	0.169 ± 0.002	$0.091^{+0.008}_{-0.007}$
J1445–1629	Y	···	···	···	115	0.545 ± 0.010	$0.186^{+0.014}_{-0.013}$
J1446+1721	Y	···	FSRQ	1.026	301	0.733 ± 0.007	0.169 ± 0.007
J1448+7601	Y	···	FSRQ	0.899	116	0.687 ± 0.008	$0.125^{+0.010}_{-0.009}$
J1450+0910	Y	···	FSRQ	2.612	292	0.117 ± 0.005	$0.656^{+0.040}_{-0.036}$
J1451–0127	Y	···	FSRQ	1.314	196	0.239 ± 0.002	$0.086^{+0.006}_{-0.005}$
J1453+1036	Y	···	FSRQ	2.27	304	0.118 ± 0.001	0.103 ± 0.006
J1453+2648	Y	···	BLL	···	128	0.475 ± 0.012	$0.288^{+0.021}_{-0.019}$
J1453+3505	Y	···	FSRQ	0.721	134	0.166 ± 0.003	$0.200^{+0.014}_{-0.013}$
J1456+5048	Y	···	···	···	204	0.073 ± 0.001	$0.118^{+0.010}_{-0.009}$
J1457+0749	Y	···	···	···	207	0.072 ± 0.001	$0.131^{+0.011}_{-0.010}$
J1458+3720	Y	···	BLL	0.333	128	0.149 ± 0.002	$0.127^{+0.010}_{-0.009}$
J1459+4442	Y	···	FSRQ	3.402	138	0.176 ± 0.002	$0.136^{+0.011}_{-0.010}$
J1459–1810	Y	···	FSRQ	0.235	110	0.130 ± 0.001	$0.062^{+0.009}_{-0.008}$
J1500+4751	Y	···	BLL	···	118	0.408 ± 0.004	$0.107^{+0.008}_{-0.007}$
J1502–1508	Y	···	···	···	110	0.129 ± 0.002	$0.133^{+0.015}_{-0.013}$
J1504+0813	Y	···	FSRQ	2.832	212	0.147 ± 0.001	0.039 ± 0.005
J1504+1029	Y	J1504.4+1029	FSRQ	1.839	305	1.910 ± 0.036	$0.325^{+0.015}_{-0.014}$
J1505+0326	Y	J1505.0+0328	AGN	0.409	225	0.532 ± 0.004	0.097 ± 0.005
J1506+3730	Y	J1505.8+3725	FSRQ	0.674	135	0.652 ± 0.007	$0.129^{+0.009}_{-0.008}$
J1506+4239	Y	···	FSRQ	0.587	129	0.672 ± 0.003	0.045 ± 0.004
J1506+4933	Y	···	FSRQ	1.395	121	0.377 ± 0.007	$0.188^{+0.014}_{-0.013}$
J1506+8319	Y	···	FSRQ	2.577	194	0.249 ± 0.001	0.044 ± 0.003
J1507–1652	Y	···	FSRQ	0.876	192	0.928 ± 0.007	0.093 ± 0.005
J1508–1548	Y	···	FSRQ	2.499	186	0.106 ± 0.001	$0.136^{+0.009}_{-0.008}$
J1510–1121	Y	···	···	···	194	0.133 ± 0.001	0.064 ± 0.007

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
J1511+0518	Y	···	AGN	0.084	216	0.599 ± 0.002	$0.044^{+0.003}_{-0.002}$
J1513–1012	Y	···	FSRQ	1.513	199	1.134 ± 0.009	$0.107^{+0.006}_{-0.005}$
J1516+0015	Y	···	NLRG	0.052	216	1.050 ± 0.003	$0.043^{+0.003}_{-0.002}$
J1516+1932	Y	J1516.9+1928	BLL	···	336	1.149 ± 0.018	0.284 ± 0.012
J1517+1332	Y	···	FSRQ	1.499	309	0.104 ± 0.001	$0.174^{+0.011}_{-0.010}$
J1521+4336	Y	···	FSRQ	2.18	130	0.453 ± 0.004	$0.082^{+0.006}_{-0.005}$
J1522+3144	Y	J1522.1+3143	FSRQ	1.487	176	0.428 ± 0.003	$0.091^{+0.006}_{-0.005}$
J1526+6650	Y	···	FSRQ	3.02	230	0.238 ± 0.002	$0.095^{+0.006}_{-0.005}$
J1526–0425	Y	···	FSRQ	1.492	222	0.083 ± 0.002	$0.285^{+0.019}_{-0.017}$
J1532–1319	Y	···	···	···	124	0.337 ± 0.011	$0.362^{+0.028}_{-0.025}$
J1533–0421	Y	···	FSRQ	0.84	222	0.213 ± 0.003	$0.178^{+0.011}_{-0.010}$
J1534+0131	Y	···	FSRQ	1.42	226	0.693 ± 0.007	0.140 ± 0.007
J1535+4957	Y	···	FSRQ	1.119	119	0.222 ± 0.002	$0.105^{+0.008}_{-0.007}$
J1537–1527	Y	···	FSRQ	1.766	194	0.176 ± 0.001	0.067 ± 0.006
J1538+1444	Y	···	···	···	171	0.104 ± 0.001	0.037 ± 0.006
J1538–1655	Y	···	···	···	183	0.088 ± 0.001	$0.080^{+0.009}_{-0.008}$
J1539+2744	Y	J1539.7+2747	FSRQ	2.19	277	0.204 ± 0.004	$0.306^{+0.016}_{-0.014}$
J1539+3104	Y	···	FSRQ	1.211	168	0.127 ± 0.003	$0.259^{+0.017}_{-0.016}$
J1540+1447	Y	···	FSRQ	0.605	179	0.984 ± 0.009	0.123 ± 0.007
J1544+3240	Y	···	FSRQ	1.05	176	0.128 ± 0.002	$0.166^{+0.012}_{-0.011}$
J1545+5135	Y	···	FSRQ	1.93	210	0.232 ± 0.003	$0.200^{+0.011}_{-0.010}$
J1548+1727	Y	···	FSRQ	1.874	195	0.096 ± 0.002	$0.233^{+0.015}_{-0.014}$
J1549+0237	Y	J1549.3+0235	FSRQ	0.414	226	1.924 ± 0.022	$0.168^{+0.009}_{-0.008}$
J1549+5038	Y	···	FSRQ	2.175	216	0.833 ± 0.006	0.095 ± 0.005
J1550+0527	Y	J1550.7+0527	FSRQ	1.422	178	2.816 ± 0.006	0.025 ± 0.002
J1551+5806	Y	···	FSRQ	1.32	207	0.316 ± 0.001	0.054 ± 0.003
J1555+1111	Y	J1555.7+1111	BLL	···	185	0.171 ± 0.001	0.085 ± 0.006
J1557–0001	Y	···	FSRQ	1.772	226	0.496 ± 0.005	0.154 ± 0.008
J1558+5625	···	J1558.9+5627	BLL	0.3	98	0.227 ± 0.002	$0.053^{+0.007}_{-0.006}$
J1558–1409	Y	···	AGN	0.097	192	0.206 ± 0.001	0.034 ± 0.004
J1559+0304	Y	···	FSRQ	3.891	227	0.294 ± 0.001	0.057 ± 0.004
J1602+2646	Y	···	FSRQ	0.373	259	0.268 ± 0.001	0.052 ± 0.003
J1602+3326	Y	···	NLRG	1.1	189	0.852 ± 0.003	$0.042^{+0.003}_{-0.002}$

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
J1603+1105	Y	···	BLL	0.143	193	0.391 ± 0.008	$0.291^{+0.017}_{-0.016}$
J1603+1554	Y	···	FSRQ	0.109	196	0.228 ± 0.001	$0.048^{+0.004}_{-0.003}$
J1603+5730	Y	···	FSRQ	2.858	209	0.221 ± 0.001	0.085 ± 0.005
J1603–1007	Y	···	FSRQ	0.959	144	0.144 ± 0.003	$0.201^{+0.016}_{-0.014}$
J1604+5714	Y	J1604.3+5710	FSRQ	0.72	213	0.468 ± 0.005	$0.165^{+0.009}_{-0.008}$
J1605+3001	Y	···	FSRQ	2.404	185	0.317 ± 0.003	0.137 ± 0.008
J1605–1139	Y	···	···	···	145	0.447 ± 0.006	$0.146^{+0.010}_{-0.009}$
J1606+2717	Y	···	FSRQ	0.934	270	0.423 ± 0.002	$0.064^{+0.004}_{-0.003}$
J1606+3124	Y	···	FSRQ	4.56	186	0.427 ± 0.001	0.017 ± 0.002
J1608+1029	Y	J1609.0+1031	FSRQ	1.226	193	0.927 ± 0.018	$0.273^{+0.016}_{-0.014}$
J1610+2414	Y	···	FSRQ	1.449	180	0.255 ± 0.002	0.119 ± 0.007
J1611+1856	Y	···	FSRQ	1.776	185	0.117 ± 0.002	$0.228^{+0.014}_{-0.013}$
J1613+3412	Y	J1613.5+3411	FSRQ	1.397	183	2.584 ± 0.021	0.106 ± 0.006
J1613+4223	Y	···	···	···	170	0.032 ± 0.001	$0.133^{+0.020}_{-0.019}$
J1616+0459	Y	···	FSRQ	3.199	236	0.574 ± 0.003	0.073 ± 0.004
J1616+4632	···	J1616.1+4637	FSRQ	0.95	111	0.102 ± 0.004	$0.423^{+0.038}_{-0.033}$
J1617+0246	Y	···	FSRQ	1.341	235	0.365 ± 0.002	$0.088^{+0.005}_{-0.004}$
J1617–1122	Y	···	···	···	142	0.211 ± 0.004	$0.223^{+0.017}_{-0.016}$
J1617–1941	Y	···	···	···	198	0.243 ± 0.002	0.073 ± 0.005
J1618+0819	Y	···	FSRQ	0.445	225	0.216 ± 0.001	0.044 ± 0.003
J1619+2247	Y	···	FSRQ	1.987	189	0.613 ± 0.002	$0.033^{+0.003}_{-0.002}$
J1619–1817	Y	···	···	···	206	0.285 ± 0.002	0.074 ± 0.005
J1620+4901	Y	···	FSRQ	1.513	156	0.366 ± 0.005	$0.177^{+0.012}_{-0.011}$
J1623+0741	Y	···	FSRQ	1.301	237	0.250 ± 0.001	0.055 ± 0.004
J1623+3909	Y	···	FSRQ	1.975	173	0.058 ± 0.001	$0.087^{+0.013}_{-0.012}$
J1623+6624	Y	···	FSRQ	0.203	232	0.189 ± 0.001	$0.075^{+0.005}_{-0.004}$
J1624+2748	Y	···	···	···	189	0.173 ± 0.001	0.042 ± 0.004
J1624+5652	Y	···	BLL	···	161	0.259 ± 0.003	$0.132^{+0.009}_{-0.008}$
J1624+5741	Y	···	FSRQ	0.789	139	0.438 ± 0.006	$0.150^{+0.011}_{-0.010}$
J1625+4134	Y	···	FSRQ	2.55	177	0.608 ± 0.008	$0.173^{+0.011}_{-0.010}$
J1628–1415	Y	···	FSRQ	1.026	209	0.231 ± 0.003	$0.153^{+0.010}_{-0.009}$
J1630+0701	Y	···	FSRQ	0.736	232	0.217 ± 0.003	$0.210^{+0.012}_{-0.011}$
J1630+5221	···	J1630.2+5220	BLL	···	95	no det.	···

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
J1631+4927	Y	···	FSRQ	0.52	158	0.398 ± 0.004	0.125 ± 0.008
J1632+8232	Y	J1635.4+8228	AGN	0.025	192	0.881 ± 0.003	0.043 ± 0.003
J1635+3808	Y	J1635.0+3808	FSRQ	1.814	141	3.391 ± 0.031	$0.109^{+0.007}_{-0.006}$
J1636+2112	Y	···	FSRQ	1.802	180	0.247 ± 0.002	$0.084^{+0.006}_{-0.005}$
J1637+4717	Y	···	FSRQ	0.74	157	1.048 ± 0.010	$0.118^{+0.008}_{-0.007}$
J1638+5720	Y	···	FSRQ	0.751	150	1.702 ± 0.025	$0.178^{+0.012}_{-0.010}$
J1639+1632	Y	···	···	···	193	0.087 ± 0.001	$0.190^{+0.013}_{-0.012}$
J1639+5357	Y	···	FSRQ	1.977	161	0.217 ± 0.001	0.061 ± 0.005
J1640+1144	···	J1641.0+1143	AGN	0.078	126	0.180 ± 0.002	$0.081^{+0.009}_{-0.008}$
J1640+3946	Y	···	FSRQ	1.666	146	1.149 ± 0.026	$0.272^{+0.018}_{-0.017}$
J1640–0011	Y	···	FSRQ	0.651	237	0.205 ± 0.003	0.186 ± 0.010
J1641+2257	Y	···	FSRQ	2.063	183	0.431 ± 0.004	0.110 ± 0.006
J1642+3948	Y	···	FSRQ	0.593	143	7.928 ± 0.065	0.096 ± 0.006
J1642+6856	Y	···	FSRQ	0.751	141	3.128 ± 0.097	$0.365^{+0.026}_{-0.023}$
J1642–0621	Y	···	BLL	1.514	149	1.074 ± 0.030	$0.341^{+0.024}_{-0.021}$
J1644–0743	Y	···	NLRG	0.139	146	0.178 ± 0.003	$0.146^{+0.012}_{-0.011}$
J1644–1804	Y	···	···	···	211	0.263 ± 0.002	$0.089^{+0.006}_{-0.005}$
J1646+4059	Y	···	FSRQ	0.835	152	0.439 ± 0.003	$0.063^{+0.005}_{-0.004}$
J1647+4950	Y	J1647.4+4948	AGN	0.047	161	0.242 ± 0.002	0.079 ± 0.006
J1648+2224	Y	···	···	···	170	0.464 ± 0.003	$0.086^{+0.006}_{-0.005}$
J1649+0412	Y	···	···	···	234	0.304 ± 0.005	$0.247^{+0.013}_{-0.012}$
J1650+0824	Y	···	FSRQ	1.965	241	0.284 ± 0.003	$0.144^{+0.008}_{-0.007}$
J1651+0129	Y	···	···	···	245	0.582 ± 0.004	$0.091^{+0.005}_{-0.004}$
J1652+0618	Y	···	···	···	244	0.226 ± 0.003	0.197 ± 0.010
J1652+3902	Y	···	FSRQ	1.299	152	0.498 ± 0.009	$0.208^{+0.014}_{-0.012}$
J1653+3107	Y	···	FSRQ	1.298	142	0.502 ± 0.003	0.062 ± 0.004
J1653+3945	Y	J1653.9+3945	BLL	0.034	197	1.176 ± 0.003	0.032 ± 0.002
J1653–1551	Y	···	···	···	221	0.149 ± 0.002	$0.188^{+0.012}_{-0.011}$
J1655+4233	Y	···	···	···	149	0.192 ± 0.002	$0.121^{+0.009}_{-0.008}$
J1656+1826	Y	···	FSRQ	2.551	223	0.213 ± 0.001	0.059 ± 0.004
J1656+6012	···	J1656.9+6017	FSRQ	0.623	97	0.303 ± 0.006	$0.186^{+0.016}_{-0.014}$
J1657+4808	Y	···	BLL	···	197	0.572 ± 0.009	$0.206^{+0.012}_{-0.011}$
J1658+0515	Y	···	FSRQ	0.879	243	0.869 ± 0.003	0.043 ± 0.002

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
J1658+3443	Y	···	FSRQ	1.939	141	0.319 ± 0.001	$0.044^{+0.004}_{-0.003}$
J1658+4737	Y	···	FSRQ	1.622	241	0.836 ± 0.002	0.034 ± 0.002
J1658–0739	Y	···	FSRQ	3.742	152	0.453 ± 0.002	$0.041^{+0.004}_{-0.003}$
J1700+6830	Y	J1700.1+6830	FSRQ	0.301	146	0.409 ± 0.012	$0.342^{+0.024}_{-0.022}$
J1701+3954	Y	···	BLL	0.507	153	0.464 ± 0.007	$0.173^{+0.011}_{-0.010}$
J1701–1903	Y	···	···	···	217	0.201 ± 0.002	0.087 ± 0.006
J1702+1502	Y	···	FSRQ	1.041	217	0.107 ± 0.001	$0.126^{+0.008}_{-0.007}$
J1707+0148	Y	···	FSRQ	2.576	246	0.772 ± 0.009	$0.179^{+0.009}_{-0.008}$
J1707+1122	Y	···	FSRQ	2.406	209	0.121 ± 0.001	0.063 ± 0.005
J1707+1331	Y	···	FSRQ	0.936	207	0.420 ± 0.004	$0.141^{+0.008}_{-0.007}$
J1707–1415	Y	···	···	···	215	0.454 ± 0.005	0.147 ± 0.008
J1709+4318	Y	J1709.6+4320	FSRQ	1.027	151	0.283 ± 0.007	$0.299^{+0.021}_{-0.018}$
J1709–1728	Y	···	···	···	134	0.490 ± 0.009	$0.206^{+0.014}_{-0.013}$
J1712+6053	Y	···	FSRQ	1.684	120	0.082 ± 0.001	0.040 ± 0.009
J1712–1820	Y	···	···	···	142	0.207 ± 0.004	$0.215^{+0.017}_{-0.015}$
J1713+4916	Y	···	FSRQ	1.552	198	0.267 ± 0.003	$0.126^{+0.008}_{-0.007}$
J1716+2152	Y	···	FSRQ	0.359	294	0.425 ± 0.001	0.036 ± 0.002
J1716+6836	Y	···	FSRQ	0.777	146	$0.638^{+0.009}_{-0.008}$	$0.156^{+0.010}_{-0.009}$
J1716–0452	Y	···	FSRQ	1.026	238	0.430 ± 0.002	0.050 ± 0.003
J1719+0817	Y	···	FSRQ	1.185	235	0.314 ± 0.001	0.017 ± 0.002
J1719+1745	Y	J1719.2+1745	BLL	0.137	296	0.640 ± 0.004	0.106 ± 0.005
J1719+4804	Y	···	FSRQ	1.083	152	0.036 ± 0.001	$0.205^{+0.018}_{-0.017}$
J1719+4858	Y	···	AGN	0.025	191	0.159 ± 0.002	$0.140^{+0.009}_{-0.008}$
J1719–1420	Y	···	FSRQ	0.64	204	0.544 ± 0.002	0.041 ± 0.003
J1721+3542	Y	···	FSRQ	0.283	158	0.298 ± 0.002	0.060 ± 0.005
J1722+1013	···	J1722.5+1012	FSRQ	0.732	119	0.533 ± 0.013	$0.261^{+0.020}_{-0.018}$
J1722+2815	Y	···	FSRQ	0.945	225	0.421 ± 0.007	$0.243^{+0.013}_{-0.012}$
J1722+5856	Y	···	FSRQ	1.979	200	0.171 ± 0.002	$0.137^{+0.009}_{-0.008}$
J1722+6105	Y	···	FSRQ	2.058	192	0.138 ± 0.006	$0.546^{+0.038}_{-0.034}$
J1724+1648	Y	···	···	···	299	0.110 ± 0.002	$0.264^{+0.014}_{-0.013}$
J1724+3303	Y	···	FSRQ	1.87	227	0.234 ± 0.001	$0.069^{+0.005}_{-0.004}$
J1724–1443	Y	···	FSRQ	0.899	195	0.500 ± 0.005	$0.138^{+0.008}_{-0.007}$
J1725+1152	Y	J1725.0+1151	BLL	···	216	0.062 ± 0.001	0.101 ± 0.010

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
J1725+3026	Y	···	FSRQ	0.978	231	0.180 ± 0.001	0.102 ± 0.006
J1726+2717	Y	···	FSRQ	0.535	315	0.169 ± 0.003	$0.247^{+0.012}_{-0.011}$
J1726+3213	Y	···	FSRQ	1.094	229	0.251 ± 0.001	0.058 ± 0.004
J1727+4530	Y	J1727.3+4525	FSRQ	0.714	152	1.416 ± 0.023	$0.199^{+0.013}_{-0.012}$
J1727+5510	Y	···	FSRQ	0.247	200	0.163 ± 0.001	0.091 ± 0.006
J1728+0427	Y	J1728.2+0431	FSRQ	0.293	243	0.838 ± 0.014	$0.266^{+0.014}_{-0.013}$
J1728+1215	Y	···	FSRQ	0.583	220	0.402 ± 0.005	$0.156^{+0.009}_{-0.008}$
J1728+5013	Y	···	BLL	0.055	201	0.137 ± 0.001	0.047 ± 0.004
J1733–0456	Y	···	···	···	154	0.241 ± 0.004	$0.174^{+0.012}_{-0.011}$
J1733–1304	Y	J1733.0–1308	FSRQ	0.902	242	4.452 ± 0.032	0.110 ± 0.005
J1734+3857	Y	J1734.4+3859	FSRQ	0.976	154	0.950 ± 0.012	0.147 ± 0.009
J1735+3616	Y	···	FSRQ	0.893	240	0.661 ± 0.005	$0.112^{+0.006}_{-0.005}$
J1735+5049	Y	···	···	···	197	0.745 ± 0.003	0.055 ± 0.003
J1735–1117	···	J1735.4–1118	···	···	120	0.059 ± 0.001	$0.174^{+0.018}_{-0.016}$
J1736+0631	Y	···	FSRQ	2.388	229	0.244 ± 0.005	$0.282^{+0.016}_{-0.014}$
J1738+3224	Y	···	FSRQ	0.126	223	0.080 ± 0.001	0.149 ± 0.011
J1738+4008	Y	···	FSRQ	3.591	234	0.174 ± 0.002	$0.162^{+0.009}_{-0.008}$
J1739+3358	Y	···	FSRQ	1.626	233	0.154 ± 0.001	0.098 ± 0.006
J1739+4737	Y	···	BLL	···	151	0.681 ± 0.004	0.074 ± 0.005
J1739+4955	Y	···	FSRQ	1.545	202	0.533 ± 0.008	$0.200^{+0.011}_{-0.010}$
J1740+2211	Y	···	FSRQ	1.406	301	0.555 ± 0.003	0.093 ± 0.004
J1740+4348	Y	···	FSRQ	2.246	155	0.164 ± 0.001	$0.085^{+0.007}_{-0.006}$
J1740+4506	Y	···	FSRQ	2.788	156	0.157 ± 0.001	0.038 ± 0.004
J1740+5211	Y	···	FSRQ	1.379	196	1.115 ± 0.017	$0.213^{+0.012}_{-0.011}$
J1740–0811	Y	···	···	···	156	0.389 ± 0.007	$0.215^{+0.015}_{-0.013}$
J1741+4751	Y	···	···	···	156	0.152 ± 0.001	0.040 ± 0.004
J1742+5945	···	J1742.1+5947	BLL	···	96	0.153 ± 0.001	$0.061^{+0.008}_{-0.007}$
J1743+3747	Y	···	FSRQ	1.958	236	0.132 ± 0.002	$0.197^{+0.011}_{-0.010}$
J1743–0350	Y	···	FSRQ	1.057	144	3.525 ± 0.042	$0.142^{+0.009}_{-0.008}$
J1745+1720	Y	···	FSRQ	1.702	301	0.486 ± 0.002	0.080 ± 0.004
J1745+2252	Y	···	FSRQ	1.884	299	0.266 ± 0.001	0.082 ± 0.004
J1745–0753	Y	J1745.6–0751	BLL	···	157	0.740 ± 0.013	$0.224^{+0.014}_{-0.013}$
J1747+4658	Y	···	BLL	1.484	153	0.310 ± 0.003	$0.094^{+0.007}_{-0.006}$

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
J1748+3404	Y	···	FSRQ	2.764	229	0.331 ± 0.004	0.163 ± 0.009
J1748+7005	Y	J1748.5+7004	BLL	0.77	226	0.558 ± 0.007	$0.190^{+0.010}_{-0.009}$
J1749+4321	Y	J1749.0+4323	BLL	···	231	0.413 ± 0.002	0.076 ± 0.004
J1751+0939	Y	J1751.5+0937	BLL	0.322	277	4.161 ± 0.091	$0.362^{+0.018}_{-0.017}$
J1752+1734	Y	···	FSRQ	0.507	305	0.698 ± 0.006	$0.144^{+0.007}_{-0.006}$
J1753+2848	Y	···	FSRQ	1.118	233	1.254 ± 0.028	$0.332^{+0.018}_{-0.017}$
J1753+4409	Y	···	FSRQ	0.871	227	0.549 ± 0.004	0.093 ± 0.005
J1754+6452	Y	···	FSRQ	0.977	218	0.137 ± 0.002	$0.196^{+0.011}_{-0.010}$
J1756+1535	Y	···	···	···	294	0.469 ± 0.006	$0.211^{+0.010}_{-0.009}$
J1756+1553	Y	···	FSRQ	0.547	303	0.323 ± 0.002	0.083 ± 0.004
J1756+3046	Y	···	FSRQ	1.983	233	0.064 ± 0.001	$0.247^{+0.015}_{-0.013}$
J1759+2343	Y	···	FSRQ	1.721	310	0.447 ± 0.003	0.093 ± 0.004
J1800+3848	Y	···	FSRQ	2.092	235	1.075 ± 0.004	0.055 ± 0.003
J1800+7828	Y	J1800.4+7827	BLL	0.68	138	2.617 ± 0.028	$0.121^{+0.008}_{-0.007}$
J1801+4404	Y	···	FSRQ	0.663	223	1.141 ± 0.021	$0.266^{+0.014}_{-0.013}$
J1803+0934	Y	···	FSRQ	0.683	277	0.131 ± 0.001	$0.163^{+0.009}_{-0.008}$
J1804+0042	Y	···	GALAXY	0.07	149	0.091 ± 0.001	$0.103^{+0.010}_{-0.009}$
J1804+0101	Y	···	FSRQ	1.522	143	0.520 ± 0.005	0.097 ± 0.006
J1805+1714	Y	···	···	···	311	0.134 ± 0.001	0.114 ± 0.006
J1806+6949	Y	J1807.0+6945	BLL	0.051	221	1.588 ± 0.008	0.069 ± 0.004
J1807+2204	Y	···	FSRQ	0.798	297	0.167 ± 0.001	0.065 ± 0.004
J1808+4542	Y	···	FSRQ	0.83	220	0.689 ± 0.004	0.078 ± 0.004
J1809+1849	Y	···	FSRQ	0.928	309	0.159 ± 0.003	$0.267^{+0.013}_{-0.012}$
J1809+2910	···	J1809.6+2908	BLL	···	123	0.082 ± 0.001	$0.136^{+0.013}_{-0.012}$
J1811+1704	Y	···	BLL	···	313	0.466 ± 0.009	$0.345^{+0.016}_{-0.015}$
J1813+0615	Y	···	BLL	···	148	0.330 ± 0.007	$0.260^{+0.018}_{-0.016}$
J1813+2952	Y	···	FSRQ	1.351	233	0.163 ± 0.002	0.187 ± 0.010
J1813+3144	···	J1813.4+3141	BLL	0.117	114	0.093 ± 0.001	$0.073^{+0.010}_{-0.009}$
J1815+1623	Y	···	FSRQ	0.742	323	0.355 ± 0.007	$0.325^{+0.015}_{-0.014}$
J1816+5307	Y	···	···	···	210	0.100 ± 0.001	0.063 ± 0.006
J1818+0903	···	J1818.1+0905	FSRQ	0.354	118	0.118 ± 0.004	$0.310^{+0.027}_{-0.024}$
J1818+5017	Y	···	FSRQ	1.395	214	0.152 ± 0.001	0.069 ± 0.005
J1823+7938	Y	···	FSRQ	0.224	143	0.358 ± 0.001	0.020 ± 0.003

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
J1824+1044	Y	···	FSRQ	1.364	314	0.611 ± 0.002	0.061 ± 0.003
J1824+5651	Y	J1824.0+5651	BLL	0.664	207	1.541 ± 0.009	$0.084^{+0.005}_{-0.004}$
J1827+2658	Y	···	···	···	244	0.153 ± 0.002	$0.141^{+0.009}_{-0.008}$
J1832+1357	Y	···	FSRQ	2.83	318	0.357 ± 0.005	$0.253^{+0.012}_{-0.011}$
J1835+2506	Y	···	FSRQ	1.973	255	no det.	···
J1835+3241	Y	···	AGN	0.058	228	0.810 ± 0.005	0.090 ± 0.005
J1835+6119	Y	···	FSRQ	2.274	207	0.175 ± 0.001	0.052 ± 0.004
J1840+2457	Y	···	FSRQ	1.635	258	0.112 ± 0.001	$0.137^{+0.009}_{-0.008}$
J1840+3900	Y	···	FSRQ	3.095	226	0.166 ± 0.002	0.133 ± 0.008
J1842+6809	Y	···	FSRQ	0.475	221	0.868 ± 0.027	$0.456^{+0.028}_{-0.025}$
J1846+3747	Y	···	FSRQ	2.441	226	0.092 ± 0.001	0.093 ± 0.007
J1848+3219	Y	J1848.5+3224	FSRQ	0.798	230	0.529 ± 0.009	$0.241^{+0.013}_{-0.012}$
J1849+3024	Y	···	FSRQ	0.672	235	0.295 ± 0.001	0.059 ± 0.004
J1849+6705	Y	J1849.3+6705	FSRQ	0.657	215	2.407 ± 0.024	$0.142^{+0.008}_{-0.007}$
J1850+2825	Y	···	FSRQ	2.56	236	1.150 ± 0.003	0.028 ± 0.002
J1852+4019	Y	···	FSRQ	2.12	189	0.503 ± 0.003	$0.074^{+0.005}_{-0.004}$
J1852+4855	Y	J1852.5+4853	FSRQ	1.25	209	0.395 ± 0.008	$0.285^{+0.016}_{-0.015}$
J1854+7351	Y	···	FSRQ	0.46	198	0.332 ± 0.003	0.105 ± 0.006
J1855+3742	Y	···	FSRQ	1.12	229	0.114 ± 0.001	0.086 ± 0.007
J1900+2701	Y	···	···	···	237	0.160 ± 0.001	0.111 ± 0.007
J1900+2722	Y	···	···	···	240	0.184 ± 0.002	$0.181^{+0.010}_{-0.009}$
J1912+3740	Y	···	FSRQ	1.104	227	0.361 ± 0.004	0.152 ± 0.008
J1916–1519	Y	···	···	···	146	0.224 ± 0.003	$0.125^{+0.010}_{-0.009}$
J1917–1921	Y	J1917.7–1922	BLL	0.137	149	0.303 ± 0.003	$0.088^{+0.007}_{-0.006}$
J1918+4937	Y	···	FSRQ	0.926	211	0.125 ± 0.001	0.094 ± 0.006
J1918+5520	Y	···	FSRQ	1.734	227	0.516 ± 0.003	$0.079^{+0.005}_{-0.004}$
J1921–1231	···	J1921.1–1234	···	···	124	0.089 ± 0.002	$0.135^{+0.014}_{-0.012}$
J1921–1607	···	J1922.0–1608	BLL	···	116	0.053 ± 0.001	$0.136^{+0.017}_{-0.015}$
J1923+3941	Y	···	···	···	185	0.095 ± 0.001	$0.158^{+0.011}_{-0.010}$
J1927+6117	Y	···	BLL	···	231	0.767 ± 0.009	$0.175^{+0.009}_{-0.008}$
J1927+7358	Y	···	FSRQ	0.303	198	3.718 ± 0.040	0.151 ± 0.008
J1933+6540	Y	···	FSRQ	1.687	231	0.333 ± 0.004	$0.168^{+0.009}_{-0.008}$
J1934+6138	Y	···	FSRQ	1.749	231	0.200 ± 0.001	$0.092^{+0.006}_{-0.005}$

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
J1935+8130	Y	···	···	···	178	0.207 ± 0.001	0.019 ± 0.003
J1936+7131	Y	···	FSRQ	1.864	141	0.409 ± 0.003	0.092 ± 0.006
J1936–0402	Y	···	FSRQ	0.49	173	0.106 ± 0.001	$0.136^{+0.010}_{-0.009}$
J1938–1749	Y	···	FSRQ	0.903	149	0.305 ± 0.002	0.066 ± 0.005
J1939–1002	Y	···	FSRQ	3.787	203	0.338 ± 0.002	0.068 ± 0.004
J1939–1525	Y	···	FSRQ	1.657	165	0.667 ± 0.005	$0.087^{+0.006}_{-0.005}$
J1941+7221	···	J1941.6+7214	···	···	118	0.100 ± 0.001	$0.081^{+0.009}_{-0.008}$
J1941–0211	Y	···	FSRQ	0.202	202	0.070 ± 0.001	$0.184^{+0.013}_{-0.012}$
J1947–0103	Y	···	···	···	202	0.182 ± 0.001	0.047 ± 0.006
J1949+5041	Y	···	FSRQ	1.927	209	0.135 ± 0.001	0.059 ± 0.006
J1949–1957	Y	···	FSRQ	2.652	168	0.609 ± 0.002	0.040 ± 0.003
J1951+0134	Y	···	FSRQ	4.114	166	0.133 ± 0.001	$0.111^{+0.009}_{-0.008}$
J1951–0509	Y	···	FSRQ	1.083	193	0.245 ± 0.004	$0.212^{+0.013}_{-0.012}$
J1954+6153	Y	···	···	···	229	no det.	···
J1954–1123	Y	J1954.8–1124	FSRQ	0.683	149	0.533 ± 0.022	$0.507^{+0.039}_{-0.035}$
J1955+0618	Y	···	···	···	65	no det.	···
J1955+5131	Y	···	FSRQ	1.223	220	1.403 ± 0.010	0.098 ± 0.005
J1959+6508	Y	J2000.0+6508	BLL	0.049	336	0.206 ± 0.001	$0.086^{+0.005}_{-0.004}$
J2000–1325	Y	···	FSRQ	0.222	162	$0.714^{+0.002}_{-0.003}$	0.034 ± 0.004
J2000–1748	Y	J2000.9–1749	FSRQ	0.652	183	1.797 ± 0.052	$0.390^{+0.025}_{-0.022}$
J2003–0421	Y	···	···	···	168	0.112 ± 0.001	$0.039^{+0.006}_{-0.005}$
J2004+7355	Y	···	···	···	141	0.206 ± 0.002	$0.079^{+0.007}_{-0.006}$
J2005+7752	Y	J2006.0+7751	BLL	0.342	236	1.101 ± 0.016	0.223 ± 0.011
J2006+6424	Y	···	FSRQ	1.574	233	1.018 ± 0.006	$0.088^{+0.005}_{-0.004}$
J2007+0636	Y	···	FSRQ	2.864	163	0.299 ± 0.004	$0.186^{+0.012}_{-0.011}$
J2007+6607	Y	···	FSRQ	1.325	225	0.570 ± 0.007	$0.172^{+0.009}_{-0.008}$
J2009+0727	Y	···	FSRQ	0.763	166	0.091 ± 0.001	$0.143^{+0.012}_{-0.011}$
J2009+7229	Y	J2009.1+7228	BLL	···	143	0.805 ± 0.012	$0.179^{+0.012}_{-0.011}$
J2010+6116	Y	···	FSRQ	0.87	227	0.120 ± 0.002	$0.189^{+0.011}_{-0.010}$
J2011–1546	Y	···	FSRQ	1.18	205	2.093 ± 0.009	0.059 ± 0.003
J2015+6554	Y	···	FSRQ	2.845	232	0.547 ± 0.006	$0.170^{+0.009}_{-0.008}$
J2015–0137	···	J2015.3–0129	BLL	···	127	0.372 ± 0.002	$0.044^{+0.005}_{-0.004}$
J2015–1252	Y	···	FSRQ	0.614	189	0.168 ± 0.002	$0.152^{+0.010}_{-0.009}$

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
J2016+1632	Y	···	···	···	219	0.487 ± 0.004	0.110 ± 0.006
J2018–0509	Y	···	FSRQ	0.905	169	0.213 ± 0.004	$0.233^{+0.016}_{-0.015}$
J2018–1109	Y	···	FSRQ	0.9	159	0.145 ± 0.001	$0.095^{+0.007}_{-0.006}$
J2020+6747	Y	···	FSRQ	2.571	230	0.163 ± 0.002	$0.120^{+0.008}_{-0.007}$
J2021+0515	Y	···	···	···	167	0.259 ± 0.001	$0.033^{+0.004}_{-0.003}$
J2022+6136	Y	···	AGN	0.227	221	2.296 ± 0.004	0.018 ± 0.001
J2022+7611	Y	J2020.4+7608	BLL	···	139	0.798 ± 0.021	$0.299^{+0.021}_{-0.019}$
J2023–0123	Y	···	···	···	164	0.175 ± 0.003	$0.219^{+0.016}_{-0.015}$
J2023–1139	···	J2023.7–1141	FSRQ	0.698	126	0.054 ± 0.001	$0.143^{+0.017}_{-0.015}$
J2024+1718	Y	···	FSRQ	1.05	220	0.655 ± 0.005	0.108 ± 0.006
J2025+0316	Y	···	FSRQ	2.21	166	0.346 ± 0.004	0.147 ± 0.009
J2027–0831	Y	···	FSRQ	1.697	158	0.262 ± 0.001	0.028 ± 0.003
J2030–0503	Y	···	FSRQ	0.543	250	0.277 ± 0.002	$0.090^{+0.006}_{-0.005}$
J2030–0622	···	J2030.3–0617	FSRQ	0.667	131	0.208 ± 0.004	$0.213^{+0.018}_{-0.016}$
J2031+0239	Y	···	FSRQ	0.858	168	0.217 ± 0.003	$0.158^{+0.011}_{-0.010}$
J2031+1219	Y	J2031.5+1219	BLL	···	302	1.402 ± 0.011	0.136 ± 0.006
J2033+2146	Y	···	FSRQ	0.173	202	0.448 ± 0.003	$0.073^{+0.005}_{-0.004}$
J2035+1056	···	J2035.4+1100	FSRQ	0.601	137	0.397 ± 0.015	$0.426^{+0.033}_{-0.029}$
J2036–0629	Y	···	FSRQ	1.636	238	0.516 ± 0.002	$0.062^{+0.004}_{-0.003}$
J2037–1522	Y	···	FSRQ	1.802	198	0.113 ± 0.001	$0.125^{+0.009}_{-0.008}$
J2039–1046	···	J2039.0–1047	BLL	···	125	0.328 ± 0.008	$0.250^{+0.019}_{-0.017}$
J2042+7508	Y	···	FSRQ	0.104	138	0.189 ± 0.002	$0.079^{+0.007}_{-0.006}$
J2043+1255	Y	···	FSRQ	3.277	296	0.099 ± 0.001	0.100 ± 0.007
J2045–1858	Y	···	···	···	195	0.189 ± 0.002	$0.105^{+0.008}_{-0.007}$
J2049+1003	Y	···	···	···	168	0.588 ± 0.005	0.115 ± 0.007
J2050+0407	Y	J2050.1+0407	BLL	···	168	0.576 ± 0.001	0.024 ± 0.002
J2051+1743	Y	···	AGN	0.195	339	0.639 ± 0.005	0.124 ± 0.005
J2101+0341	Y	···	FSRQ	1.015	256	0.733 ± 0.003	$0.066^{+0.004}_{-0.003}$
J2102+6758	Y	···	···	···	230	0.063 ± 0.001	$0.271^{+0.018}_{-0.017}$
J2106+0231	Y	···	FSRQ	2.942	250	0.125 ± 0.001	0.052 ± 0.004
J2106+2135	Y	···	FSRQ	0.647	306	0.184 ± 0.002	$0.177^{+0.009}_{-0.008}$
J2108+1430	Y	···	FSRQ	2.017	336	0.362 ± 0.003	0.160 ± 0.007
J2110+0809	···	J2110.0+0811	FSRQ	1.58	136	0.200 ± 0.005	$0.243^{+0.019}_{-0.017}$

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
J2110–1020	Y	···	···	···	222	0.548 ± 0.001	0.027 ± 0.002
J2114+1714	Y	···	···	···	354	0.089 ± 0.001	0.110 ± 0.006
J2114+2832	Y	···	FSRQ	2.345	219	0.377 ± 0.002	0.066 ± 0.004
J2115+2933	Y	J2115.5+2937	FSRQ	1.514	223	0.747 ± 0.007	$0.129^{+0.007}_{-0.006}$
J2115–1416	Y	···	FSRQ	1.7	150	0.281 ± 0.002	0.066 ± 0.005
J2117+0503	Y	···	FSRQ	1.794	256	0.309 ± 0.002	0.063 ± 0.004
J2118+0013	···	J2117.8+0016	FSRQ	0.463	134	0.120 ± 0.004	$0.309^{+0.025}_{-0.022}$
J2118–0636	Y	···	FSRQ	0.328	223	0.124 ± 0.001	$0.087^{+0.007}_{-0.006}$
J2120+0533	Y	···	NLRG	0.535	259	0.132 ± 0.001	$0.145^{+0.009}_{-0.008}$
J2123+0535	Y	···	FSRQ	1.941	257	1.699 ± 0.008	0.069 ± 0.003
J2124–1941	Y	···	···	···	143	0.136 ± 0.001	$0.099^{+0.009}_{-0.008}$
J2125+0441	Y	···	FSRQ	1.394	264	0.256 ± 0.003	$0.179^{+0.010}_{-0.009}$
J2128–0244	Y	···	FSRQ	1.812	262	0.148 ± 0.003	$0.314^{+0.018}_{-0.016}$
J2129–1538	Y	···	FSRQ	3.28	155	0.791 ± 0.003	0.043 ± 0.003
J2130–0927	Y	···	FSRQ	0.78	220	0.757 ± 0.012	$0.221^{+0.012}_{-0.011}$
J2131–1207	Y	···	FSRQ	0.501	145	2.095 ± 0.013	$0.071^{+0.005}_{-0.004}$
J2133+1443	Y	···	···	···	353	0.220 ± 0.003	$0.265^{+0.012}_{-0.011}$
J2134–0153	Y	J2134.0–0203	FSRQ	1.284	249	2.373 ± 0.013	0.082 ± 0.004
J2136+0041	Y	···	FSRQ	1.941	264	6.744 ± 0.020	0.047 ± 0.002
J2139+0122	Y	···	FSRQ	1.401	265	0.180 ± 0.002	0.125 ± 0.007
J2139+1423	Y	···	FSRQ	2.427	298	2.601 ± 0.011	0.071 ± 0.003
J2142–0437	Y	···	FSRQ	0.344	239	0.335 ± 0.004	0.176 ± 0.009
J2143+1743	Y	J2143.4+1742	FSRQ	0.211	309	0.823 ± 0.009	$0.197^{+0.009}_{-0.008}$
J2145+1115	Y	···	FSRQ	0.548	312	0.424 ± 0.003	$0.092^{+0.005}_{-0.004}$
J2146–1525	Y	···	FSRQ	0.698	168	0.600 ± 0.005	$0.106^{+0.007}_{-0.006}$
J2147+0929	Y	J2147.2+0929	FSRQ	1.113	316	0.859 ± 0.014	$0.283^{+0.013}_{-0.012}$
J2148+0211	Y	···	···	···	267	0.061 ± 0.001	0.135 ± 0.011
J2148+0657	Y	J2148.5+0654	FSRQ	0.999	263	5.407 ± 0.017	0.048 ± 0.002
J2148–1723	Y	···	FSRQ	2.13	168	0.517 ± 0.006	$0.150^{+0.009}_{-0.008}$
J2149+0322	···	J2149.7+0327	BLL	···	141	0.080 ± 0.001	$0.069^{+0.012}_{-0.011}$
J2151+0552	Y	···	FSRQ	0.74	263	0.421 ± 0.001	0.034 ± 0.002
J2151+0709	Y	···	FSRQ	1.364	252	0.860 ± 0.006	0.104 ± 0.005
J2152+1734	Y	J2152.5+1734	BLL	0.871	277	0.508 ± 0.002	0.048 ± 0.003

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
J2153–1136	Y	···	···	···	149	0.271 ± 0.008	$0.345^{+0.025}_{-0.022}$
J2156–0037	Y	···	BLL	0.495	234	0.525 ± 0.006	0.157 ± 0.008
J2157+3127	Y	J2157.4+3129	FSRQ	1.486	215	0.653 ± 0.007	$0.148^{+0.008}_{-0.007}$
J2158–1501	Y	J2157.9–1503	FSRQ	0.672	172	1.969 ± 0.021	0.140 ± 0.008
J2200+0234	Y	···	FSRQ	1.323	229	0.094 ± 0.001	0.156 ± 0.010
J2200+2137	Y	···	BLL	···	337	0.160 ± 0.001	0.073 ± 0.004
J2203+1007	Y	···	NLRG	1.005	311	0.130 ± 0.001	0.057 ± 0.005
J2203+1725	Y	J2203.5+1726	FSRQ	1.076	306	1.198 ± 0.011	0.160 ± 0.007
J2203+3145	Y	···	FSRQ	0.298	206	2.815 ± 0.018	$0.089^{+0.005}_{-0.004}$
J2204+0440	···	J2204.6+0442	AGN	0.027	144	0.343 ± 0.003	$0.096^{+0.007}_{-0.006}$
J2204+3632	Y	···	GALAXY	0.073	162	0.140 ± 0.001	$0.080^{+0.007}_{-0.006}$
J2206–0031	Y	···	···	···	222	0.292 ± 0.007	$0.323^{+0.019}_{-0.017}$
J2207+1652	Y	···	FSRQ	1.639	318	0.243 ± 0.002	0.134 ± 0.006
J2210+2013	Y	···	FSRQ	0.282	336	0.468 ± 0.002	0.057 ± 0.003
J2211+1841	Y	···	AGN	0.07	344	0.647 ± 0.008	$0.229^{+0.010}_{-0.009}$
J2212+2355	Y	J2212.1+2358	FSRQ	1.125	340	0.931 ± 0.007	$0.133^{+0.006}_{-0.005}$
J2212+2843	Y	···	···	···	221	0.056 ± 0.001	$0.072^{+0.011}_{-0.010}$
J2214+3739	Y	···	FSRQ	2.249	150	0.334 ± 0.001	0.035 ± 0.003
J2216+3102	Y	···	FSRQ	2.462	176	0.211 ± 0.001	0.060 ± 0.004
J2216+3518	Y	···	FSRQ	0.51	173	0.401 ± 0.004	$0.142^{+0.009}_{-0.008}$
J2217+2421	Y	J2217.1+2423	BLL	0.505	337	0.791 ± 0.007	$0.162^{+0.007}_{-0.006}$
J2218+1520	Y	···	FSRQ	2.335	317	1.032 ± 0.007	0.121 ± 0.005
J2218–0335	Y	···	FSRQ	0.901	202	1.666 ± 0.010	$0.082^{+0.005}_{-0.004}$
J2219+1806	Y	J2219.3+1804	FSRQ	1.071	325	0.199 ± 0.002	0.165 ± 0.008
J2219+2613	Y	···	AGN	0.085	345	0.128 ± 0.001	0.085 ± 0.005
J2225+2118	Y	···	FSRQ	1.953	333	1.480 ± 0.012	0.145 ± 0.006
J2226+0052	Y	···	FSRQ	2.248	200	0.407 ± 0.007	$0.252^{+0.014}_{-0.013}$
J2228+2503	Y	···	FSRQ	0.558	340	0.133 ± 0.001	0.133 ± 0.006
J2229–0832	Y	J2229.7–0832	FSRQ	1.559	165	2.256 ± 0.035	$0.194^{+0.012}_{-0.011}$
J2230+6946	Y	···	FSRQ	1.426	194	0.664 ± 0.003	0.050 ± 0.003
J2230–1325	Y	···	FSRQ	1.42	172	0.685 ± 0.010	$0.180^{+0.011}_{-0.010}$
J2235–1826	Y	···	FSRQ	2.153	212	0.118 ± 0.001	0.129 ± 0.009
J2236+2828	Y	J2236.2+2828	FSRQ	0.795	277	1.232 ± 0.008	0.105 ± 0.005

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
J2236+7322	Y	···	FSRQ	1.345	192	0.177 ± 0.001	0.057 ± 0.004
J2236–1433	Y	J2236.4–1432	BLL	···	176	0.538 ± 0.007	$0.167^{+0.010}_{-0.009}$
J2238+0724	Y	···	FSRQ	1.012	215	0.158 ± 0.001	0.067 ± 0.006
J2238+2749	Y	···	FSRQ	0.836	337	0.177 ± 0.002	0.224 ± 0.010
J2241+0953	Y	···	FSRQ	1.707	292	0.500 ± 0.005	0.153 ± 0.007
J2241+4120	Y	···	BLL	···	265	0.219 ± 0.002	0.154 ± 0.008
J2244+2600	Y	···	FSRQ	2.043	347	0.110 ± 0.001	0.059 ± 0.005
J2245+0324	Y	···	FSRQ	1.34	194	0.331 ± 0.002	0.061 ± 0.004
J2245+0500	Y	···	FSRQ	1.091	207	0.276 ± 0.005	$0.243^{+0.014}_{-0.013}$
J2246–1206	Y	···	FSRQ	0.63	171	2.367 ± 0.013	0.070 ± 0.004
J2247+0310	Y	···	FSRQ	1.571	174	0.468 ± 0.004	0.101 ± 0.006
J2247–1237	Y	···	FSRQ	1.892	180	0.524 ± 0.003	0.057 ± 0.004
J2249+2107	Y	···	FSRQ	1.274	325	0.552 ± 0.004	0.114 ± 0.005
J2250+3824	···	J2250.1+3825	BLL	0.119	132	0.062 ± 0.001	0.054 ± 0.012
J2251+4030	···	J2251.7+4030	BLL	···	131	0.069 ± 0.001	$0.112^{+0.013}_{-0.012}$
J2253+1608	Y	J2253.9+1608	FSRQ	0.859	294	13.347 ± 0.379	$0.484^{+0.025}_{-0.023}$
J2253+1942	Y	···	FSRQ	0.284	279	0.348 ± 0.005	$0.229^{+0.011}_{-0.010}$
J2257+0243	Y	···	FSRQ	2.089	140	0.329 ± 0.001	0.033 ± 0.003
J2259–0811	Y	···	FSRQ	1.376	163	0.128 ± 0.001	$0.044^{+0.007}_{-0.006}$
J2300+1655	Y	···	FSRQ	1.283	529	0.547 ± 0.002	0.079 ± 0.003
J2301–0158	Y	···	FSRQ	0.778	132	0.805 ± 0.015	$0.208^{+0.015}_{-0.013}$
J2303+3853	Y	···	···	···	270	0.065 ± 0.000	0.062 ± 0.009
J2305+8242	Y	···	BLL	···	219	0.210 ± 0.003	$0.178^{+0.010}_{-0.009}$
J2307+1450	Y	J2307.3+1452	BLL	···	268	0.081 ± 0.002	$0.281^{+0.017}_{-0.016}$
J2307+3230	Y	···	FSRQ	1.937	268	0.301 ± 0.001	0.046 ± 0.003
J2308+2008	Y	···	FSRQ	0.25	316	0.252 ± 0.001	$0.070^{+0.004}_{-0.003}$
J2310+1055	Y	···	FSRQ	0.494	278	0.397 ± 0.003	0.130 ± 0.006
J2311+4543	Y	···	FSRQ	1.447	269	0.521 ± 0.007	$0.227^{+0.011}_{-0.010}$
J2321+2732	···	J2321.6+2726	FSRQ	1.253	133	0.878 ± 0.007	$0.086^{+0.006}_{-0.005}$
J2321+3204	···	J2322.0+3208	FSRQ	1.489	130	0.479 ± 0.004	0.076 ± 0.006
J2322+1843	Y	···	FSRQ	1.725	522	0.228 ± 0.001	$0.054^{+0.003}_{-0.002}$
J2322+4445	Y	···	···	···	270	0.413 ± 0.003	$0.114^{+0.006}_{-0.005}$
J2323–0317	Y	J2323.5–0315	FSRQ	1.41	141	0.974 ± 0.008	$0.089^{+0.006}_{-0.005}$

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
J2327+0940	Y	J2327.7+0943	FSRQ	1.843	215	1.686 ± 0.038	$0.328^{+0.019}_{-0.017}$
J2327+1524	Y	···	FSRQ	0.046	287	0.270 ± 0.001	0.071 ± 0.004
J2327+1533	Y	···	FSRQ	0.989	286	0.206 ± 0.001	0.098 ± 0.005
J2329+0834	Y	···	FSRQ	0.948	199	0.215 ± 0.006	$0.381^{+0.024}_{-0.022}$
J2330+1100	Y	···	FSRQ	1.489	224	0.844 ± 0.003	0.054 ± 0.003
J2330+3348	Y	···	FSRQ	1.809	214	0.450 ± 0.003	0.087 ± 0.005
J2331–1556	Y	···	FSRQ	1.153	204	0.843 ± 0.008	0.130 ± 0.007
J2334+0736	Y	J2334.3+0735	FSRQ	0.401	209	1.111 ± 0.011	0.137 ± 0.007
J2335–0131	Y	···	FSRQ	1.184	148	0.402 ± 0.007	$0.197^{+0.013}_{-0.012}$
J2336–1451	Y	···	···	···	179	0.065 ± 0.001	$0.155^{+0.013}_{-0.012}$
J2337+2617	Y	···	···	···	240	0.123 ± 0.001	$0.160^{+0.009}_{-0.008}$
J2337–0230	Y	J2338.3–0231	FSRQ	1.072	142	0.721 ± 0.004	$0.067^{+0.005}_{-0.004}$
J2339–1206	Y	···	···	···	228	0.200 ± 0.005	$0.371^{+0.021}_{-0.020}$
J2340+2641	Y	···	BLL	0.372	234	0.489 ± 0.001	0.038 ± 0.003
J2343+1543	Y	···	FSRQ	1.445	239	0.119 ± 0.001	$0.144^{+0.009}_{-0.008}$
J2343+2339	Y	···	FSRQ	1.328	237	0.379 ± 0.003	$0.125^{+0.007}_{-0.006}$
J2345–1555	Y	J2344.6–1554	FSRQ	0.621	213	0.581 ± 0.009	$0.219^{+0.012}_{-0.011}$
J2346+0930	Y	···	FSRQ	0.677	219	1.009 ± 0.006	0.078 ± 0.004
J2346+8007	Y	···	BLL	···	225	0.217 ± 0.006	$0.407^{+0.024}_{-0.022}$
J2348–0425	Y	···	FSRQ	1.106	142	0.295 ± 0.003	$0.122^{+0.009}_{-0.008}$
J2348–1631	Y	J2348.0–1629	FSRQ	0.576	218	$1.934^{+0.021}_{-0.022}$	0.162 ± 0.008
J2350+1106	Y	···	FSRQ	0.921	232	0.340 ± 0.003	0.111 ± 0.006
J2352+3947	Y	···	FSRQ	0.858	226	0.374 ± 0.003	0.100 ± 0.006
J2354–1513	Y	···	FSRQ	2.675	228	0.690 ± 0.006	0.114 ± 0.006
J2357–0152	Y	···	BLL	0.812	137	0.213 ± 0.002	$0.115^{+0.009}_{-0.008}$
J2357–1125	Y	···	FSRQ	0.96	241	0.889 ± 0.002	0.026 ± 0.002
J2358+1955	Y	···	FSRQ	1.066	268	0.499 ± 0.005	$0.146^{+0.007}_{-0.006}$
J2358–1020	Y	···	FSRQ	1.639	227	0.853 ± 0.012	$0.205^{+0.011}_{-0.010}$
M82	···	J0956.5+6938	AGN	0.001	47	1.639 ± 0.007	$0.023^{+0.004}_{-0.003}$
MS14588+2249	···	J1501.1+2237	BLL	0.235	56	0.062 ± 0.002	$0.171^{+0.029}_{-0.024}$
NGC1218	···	J0308.3+0403	AGN	0.029	108	1.479 ± 0.005	$0.031^{+0.003}_{-0.002}$
PG1437+398	···	J1439.2+3930	BLL	0.349	63	0.045 ± 0.000	< 0.128
PKS0214–085	···	J0217.0–0829	FSRQ	0.607	63	0.560 ± 0.006	$0.082^{+0.009}_{-0.008}$

continued...

Table C.1 (continued)

Source	CGR?	1FGL Name	Class	z	$N_{\rm obs}$	S_0 (Jy)	$\overline{m}$
PKS0336–177	···	J0339.1–1734	AGN	0.065	60	0.097 ± 0.002	$0.085^{+0.018}_{-0.016}$
PKS1508–05	···	J1511.1–0545	FSRQ	1.185	90	$1.280^{+0.005}_{-0.006}$	0.032 ± 0.004
PKS1509+022	···	J1512.3+0201	FSRQ	0.219	105	0.326 ± 0.003	$0.080^{+0.008}_{-0.007}$
PKS1510–089	···	J1512.8–0906	FSRQ	0.36	111	$2.341^{+0.055}_{-0.054}$	$0.242^{+0.019}_{-0.017}$
PKS1728+004	···	J1730.4+0008	FSRQ	1.335	73	0.306 ± 0.008	$0.211^{+0.021}_{-0.019}$
PKS2320–021	···	J2322.3–0153	FSRQ	1.774	61	0.188 ± 0.001	0.043 ± 0.007
RBS1752	···	J2131.7–0914	BLL	0.449	77	no det.	···
RBS421	···	J0325.9–1649	BLL	0.291	57	no det.	···
RBS76	···	J0033.5–1921	BLL	0.61	57	no det.	···
RXJ0035.2+15	···	J0035.1+1516	BLL	···	66	no det.	···
RXJ0202.4+08	···	J0202.1+0849	BLL	···	63	no det.	···
RXJ0805.4+75	···	J0804.7+7534	BLL	0.121	44	no det.	···
RXJ0850.6+34	···	J0850.2+3457	BLL	0.149	52	no det.	···
RXJ0909.0+23	···	J0909.2+2310	BLL	0.223	52	no det.	···
RXJ1448.0+36	···	J1447.9+3608	BLL	···	55	no det.	···
SDSSJ141826.	···	J1418.3–0235	BLL	···	56	no det.	···
SDSSJ205528.	···	J2055.5–0023	BLL	···	88	no det.	···
TXS1029+378	···	J1032.7+3737	BLL	···	52	no det.	···

Appendix D

QUIET

This appendix contains documents describing some of my contributions the QUIET project. Most of this work is presented as reproductions of internal specification documents and memoranda. In section D.6, we present a preprint manuscript of the paper describing the first results from the QUIET program.

D.1 Bias Electronics

Here, we present interface specifications for the bias electronics boards to provide bias to the polarimeter modules. Two boards were designed. The *MMIC bias* electronics were used to generate bias signals for the monolithic microwave integrated circuit (MMIC) amplifiers in the modules. The *phase switch bias* electronics were used to provide bias for the phase switches in the modules. The electronics design for the bias boards was developed chiefly by Dr. Michael D. Seiffert and Steve Smith, and the hardware was built and basic functions were tested by Steve Smith. In addition to developing the interface specification, I was responsible for detailed testing, characterization, and debugging of the electronics. The interface specifications are presented here with their original formatting and pagination. In each, a lengthy appendix containing pin assignment tables is not shown here.

QUIET MMIC Bias Interface Specification

Joey Richards*, Mike Seiffert, Steve Smith
Revision 1.5

January 11, 2006

1 Introduction

This document specifies the various interfaces between the MMIC bias cards and the rest of the QUIET system. The MMIC bias card design described in this document is intended to support both 40 GHz (Q-band)and 90 GHz (W-band) receiver modules with at most minor component value changes between versions.

1.1 Schedule

Key dates in the design and production schedule for the MMIC bias cards are listed in Table 1.

Item	Start	Finish
Interface definition	—	6/17/2005
Prototype run (qty 2)	In progress	July
Prototype test/debug	July	7/31/2005
Deliver board to Chicago (qty 1)	—	7/31/2005
Production run (for 19 Q-band elements, 91 W-band elements, plus spares)	TBD	TBD

Table 1: MMIC bias card design/construction schedule.

2 Mode of Operation

In this section, we describe generally how to interact with the MMIC bias card. Details of the operations required are in later sections of this document. Please take note of the warnings regarding circuit operation in this section.

* *joey@caltech.edu*

2.1 Power-Up Sequence

The MMIC bias card will power-on into a safe "off" state with the MMIC bias is in a low-current mode. The bias can be left in this state indefinitely without risk of module damage.

Incomplete list of Concerns:

- Current balance between MMICs in pre-amp (common drain)
  - Turn both gates on first, then apply drain current

2.1.1 Reset Delay

At power-up or after the DAC /CLR is asserted and de-asserted, an RC delay circuit will hold the /CLR input to the DACs low. A 1 ms delay should be observed after power-up or de-asserting /CLR before attempting to program the DACs.

2.2 Warnings

Please be aware of the following important warnings about MMIC bias card operation.

2.2.1 Limitations of Hardware Failsafes

As much as possible, hardware failsafes (current limits, voltage clamps, etc) have been designed into the MMIC bias circuits. However, it appears that it may not be possible to design the circuits to be completely incapable of damaging the receiver modules. In these cases, software protection will be required to prevent damage to the modules due to user error or software defects. It is strongly advised that this protection be implemented at the lowest level possible in the software drivers to provide maximum protection against software defects.

A detailed list of limitations to hardware failsafes will be included here. These will not be known until the bias cards are designed.

2.2.2 Bias Programming Timing

To prevent contamination of science data, bias values should not be changed while observations are being made. The bias DACs should be updated only during receiver dead time.

2.2.3 Multiplexer Switching Timing

To prevent contamination of science data, the monitor multiplexer address should be changed only during receiver dead time.

2.2.4 Multiplexer Switching Restrictions

In addition to the timing restrictions above, during science data collection, the multiplexer address should only be changed after an even number of phase switch toggles (i.e., synchronously with the 4 kHz phase switch clock). This is to minimize the effect of any bias level offset caused by an interaction with the multiplexer by ensuring such an offset is common mode between the two phase switch states.

Alternatively, sampling may be restricted to the receiver dead time. In this mode, during the dead-time, the multiplexer would switch to the desired signal, the housekeeping board would sample, and the multiplexer would switch back to an isolated state before the end of the dead time.

Interface	Signal Type	Signal Count	Total
Bias Output	Analog	11 per module	154
Bias Control Input	LVDS	5 x 2	10
MAB Ground Sense	Analog	2	2
Monitor/Mux Control Input	LVDS	8 x 2	16
Monitor/Mux Output	Analog	1	1
Power	Power	7	7
Ground	Ground	6	6
Totals:		42 + 11 per mod	196

Table 2: Summary of interface signal counts for the MMIC bias card.

2.2.5 Interactions Between Module Bias Settings

The various bias settings within a MAB will not be entirely independent due to a shared ground return. A ground voltage sense signal is provided to allow software to monitor the voltage of the MAB ground (which will differ from the voltage at the MMIC bias card due to IR drop across the ground wire). It is not expected that this could cause hardware damage, but the bias control software must account for such interactions.

2.2.6 Assumptions About MAB

The gate bias circuit assumes a 1:7.19 ratio voltage divider is installed on the MAB. Use without this divider will damage the module.

3 Electrical Interface

The electrical interface consists of the digital inputs to control the MMIC bias circuit and multiplexers, analog outputs to the MABs, and analog monitor outputs to the housekeeping board. Additionally, the grounding strategy and transient protection requirements are described in this section.

Each MMIC bias card will provide bias for the receiver modules on either two MABs (14 modules). The interface signal requirements are summarized in Table 2.

3.1 Digital Signal Characteristics

All digital inputs to the MMIC bias card are LVDS differential signals. These signals are translated to single-ended digital signals and optoisolated on the MMIC bias card. A power supply and ground must be provided for the LVDS level converter and the input side of the optoisolators (see Section 3.5.2). This power supply and ground must be DC-isolated from the power supplies for the analog sections of the MMIC bias cards (see Section 3.5.3), though it may be common between cards.

3.2 Bias Output

For each 90 GHz receiver module, 10 bias signals must be provided. These are listed in Table 3. The 40 GHz modules will use the same outputs, although they may be connected differently inside the modules.

The descriptions in Table 3 refer to two amps (1 and 2). These correspond to the LNA section (amp 1) and the second gain stage (amp 2) in the module. As seen in the module block diagram in the QUIET

Name	DAC	A3	A2	A1	A0	Description
GX11	0	0	0	0	1	Gate voltage, side X, amp 1, gate 1
GX12	0	0	0	1	1	Gate voltage, side X, amp 1, gate 2
GX2	1	0	0	0	1	Gate voltage, side X, amp 2
DX1	0	0	0	1	0	Drain current, side X, amp 1
DX2	1	0	0	1	0	Drain current, side X, amp 2
GY11	1	0	0	1	1	Gate voltage, side Y, amp 1, gate 1
GY12	0	0	1	0	0	Gate voltage, side Y, amp 1, gate 2
GY2	1	0	1	0	1	Gate voltage, side Y, amp 2
DY1	1	0	1	0	0	Drain current, side Y, amp 1
DY2	0	0	1	0	1	Drain current, side Y, amp 2
RET	—	—	—	—	—	Ground return for module

Table 3: 90 GHz module bias control signals and DAC channel assignments.

Signal	Absolute Maximum			Operating Range			
	V_{MIN}	V_{MAX}	I_{MAX}	V_{MIN}	V_{MAX}	I_{MAX}	ΔV
Gx11	–0.2 V	0.4 V	0.2 mA	–0.1 V	0.3 V	0.1 mA	1 mV
Gx2	–0.2 V	0.4 V	0.2 mA	–0.1 V	0.3 V	0.4 mA	1 mV

Table 4: 40 GHz and 90 GHz module gate bias requirements. ΔV is the gate voltage resolution required to ensure a suitable bias point can be found. These voltages are measured *after* the 1:7.19 voltage divider on the MAB.

90 GHz Module Development memo, the LNA actually consists of two MMIC amplifiers.[1] These each have a separately controlled gate, but their drains are connected to a single pin as indicated.

3.2.1 Module Bias Requirements

The 40 GHz and 90 GHz module bias requirements consist of absolute maximum ratings and operating range ratings. The absolute maximum ratings reflect damage thresholds that must not be exceeded. The operating range ratings reflect the minimum bias range needed to allow proper operation of the module. Ideally, the MMIC bias card hardware will provide bias up to a limit that exceeds the operating range rating but is no greater than the absolute maximum rating.

The gate bias requirements are described in Table 4 and the drain bias requirements are described in Table 5. These tables are based on the QUIET 90 GHz Module Development memo[2] and personal communications with Todd Gaier. Note that the drain current resolution of 0.1 mA corresponds to a drain voltage resolution of approximately 3 mV or finer.

The MAB circuit design will implement a voltage divider on the gate input with a ratio of 1/7.19. The above specifications refer to the voltages *after* the voltage divider.

[1] T. Gaier, "Quiet 90 GHz Module Development," March 16, 2005, internal memo, 1 (Figure 1).
[2] *Ibid.*

Signal	Absolute Maximum	Operating Range			
	I_{MAX}	I_{MIN}	I_{MAX}	V_{MAX}	ΔI
Dx1	30 mA	0	25 mA	1.2 V	0.1 mA
Dx2	60 mA	0	50 mA	1.2 V	0.1 mA

Table 5: 40 GHz and 90 GHz module drain bias requirements. ΔI is the drain current resolution required to ensure a suitable bias point can be found.

Signal	Type	V_{MIN}	V_{MAX}	I_{MIN}	I_{MAX}	Output Impedance
Gx11	voltage source	–0.20 V	0.35 V	—	—	TBD
Gx12	voltage source	–0.20 V	0.35 V	—	—	TBD
Gx2	voltage source	–0.20 V	0.35 V	—	—	TBD
Dx1	current source	—	—	1.5 mA	30 mA	—
Dx2	current source	—	—	1.5 mA	30 mA	—

Table 6: MMIC bias card output capabilities. This table reflects the actual capabilities of the MMIC bias card. Circuits should not be designed to these specs; rather, design requirements are listed in Section 3.2.1. V_{MIN} and V_{MAX} capabilities listed are as measured at the module gate input, after a 1:7.19 voltage divider on the MAB.

3.2.2 Bias Output Capabilities

The MMIC bias card's output capabilities are described in Table 6. Transient suppression / over-voltage protection is discussed in Section 3.7.

The MMIC bias card gate output is designed for use with an MAB gate circuit including a 1:7.19 voltage divider (constructed from a 1.00 kΩ 1% shunt resistor and a 6.19 kΩ 1% series resistor) on the gate input. Without this divider, the MMIC bias card will damage the MMICs.

3.2.3 Bias Output Noise Requirements

Noise on the MMIC bias outputs will contribute to the overall noise of the receiver system. This noise consists of two components. The first is common mode on all the gates and drains, such as would result from variations in the module ground potential. This will be referred to as *ripple*. The second component is uncorrelated between control signals and will be referred to simply as *noise* in the following.

For a complete analysis, interactions between bias noise and the scan strategy should be considered. A ripple component synchronous to the scan could be particularly harmful. We therefore distinguish limits for asynchronous ripple and synchronous ripple.

Bias output noise requirements are listed in Table 7. These are rough estimates extrapolated from the Planck-LFI requirements.[3]

3.3 Bias Control Interface

The bias control interface consists of the digital inputs used to program the DACs that set the bias levels to the modules. Here we describe the protocol for setting the DACs.

[3]T. Gaier, "Planck-LFI Bias Electronics Requirements– Technical Note," December, 1999.

Signal	Noise	Async. Ripple	Sync. Ripple
Gx11	0.58 uV/$\sqrt{\text{Hz}}$	0.58 uV	5.8 nV
Gx12	0.58 uV/$\sqrt{\text{Hz}}$	0.58 uV	5.8 nV
Gx2	0.58 uV/$\sqrt{\text{Hz}}$	0.58 uV	5.8 nV
Dxx	0.58 uV/$\sqrt{\text{Hz}}$	0.58 uV	5.8 nV

Table 7: Bias output noise requirements. *Noise* refers to uncorrelated noise on each line, assumed Gaussian. *Ripple* specs refer to noise signals correlated between output lines, either synchronous or asynchronous to the scan period.

Name	Description
SCK	Serial clock input. Rising-edge active.
DIN	Serial data input.
/CS	Active-low chip select.
/CLR	Active-low asynchronous clear (resets internal DAC registers and output voltages to 0). See note in text regarding /CLR timing.

Table 8: DAC interface signals.

3.3.1 DAC Control Logic

The bias outputs of the MMIC bias circuits are set by an array of Linear Technologies LTC1660[4] 10-bit 8-channel DACs on the card. The digital inputs of the DACs are daisy-chained so all DACs on each card are controlled via a single serial interface with chip select plus an asynchronous clear input. The interface signals are summarized in Table 8. These inputs are optoisolated on the MMIC bias card as necessary. The electrical properties of these signals are described in Section 3.1.

Two LTC1660 DACs are used for each channel for a total of 28 DACs. The two DACs for each module are numbered 0 and 1. The bias settings assigned to each DAC are indicated in Table 3. The bias DACs are divided into two sets of 7 (A and B) corresponding to the two MABs biased by the card. A single chip select controls the entire MMIC bias card.

The DACs are programmed by writing 28 16-bit DAC control words to the serial interface. The ordering of the words is shown in Table 9. To load the DACs, begin by raising /CS, then lowering SCK. Next, lower /CS to enable data input to the DACs. Place the first data bit on DIN and raise SCK to clock in the first bit. Lower SCK and continue clocking in the rest of the bits on rising edges of SCK. After all 448 bits, lower SCK and raise /CS to complete the write operation. For a detailed timing diagram, see page 8 of the LTC1660 data sheet.

[4]Data sheet available at http://www.linear-tech.com.

MAB B					MAB A				
Mod 6 DAC 1	Mod 6 DAC 0	...	Mod 0 DAC 1	Mod 0 DAC 0	Mod 6 DAC 1	Mod 6 DAC 0	...	Mod 0 DAC 1	Mod 0 DAC 0

Table 9: DAC control word ordering. Words transmitted in left-to-right order.

Address/Control				**Input Code**										**Don't Care**	
A3	A2	A1	A0	D9	D8	D7	D6	D5	D4	D3	D2	D1	D0	X1	X0

Table 10: DAC control word format. Bits transmitted in left-to-right order (A3 first, X0 last).

Clock Rate	$\mathbf{t_{SINGLE}}$ (usec)	$\mathbf{t_{TOTAL}}$ (usec)
10 kHz (t_{CLK}=100 usec)	22,500	112,500
100 kHz (t_{CLK}=10 usec)	2,250	11,250
1 MHz (t_{CLK}=1 usec)	225	1,125
8 MHz (t_{CLK}=125 nsec)	28	141

Table 11: Time required to set DAC control words at various serial clock rates.

The format of each 16-bit data word is described in Table 10. The Address/Control input (A3-A0) specifies which of the 8 DAC outputs should be updated with the data in the Input Code field. The address corresponding to each bias control line is indicated in Table 3. Additionally, the Address/Control value "0000" indicates that the Input Code field should be ignored. By setting Address/Control = "0000" in all but one of the 28 words, it is possible to update only a single DAC . For more information about the Address/Control field, see page 10 of the LTC1660 data sheet.

The Input Code (D9-D0) field is the 10-bit value to write to the DAC output specified by the Address/Control field. D9 is the MSB of the value, D0 the LSB. The last two bits of the DAC control word, X1 and X0, are don't-care bits. Their value is ignored.

In addition to the synchronous serial control interface, there is an asynchronous clear input, /CLR. As described in the LTC1660 data sheet, asserting this signal resets the DAC registers to 0 and immediately sets all DAC outputs to 0 Volts. To prevent DAC glitches at power up, the /CLR input is connected to an RC circuit on the MMIC bias card. To accommodate this, a delay of at least 1 ms following power up or assertion of the /CLR must be observed before beginning to load the DAC.

3.3.2 DAC Control Timing

For detailed timing parameters of the DAC interface, please consult the LTC1660 data sheet. Timing values used in this section are from the table on page 3 of the data sheet.

The maximum serial clock rate supported by the LTC1660 is 16.7 MHz ($t_{CLK} = t_3 + t_4 = 60$ ns). However, because we are daisy-chaining their outputs, we are limited to a maximum serial clock rate of 8.3 MHz ($t_{CLK} = t_1 + t_8 = 120$ ns) to guarantee valid data to all DACs. In our calculations below, we will also consider slower clock rates.

To write one set of DAC control words to one MAB on the MMIC bias card requires time $t_{SINGLE} = t_5 + t_6 + t_7 + 14 \times 16 \times t_{CLK}$. To update every DAC on the bias card requires that this be repeated 5 times, so $t_{TOTAL} = 5 \times t_{SINGLE}$. The results for various clock rates are shown in Table 11. Note that in these calculations, t_6 and t_7 were assumed to be 0, and t_5 was set to one clock period.

Note that the above calculations represent only the time to program the digital interfaces of the MMIC bias DACs. Once these are set, the analog outputs will slew to their new voltage settings at the rates indicated on page 3 of the LTC1660 data sheet. A full-scale slew followed by typical settling time adds less than 50 usec to the above times, however, so it is clear that even a 10 kHz serial clock will allow the MMIC biases for all 91 modules to be updated at about 1 Hz, even if the MMIC bias cards must be programmed in series.

Module	Multiplexer Sub-address	Module	Multiplexer Sub-address
A0	1111 xxxx	B0	0111 xxxx
A1	1110 xxxx	B1	0110 xxxx
A2	1101 xxxx	B2	0101 xxxx
A3	1100 xxxx	B3	0100 xxxx
A4	1011 xxxx	B4	0011 xxxx
A5	1010 xxxx	B5	0010 xxxx
A6	1001 xxxx	B6	0001 xxxx
GND SENSE A	1000 xxxx	GND SENSE B	0000 xxxx

Table 12: Receiver module multiplexer sub-addresses. Note that the ground sense signals for the MABs are assigned here.

3.3.3 DAC Transfer Functions

This section will describe the functions that determine the analog output level corresponding to a particular DAC section. This information will not be available until the bias card design is complete.

3.4 Monitor / Multiplexer Interface

For each receiver module, the MMIC bias card provides several signals to be monitored as part of housekeeping. Multiplexers on each MMIC bias card allow selection of one of the monitor signals on the card to be connected to a single analog output to the backplane. This output should be routed to one of the housekeeping boards for digitization.

3.4.1 Monitor Signal Addressing

The MMIC bias card multiplexer is controlled by 8 digital address lines . The 4 MSBs of the address specify from which receiver module to select a monitor signal and the 4 LSBs specify the sub-address of a particular signal. Table 12 and Table 13 list the sub-addresses of the modules and monitor signals.

For example, suppose IDY2MON from receiver module A5 is desired. From Table 12, for module A5 the MSBs of the address should be binary 1010. From Table 13, the sub-address for IDY2MON is 0010, which specifies the LSBs of the address. Thus, the complete multiplexer address for the desired signal is 10100010.

An unused module address (probably 1111xxxx) may be reserved to monitor MMIC bias signals not related to modules. For example, this may be used to allow MMIC bias card power and ground voltages to be monitored.

3.4.2 Multiplexer Output

The output of the multiplexer is passed through an analog optoisolator before being routed to the connector pin. Power and ground for the output side of the optoisolator must be provided separately from the MMIC bias card main power supplies. Power supply requirements are listed in Section 3.5.2.

The housekeeping output signal range is –2 V to +2 V.

3.4.3 Interpreting Monitor Outputs

This section will describe the functions needed to convert the raw monitor signal voltages into operating point measurements. This information will not be available until the bias card design is complete.

Signal Name	Multiplexer Sub-address	Description
IDX1MON	xxxx 1000	Drain current, side X, amp 1
IDX2MON	xxxx 1110	Drain current, side X, amp 2
IDY1MON	xxxx 1011	Drain current, side Y, amp 1
IDY2MON	xxxx 0010	Drain current, side Y, amp 2
VDX1MON	xxxx 1001	Drain voltage, side X, amp 1
VDX2MON	xxxx 1111	Drain voltage, side X, amp 2
VDY1MON	xxxx 1100	Drain voltage, side Y, amp 1
VDY2MON	xxxx 0011	Drain voltage, side Y, amp 2
VGX11MON	xxxx 1010	Gate voltage, side X, amp 1, gate 1
VGX12MON	xxxx 0110	Gate voltage, side X, amp 1, gate 2
VGX2MON	xxxx 0001	Gate voltage, side X, amp 2
VGY11MON	xxxx 1101	Gate voltage, side Y, amp 1, gate 1
VGY12MON	xxxx 0101	Gate voltage, side Y, amp 1, gate 2
VGY2MON	xxxx 0100	Gate voltage, side Y, amp 2
+2.5V	xxxx ????	+2.5 V rail
None	xxxx ????	Unallocated

Table 13: MMIC bias monitor signals for each receiver module.

Supply	Current	Max Ripple
+3.3 V	500 mA	5 mV
–5 V	25 mA	5 mV

Table 14: Main power supplies for each MAB. Each supply must be DC-isolated from all other supplies in the system. Two sets are required per MMIC bias card.

3.5 Power and Ground

3.5.1 Main Power Supplies

The MMIC bias card design will completely isolate the power and ground for the two MABs it biases. Separate power supply and ground return pins for each MAB will be provided. These should be powered by a separate, DC-isolated set of power supplies for each MAB (i.e., two sets of power supplies per MMIC bias card).

Each MAB requires two power supplies, one positive and one negative. These are summarized in Table 14.

3.5.2 Optoisolator Power Supplies

In addition to the main power supplies, a separate power supply is required for the input level converters and optoisolators and for the output analog optoisolator. These power supplies may be shared between different MMIC and phase switch bias cards. The specifications are listed in Table 15.

Supply	Voltage	Current	Max Ripple.
Input Opto	+5 V	50 mA	5 mV
Output Opto (+)	+5 V	50 mA	5 mV
Output Opto (–)	–5 V	50 mA	5 mV

Table 15: Optoisolator power supplies for each MMIC bias card. Power supplies may be shared between cards.

3.5.3 Grounding Strategy

Because they provide power and ground connections to the receiver modules, the MMIC bias cards are a critical piece of the QUIET instrument grounding scheme. The grounding scheme must address two issues. First, radiometer signal contamination due to ground loops or other effects must be minimized. Second, interactions between various bias settings should be minimized.

Each MMIC bias card has two separate ground sections, corresponding to the two MABs it biases. Each ground will be connected to all the modules in the corresponding MAB. Because the receiver modules are case-grounded to the feed array, all 91 module grounds will be connected at the top of the dewar. Therefore, unless an isolated supply is used for each module, some ground loops will be introduced. As a compromise, one power supply per MAB bias card allows the area of the ground-loop to be minimized without requiring the expense (power, pin count, financial, etc.) of 91 isolated power supplies.

Because the ground loops are contained to a single MAB, effect of a shared ground at the bias card can be minimized. This should be done by keeping the signal and ground connections between the bias card and its MAB or MABs as close together as possible as they are routed through the backplane and the FPC connector. All connections between a MAB and its bias card should be routed on a single FPC if possible. If multiple FPCs are required per MAB, care should be taken to route the FPCs destined for a particular MAB in close proximity along their entire lengths.

Another concern is the interaction between bias settings on different modules. Because ground traces and wires have non-zero resistance, return current in the ground path between the module and the bias card causes the ground voltage at the module to differ from that at the bias card. If two modules share a ground return, changes in the bias current of one module will cause changes in the bias card's ground voltage relative to the shared ground at the feed array. This will change the bias settings of the other module.

This interaction will be minimized in two ways. First, simply by limiting the number of modules with a shared ground return, the number of possible interactions between modules is limited. Second, a ground voltage sense wire will be connected to the common ground at each MAB. This sense wire is available to the housekeeping system as a monitor output signal at the address listed in Table 12.

3.6 FPC Connection Requirements

The FPC connections to the MABs should be designed to accommodate the requirements in Table 16.

3.7 Transient Protection

Because the transient protection diodes on the MABs will be not provide protection when cooled in the cryostat, transient protection during cooled operation will be provided on the MMIC bias card as follows.

- Gates
  - Zener diode protection on outputs
- Drains
  - Shunt capacitors at outputs to roll off transients

Signal	Max Current	Resistance
Gates	0.1 mA	20 Ω
Dx1	25 mA	2 Ω
Dx2	50 mA	2 Ω
GND	150 mA	0.5 Ω
GND Sense	< 0.1 mA	20 Ω

Table 16: FPC connection requirements.

  - No protection diodes — would interfere with bias circuit

- Overall
  - Zener diode protection on regulator output in case of regulator failure

4 Mechanical Interface

The mechanical interface includes the physical size of the cards, the connector type and placement, and the location of any mounting holes or brackets on the board.

4.1 Card Dimensions

The MMIC bias card is 207.95 mm high by 160 mm deep (8.187 inches by 6.299 inches). The board dimensions, connector locations, etc., are based on the IEEE 1101.2-1992 conduction-cooled Eurocard specification, modified to reduce the board height. A mechanical drawing is shown in Figure 1.

4.2 Connectors

There are two 160 pin VME connectors (P1 and P2) along one long edge of the PCB. All connections to the MABs are routed to P2. All other connections (digital controls, power inputs, housekeeping outputs) are routed to P1. Connector pin outs are listed in Appendix A.1

5 Thermal Specifications

5.1 Power Budget

The Complete QUIET Electronics memo estimates 24 mW of "overhead" power will be dissipated by the MMIC bias cards for each module.[5] This did not account for linear power regulation on the MMIC bias cards, however. Measurements of early prototype MMIC bias circuits showed an actual power overhead of approximately 350 mW per module. It is believed this can be reduced to approximately 220 mW in the final design. This is in addition to power delivered to (and dissipated in) the receiver module itself.

5.2 Heat Sinking

Heat will be conducted out of the MMIC bias card through wedgelock connectors attached to the bare copper regions along the top and bottom (160 mm edges).

[5] C. Bischoff, et al, "Complete QUIET Electronics, 91-Element Prototype Edition, V.2," March 11, 2005, internal memo, 15-16.

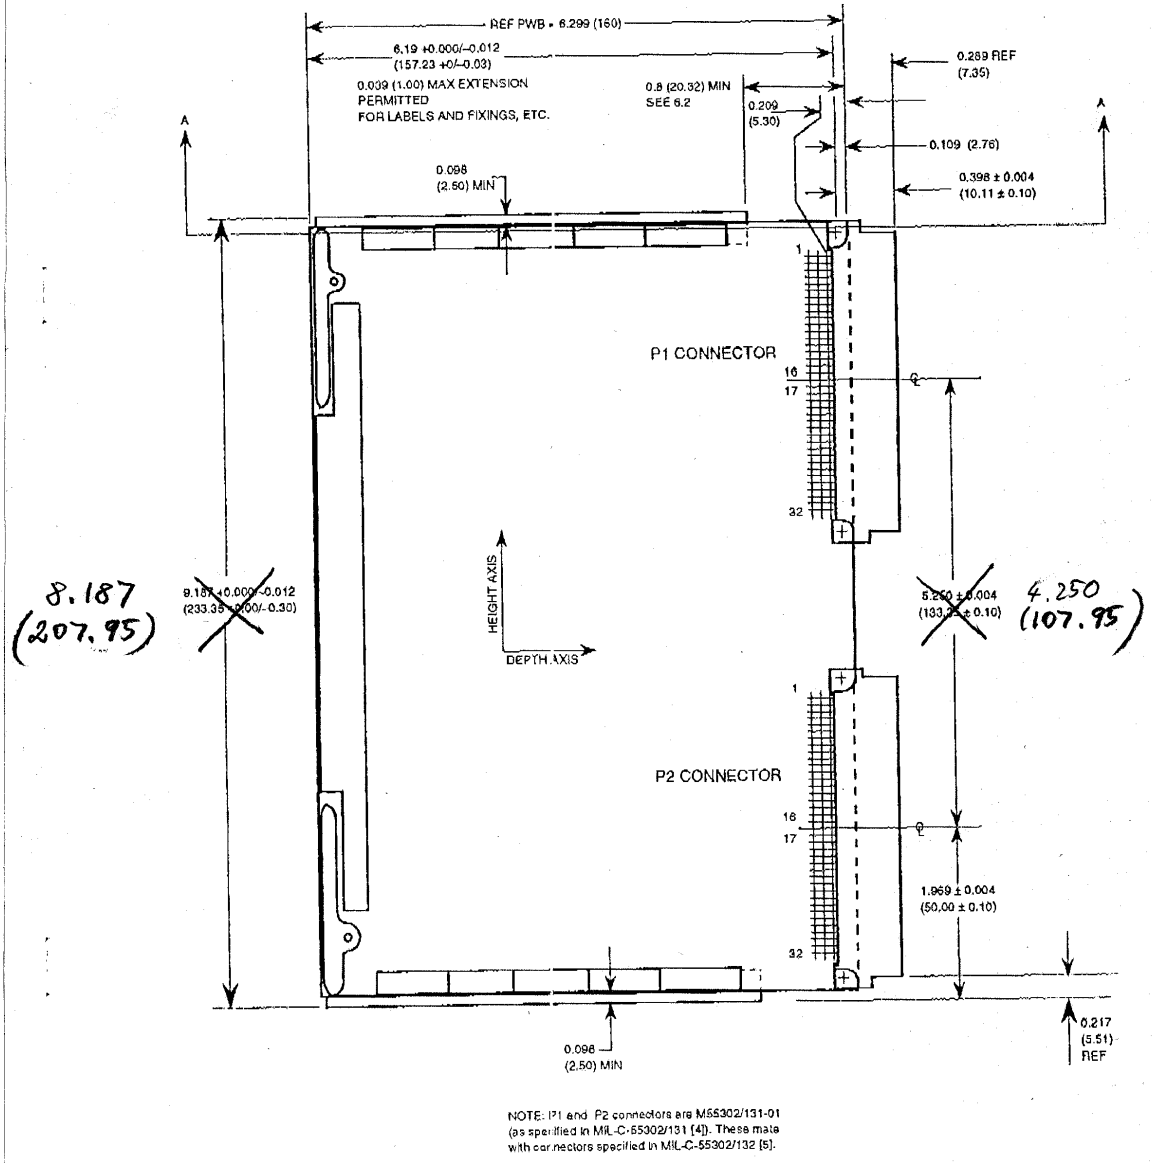

NOTE: P1 and P2 connectors are M55302/131-01 (as specified in MIL-C-55302/131 [4]). These mate with connectors specified in MIL-C-55302/132 [5].

Fig 1
Double-Height Board, Component Side View

Figure 1: Mechanical drawing of board outline and keep-outs.

Modules/Card	**Est. Power Dissipation** (mW)
7	1,540
14	3,080

Table 17: Per-card overhead power dissipation estimates.

5.3 Thermal Stability

To maintain bias stability, the temperature of the MMIC bias card must be maintained. The target gate output voltage stability is a variation of less than 10 uV. To maintain this output stability, the card temperature must be maintained within 1 degree C.

A Appendices

A.1 Connector Pin Outs

A.1.1 P1 Pin Out

The P1 connector (see Figure 1) carries all signals not connected to the MABs. The pin out is in Tables 18 – 22.

A.1.2 P2 Pin Out

The P2 connector (see Figure 1) carries all signals destined for the MABs. The pin out is in Tables 23 – 27.

A.2 Revision History

- Revision 1.5 (December 14, 2005)
  - Fixed DAC and MUX address tables to match actual hardware.
  - Changed MMIC leg designation from A/B to X/Y.
  - Fixed typos in P2 pinout (AORET→A0RET, second pin labelled A5DA1 corrected to A5DA2).
  - Began this revision history section.

QUIET Phase Switch Bias Interface Specification

Joey Richards*, Mike Seiffert, Steve Smith
Revision 1.4

January 11, 2006

1 Introduction

This document specifies the various interfaces between the phase switch bias cards and the rest of the QUIET system. This phase switch bias design is intended for use in both the 40 GHz (Q-band) and 90 GHz (W-band) systems.

1.1 Schedule

Key dates in the design and production schedule for the MMIC bias card are listed in Table 1.

Item	Start	Finish
Interface definition	—	6/17/2005
Prototype run (qty 2)	In Progress	July
Prototype test/debug	July	7/31/2005
Deliver board to Chicago (qty 1)	—	7/31/2005
Production run (for 19 Q-band elements, 91 W-band elements, plus spares)	TBD	TBD

Table 1: Phase switch bias card design/construction schedule.

2 Mode of Operation

In this section, we describe generally how to interact with the phase switch bias card. Details of the operations required are in later sections of this document. Please take note of the warnings regarding circuit operation in this section.

2.1 Power-Up Sequence

This section will describe any special procedures or precautions necessary to safely power up the phase switch bias card. This is not expected to be particularly sensitive.

* *joey@caltech.edu*

Interface	Signal Type	Signal Count	Total
Bias Output	Analog	5 per module	105
Bias Control Input	LVDS	7 x 2	14
Monitor/Mux Control Input	LVDS	6 x 2	12
Phase Toggle	LVDS	4 x 2	8
Housekeeping	Analog	1	1
Power	Power	9	9
Ground	Ground	8	8
		Total:	157

Table 2: Summary of interface signals for the phase switch bias card.

2.1.1 Reset Delay

At power-up or after the DAC /CLR is asserted and de-asserted, an RC delay circuit will hold the /CLR input to the DACs low. A 1 ms delay should be observed after power-up or de-asserting /CLR before attempting to program the DACs.

2.2 Warnings

2.2.1 Receiver Dead Time

Because the phase switches will produce spurious output during switching events, data collected during switching events should be discarded.

2.2.2 Multiplexer Switching Timing

To prevent contamination of science data, the housekeeping monitor multiplexer address should be changed only during receiver dead time.

2.2.3 Multiplexer Switching Restrictions

In addition to the timing restrictions above, during science data collection the multiplexer address should only be changed after an even number of phase switch toggles (i.e., synchronously with the 4 kHz phase switch clock). This is to minimize the effect of any bias level offset caused by an interaction with the multiplexer by ensuring such an offset is common mode between the two phase switch states.

3 Electrical Interface

The electrical interface consists of digital inputs to control the phase switch bias levels and polarity, a set of analog bias outputs destined for the MABs, and digital inputs and an analog output to select and monitor housekeeping signals. Additionally, we describe the grounding strategy and transient protection in this section.

Each phase switch card will provide bias for the phase switches in the modules on three MABs (21 modules per phase switch card). The interface signal requirements are summarized in Table 2.

Name	Description
PA1	Side A phase switch bias 1
PA2	Side A phase switch bias 2
PB1	Side B phase switch bias 1
PB2	Side B phase switch bias 2
GND	Ground return for module

Table 3: Phase switch bias control signals.

Signal	Absolute Maximum		Operating Range		
	V_{MIN}	I_{MAX}	V_{MIN}	V_{MAX}	I_{MAX}
Pxx	-2.5 V	2 mA	-1.8 V	0.85 V	820 uA

Table 4: Phase switch bias requirements, from personal communication with T. Gaier.

3.1 Digital Signal Characteristics

All digital inputs to the phase switch bias card are LVDS differential signals. These signals are translated to single-ended digital signals and optoisolated on the phase switch bias card. A power supply and ground must be provided for the LVDS level converter and the input side of the optoisolators (see Section 3.6.2). This power supply and ground must be DC-isolated from the power supplies for the analog sections of the phase switch bias cards (see Section 3.6.3), though it may be common between phase switch cards.

3.2 Bias Output

Each of the two phase switches on each receiver module requires two bias connections. The phase switch bias card outputs for each module are listed in Table 3.

3.2.1 Phase Switch Bias Requirements

The phase switch bias requirements consist of absolute maximum ratings and operating range ratings. The absolute maximum ratings reflect damage thresholds that must not be exceeded. The operating range ratings reflect the minimum bias range needed to allow proper operation of the module. Ideally, the phase switch bias card hardware will provide bias up to a limit that exceeds the operating range rating but is no greater than the absolute maximum rating. The ratings are summarized in Table 4.

3.2.2 Bias Output Capabilities

The phase switch bias card's output capabilities are described in Table 5. These will be specified to meet the phase switch bias requirements when connected to an MAB through a flexible printed circuit (FPC) that meets the specifications in Section 3.7.

Signal	V_{MIN}	V_{MAX}	I_{MAX}
Pxx	?	?	1.6 mA

Table 5: Phase switch bias card output capabilities.

Name	Description
SCK	Serial clock input. Rising-edge active.
DIN	Serial data input.
/CS	Active-low chip select.
/CLR	Active-low asynchronous clear (resets internal DAC registers and output voltages to 0). See note in text regarding /CLR timing.

Table 6: DAC interface signals.

3.2.3 Bias Output Noise Estimates

TBD

3.3 Bias Control Interface

The bias control interface consists of the digital inputs used to program the DACs that set the bias levels to the phase switches. Here we describe the protocol for setting the DACs.

3.3.1 DAC Control Logic

The bias outputs of the phase switch bias circuits are set by an array of Linear Technologies LTC1660 10-bit 8-channel DACs on the card. The digital inputs of the DACs are daisy-chained so all DACs on each card are controlled via a single serial interface with chip select plus an asynchronous clear input. The interface signals are summarized in Table 6. These inputs are optoisolated on the phase switch bias card as necessary. The electrical properties of these signals are described in Section 3.1

Each phase switch requires only one programmable bias setting so only one DAC channel is required for each module. Thus, there are three eight-channel DACs to support 21 modules per card. The DACs are numbered 0–2. DAC channel assignments are listed in Table 7. The DACs are programmed by transmitting three 16-bit DAC control words to the serial interface. The DAC serial interfaces are daisy-chained so that the first word transmitted sets DAC 2, the second DAC 1, and the third DAC 0. Any time the DACs are updated, three words must be transmitted, even if only one DAC is to be updated.

To load the DACs, begin by raising /CS, then lowering SCK. Next, lower /CS to enable data input to the DACs. Place the first data bit on DIN and raise SCK to clock in the first bit. Lower SCK and continue clocking in the rest of the bits on rising edges of SCK. After all bits have been transmitted, lower SCK and raise /CS to complete the write operation. For a detailed timing diagram, see page 8 of the LTC1660 data sheet.

The format of each 16-bit control word is described in Table 8. The Address/Control input (A3-A0) specifies which of the 8 DAC outputs should be updated with the data in the Input Code field. The address corresponding to each bias control line is indicated in Table 7. Additionally, the Address/Control value “0000” indicates that the Input Code field should be ignored. By setting Address/Control = “0000” in all but one word, it is possible to update only a single DAC . For more information about the Address/Control field, see page 10 of the LTC1660 data sheet.

The Input Code (D9-D0) field is the 10-bit value to write to the DAC output specified by the Address/Control field. D9 is the MSB of the value, D0 the LSB. The last two bits of the DAC control word, X1 and X0, are don’t-care bits. Their value is ignored.

In addition to the synchronous serial control interface, there is an asynchronous clear input, /CLR. As described in the LTC1660 data sheet, asserting this signal resets the DAC registers to 0 and immediately sets all DAC outputs to 0 Volts. To prevent DAC glitches at power up, the /CLR input is connected to an

Module	DAC	Addr	Module	DAC	Addr	Module	DAC	Addr
A0	0	0001	B0	1	0001	C0	2	0001
A1	0	0010	B1	1	0010	C1	2	0010
A2	0	0011	B2	1	0011	C2	2	0011
A3	0	0100	B3	1	0100	C3	2	0100
A4	0	0101	B4	1	0101	C4	2	0101
A5	0	0110	B5	1	0110	C5	2	0110
A6	0	0111	B6	1	0111	C6	2	0111

Table 7: DAC interface signals. Each DAC controls the modules on one MAB (labeled A–C).

Address/Control				Input Code										Don't Care	
A3	A2	A1	A0	D9	D8	D7	D6	D5	D4	D3	D2	D1	D0	X1	X0

Table 8: DAC control word format. Bits transmitted in left-to-right order (A3 first, X0 last).

RC circuit on the MMIC bias card. To accommodate this, a delay of at least 1 ms following power up or assertion of the /CLR must be observed before beginning to load the DAC.

3.3.2 DAC Control Timing

For detailed timing parameters of the DAC interface, please consult the LTC1660 data sheet. Timing values used in this section are from the table on page 3 of the data sheet.

The maximum serial clock rate supported by the LTC1660 is 16.7 MHz ($t_{CLK} = t_3 + t_4$ =60 ns). However, because we are daisy-chaining their outputs, we are limited to a maximum serial clock rate of 8.3 MHz ($t_{CLK} = t_1 + t_8$ =120 ns) to guarantee valid data to all DACs. Writing one set of DAC control words to the phase shift cards requires time $t_{SINGLE} = t_5 + t_6 + t7 + 3 \times 16 \times t_{CLK}$ (assuming 21 modules per card). Updating all DAC channels requires time $t_{TOTAL} = 7 \times t_{SINGLE}$.

For example, even with a very slow 10 kHz clock, t_{CLK} = 100 usec. Neglecting t_6 and t_7 and assuming $t_5 = t_{CLK}$, we find t_{SINGLE} = 4.9 msec and t_{TOTAL} = 34.3 msec. In the worst case, even if the phase shifter cards must be programmed in series, programming all five phase shift bias cards requires less than 180 msec.

3.3.3 DAC Transfer Functions

This section will describe a function to predict the bias current that will result from a particular DAC setting.

3.4 Phase Switch Toggle Input

Each phase switch has two inputs that are used to bias diodes inside the switch. At any time, one bias is forward-biased and the other reverse-biased. By switching which diode is forward-biased, the phase switch is flipped by 180 degrees.

This is accomplished in the phase switch bias circuit by switching the two bias outputs between a fixed negative reference voltage (-1.9 V) and a forward bias current set by the bias DAC. Which of the two inputs of the phase switch is forward-biased is controlled using the PCLK toggle input. Table 9 describes the effect of the PCLK input on the phase switch bias.

Each receiver module has two independently controlled phase switches. During ordinary QUIET operation, one phase switch in each module will be kept in a constant state (i.e., its PCLK line will not toggle)

PCLKA/PCLKB	PA1/PB1	PA2/PB2
HI	Fwd biased	Rev biased
LO	Rev biased	Fwd biased

Table 9: PCLK truth table. PCLKA controls PA1 and PA2 on the module, PCLKB controls PB1 and PB2.

Module	Multiplexer Sub-address	Module	Multiplexer Sub-address	Module	Multiplexer Sub-address
A0	11111x	B0	10111x	C0	01111x
A1	11110x	B1	10110x	C1	01110x
A2	11101x	B2	10101x	C2	10010x
A3	11100x	B3	10100x	C3	01100x
A4	11011x	B4	10011x	C4	01011x
A5	11010x	B5	10010x	C5	01010x
A6	11001x	B6	10001x	C6	01001x

Table 10: Receiver module multiplexer sub-addresses.

while the other is switched at 4 kHz. During testing and calibration, it may be necessary to toggle the state of both switches on each module, so this capability is provided.

Each phase switch bias board provides 2 pairs of PCLK lines (4 total), labeled PCLK0a/b and PCLK1a/b. PCLK0 controls even-numbered modules on all three MABs (0, 2, 4, 6) and PCLK1 controls odd-numbered modules on all three MABs (1, 3, 5). The -a and -b designation indicates which phase switch (side a or side b) in each module is controlled.

3.4.1 Phase Switch Toggle Timing

During ordinary QUIET operation, one phase switch in each receiver module will be toggled at 4 kHz. Because the phase switch will generate spurious output during each transition, no data can be collected during the toggle. For reasonable operation, the phase switch bias circuit switching time, t_{SWITCH}, will not exceed 10% of the switching period. There are two transitions in each period of the 4 kHz clock, so $t_{SWITCH} \leq 12.5$ usec.

3.5 Monitor / Multiplexer Interface

The bias current to the forward-biased diode in each phase switch is available as an analog output. Multiplexers on each phase switch bias card allow one phase switch bias current monitor signal on the card to be connected to a single analog output to the backplane. This output should be routed to one of the housekeeping boards for digitization.

3.5.1 Monitor Signal Addressing

The phase switch bias card multiplexer is controlled by 6 digital address lines. The 5 MSBs of the address specify from which receiver module to select a monitor signal and the LSB specifies the sub-address of the particular signal. Table 10 and Table 11 list the sub-addresses of the monitor signals and the modules.

For example, suppose Ib from receiver module B5 is desired. From Table 10, the MSBs of the address for module B5 are binary 01101. From Table 11, the sub-address for Ib is 1, which specifies the LSB of the address. Thus, the complete multiplexer address for the desired signal is 011011.

Signal Name	Multiplexer Sub-address	Description
Ia	xxxxx1	Forward bias current, phase switch A
Ib	xxxxx0	Forward bias current, phase switch B

Table 11: Phase switch bias monitor signals for each receiver module.

Supply	Current	Max Ripple
+5 V	50 mA	5 mV
–5 V	25 mA	5 mV

Table 12: Main power supplies for each phase switch bias card. Use of a single DC-isolated supply for each phase switch bias card is allowed.

3.5.2 Multiplexer Output

The output of the multiplexer is passed through an analog optoisolator before being routed to the connector pin. Power and ground for the output side of the optoisolator must be provided separately from the phase switch bias card main power supplies. Power supply requirements are listed in Section 3.6.2.

The housekeeping output signal range is –2 V to +2 V.

3.5.3 Interpreting Monitor Outputs

This section will describe the function to calculate bias current given the measured housekeeping voltage.

3.6 Power and Ground

3.6.1 Main Power Supplies

The phase switch bias card design will completely isolate the power and ground for the three MABs it biases. Separate power supply and ground return pins for each MAB will be provided.

It is allowable to power all three MABs on each phase switch bias card with a single set of power supplies. Each phase switch bias card must have its own set of DC-isolated power supplies, however, and all phase switch bias card power supplies must be DC-isolated from all other power supplies in the system. If a single set of supplies is used, the backplane must route these supplies to the input pins for all three MABs.

Each MAB requires two power supplies, one positive and one negative. These are summarized in Table 12.

3.6.2 Optoisolator Power Supplies

In addition to the main power supplies, a separate power supply is required for the input level converters and optoisolators and for the output analog optoisolator. These power supplies may be shared between different phase switch and MMIC bias cards. The specifications are listed in Table 13.

3.6.3 Grounding Strategy

Because they provide power and ground connections to the receiver modules, the phase switch bias cards are a critical piece of the QUIET instrument grounding scheme. Because there are fewer settings for each module and the phase switch bias current to each module will be very nearly constant, interactions between modules are less of a concern than for the MMIC bias cards. However, preventing noise pick-up or coupling to other circuit elements is still important.

Supply	Voltage	Current	Max Ripple
Input Opto	+5 V	50 mA	5 mV
Output Opto (+)	+5 V	75 mA	5 mV
Output Opto (–)	–5 V	75 mA	5 mV

Table 13: Optoisolator power supplies for each phase switch bias card. Power supplies may be shared between cards.

Signal	Max Current	Resistance
GND	4 mA	2 Ω
Pxx	1 mA	2 Ω

Table 14: FPC connection requirements.

Each phase switch bias card should be powered from a dedicated, DC-isolated power supply with an isolated ground. This ground will be connected to all the modules biased by the card. Because the receiver modules are case-grounded to the feed array, all 91 module grounds will be connected at the top of the dewar. Therefore, unless an isolated supply is used for each module, some ground loops will be introduced. As a compromise, one power supply per phase switch bias card allows the area of the ground-loop to be minimized without requiring the expense (power, pin count, financial, etc.) of 91 isolated power supplies.

Because each phase switch bias card biases the modules on only three MABs, the effect of a shared ground at the bias card can be minimized. This should be done by keeping the signal and ground connections between the bias card and its MABs as close together as possible as they are routed through the backplane and the FPC connector. The three MABs biased by each phase switch bias card should be neighbors in the array to minimize loop area.

3.7 FPC Connection Requirements

The FPC connections to the MABs should be designed to accommodate the requirements in Table 14.

3.8 Transient Protection

The following transient protection will be included on the phase switch bias card. This is intended to protect the phase switches from damage due to transients.

- Forward bias circuit
  - Current-limiting resistor
- Reverse bias circuit
  - Zener diode voltage protection
- Overall
  - Zener diode protection on regulator output in case of regulator failure

4 Mechanical Interface

The mechanical interface includes the physical size of the cards, the connector type and placement, and the location of any mounting holes or brackets on the board.

4.1 Card Dimensions

The MMIC bias card is 207.95 mm high by 160 mm deep (8.187 inches by 6.299 inches). The board dimensions, connector locations, etc., are based on the IEEE 1101.2-1992 conduction-cooled Eurocard specification, modified to reduce the board height. A mechanical drawing is shown in Figure 1.

4.2 Connectors

There are two 160 pin VME connectors (P1 and P2) along one long edge of the PCB. All connections to the MABs are routed to P2. All other connections (digital controls, power inputs, housekeeping outputs) are routed to P1. Pin outs for these connectors are listed in Appendix A.1.

5 Thermal Specifications

5.1 Power Budget

TBD, but small.

5.2 Heat Sinking

Heat will be conducted out of the MMIC bias card through wedgelock connectors attached to the bare copper regions along the top and bottom (160 mm edges).

5.3 Thermal Stability

To maintain bias stability, the temperature of the phase switch bias card must be maintained. The allowable temperature swing is to be determined.

A Appendices

A.1 Connector Pin Outs

A.1.1 P1 Pin Out

The P1 connector (see Figure 1) carries all signals not connected to the MABs. The pin out is in Table 15 – Table 19.

A.1.2 P2 Pin Out

The P2 connector (see Figure 1) carries all signals destined for the MABs. The pin out is in Tables 20 – 24.

D.2 Module Protection Circuitry

In this section we present a memorandum describing a recommended design for protection circuitry to prevent damage to the QUIET polarimeter modules when the cryostat was cold. The polarimeter modules are installed in the cold section of the cryostat and installed on a printed circuit board called a module attachment board (MAB) with connectors for the flexible printed circuit (FPC) cabling that extended out to the cryostat interface plate. Circuitry to protect the polarimeter module from static discharge was installed on the MAB, but much of the protection was provided by diodes that were not functional at cryogenic temperatures. Thus, if the connectors on the warm side of the interface plate were exposed to a static discharge during maintenance, polarimeter module damage could result. In this memorandum, Robert Dumoulin and I propose an additional set of "warm protection" circuitry to be installed on the FPC cabling to address this risk. It is presented with its original formatting and pagination.

Static Protection for Cold QUIET Modules

Robert Dumoulin*and Joey Richards†

July 11, 2006

Abstract

We propose that static protection circuits be integrated into a modified FPC design to protect the QUIET modules from damage due to static discharges when the MAB protection circuits are cooled.

1 Introduction

The QUIET modules contain several circuit elements that are susceptible to damage due to static shocks. The module attachment boards (MABs) and bias cards each contain protection circuitry to mitigate the risk of damage. However, the protection in these areas is insufficient to fully protect the modules during operation. The protection diodes on the MABs are not rated for cryogenic operation and are expected to fail when the modules are cooled. The redundant protection on the bias cards will continue to function, but will not provide any protection if the flexible printed circuit (FPC) connections are unplugged for maintenance.

We propose that the FPCs be modified to include protection circuitry that will be kept in the warm region of the cryostat. This will guarantee protection of the modules at all times, even when cooled for operation.

2 Existing Protection Circuits

Currently, protection is installed on both the MAB and on the bias cards. We propose that this protection be kept in place. For clarity, we summarize here the existing protection as we understand it.

2.1 MAB Protection

Figures 1–3 show the protection circuits to be included on the MAB for the gate, drain, and phase switch pins. The diode types are not specified here, we assume they will be selected appropriately to provide protection without impeding the operating range of the bias signals. This information is based on the schematics in the July 7 QUIET MAB Interface Specification[1] and recent communication with Dan Kapner about the voltage divider ratio.

It should be noted that the protection diodes on the gate circuits are on the input side of the 5.1:1 voltage dividers on those lines. This is a change from the initial MAB implementation which placed the diodes directly on the module gate. This change was previously suggested in order to reduce the size of voltage spikes that could appear on the module gate.

*robert@phys.columbia.edu

†joey@caltech.edu

[1]The QUIET Mab Interface Specification indicates that BAV99 diodes will be used. However, it's not clear that these will provide sufficient gate voltage range in the circuit as drawn.

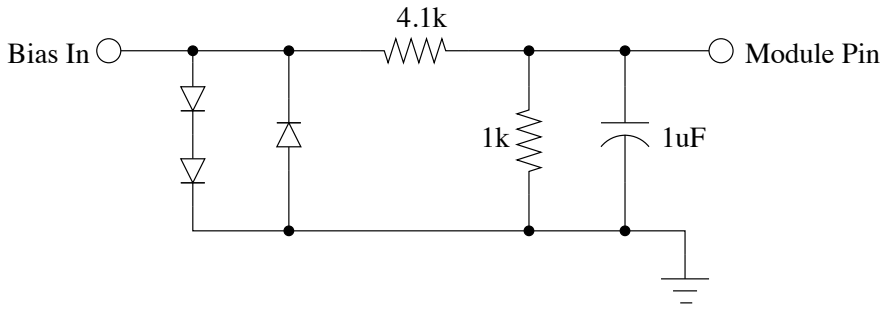

Figure 1: Schematic of gate protection to be installed on MAB. Note that the diode protection appears on the input side of the 5.1:1 voltage divider.

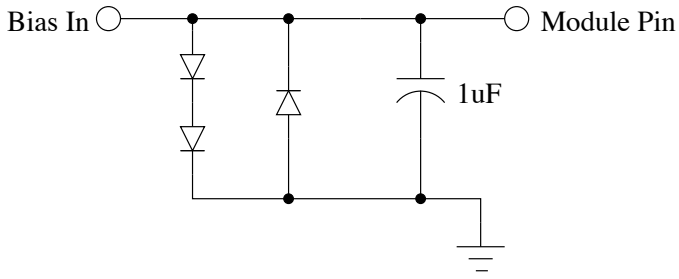

Figure 2: Schematic of drain protection to be installed on MAB.

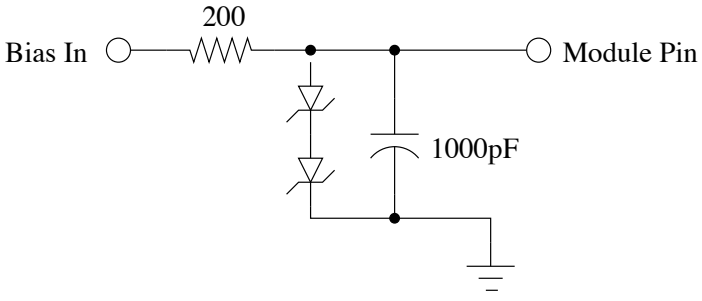

Figure 3: Schematic of phase switch protection to be installed on MAB.

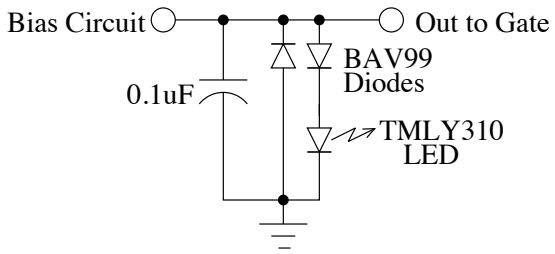

Figure 4: Schematic of gate protection on the MMIC bias card. The LED is used to provide a sharper turn-on knee when the circuit begins limiting the voltage.

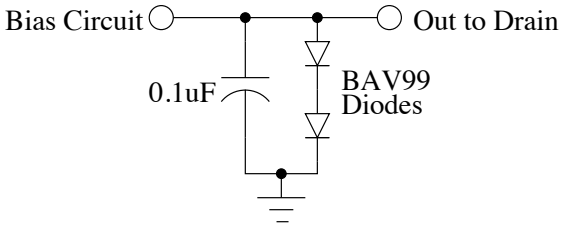

Figure 5: Schematic of drain protection circuit on the MMIC bias card.

In our proposal, we add an additional reason to place these diodes in the position shown in Figure 1. Diode protection placed upstream of the MAB will necessarily occur before the voltage divider[2]. Thus, diodes compatible with this position must be specified for the added protection and it is sensible to use the same components on the MAB itself.

2.2 Bias Card Protection

The protection circuits on the MMIC bias card are shown in Figures 4 and 5. These circuits come from the June 26 schematics used in the recent design review.

Figure 6 shows the protection circuit on the phase switch bias card. This circuit is from the schematics used for the 2005 phase switch bias card prototype build.

[2]It is undesirable to move the voltage divider on to the FPC because this will increase the impact of noise picked up by the portion of the FPC following the divider. The divider should be placed as near to the module gate pin as possible to minimize noise susceptibility.

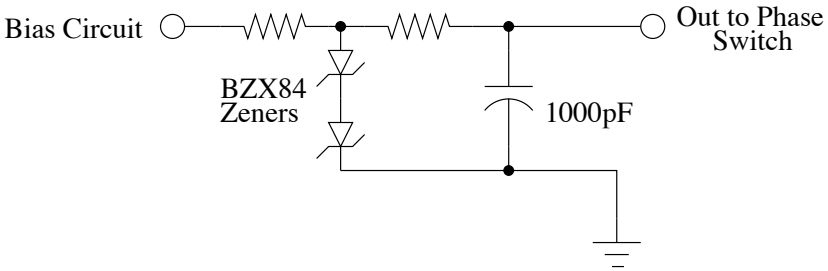

Figure 6: Schematic of protection circuit on the phase switch bias card. The resistors are both 1 kΩ.

3 Proposal for Additional Protection

We propose incorporating an extra set of "warm-protection" circuits for the modules integrated into the FPCs themselves. These circuits should be located on the portion of the FPC that is inside the cryostat to minimize flexure of the component-bearing portion of the FPC. Because we do not have very many options for small protection diodes, we will most likely have to widen the cable to accommodate all of the components.

3.1 Widening the Cable

Each trace for the FPC is currently about 0.5 mm. In order to accommodate a diode in a SOT-363 package, we would have to widen each trace (at a local region/point) to about 2.5 mm to accommodate them. We would also have to put resistors/capacitors and ground links for the cable to function as protection circuitry. According to engineers at Cirexx, our FPC manufacturer, it is not a problem to widen the cable for any region/length (leaving the end connectors the same size). This can probably be done neatly, leaving most of the cable unchanged with a small thick region where the components are.

3.2 Rigidity of the Cable

Since the cable will have quite a few components in a particular region, the cable must to be fairly rigid for this small component bearing of the cable (perhaps several inches). This should not present too many problems, as we should be able to find a small warm section where the cable does not need very much flexibility.

3.3 Additional Cost of Protection on FPC

Until we submit designs to Cirexx, we cannot get a specific quote for the cost of this additional work. According to two engineers at Cirexx, widening the cable and adding components (we would have to stuff the FPCs ourselves) will probably cost an additional $15 - 20\%$, so long as we do not change the number of layers on the FPC. The maximum additional cost of this would be about an additional 25% over the present cost. According to the November 15, 2005, FPC Memo, the cost of 5 sets of 5 FPCs is about $3400, so the additional cost for a component-bearing FPC will not likely exceed $850 for this quantity.

If additional FPC layers must be added, the cost increase will be much greater, probably resulting in a 100% increase in price. It is our understanding that the existing FPC uses two routing layers with signals on one layer and ground connections on the other. This is convenient for routing because all protection circuits will require a connection to ground.

3.4 FPC Protection Circuits

We suggest that the same protection circuits (minus resistors) used on the MAB be replicated on the FPC. During cold operation, this should pose no problems. During warm operation, the parallel protection diodes on the FPC and on the MAB will draw extra current (because the current through the "off" diodes is not exactly zero). This may be inconvenient for the drains in particular, but should only affect warm operation.

4 Rejected Option: Additional PCB

Initially, we considered the option of introducing a new PCB that would contain the new warm protection circuitry. This PCB would sit inside the cryostat and connect to a shortened FPC from the MAB on one side and connect either to another FPC or to another type of high-density cable on the other. This introduces either two extra sets of FPC connections or the challenge of identifying a more convenient, mechanically robust high-density cable. Because of this and because we believe the cost of integrating the protection circuits directly on the FPC to be similar to (or less than) that of an additional PCB with connectors, we

have rejected this option. However, in the unlikely event that routing the protection circuitry will require additional FPC layers, this option may be worth reconsidering.

5 Conclusions

We believe that additional protection circuitry is necessary to guard against static damage to the QUIET modules when the modules are cooled in the cryostat. It appears that a duplicate set of protection circuits can be integrated into the FPC "cables" at a reasonable cost. We therefore propose that the FPC design be modified to include these protection circuits.

D.3 Housekeeping Measurement Procedures

This section contains a memorandum describing a recommended set of procedures for measuring and calibrating the housekeeping data outputs of the bias electronics. This work was based on experience from testing prototype bias electronics cards in the lab. This memorandum is presented with its original formatting and pagination.

QUIET MMIC Bias Housekeeping Calibration Procedure

Joey Richards*
Revision 1.1

February 1, 2008

1 Introduction

This note describes a basic procedure for calibrating the QUIET MMIC bias card housekeeping output. At this time, careful analysis of unit-to-unit variation, temperature sensitivity, and other potentially important effects has not been done. It is not entirely clear what degree of absolute accuracy is required for these outputs. Module performance is quite likely sensitive to variation within the errors that will result from the calibration procedure described here. Based on testing with one module site, it appears that this method will result in measurement accuracy of approximately 2-3 mV for gate voltages, 5 mV for drain voltages, and 1 mA for drain currents (up to 1 V drain voltage).

2 Calibration Procedure

For quick reference, the following calibration procedure is suggested.

1. Measure analog optoisolator voltage transfer function, f_{opto} by injecting a signal into the GNDSENSE input.
2. Measure impedance of each MAB's ground connection and each drain bias line. For the accuracy described here, 1 Ω precision is sufficient—it may be possible to achieve this without measurement. If it's convenient to measure gate bias lines as well, these may be useful for improving precision in the future.
3. Use f_{opto}^{-1} and equations (1), (8), and (9) to convert housekeeping measurements into actual bias parameters.

3 Circuit Analysis

The key component to be calibrated is the analog optoisolator. This is common to all modules on a MAB, so only two sets of measurements per MMIC bias card are required. Once this curve is known, it is straightforward to compute estimates of the actual gate and drain bias parameters from the housekeeping measurements. For the level of precision described here, this measurement need only be made once per MAB-worth of circuits; no measurements of individual bias circuits are required.

* *joey@caltech.edu*

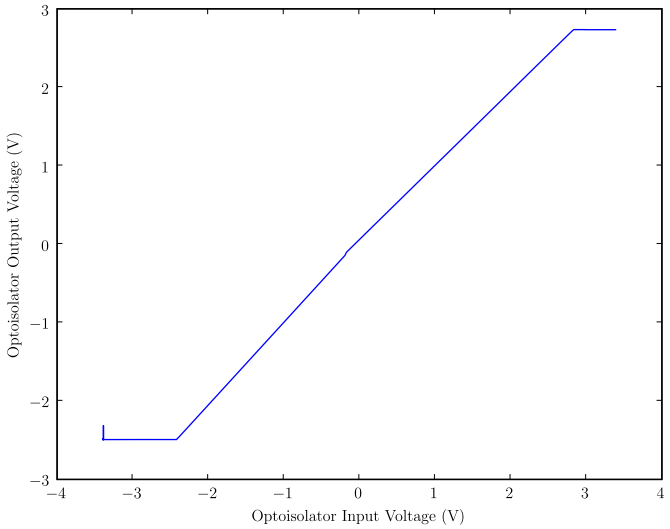

Figure 1: Housekeeping analog optoisolator voltage transfer curve, f_{opto} measured by injecting a signal into the GNDSENSE input. The Y axis indicates the voltage on the housekeeping output line.

3.1 Optoisolator

As of revision 2 of the MMIC bias card design, the optoisolator circuit's voltage response is nonlinear. The circuit consists of a separate LED/photo-diode circuits for positive and negative voltage swings. These circuits are not perfectly matched, so the positive and negative regions have somewhat different slopes and offsets, with a non-linear transition region slightly below zero output. Figure 1 shows an example of the curve.

This curve should be measured twice for each MMIC bias board: once for each MAB. This can most easily be done by injecting signal on the GNDSENSE input to the bias board. This signal should be referenced to the bias board ground, not the MAB ground. This ground should be one of the bias circuit grounds, labeled as A0–A9 on the MMIC bias board schematic. A connection can be obtained by soldering a lead to the bias board or by connecting to, e.g., pin Z15 of P2 (A0GND) for MAB A or pin Z18 of P2 (B0GND) for MAB B. To cover the entire range of the optoisolator, the input should be swept from -3 V to +3 V . Measuring with small step size is suggested, particularly in the transition region near zero output.

At JPL, we swept the input with a slow triangle wave and recorded the housekeeping output over several periods while recording the actual input voltage on another ADC channel. A table was constructed from these data. Samples are linearly interpolated between nearest neighbors to fill in missing points. This appears to be sufficient.

3.2 Gate Calibration

Figure 2 shows a simplified schematic of a single gate output. R1 and R2 are on the MMIC bias board, R_{lead} and R_{gnd} represent the FPC resistances, and R3 and R4 make up the voltage divider on the MAB. Table 1 lists the circuit component values. R_{lead} and R_{gnd} will depend on the test setup and should ideally be measured for each channel. R_{gnd} is common to an entire MAB.

Circuit analysis yields the following relationship between the housekeeping output and the actual gate

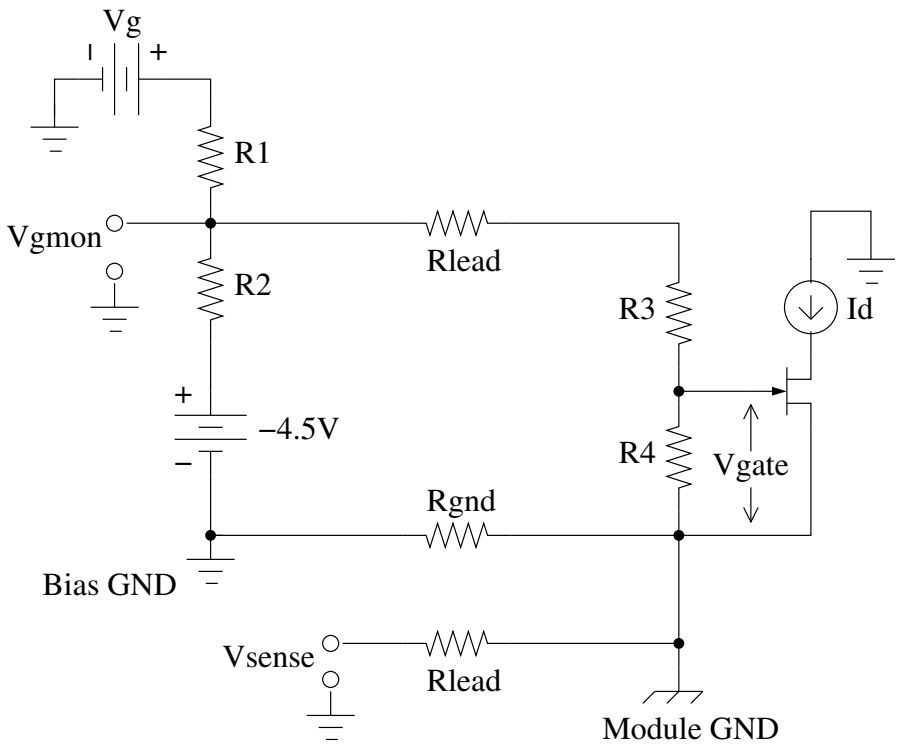

Figure 2: Schematic of a single gate control and housekeeping circuit, including lead resistances. V_g represents the DAC output that sets the gate voltage. The labeled −4.5 V power supply value is actually −4.56 V for W-band and −2.81 V for Q-band.

Resistance	**W Band**	**Q Band**
R1	499 Ω	1.00 kΩ
R2	3.74 kΩ	2.00 kΩ
R3	4.12 kΩ	499 Ω
R4	1.00 kΩ	1.00 kΩ

Table 1: Gate circuit resistor values for W- and Q-band. R_1 and R_2 values are quoted from JPL schematics. R_3 and R_4 values for W-band are from the March 1, 2007, Version 2 W-Band MAB schematic posted on the QUIET Wiki. For Q-band, these are from *Revised Q-band MAB Design*, R. Dumoulin, January 31, 2008.

Quantity	Value
$\mathbf{R_{ee}}$ (Drain 1)	24.9 Ω
$\mathbf{R_{ee}}$ (Drain 2)	11.3 Ω
$\mathbf{V_{supply}}$	2.5 V

Table 2: Drain circuit values for both W- and Q-bands. V_{supply} can be monitored on a housekeeping channel if desired.

voltage.

$$V_{gate} = G \times (V_{gmon} - V_{sense}) - \epsilon \times V_{sense} \ , \tag{1}$$

where

$$G = \frac{R_4}{R_{lead} + R_3 + R_4 + R_{gnd}} \ , \tag{2}$$

and

$$\epsilon = G \times \frac{R_{gnd}}{R_{lead} + R_3 + R_4 + R_1 \parallel R_2} \ . \tag{3}$$

Remember that both V_{gmon} and V_{sense} were measured through the optoisolator, so must be corrected using the inverse of the optoisolator transfer function described in Section 3.1. That is,

$$V_{gmon} = f_{opto}^{-1}(V_{gmon,measured}) \tag{4}$$
$$V_{sense} = f_{opto}^{-1}(V_{sense,measured}) \ . \tag{5}$$

Figure 3 shows the error in gate voltage measured using this approach. Each gate on one output was swept several times with varying MAB ground offsets injected with a floating lab power supply. The effect of ground offset has clearly been removed quite effectively.

3.3 Drain Calibration

Figure 4 shows a simplified drain circuit. Again, the optoisolator transfer function should be used to correct the measured housekeeping outputs.

$$V_{Idmon} = f_{opto}^{-1}(V_{Idmon,measured}) \tag{6}$$
$$V_{Vdmon} = f_{opto}^{-1}(V_{Vdmon,measured}) \ . \tag{7}$$

The drain current should be estimated first because it is needed to estimate the drain voltage. It can be found from the following equation.

$$I_{drain} = \frac{V_{supply} - V_{Idmon}}{\left(1 + {h_{FE}}^{-1}\right) \times R_{ee}} \ , \tag{8}$$

where h_{FE} is the current source transistor's forward current transfer ratio. For the 2N2907A resistor used in the revision 2 bias circuits, h_{FE} is at least 100, so the ${h_{FE}}^{-1}$ term can be safely ignored.

The drain voltage is simply V_{Vdmon} corrected for the voltage drops across R_{lead} and R_{gnd}. V_{sense} allows us to directly measure the R_{gnd} drop, and we know the current through R_{lead} is the same as the drain

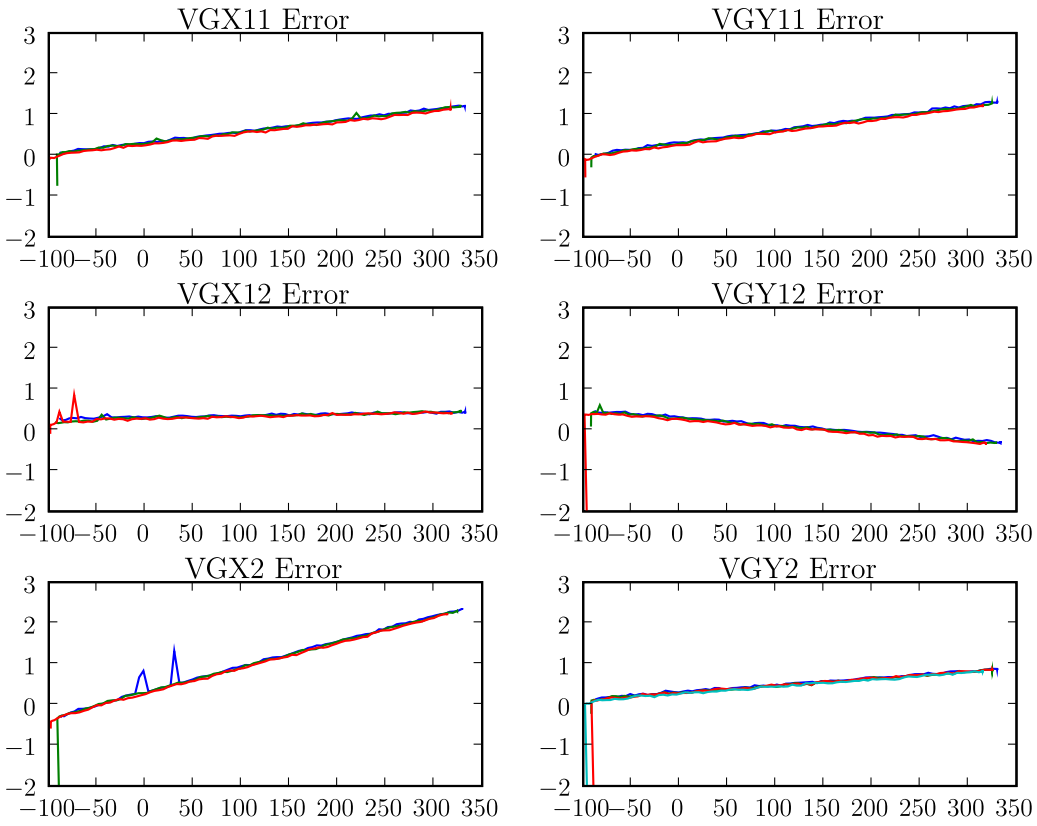

Figure 3: Gate measurement errors for a W-band single-board MAB with no module installed. The X axes indicate the actual gate voltage in mV. The Y axes indicate the error, also in mV. Each plot contains three sweeps, corresponding to different MAB ground voltage offsets. These were, approximately, 17 mV for blue, 50 mV for green, and 120 mV for red.

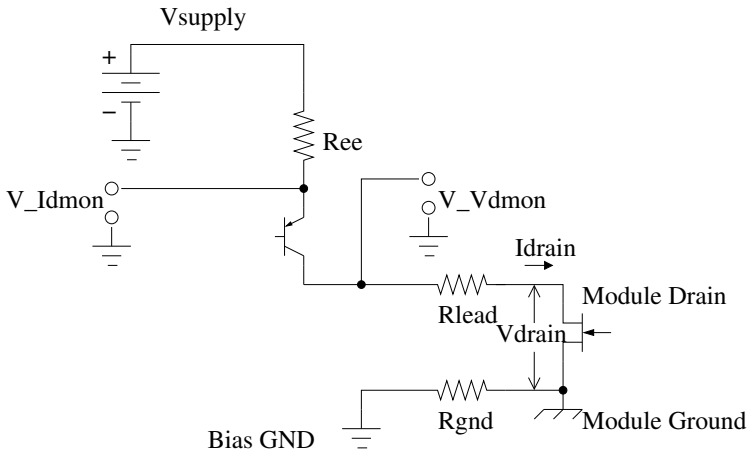

Figure 4: Simplified drain schematic. V_{Idmon} and V_{Vdmon} are the housekeeping outputs. The DAC output that controls the current is not shown; it sets the voltage at the base of the transistor shown here.

current, ignoring current through the protection diodes (not shown). For drain voltages below about 1 V , this is a reasonable assumption. The drain voltage estimate is then

$$V_{drain} = V_{Vdmon} - V_{sense} - I_{drain} \times R_{lead} \ . \tag{9}$$

A reasonable estimate of R_{lead} is necessary since significant current flows through the lead wire. However, since R_{lead} should only be 2 Ω according to the MMIC bias card specification, it should not be difficult to know this to within, say, 1 Ω . That would result in an error of 1 mV per mA of drain current, which should be acceptable.

Figure 5 shows measured errors between the actual drain current and voltages and the estimates from the housekeeping output. The large deviations at high drain voltage result from the protection diodes turning on. These cause overestimates of the drain current, which result in errors in the drain voltage as well. Several ground offset voltages were used in the test, which is why there are several different protection turn-on curves in each plot.

This test was run with a MMIC bias card connected to a warm MAB with protection via 2 Ω cryowires. Because of the drops in the cryowires, the voltage seen by the protection on the MMIC bias card is higher than the drain voltage. As a result, current through that protection will be much greater than that through the on-MAB protection. This will be even more true when the MAB's protection is cold.

Figure 6 plots the drain errors as a function of V_{Vdmon}. A single function of V_{Vdmon} can describe these error curves quite well, which confirms the expectation that the MMIC board protection is the dominant current "thief."

If operation above a drain voltage of 1 V is necessary, it may be necessary to correct for the current lost to the protection circuits. Fortunately, this should be straightforward as long as all warm protection circuits are connected to the MMIC bias card via low impedance wires. It should apparently be possible to reach a few mA accuracy on the drain current using a single correction curve with V_{Vdmon} as its input.

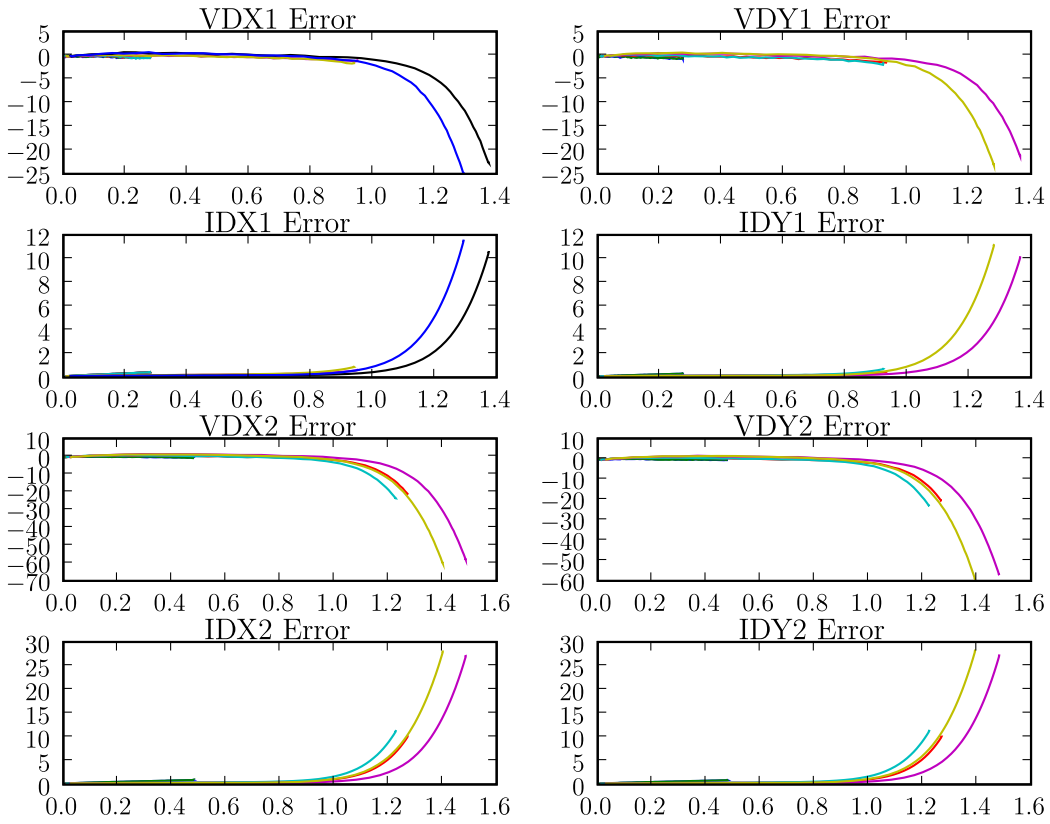

Figure 5: Drain voltage and current estimate errors as a function of actual drain voltage. Voltages are plotted in mV, currents in mA, against drain voltage in Volts. Drains were simulated with a resistor to ground. Resistors of 10.0 Ω , 34.8 Ω , and 81.2 Ω were used. Additionally, the MAB ground was offset by several values between 3 and 124 mV for each resistor value.

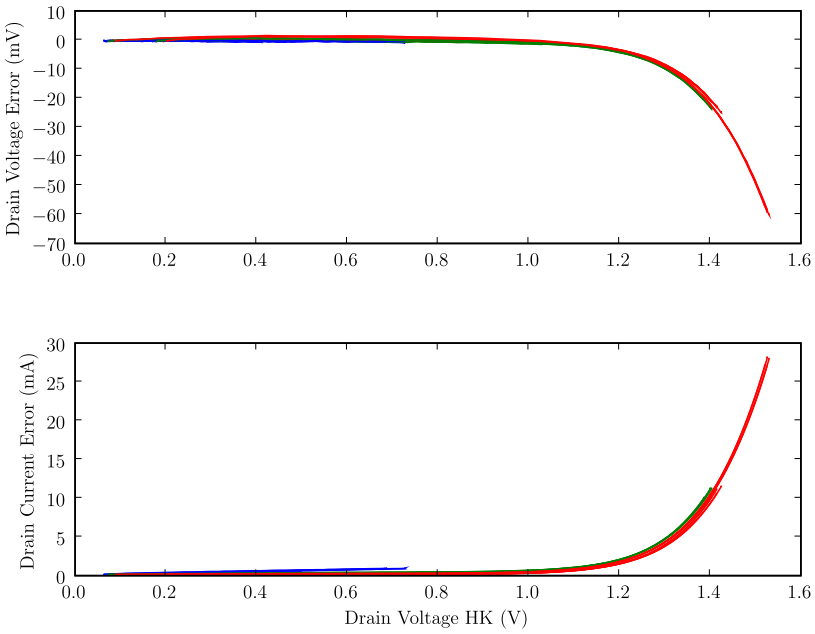

Figure 6: Same data from Figure 5, now plotted versus the housekeeping drain voltage monitor, V_{Vdmon}. It is seen that the deviation, caused by diode protection stealing current from the drains, is well-described as a single function of V_{Vmon}. This suggests that the protection on the MMIC bias board is primarily responsible for the "missing" drain current, as is expected.

D.4 Receiver Characterization

In this section we present four memoranda describing detailed study of QUIET polarimeter module performance, with emphasis on study of the $1/f$ noise produced by the QUIET polarimeter modules. These memoranda discuss measurement procedures that were carried out using a test apparatus including a cryostat, bias electronics, and digital readout electronics that I assembled with assistance from Dr. Michael D. Seiffert. Components for this test apparatus included a prototype QUIET bias electronics board set with a heavily modified interface to connect to a single computer workstation running a custom bias board control interface program I developed. A preamplifier I constructed, based on the design described in section D.5, was used for some tests. Existing preamplifiers and a manually controlled bias electronics box were used for some tests. In some cases an existing data acquisition and bias control system were used rather than one of my design and construction. The memoranda describe work that was primarily carried out by me, with guidance and consultation with Dr. Seiffert, particularly with regard to the measurement design and method. These memoranda are presented with their original formatting and pagination.

Module Noise Performance Draft Report

W-Band Module 9
Revision 1.0

Joey Richards*and Mike Seiffert†

February 23, 2006

1 Introduction

This draft report describes the results of noise performance testing on QUIET W-band receiver module 9. Demodulated time series data were collected and analyzed to characterize the performance of the module. We have measured the noise-effective bandwidth, the 1/f noise knee frequency, and we have begun to investigate processing techniques for correcting imbalances in the RF system. In these tests, the module was at room temperature with its RF inputs unconnected, looking into the room.

From our noise measurements, we find the module's noise-effective bandwidth to be 12–13 GHz, compared to an estimated 16–17 GHz calculated from its gain curve.

We estimated the 1/f noise knee frequency for several configurations. For a single diode with no differencing, the knee is at about 30 Hz. For a single diode differenced between phase switch states, the knee drops to about 300 mHz. For a weighted double-difference between two diodes, we find a knee frequency of about 100 mHz.

Weighting the two phase states for each diode substantially lowers the 1/f noise knee frequency. With a weighted difference, single-diode knee frequencies are reduced to about 100 mHz. The doubly-weighted double-difference 1/f knee is at about 50 mHz.

2 Data Acquisition System

Time series data were collected from two of the four detector diodes on the module. These were the two diodes corresponding to Q polarization measurements, designated diode 1 (D1) and diode 4 (D4). These data were collected using a data acquisition system originally designed for the Planck project.

2.1 Data Acquisition System

The data acquisition system consists of a low-noise preamp for each diode, followed by an analog integration circuit connected to an ADC in a PC. The PC stores digitized data from the integrator and controls the phase switch clock.

Each preamp consists of an Analog Devices OP37 low-noise op amp gain stage (voltage gain of 100), followed by an OP27 unity gain buffer. The preamp input is DC-coupled directly to the detector diode. No DC bias is provided to the diode for room-temperature operation.

The integrator circuit is based on a Burr-Brown (now TI) IVC102 integrator chip. The circuit integrates the two phase switch states separately. The phase switch clock is driven at 4096 kHz with 50% duty cycle. A blanking time of approximately 10 microseconds is inserted at each phase switch clock transition. Following

*joey@caltech.edu
†Michael.D.Seiffert@jpl.nasa.gov

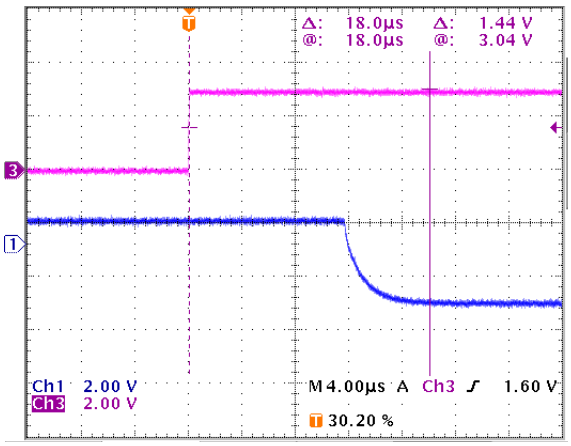

Figure 1: Rising edge of phase switch clock (top trace) followed by falling voltage on phase switch diode (bottom trace). Delay from rising edge to stable bias is approximately 18 microseconds.

this blanking interval, the active integration channel is reset and the preamp output is then integrated, held, and digitized by the data acquisition PC.

The integrator provides additional voltage gain of about 10. Additionally, each integration channel has a DC offset trim pot. Before each data run, this is adjusted to zero the DC level of each channel.

2.2 Module Bias

Module bias is set using a pair of QUIET bias cards, one for the MMICs and one for the phase switches. These are set using a Visual Basic application running on the data acquisition PC. For these tests, the bias levels were set at the start of data collection and were not monitored or adjusted during operation.

The MMIC bias values were set to approximate the recommended room-temperature bias settings provided by Todd Gaier.[1] The phase switch bias was adjusted to balance the DC detector output between the two phase states as much as possible.

2.3 Known Issues

There are several known problems with the data acquisition process used in tests so far. The effect of these problems is to reduce (worsen) the figures of merit we present in this report. The results presented here are thus lower bounds on the actual performance of the module.

2.3.1 Integrator Blanking Incorrectly Positioned

The 10 microsecond blanking interval of the integrator circuit is positioned to coincide with the phase switch clock transition. However, due to slow optoisolators on the phase switch bias card, there is a long (approximately 20 microsecond) delay between this transition and the actual phase switch event. As a result, the blanking interval probably misses the phase switch interval entirely. Figures 1 and 2 show this delay.

2.3.2 DC Level Underestimated

In some of the calculations that follow, the ratio $\frac{\Delta V}{V}$ is used to estimate the noise-effective bandwidth. Because of the incorrect position of the blanking interval, the DC voltage is underestimated by our acquisition system. This results in an underestimate of the noise-effective bandwidth.

[1] Data from *module9_02-06-2006_300K-2.xls*.

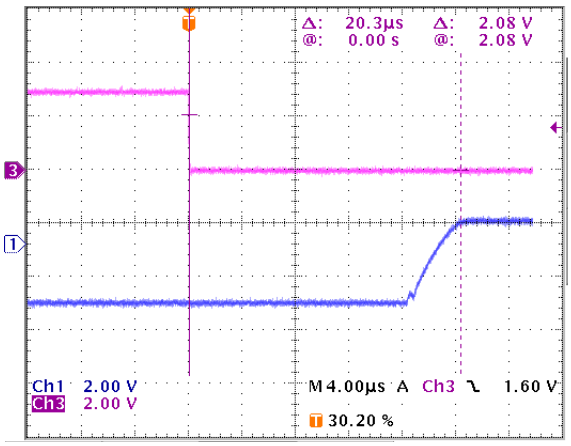

Figure 2: Falling edge of phase switch clock (top trace) followed by rising voltage on phase switch diode (bottom trace). Delay from falling edge to stable bias is approximately 20 microseconds.

2.3.3 Open-Loop Biasing

We have not connected to the housekeeping outputs of the bias cards or otherwise instrumented the bias circuits to monitor the bias settings during operation. Although care was taken to achieve the recommended bias settings for the module, it is possible that these settings have drifted with temperature.

Also, the recommended drain bias was specified in terms of both current and voltage settings. We biased the drains by first setting the gate voltage, then adjusting the bias current until the drain voltage reached the recommended level. We did not directly measure the drain current, however.

3 Noise-Effective Bandwidth

The sensitivity of the QUIET radiometer is given by

$$\frac{\Delta T}{T} = \frac{\Delta V}{V} = \frac{2}{\sqrt{\beta \tau}} \tag{1}$$

where τ is the integration time. β is the noise-effective bandwidth (also known as the reception bandwidth) of the receiver. The noise-effective bandwidth can also be estimated from the receiver's gain curve,

$$\beta = \frac{\left[\int G(f)\, df\right]^2}{\int \left[G(f)\right]^2 \, df} \tag{2}$$

where $G(f)$ is the gain of the receiver at frequency f.

We calculated the noise-effective bandwidth both from noise measurements and from a measured gain curve of the module. The results are in Table 1 and notes on the calculation methods follow.

3.1 Bandwidth from Noise Measurements

To calculate the noise-effective bandwidth from noise measurements, Equation 1 was solved for β, giving

$$\beta = \frac{4}{\tau \left(\frac{\Delta V}{V}\right)^2}. \tag{3}$$

Diode	$\beta_{\mathbf{noise}}$	$\beta_{\mathbf{gain}}$
D1	13.2 GHz	17.4 GHz
D4	12.1 GHz	15.9 GHz

Table 1: Noise-effective bandwidths. β_{noise} was calculated from noise data, β_{gain} was calculated from the gain curve.

Our data acquisition system reports the DC voltage level measured on the detector diode.[2] This was used for V in the equation. For ΔV, we computed the noise spectrum and used the white noise amplitude in $\text{nV}/\sqrt{\text{Hz}}$. Because of this choice of units, our effective total integration period is 1 second. This is divided into two phase states, so this is divided by two. Finally, approximately 10% of the integration time is lost to blanking during the phase state transition, so we use $\tau = 0.45$ seconds in our calculations.[3]. The data used for these calculations is given in Table 2.

Diode	**V**	**ΔV**	τ
D1	0.709 mV	18.4 nV	0.45 sec
D4	0.947 mV	25.7 nV	0.45 sec

Table 2: Data for noise-effective bandwidth calculation.

3.2 Bandwidth from Gain Curves

The gain curves for the diodes are shown in Figures 3 and 4. These were used to approximate the integrals in Equation 2. Only the region between 83 GHz and 103 GHz was included in the integral. For calculation, the gain curves were approximated by piecewise linear functions connecting local extrema.

4 Noise Data

The following analysis used two noise data sets. Each data set represents approximately 15 minutes of data.

For the first data set (*6_2_17_11_38.dat*), the data acquisition system was connected to D1 and D4 of module 9. The module was biased to the recommended room-temperature settings. The module's RF inputs were unconnected, looking into the room. This data set was used for most of the measurements described below.

For the second data set (*6_2_17_14_47.dat*), the data acquisition system was again connected to D1 and D4 of module 9. In this case, the MMIC and phase switch biases were set to 0. This data set was used for the baseline spectrum computations only.

4.1 Data Acquisition Baseline

Figures 5 through 7 were computed from input-free data and show the baseline noise characteristics of the data acquisition system. The single-difference noise floor of each input channel is at approximately 5 $\text{nV}/\sqrt{\text{Hz}}$.

Figure 7 illustrates that the noise on these channels is at least reasonably uncorrelated. Differencing the two input streams results in a signal with almost exactly $\sqrt{2}$ times the individual channel white noise levels.

[2]This is probably an underestimate, see Section 2.3.2.

[3]The actual integration time is actually probably a bit lower than this, see Section 2.3.1.

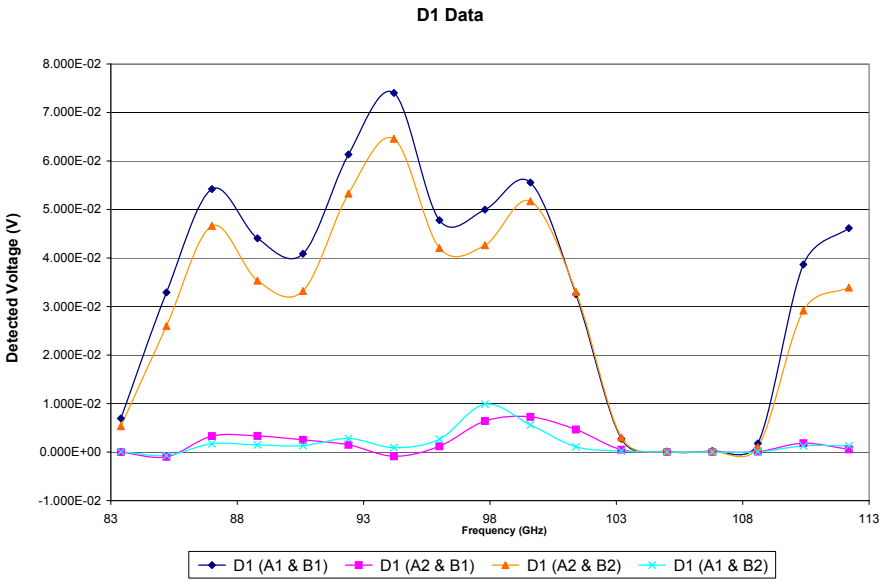

Figure 3: Gain curve for D1. Relevant curve is the dark blue plot, "A1 & B1" in the legend.

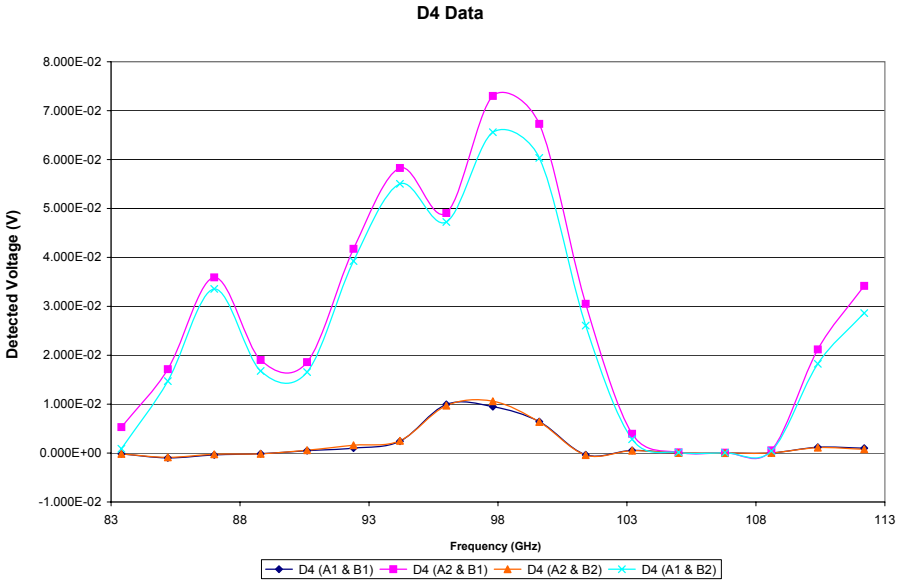

Figure 4: Gain curve for D4. Relevant curve is the light blue plot, "A1 & B2" in the legend.

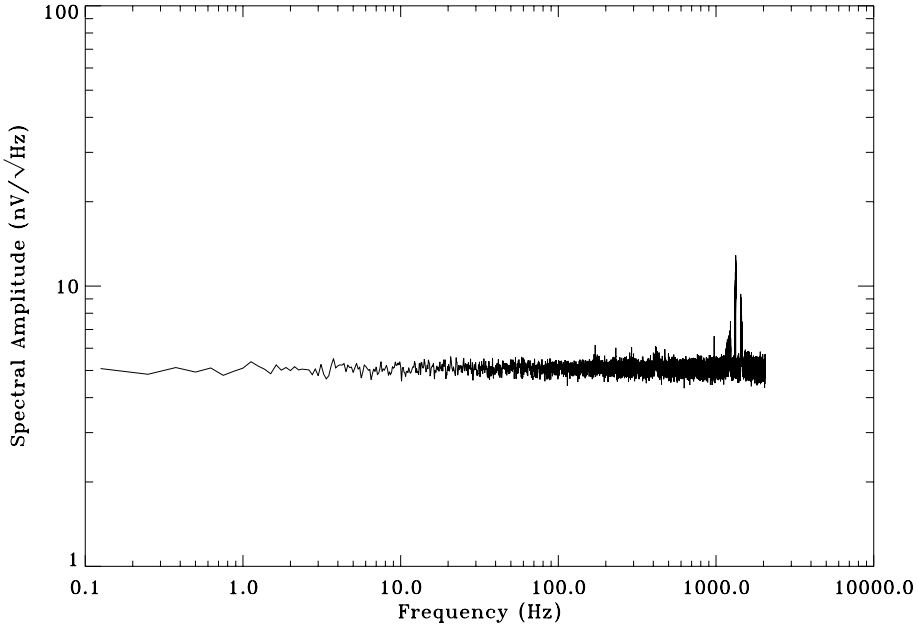

Figure 5: Channel A data acquisition system noise baseline. Input connected to D1 of unpowered module. Computed from difference of "plus" and "minus" phase states. White noise at 5.2 nV/sqrtHz.

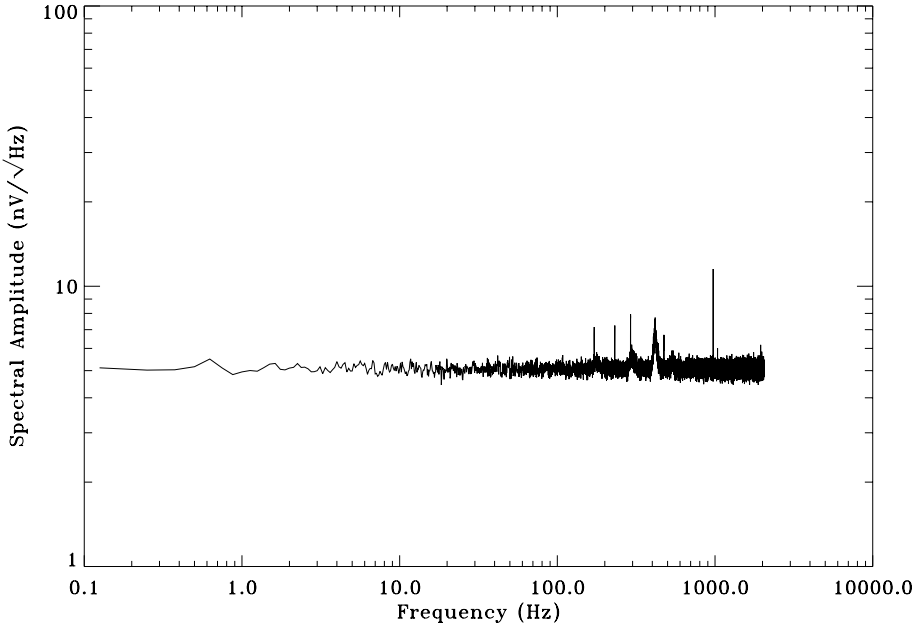

Figure 6: Channel B data acquisition system noise baseline. Input connected to D4 of unpowered module. Computed from difference of "plus" and "minus" phase states. White noise at 5.1 nV/sqrtHz.

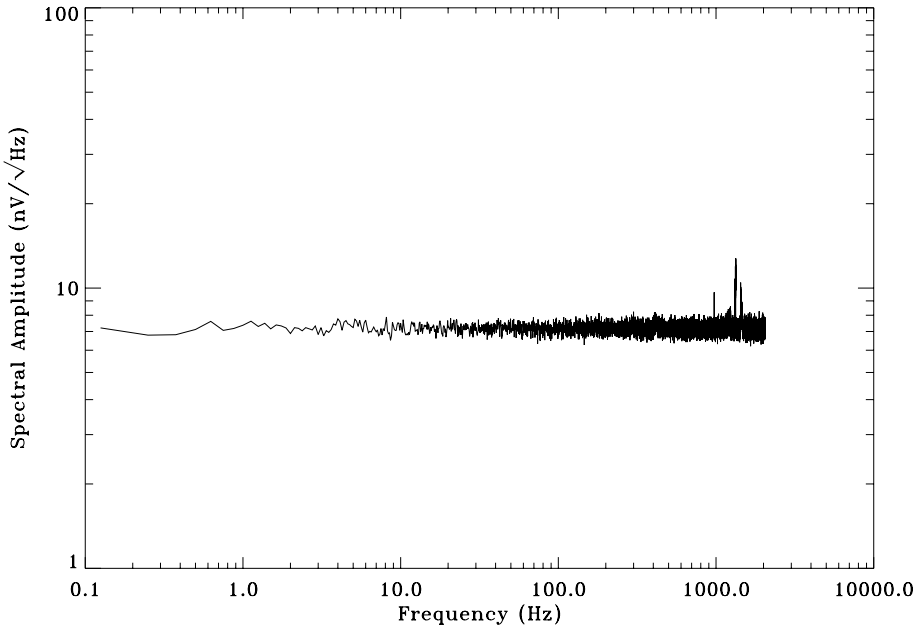

Figure 7: A-B difference noise spectrum. White noise level increased by approximately $\sqrt{2}$ from individual channels, indicating that the channel noise is uncorrelated. White noise at 7.2 nV/sqrtHz.

4.2 Time Series Comparison

Figure 8 shows the effect of single-differencing on the data streams. The D1 data shown are not significantly different from the D4 data, so only one plot is included.

4.3 Single Detector Diode Noise Spectra

Figures 9 and 10 show the noise spectra collected in the "plus" phase state on each diode. These did not differ significantly from the "minus" phase state spectra.

4.4 Single-Difference Noise Spectra

The spectra in Figures 11 and 12 were computed from singly-differenced data. No weighting was done to correct imbalances in the phase switch states.

Figures 13 to 16 show single differences between the two diodes. Weighting did not substantially improve the noise level or knee frequencies, so was not used in these plots.

4.5 Double-Difference Noise Spectra

Figure 17 was computed from a double-difference between diodes D1 and D4. A single data set for each diode was computed by differencing the two phase states. These data were then differenced between D1 and D4 to produce the spectrum.

Because the signal levels from the two diodes were not identical, the signal from each diode was weighted prior to combining into the difference signal. It seems that inversely weighting each diode by its mean white noise level yields the best combination, both in terms of double difference white noise level and 1/f knee frequency. Inversely weighting by the DC diode level gives an almost identical result.

No weighting was applied to the individual phase state signals on each diode. Only the second difference in the double-difference was weighted.

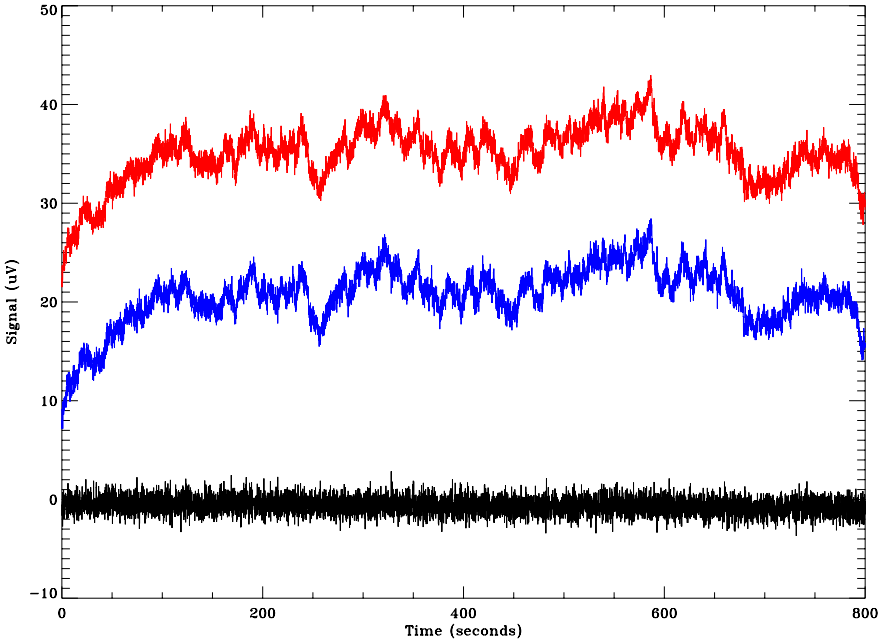

Figure 8: Comparison of time series data from D1 in the two phase states. Also shown is the single-difference data, demonstrating the reduction in 1/f noise. Red (top) is the "plus" phase state, blue (middle) is the "minus" phase state, and black (bottom) is the difference. The data are plotted with arbitrary DC offset and have been downsampled to 10 Hz.

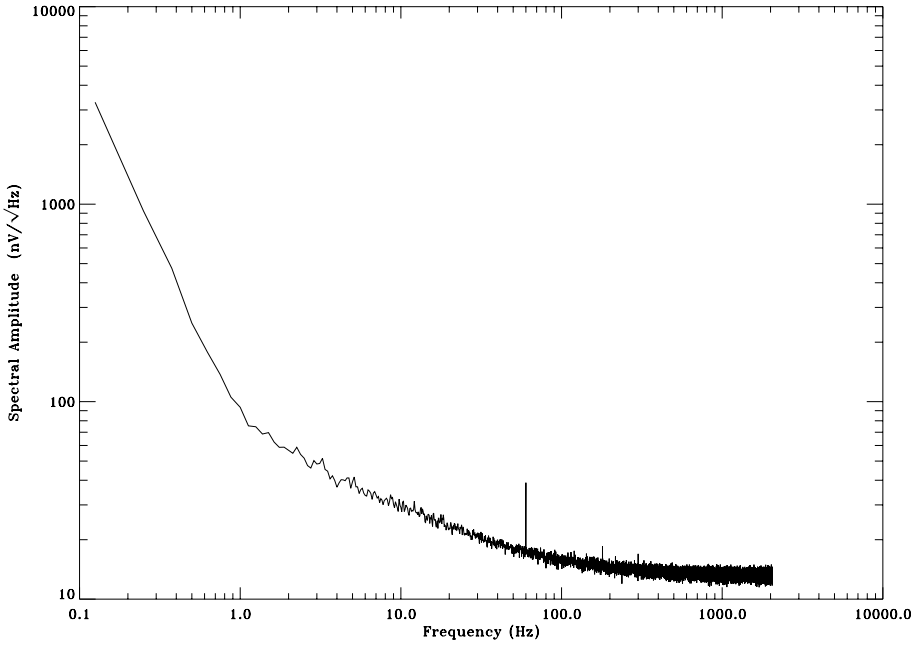

Figure 9: Noise spectrum of D1 "plus" state with no differencing. "Minus" state spectrum is essentially identical. Spectrum computed with resolution of 100 mHz.

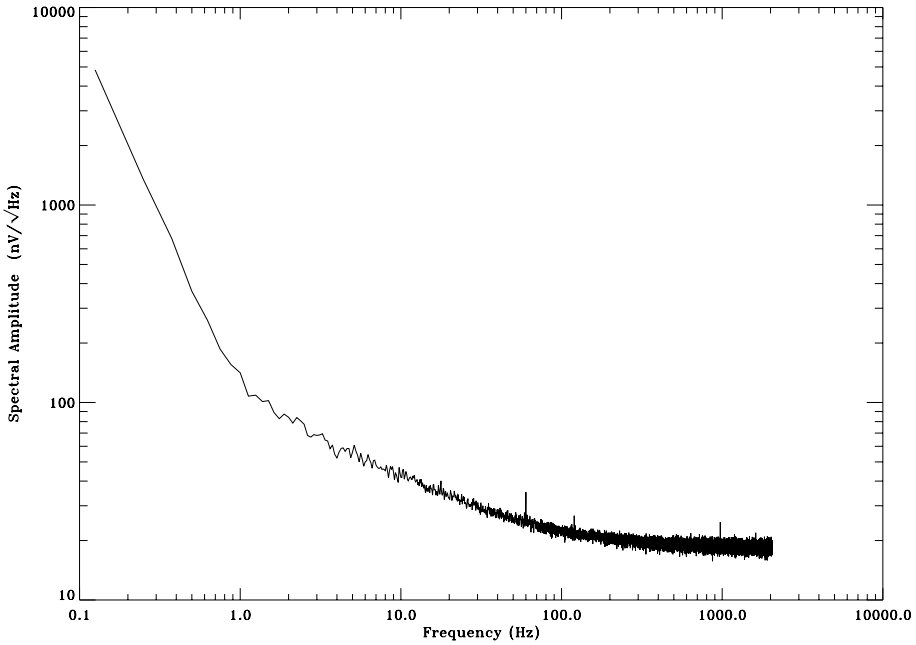

Figure 10: Noise spectrum of D4 "plus" state with no differencing. "Minus" state spectrum is essentially identical. Spectrum computed with resolution of 100 mHz.

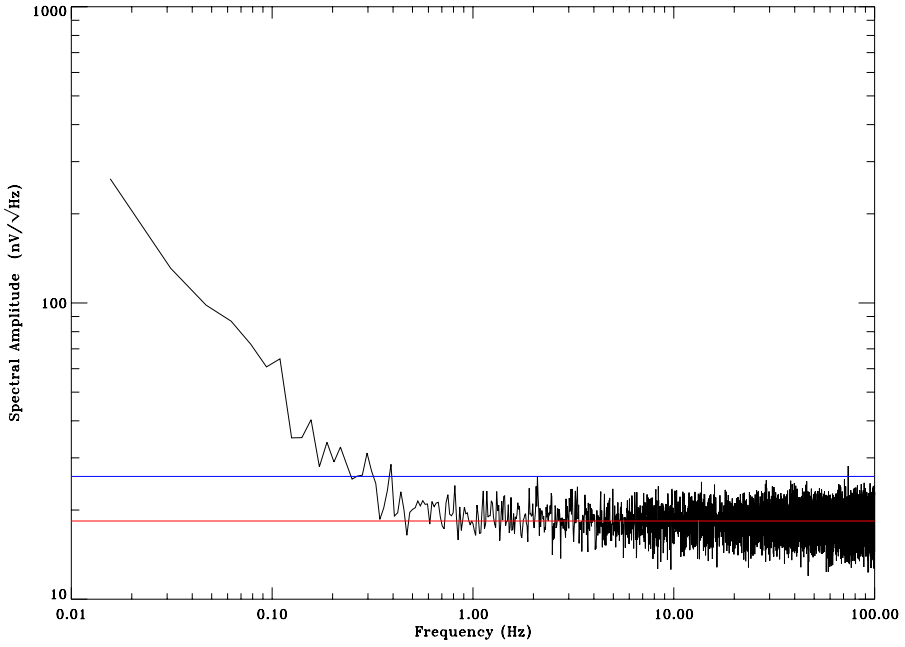

Figure 11: Single-difference noise spectrum from D1, "plus" state minus "minus" state. White noise level is 18.4 nV/$\sqrt{\text{Hz}}$, indicated by red (lower) horizontal line. Blue (upper) horizontal line indicates $\sqrt{2}$ times the white noise level. Spectrum computed with resolution of 10 mHz.

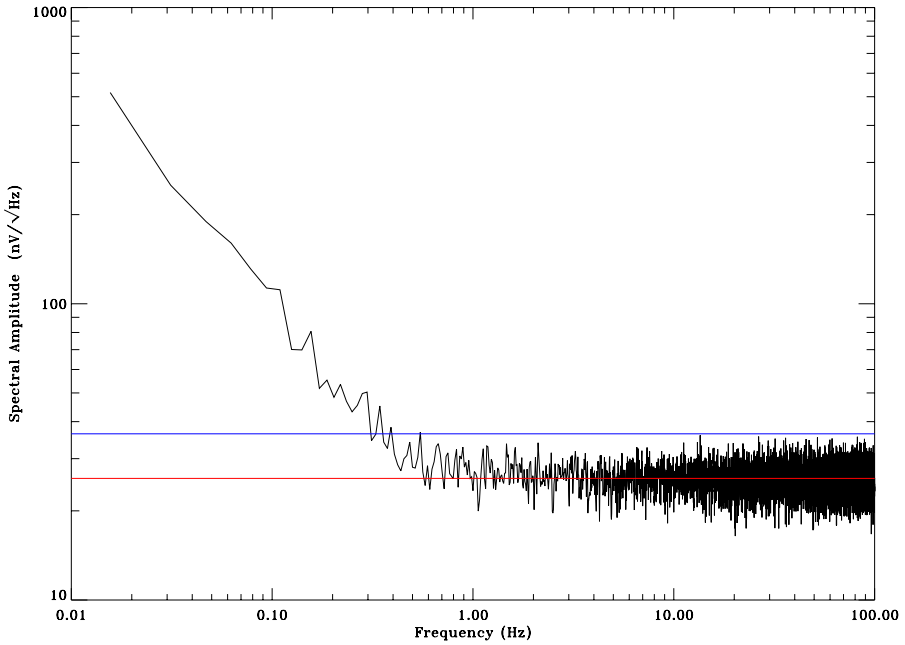

Figure 12: Single-difference noise spectrum from D4, "plus" state minus "minus" state. White noise level is 25.7 nV/$\sqrt{\text{Hz}}$, indicated by red (lower) horizontal line. Blue (upper) horizontal line indicates $\sqrt{2}$ times the white noise level. Spectrum computed with resolution of 10 mHz.

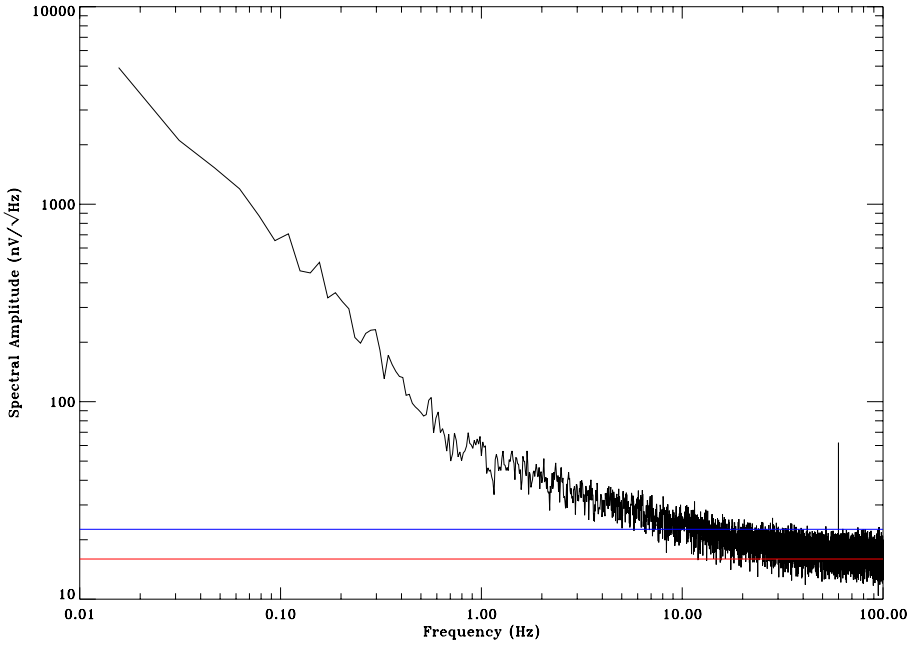

Figure 13: Single-difference noise spectrum of D1 "plus" state minus D4 "plus" state. White noise level indicated by red (lower) horizontal line. Blue (upper) horizontal line indicates $\sqrt{2}$ times the white noise level. Spectrum computed with resolution of 10 mHz.

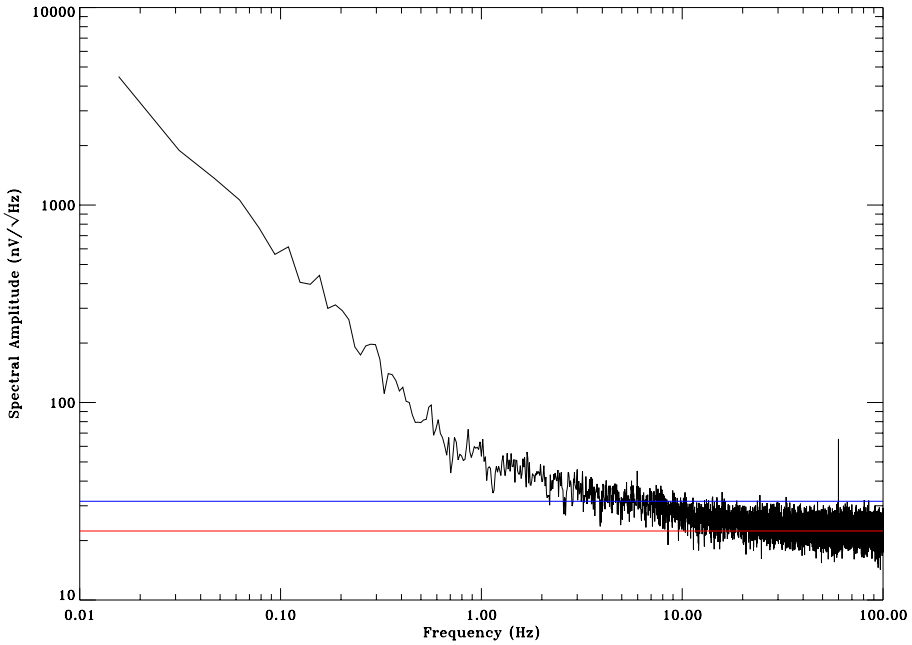

Figure 14: Single-difference noise spectrum of D1 "plus" state minus D4 "minus" state. White noise level indicated by red (lower) horizontal line. Blue (upper) horizontal line indicates $\sqrt{2}$ times the white noise level. Spectrum computed with resolution of 10 mHz.

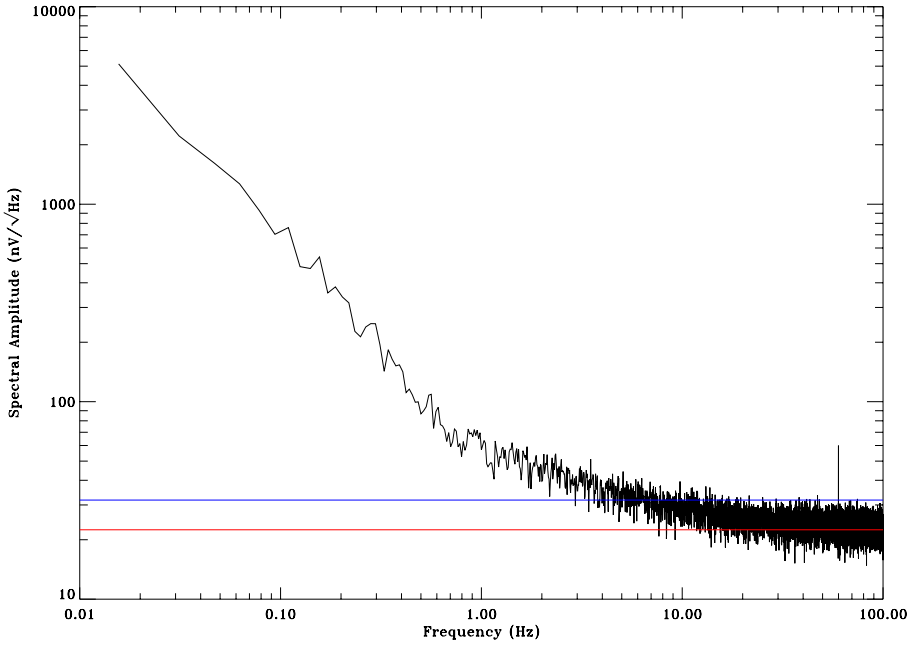

Figure 15: Single-difference noise spectrum of D1 "minus" state minus D4 "plus" state. White noise level indicated by red (lower) horizontal line. Blue (upper) horizontal line indicates $\sqrt{2}$ times the white noise level. Spectrum computed with resolution of 10 mHz.

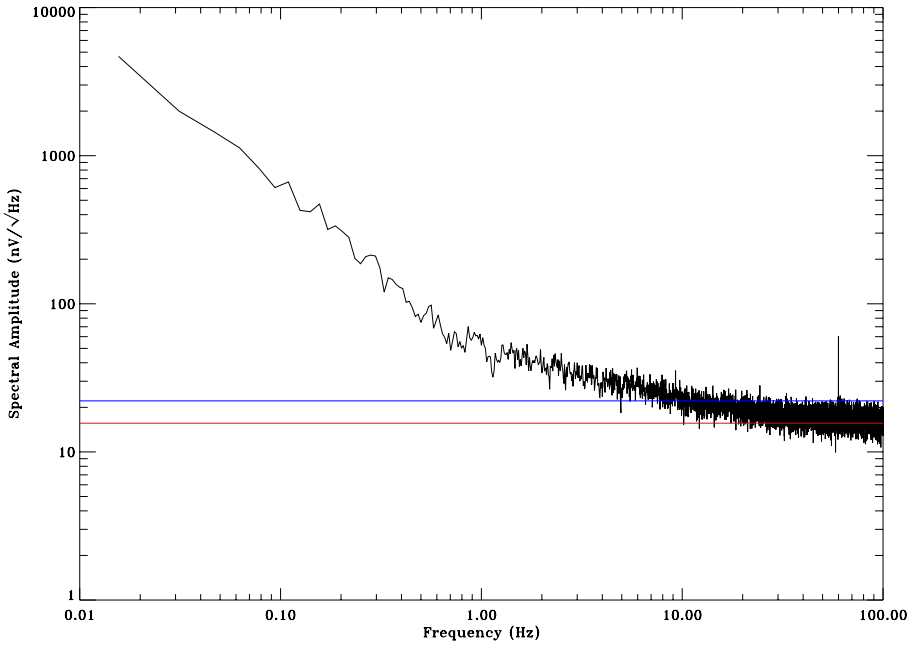

Figure 16: Single-difference noise spectrum of D1 "minus" state minus D4 "minus" state. White noise level indicated by red (lower) horizontal line. Blue (upper) horizontal line indicates $\sqrt{2}$ times the white noise level. Spectrum computed with resolution of 10 mHz.

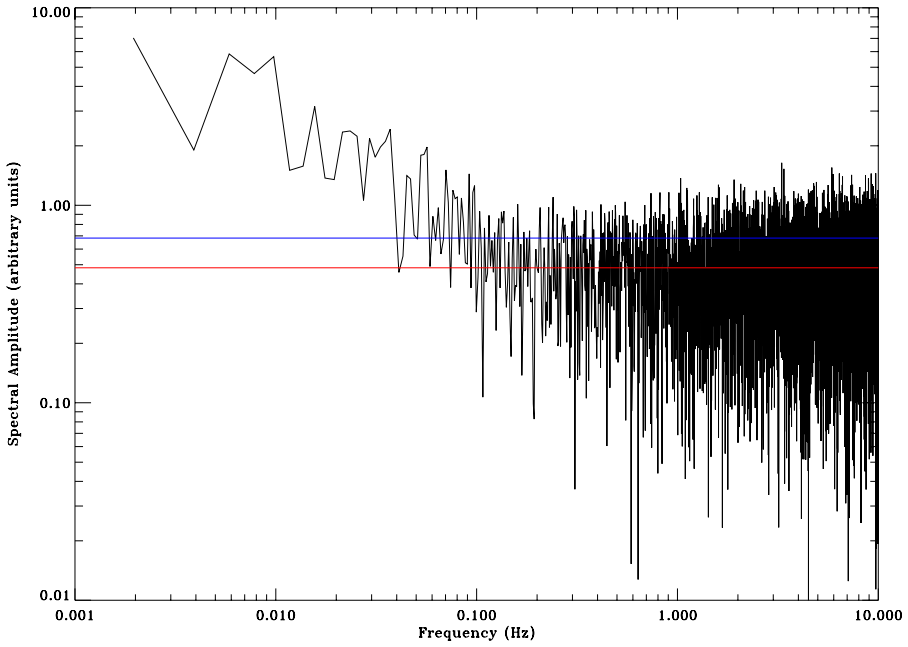

Figure 17: Double-difference noise spectrum between D1 and D4. Diode signals inversely weighted by single-difference white noise level. Y-axis scale in arbitrary units (proportional to $V/\sqrt{Hz}$). White noise level indicated by red (lower) horizontal line. Blue (upper) horizontal line indicates $\sqrt{2}$ times the white noise level. Spectrum computed with resolution of 1 mHz.

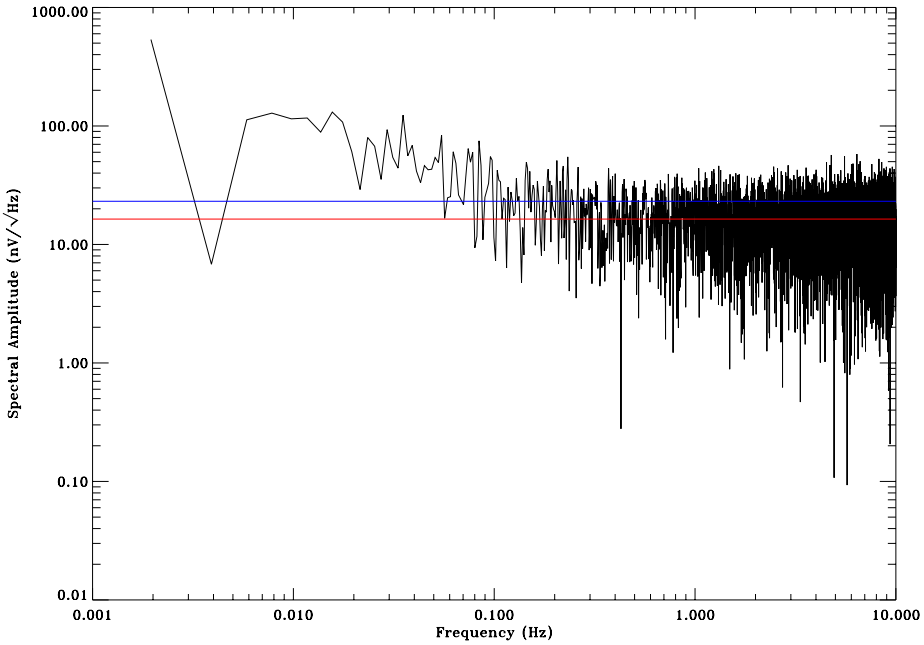

Figure 18: Weighted D1 single-difference noise spectrum. Weights adjusted to minimize white noise level and knee frequency. White noise level indicated by red (lower) horizontal line. Blue (upper) horizontal line indicates $\sqrt{2}$ times the white noise level. Spectrum computed with resolution of 1 mHz.

4.6 Weighted Phase State Differencing

By properly weighting the two phase states in the single difference, a substantial reduction in the 1/f noise knee frequency can be achieved. Figure 18 shows the noise spectrum for a weighted single-difference using the same data used in Figure 11. The weighting was adjusted to minimize the resulting white noise level — this coincided with minimizing the knee frequency. In this case, the D1 "plus" phase state was weighted by 0.9894 and the "minus" phase state was weighted by 1.0106. The 1/f knee frequency has been reduced from about 300 mHz to about 100 mHz compared to the unweighted spectrum. A similar effect was seen on the D4 diode.

When combined into a double-difference, the final knee frequency is also reduced. Figure 19 shows that the knee frequency has dropped from about 100 mHz in Figure 17 to about 50 mHz.

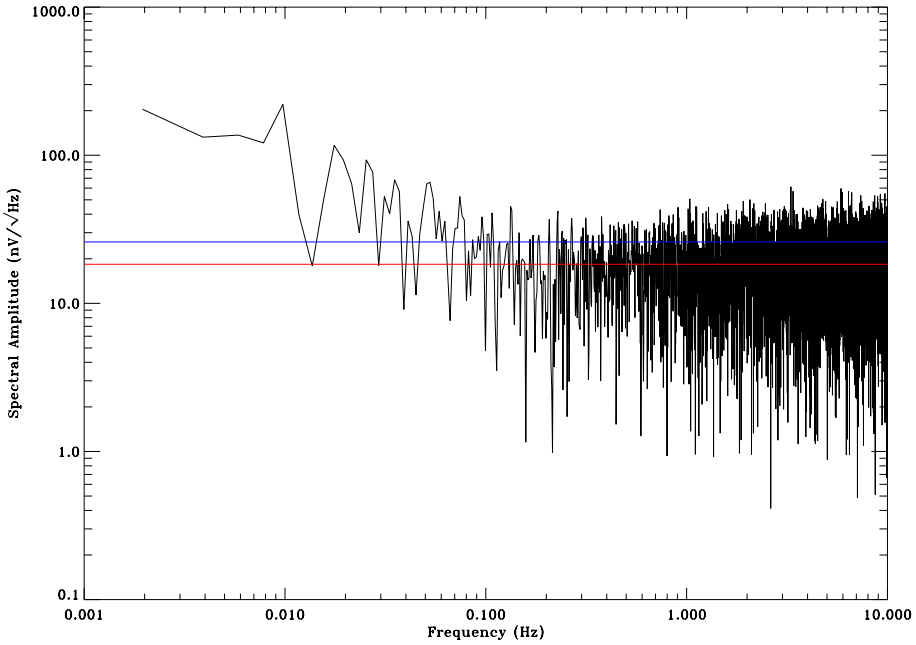

Figure 19: Double-difference noise spectrum between D1 and D4 using weighted single-differences. White noise level indicated by red (lower) horizontal line. Blue (upper) horizontal line indicates $\sqrt{2}$ times the white noise level. Spectrum computed with resolution of 1 mHz.

Module Noise Performance Draft Report

W-Band Module 9 Part 2
Revision 0.1

Joey Richards*and Mike Seiffert†

March 15, 2006

1 Introduction

This is an addendum to the previous W-Band Module 9 draft report. These documents remain a work in progress and will ultimately be reduced to a single report. For now, however, the new data are presented separately here.

2 Data Acquisition System

Since the previous report, we have made two changes to the hardware configuration. First, we have bypassed the optoisolators on the phase switch bias card. Figure 1 shows the alignment of the blanking period with the phase switch transition. Second, we have connected the 3.3 V power supply to the MMIC bias card to the same supply as the digital sections of the bias cards and increased that voltage to 3.6 V. This is a temporary measure that was necessary because of a failing power supply. This deviation from the final power supply configuration should not affect the tests described here.

We are in the process of characterizing the data acquisition system. At present, we are working to confirm the calibration used to generate absolute amplitude measurements. Measurements appear reasonable in order of magnitude, but there are some outstanding questions of $\sqrt{2}$ factors that must be resolved. This does not affect knee frequency measurements, but does affect the reception bandwidth calculation.

3 Bare Module

3.1 Corrected Blanking Period

Figure 2 through Figure 4 show the effect of correcting the blanking period. It appears that there is very little change in the undifferenced data, but both white noise and $1/f$ noise are reduced slightly in the single difference data.

3.2 Spectra and Knee Frequencies

Figure 5 through Figure 11 show spectra computed from a 19-hour run with a nominally-biased room temperature module with nothing connected to its RF input ports. The data acquisition system was used in "34x Compression" mode to reduce the size of the data set. In this mode, 34 consecutive samples are averaged and represented by a single point in the data set (each phase state is averaged separately). The resulting Nyquist frequency is just over 60 Hz.

*joey@caltech.edu
†Michael.D.Seiffert@jpl.nasa.gov

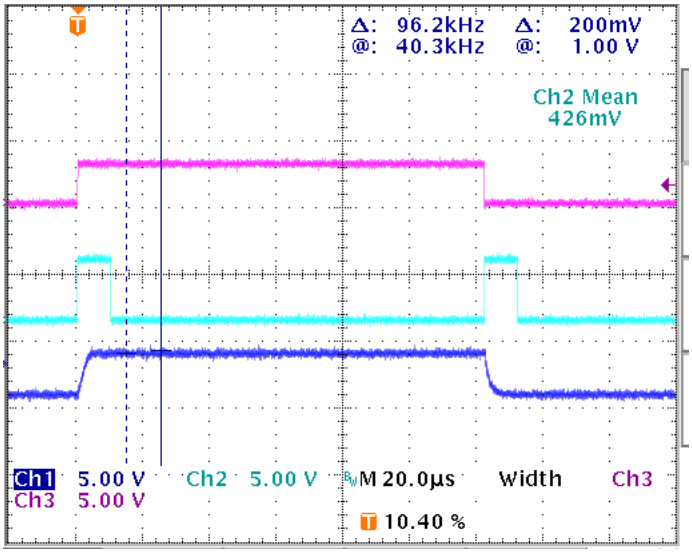

Figure 1: This scope trace shows the alignment of the blanking period with the actual phase switch transition. The magenta (top) trace is the phase switch clock, the cyan (middle) trace is the blanking pulse, and the blue (bottom) trace is the voltage on the phase switch diode. The integrator is held in reset when the blanking pulse is high. The phase switch transition is clearly within the blanking period.

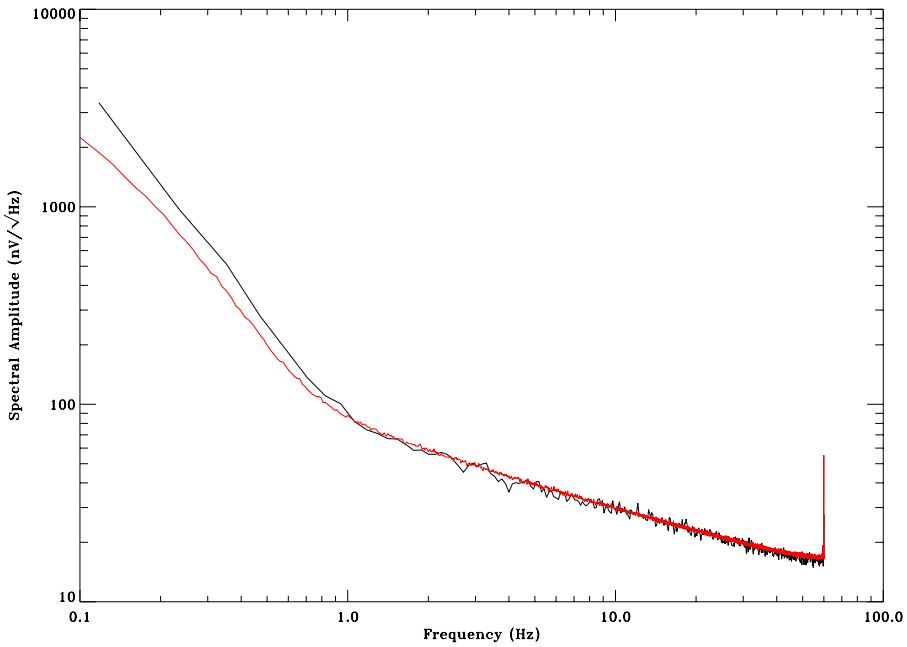

Figure 2: This plot compares the D1 "plus" phase state spectrum before and after the blanking fix. The black (noisier) plot is the "before" and the red (smoother) plot is the "after." The reduced variance on the spectrum is because the "after" data set is substantially longer.

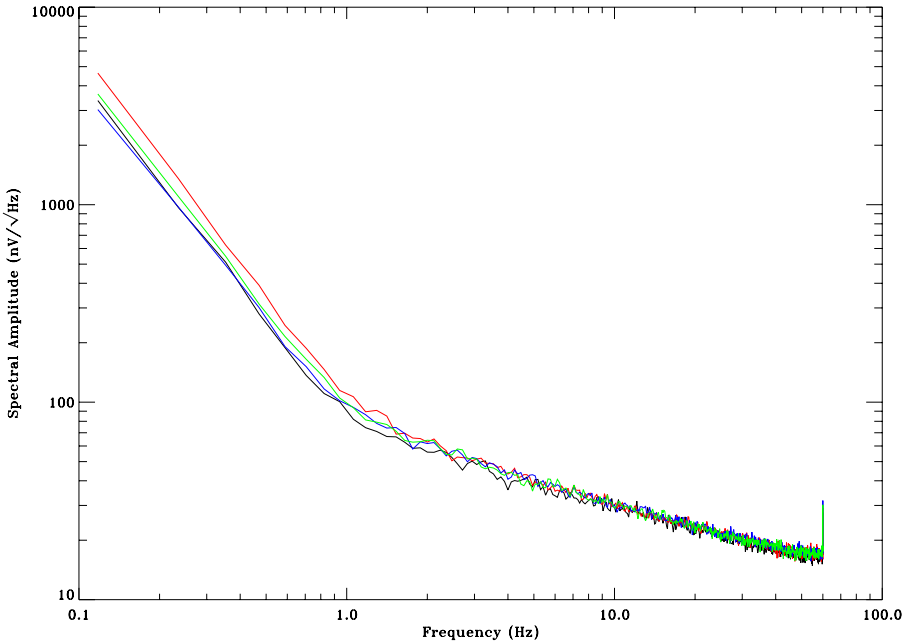

Figure 3: This plot shows the same "before" spectrum as in Figure 2 (black plot) along with three spectra computed using subsets of the same "after" data set (red, green, and blue plots). All data sets were the same length and were smoothed by a 34-point boxcar integration and downsampled by a factor of 34. This plot illustrates the similarity of the "before" and "after" undifferenced data.

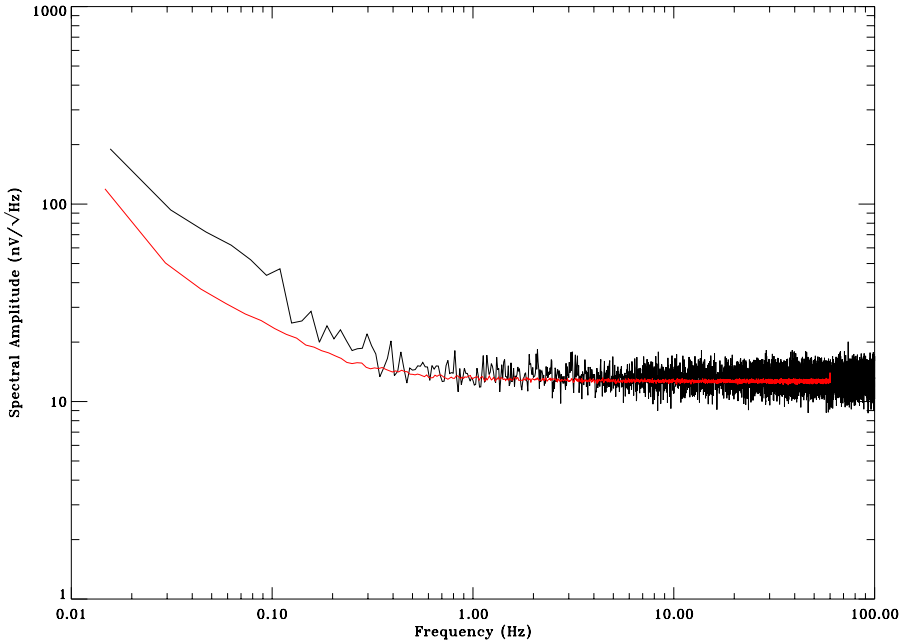

Figure 4: This plot shows the unweighted single-difference data from D1 for the same "before" and "after" data. The white noise level is about 5% lower for the "after" (red, smooth) spectrum and it appears that the $1/f$ noise level is somewhat lower as well.

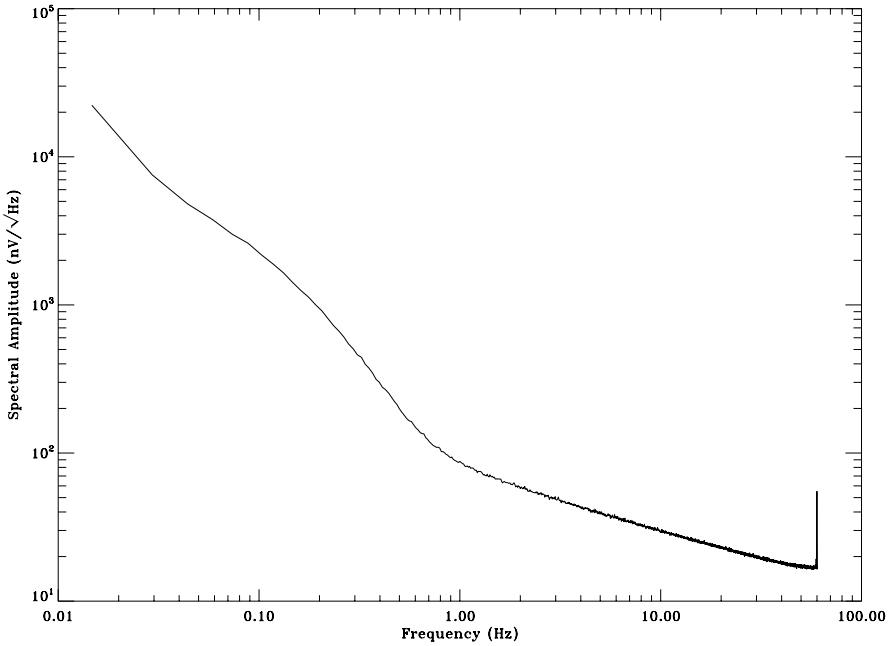

Figure 5: Undifferenced D1 "plus" state spectrum.

To reduce the size of this document, the spectra have generally been plotted over a limited frequency range. No features have been hidden by this. In most spectra, a horizontal red bar indicates the white noise level and a horizontal blue bar indicates $\sqrt{2}$ times the white noise level, the amplitude at the $1/f$ knee frequency.

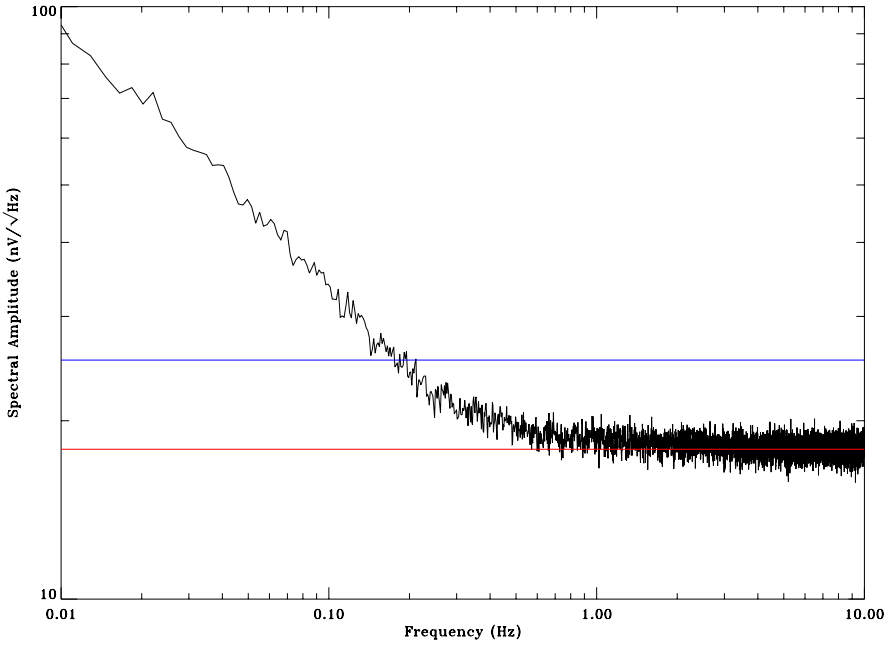

Figure 6: D1 unweighted single difference spectrum. $1/f$ knee at about 180 mHz.

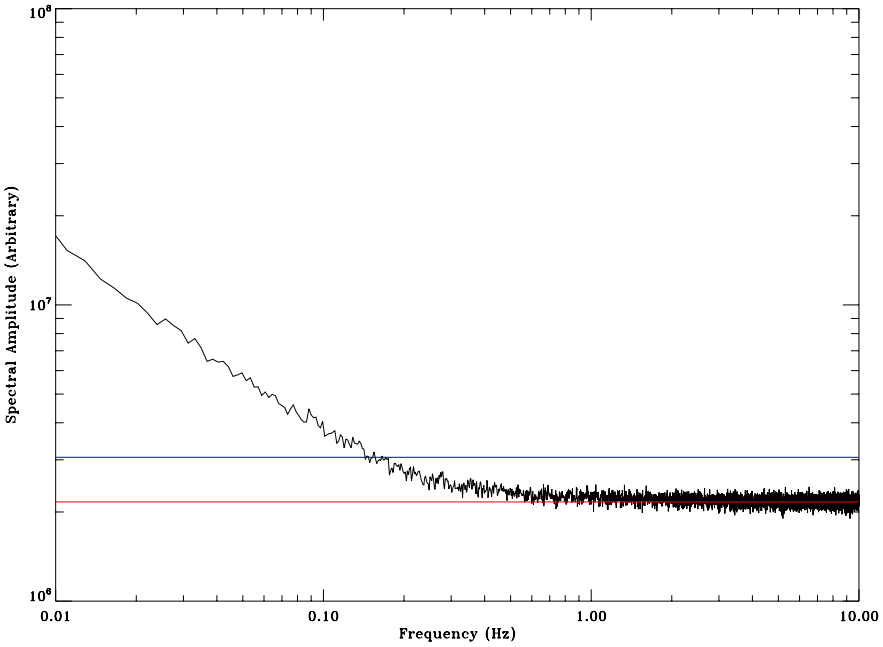

Figure 7: D1 weighted single difference spectrum. $1/f$ knee at about 150 mHz. Weighted by inverse standard deviation.

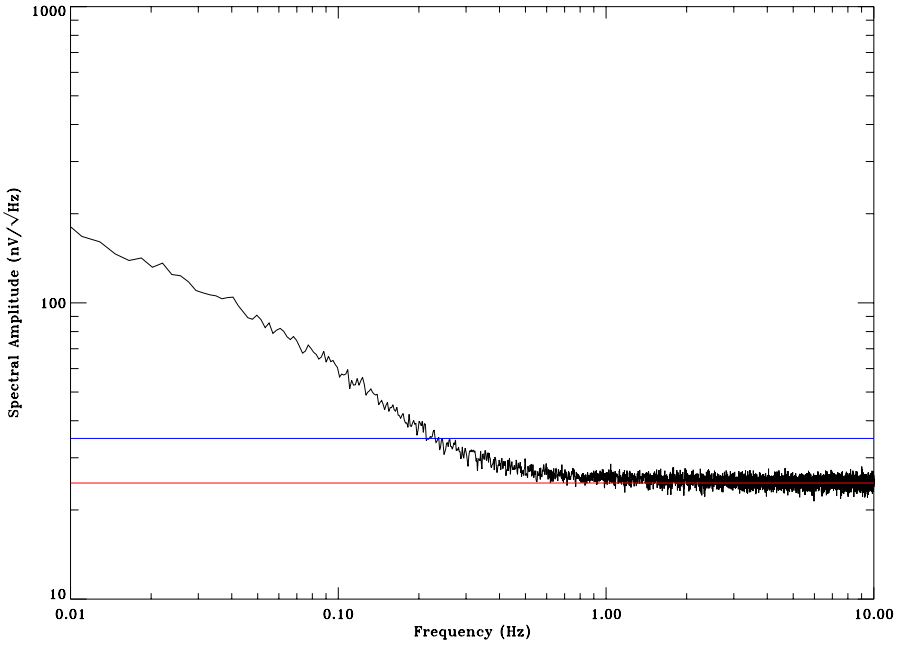

Figure 8: D4 unweighted single difference spectrum. $1/f$ knee at about 220 mHz.

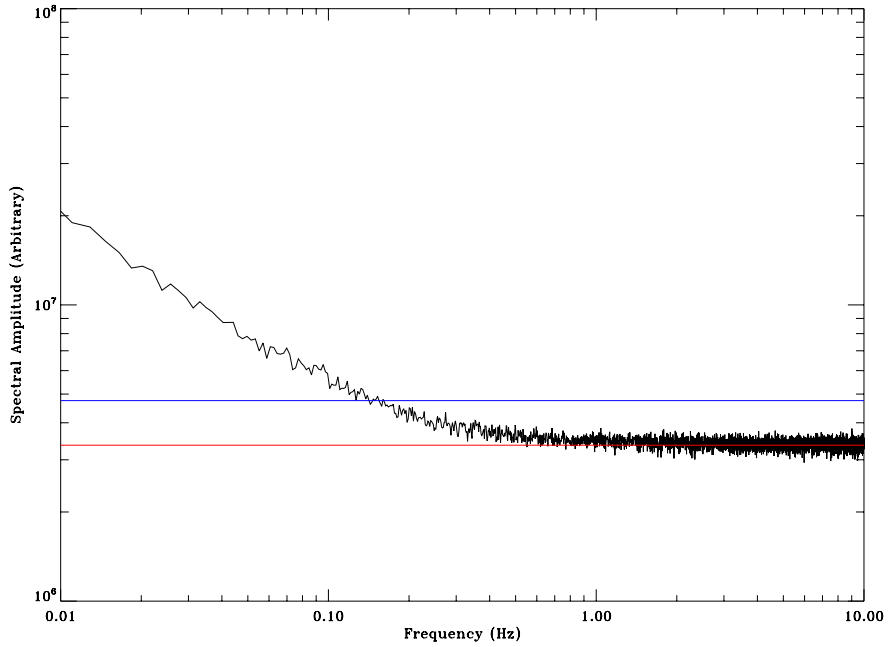

Figure 9: D4 weighted single difference spectrum. $1/f$ knee at about 150 mHz. Weighted by inverse standard deviation.

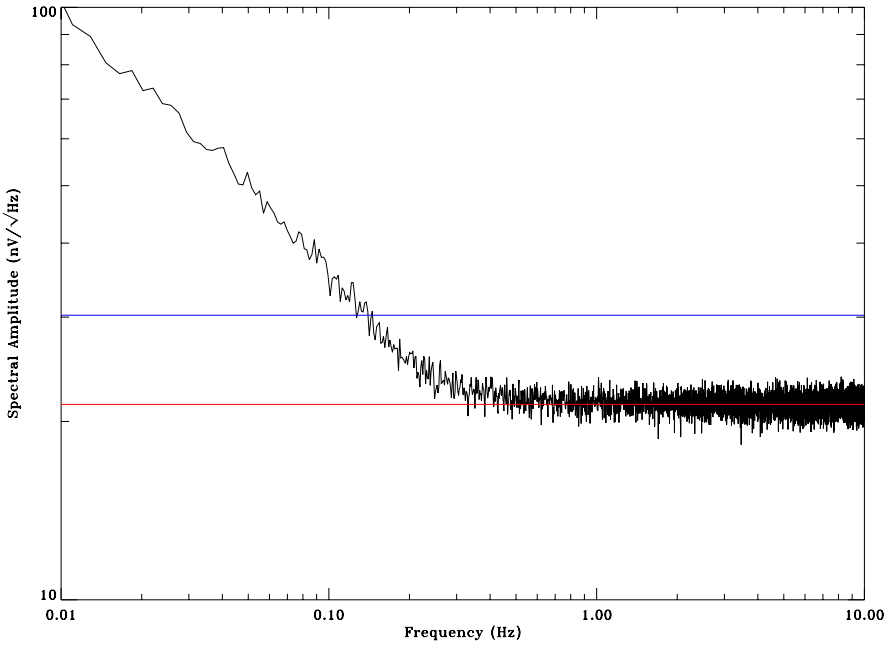

Figure 10: Unweighted double difference spectrum. $1/f$ knee at about 140 mHz.

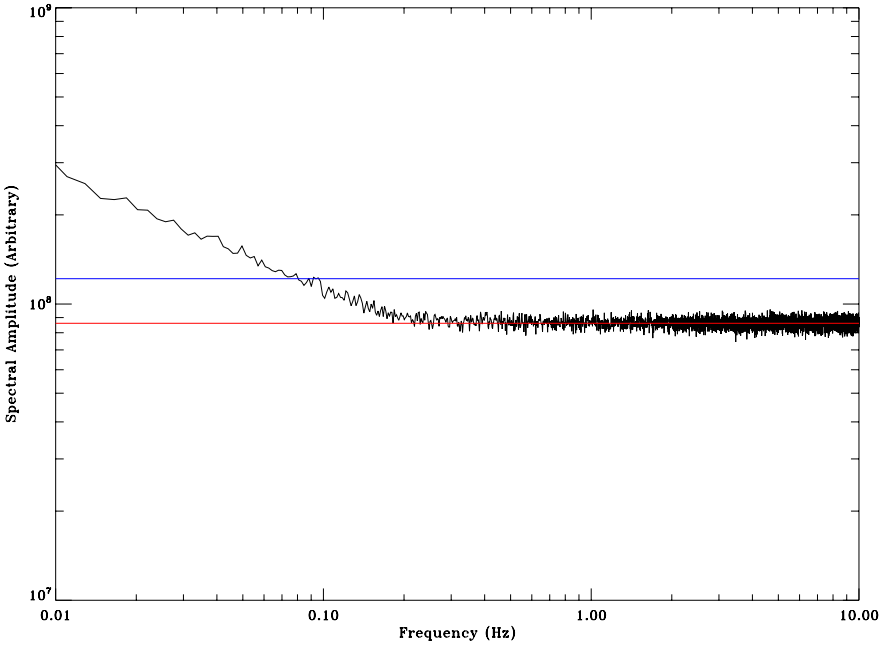

Figure 11: Doubly-weighted double difference spectrum. $1/f$ knee at about 80 mHz. This is a weighted difference of weighted single differences. All weights were inverse standard deviation.

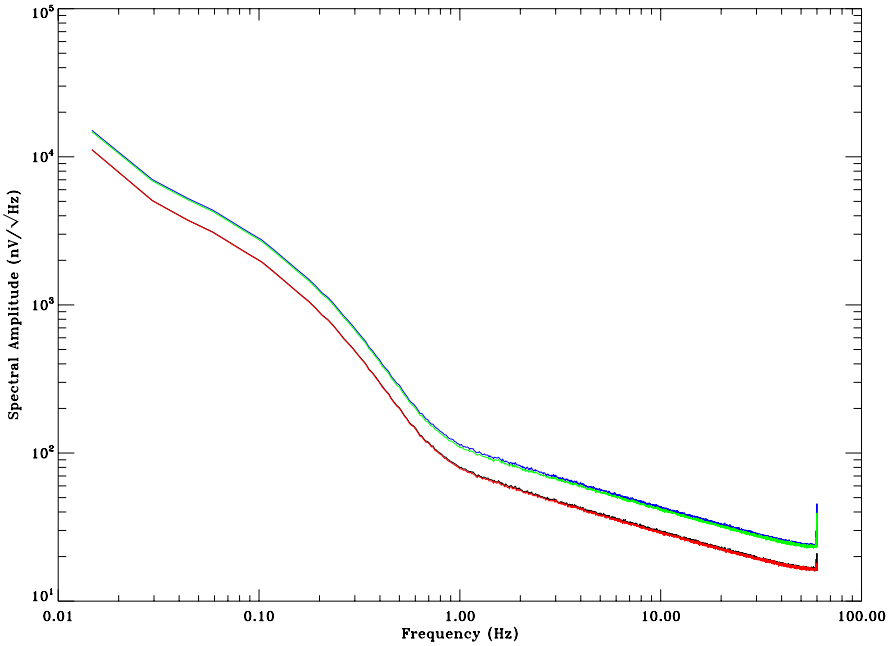

Figure 12: These are the spectra of undifferenced data plotted separately for each diode and phase switch state. Black and red plots (virtually identical, on bottom) are D1 "plus" and "minus," respectively. Blue and green plots (also virtually identical, on top) are D4, "plus" and "minus." The module was connected to an OMT and feedhorn.

4 Module with OMT

We connected an OMT and a feedhorn to the module and collected data sets. For these tests, the horn was aimed at a room-temperature (290K) Eccosorb target.

4.1 The "Bulge"

Many of the spectra, both with and without the OMT, show an unexpected "bulge" in the $1/f$-dominated region. This is generally suppressed in unweighted differences and often *not* suppressed in weighted differences. Furthermore, there is no bulge apparent in the doubly-weighted double difference spectrum Figure 18, which is the weighted difference between Figure 14 and Figure 16. Each of these spectra clearly shows the bulge. Thus, it appears that this effect is common to both phase states and both diodes. More investigation is required.

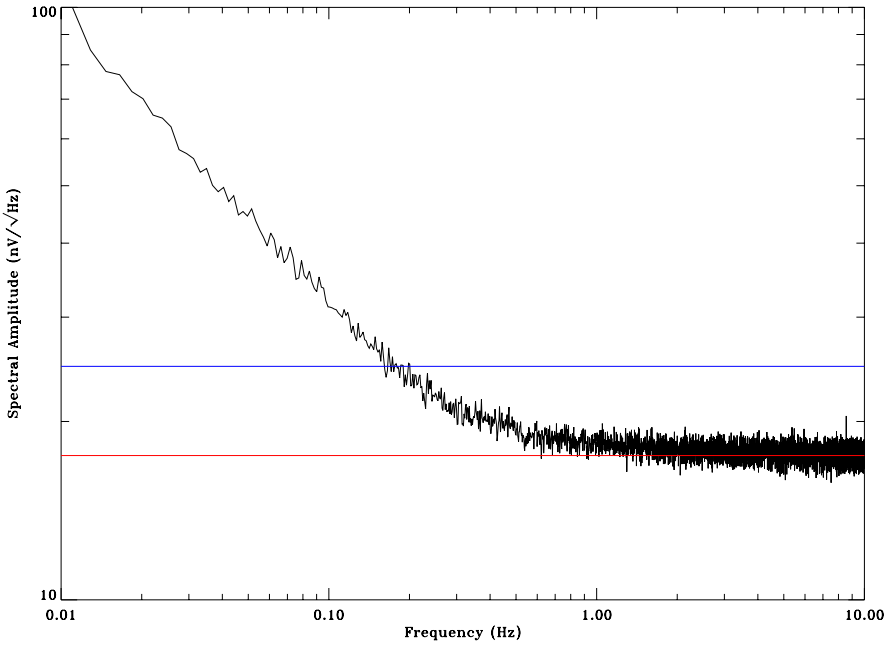

Figure 13: D1 unweighted single difference spectrum. $1/f$ knee at about 180 mHz. The module was connected to an OMT and feedhorn.

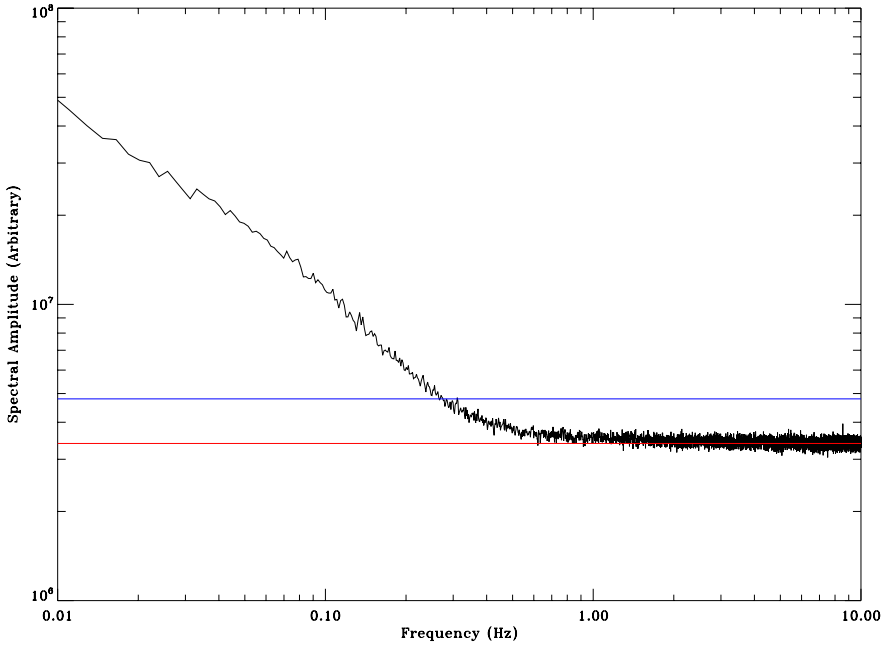

Figure 14: D1 weighted single difference spectrum. $1/f$ knee at about 260 mHz, which is higher than the unweighted knee. Weight used here was the inverse standard deviation of the spectrum data; other weights were tried, but all increased the knee frequency relative to the unweighted single difference knee. The module was connected to an OMT and feedhorn.

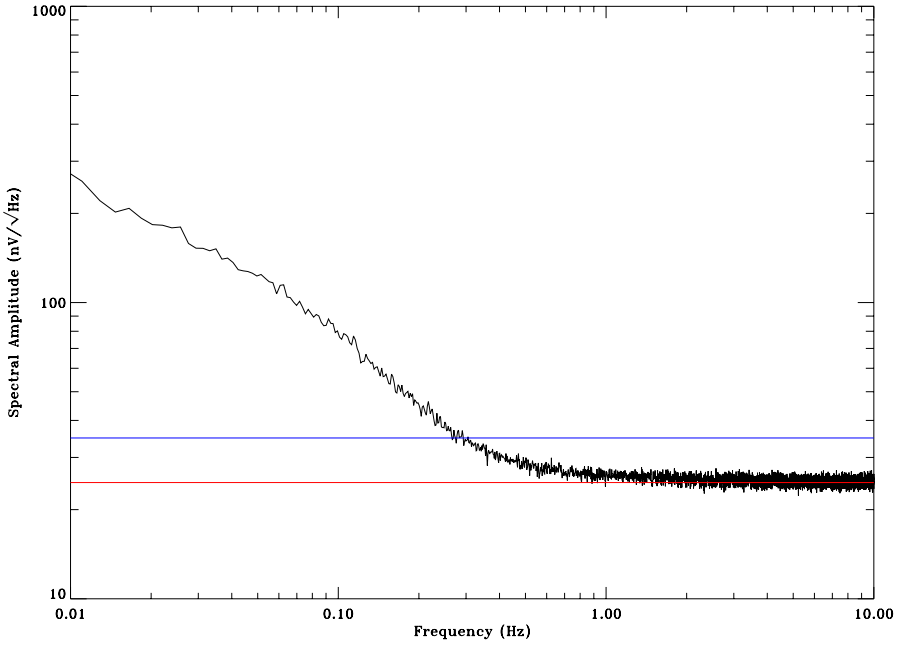

Figure 15: D4 unweighted single difference spectrum. $1/f$ knee at about 290 mHz. The module was connected to an OMT and feedhorn.

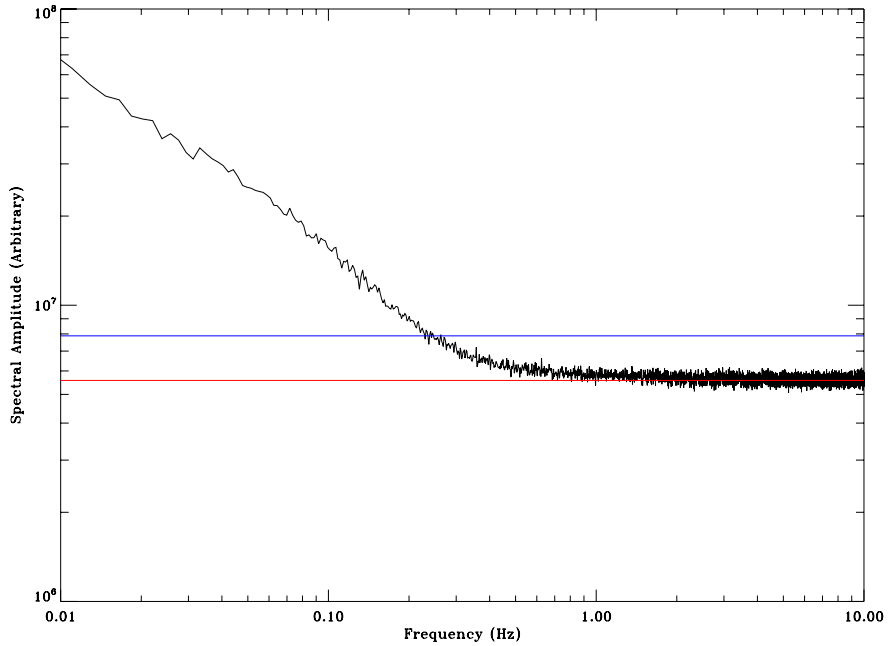

Figure 16: D4 weighted single difference spectrum. $1/f$ knee at about 240 mHz. Weighted by inverse standard deviation. The module was connected to an OMT and feedhorn.

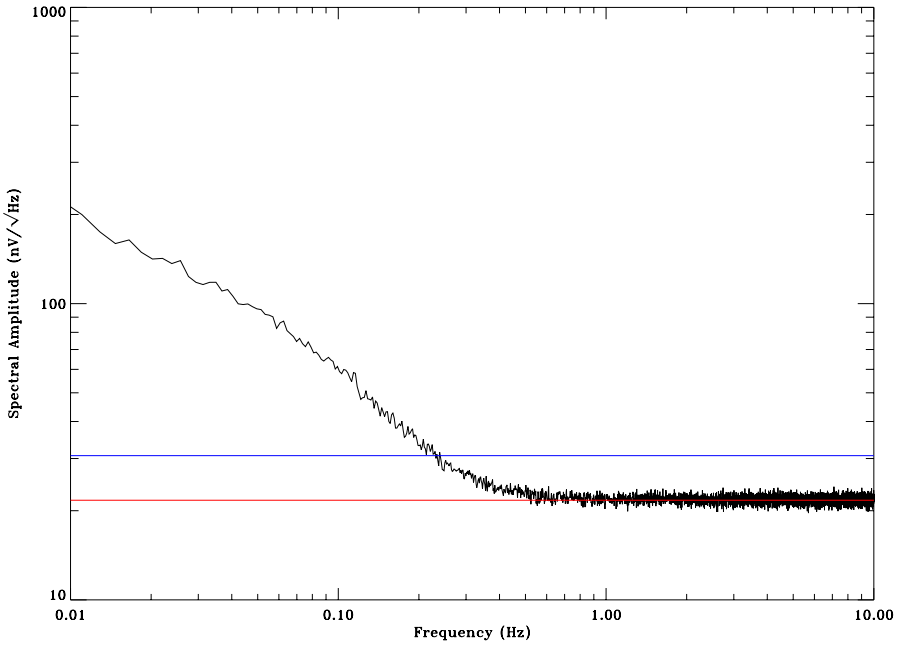

Figure 17: Unweighted double-difference spectrum. $1/f$ knee at about 230 mHz. The module was connected to an OMT and feedhorn.

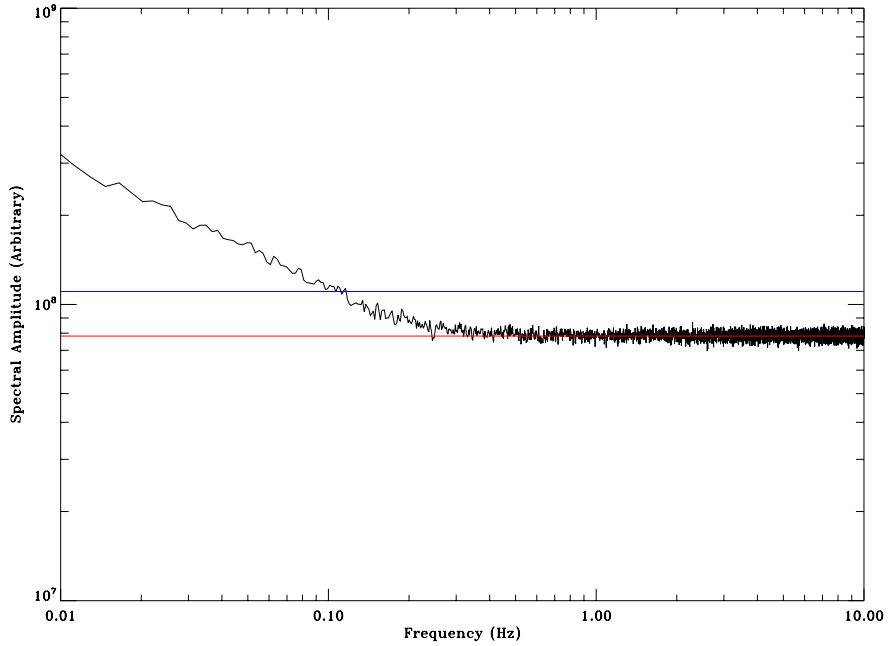

Figure 18: Doubly-weighted double-difference spectrum. $1/f$ knee at about 100 mHz. This is a weighted difference of weighted single differences. All weights were inverse standard deviation. The module was connected to an OMT and feedhorn. Also, it is not plotted, but a singly-weighted double difference was constructed from unweighted single difference data. It appears that this can reduce the $1/f$ knee frequency to between 100 and 150 mHz.

Effect of Bias on $1/f$ Gain Fluctuations in a QUIET MMIC

Joey Richards*and Mike Seiffert

April 8, 2008

Abstract

This memo describes preliminary results from measurements of the $1/f$ gain fluctuations in a single QUIET MMIC. By using a method similar to that used by Jarosik [1], we minimize the $1/f$ and white noise contributions of our test set. Our preliminary results show a correlation between the level of $1/f$ fluctuations at 1 Hz and the MMIC's drain bias voltage. These measurements were performed using MMIC W-26.

1 Test Apparatus

Our test apparatus consists of two nominally identical back-end MMIC amplifier chains, each connected to a wide-band detector diode. The configuration is shown in Figure 1. The device under test (DUT) is connected to one input a magic-tee. The other magic-tee input is terminated. Its outputs each pass through an adjustable attenuator and into one of the back-end amplifier chains.

The back-end amplifier chains, labeled A and B, each consist of two MMIC amplifier modules. MMIC A1 (B1) is the input to the chain and is followed by a band-pass filter, an isolator, a high-pass filter, then MMIC A2 (B2). The second MMIC's output passes through an isolator and into the wide-band detector diode.

The outputs of the detector diodes are amplified by a low-noise preamp with a voltage gain of 100 and an input impedance of approximately 1 kΩ.[1] Each preamp output connects to an integrator circuit and is further amplified, then integrated and sampled by a National Instruments ADC card in the test PC. The total low-frequency voltage gain from the preamp output to the ADC input is approximately 110. Unless otherwise specified, in our tests reported here, back-end chain A connects to ADC channel 1 and chain B to channel 0.

Before connecting a DUT to the system, we measured the background noise floor with the inputs terminated. Figure 2 shows a typical spectrum. The white noise floor is due to the test set—thus far, we are unable to measure the noise from a warm load above the noise floor of our data set. There are significant $1/f$ fluctuations apparent in the correlated (red) spectrum. We believe these correlations are due to noise from the integrator circuits, but more investigation is necessary. For tests with a MMIC, these fluctuations are low enough to make acceptable measurements, however.

Our DUT is a singly-packaged QUIET MMIC, serial number W-26. It is biased using a separate, DC-isolated bias channel in the same bias box providing bias to the back-end MMICs. This box provides both fixed gate and drain voltages, unlike the QUIET MMIC bias board which provides fixed gate voltage and fixed drain current. The attenuators on both chains were adusted to produce 1 mV on each detector diode with a nominal bias of the DUT. Biases for the DUT were chosen to vary the drain voltage while keeping the DC level on each detector diode close to 1 mV, without adjusting the attenuators.

*joey@caltech.edu

[1] The relatively low input impedance significantly loads the detector diode output, but this does not appear to be a limiting factor for our measurements so far.

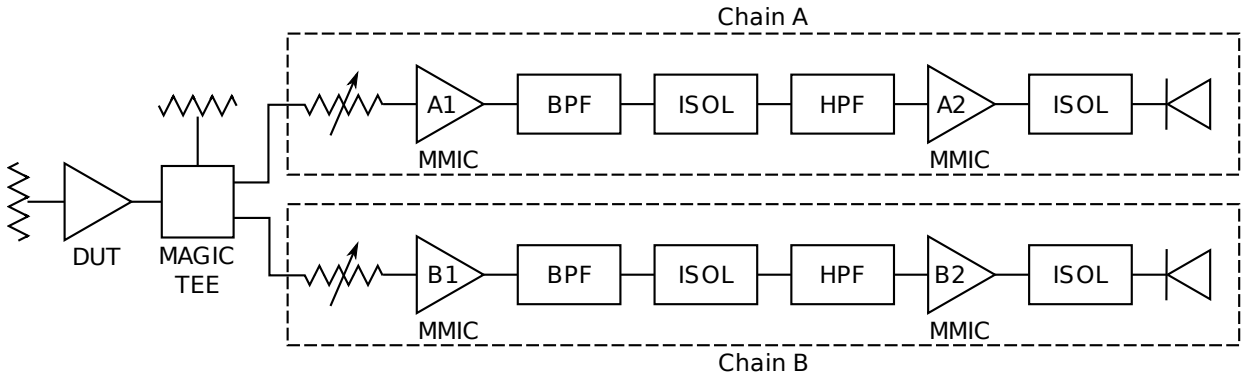

Figure 1: Schematic of the RF portion of the test apparatus. All components are at room temperature. Resistor symbols indicate terminators and potentiometers represent variable attenuators. "BPF" and "HPF" indicate band-pass and high-pass filter, respectively. "ISOL" indicates an isolator. The output of each detector diode is amplified and integrated/sampled by a base-band test set that is not shown.

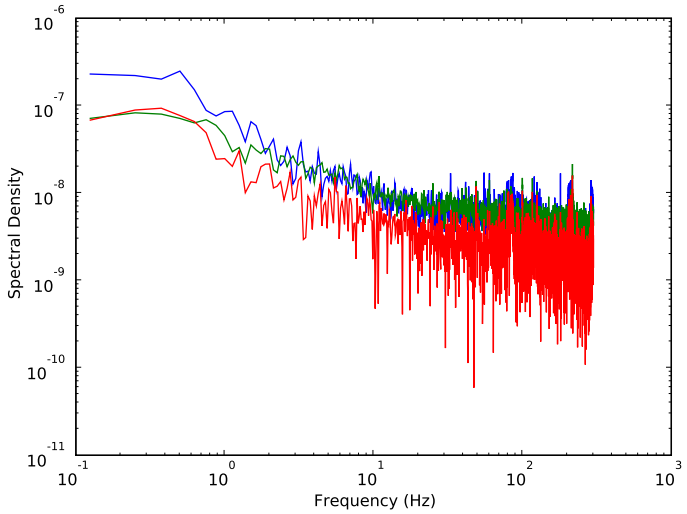

Figure 2: Typical spectrum of test set noise. The DUT input to the magic-tee is terminated. The attenuator on each chain input is set to slight attenuation, but the noise spectra do not depend on this attenuation. The blue and green data are ADC channels 0 and 1, respectively. The red data are the absolute value of the real cross-spectrum. In this test, chain A is connected to ADC channel 0 and chain B is connected to channel 1.

2 Data Analysis

2.1 Cross-Spectra

To reduce the noise from the back-end amplifiers and electronics, we compute the cross-spectrum of the two channels. An ordinary power spectrum is the Fourier transform of the auto-correlation series of a single input channel. The cross-spectrum is the generalization to two different input series: the Fourier transform of the cross-correlation series.

Unlike a power spectrum, the cross-spectrum is not a real-valued spectrum. The power spectrum is real because the auto-correlation series of a real series is a symmetric, real series, so its Fourier transform is also symmetric and real. The cross-correlation series, however, is merely real, not symmetric. As a result, the cross-spectrum has a symmetric real part and an anti-symmetric imaginary part.

We eliminate the imaginary part of the cross-spectrum by folding the two-sided cross-spectrum into a one-sided spectrum. When we add each negative-frequency component to its positive-frequency counterpart, the imaginary parts cancel (due to anti-symmetry) and we are left with a purely real one-sided spectrum.

This one-sided spectrum is not necessarily positive, since the mean value of the product of two series can be positive or negative. If there is a large correlated component between the two input series, it will dominate and the result will be positive. However, if uncorrelated noise is at a similar level to the correlated signal, the result will occasionally dip below zero. For our purposes, we simply plot the absolute value of the one-sided cross-spectrum—this will not result in significant distortions as long as we work with a signal that dominates the uncorrelated noise.

2.2 Estimating Gain Fluctuations

The cross-spectrum computed as discussed above is an unbiased estimate of the power spectrum of the correlated noise common to the two input chains. Because the correlated signal due to the DUT dominates the background correlated signal, the cross-spectrum of the two chains as an estimate of the power spectrum of the DUT.

We use the spectral density at 1 Hz as a fiducial measure of the gain fluctuations in the DUT. When divided by the DC detector value, this yields the gain fluctuations in Hz^{-1}. The background noise in our test set with the DUT powered off produces a DC level approximately 15% the DC level when the DUT is powered on. This background DC level is subtracted from the value used to divide the spectral density.

The variable attenuators were adjusted to produce equal 1.025 mV responses at both detector diodes at a nominal bias in the middle of the range tested. The attenuators were not adjusted, but as we varied the DUT bias, the two detector diodes did not always agree. At the maximum, they differed by about 0.13 mV. We used the geometric mean of the two detector diodes to normalize the gain fluctuations. Other methods of averaging differed only slightly and did not affect the qualitative features of the plots.

For each bias setting, a run of at least 45 minutes was collected. The longest run was a 20-hour overnight run. Comparisons of long runs with shorter runs at the same DUT bias did not show a significant effect on the estimated noise level. Each data run was divided into 10 second sections and the cross-spectra from these separate sections were averaged to produce a low-noise spectrum as shown in Figure 3.

From each averaged spectrum, the spectral density at 1 Hz was estimated both by eye and by fitting a $1/f^{\alpha}$ to the region between 0.5 Hz and 5 Hz. These methods produced consistent results. Error bars were crudely estimated by examining the "fur" on the spectrum. The error estimates were examined both on log-log and semi-log scales to avoid distortion by the logarithmic Y-axis scale. A more careful error analysis is intended in the future.

3 Results

See figures that follow. Only plots as a function of drain voltage and drain current bias values are included. Plots against other bias values (gate voltage and gate current) are qualitatively similar to the drain current plots here.

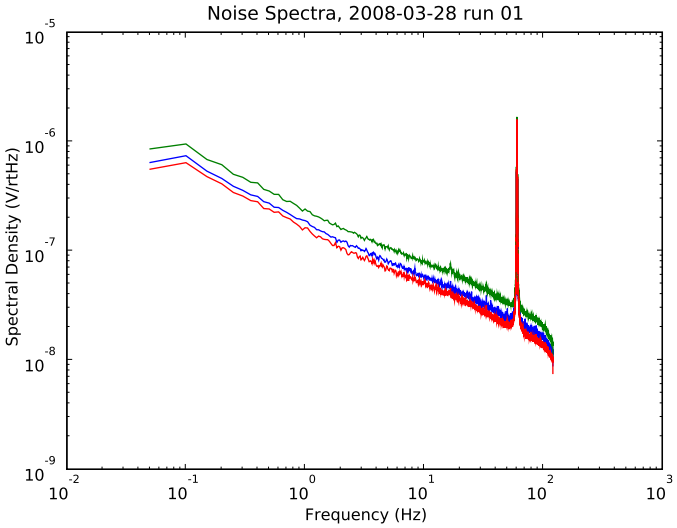

Figure 3: Typical spectrum from a 5000-second (1.4-hour) data set. The blue and green data are detector channels 0 and 1, respectively. The red data are the real cross-spectrum of the two channels. The large peak is at 60 Hz.

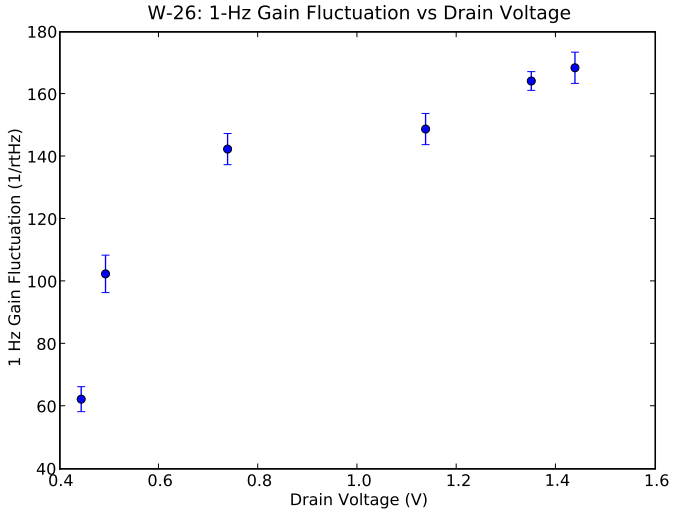

Figure 4: Gain fluctuations as a function of MMIC drain voltage. Note that the y-axis label units are incorrect. Y-axis units are arbitrary, but proportional to Hz^{-1}.

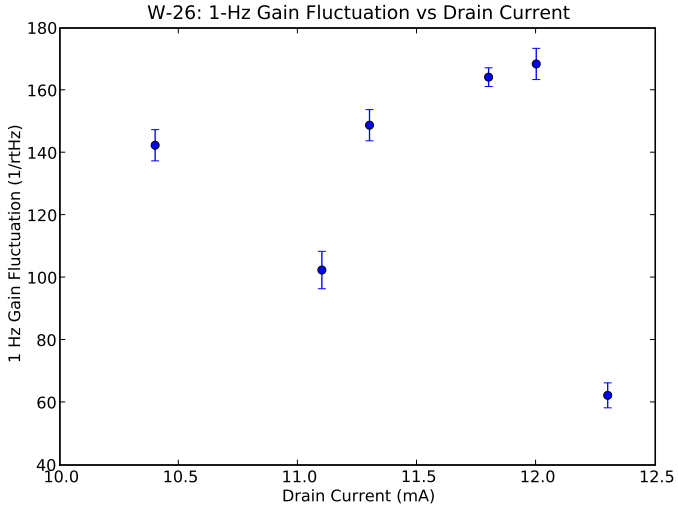

Figure 5: Gain fluctuations as a function of MMIC drain current. Note that the y-axis label units are incorrect. Y-axis units are arbitrary, but proportional to Hz^{-1}.

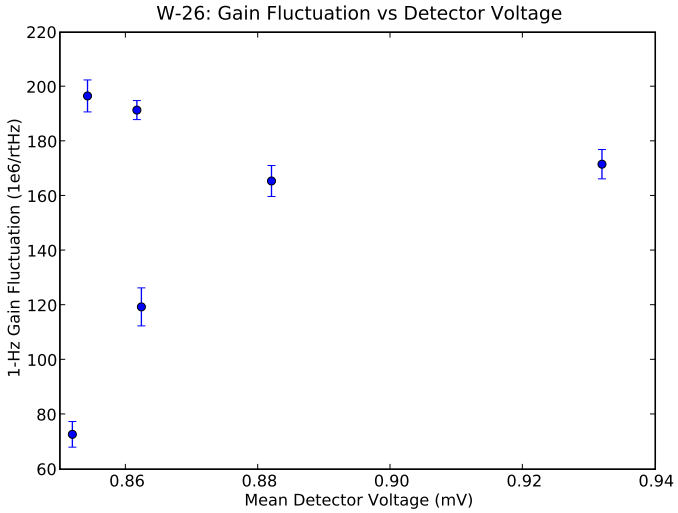

Figure 6: Gain fluctuations as a function of DC detected voltage. Here, the y-axis units are correct. The DC detector voltage is the geometric mean of the DC detected voltage at the two detector diodes. The difference between the two was less than about 0.1 mV, so the averaging method has little effect.

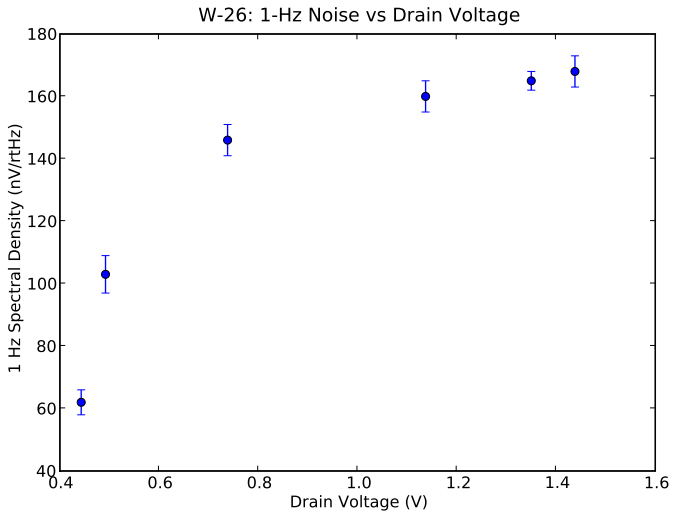

Figure 7: Spectral density at 1 Hz as a function of MMIC drain voltage. These are the same data as plotted in Figure 4 before dividing by the detected DC voltage.

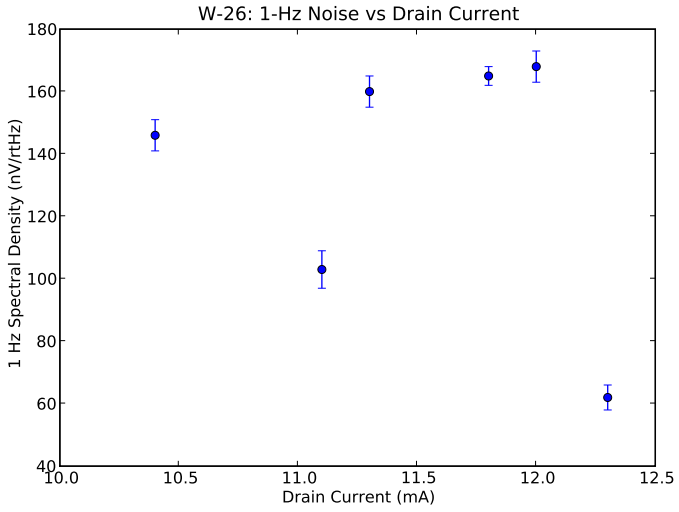

Figure 8: Spectral density at 1 Hz as a function of MMIC drain current. These are the same data as plotted in Figure 5 before dividing by the detected DC voltage.

4 Discussion / Future Plans

These results indicate a trend of lower $1/f$ gain fluctuations as the drain voltage of the DUT is decreased. It does not appear that there are systematic correlations with other bias variables. Several avenues of future measurements are immediately clear. These results come from tests with a single sample of a MMIC. A second MMIC should be tested similarly for comparison. Tests of these MMICs at cryogenic temperatures should be performed—these results are only interesting if they hold true at cryogenic temperatures as well as at room temperature.

Assuming this effect is real, it also remains to be verified that the reduction in $1/f$ gain fluctuations on a single MMIC tested in this manner translates into an effective $1/f$ noise reduction in a QUIET module in actual operation. This depends on the cause of residual $1/f$ noise after demodulation and will shed light on the origins of this residual $1/f$ whether or not the reduction is realized.

References

[1] N. C. Jarosik, "Measurements of the low-frequency-gain fluctuations of a 30-GHzhigh-electron-mobility-transistor cryogenic amplifier," *IEEE Transactions on Microwave Theory and Techniques*, vol. 44, pp. 193–197, Feb. 1996.

$1/f$ Fluctuation Studies: Uncorrelated MMIC Tests

Joey Richards*and Mike Seiffert

September 26, 2011

1 Apparatus

For these measurements we completely isolated the two legs of our test set, as shown in Figure 1. The back-end biases were set to approximately the same values as in our previously reported measurements. The MMIC devices under test (DUTs) were biased as follows.

	DUT A		DUT B	
	V	I	V	I
Gate A	0.106 V	30 uA	0.044 V	-8 uA
Gate B	0.102 V	30 uA	0.052 V	-5 uA
Drain	0.970 V	10.1 mA	0.739 V	9.9 mA

2 Spectra

In this document, we plot spectra on semilog or linear plots rather than log-log plots as we have done in the past. This is more appropriate for cross-spectra because it does allows both positive and negative values to be plotted. In particular, when a cross-spectrum is near zero-mean, plotting the absolute value (as is necessary for a log-log plot) tends to make its mean value appear significantly larger.

*joey@caltech.edu

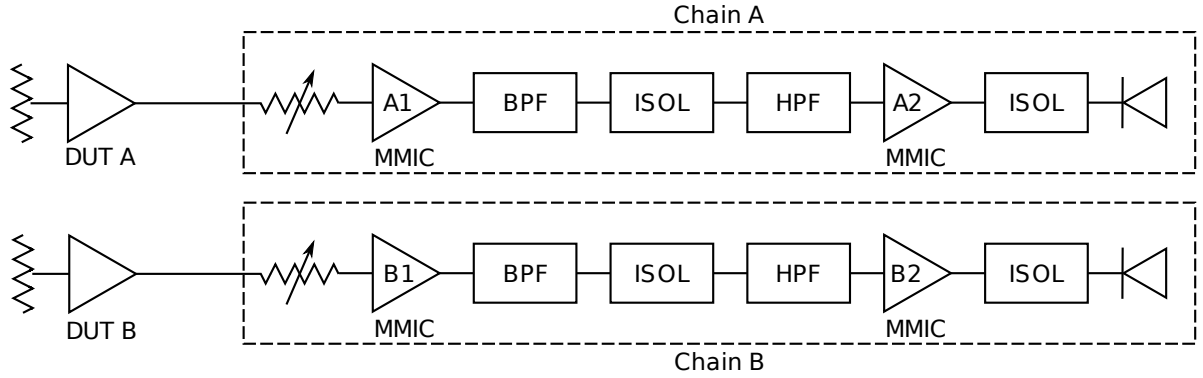

Figure 1: RF test set schematic for the uncorrelated MMIC tests. MMIC A is serial number L401. MMIC B is serial number W-26 and was used in the previous single-DUT tests. All components are room temperature.

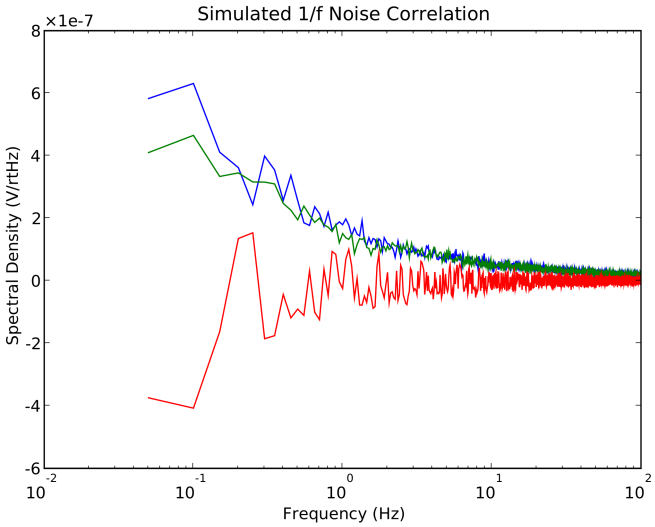

Figure 2: Simulated cross-spectrum of two independent $1/f$-noise series.

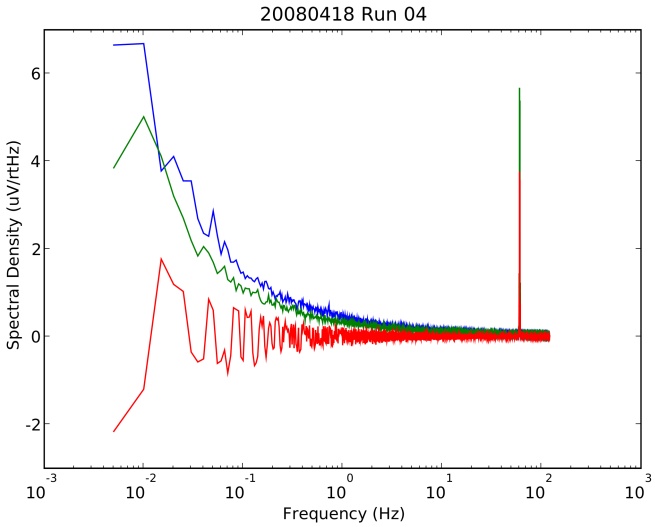

Figure 3: Power spectra of A and B signals (blue and green) and the cross-spectrum (red) using the fully independent two-DUT test set shown in Figure 1. Due to a systematic processing error, data above 60 Hz is invalid.

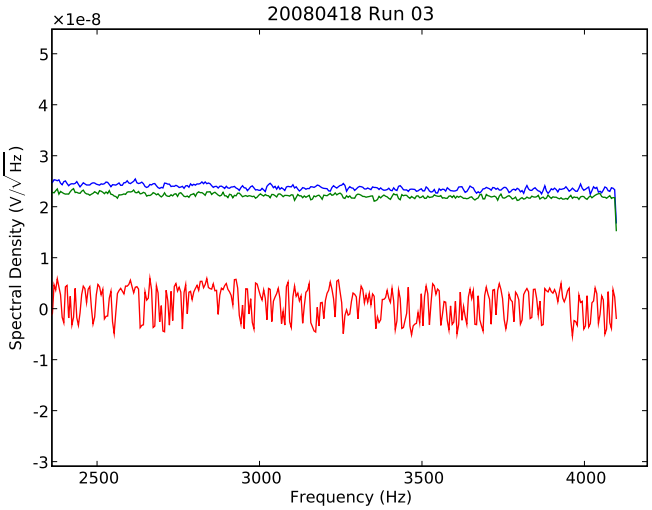

Figure 4: High-frequency spectra of A and B signals (blue and green) and the cross-spectrum (red). The RMS spectral densities over these data are 24.1 nV/$\sqrt{\text{Hz}}$, 22.5 nV/$\sqrt{\text{Hz}}$, and 2.3 nV/$\sqrt{\text{Hz}}$ for blue, green, and red, respectively.

3 Two Notes

3.1 Statistical Note

The cross-spectrum in Figure 4 appears to have a near-zero mean, but it appears to the eye that its values are clustered around non-zero positive and negative values with relatively few points very close to zero. A histogram of these data is shown in Figure 5, clearly demonstrating the absence of points near zero.

The voltage spectral density plot is calculated as the point-wise square root of the power spectral density. Because the cross-spectrum can be positive or negative, we take the absolute value of the power spectral density before the square root, then multiply by its sign to produce a real-valued voltage spectral density. Figure 6 shows the power spectral density before this processing. There is clearly no absence of points near the mean of the distribution.

The absence of points near zero in the voltage spectral density plot is an artifact of the square-rooting process we use. Let x be a Gaussian random variable with zero mean and unit standard deviation. The probability density for x is

$$p_x(x) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{x^2}{2}\right) \ . \tag{1}$$

Now consider $y = \text{sign}(x)\sqrt{|x|}$, the same transformation we use to produce our voltage spectral density plot. To find its probability density, we just need the derivative with respect to x. We'll ignore the discontinuity in the derivative at $x = y = 0$ that comes from the sign function, as it is inconsequential for our

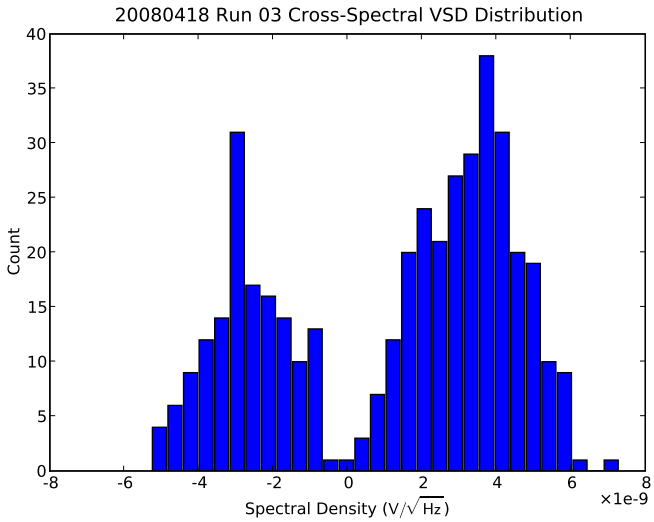

Figure 5: Histogram of the cross-spectrum data from Figure 4. Clearly, there are two separate clusters of values, one positive and one negative.

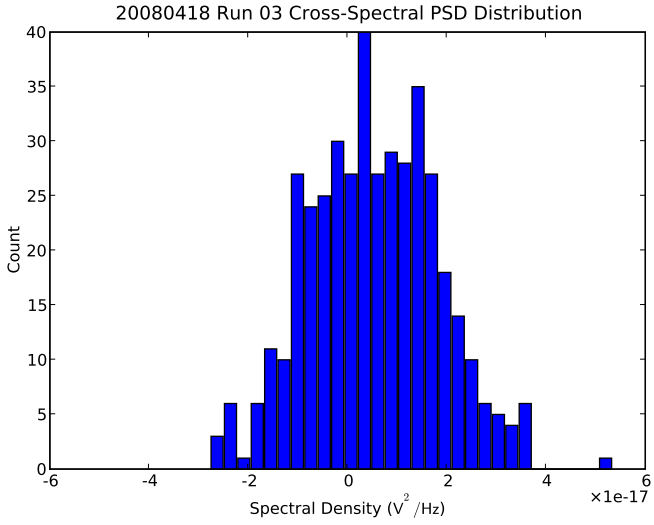

Figure 6:

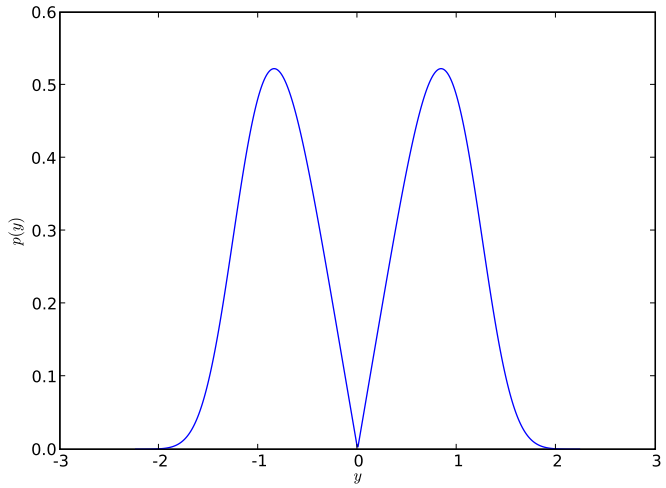

Figure 7: Probability density function of $y = \mathrm{sign}(x)\sqrt{|x|}$ for a zero-mean, unity-variance Gaussian random variable x. This qualitatively matches the features of the histogram of the voltage spectral density plot in Figure 5.

purposes.

$$p_y(y)dy = p_x(x)dx \ , \tag{2}$$

$$p_y(y) = p_x(x)\left(\frac{dx}{dy}\right) \ , \tag{3}$$

$$p_y(y) = p_x(x)\left(2|y|\right) \ , \tag{4}$$

$$p_y(y) = \frac{1}{\sqrt{2\pi}}\exp\left(-\frac{x^2}{2}\right)\left(2|y|\right) \ , \tag{5}$$

$$p_y(y) = \sqrt{\frac{2}{\pi}}\ |y|\exp\left(-\frac{y^4}{2}\right) \ . \tag{6}$$

This probability distribution is plotted in Figure 7. It features a null at zero and is qualitatively quite similar to our observed histogram. The imbalance between the positive and negative lobes is not understood, but probably reflects correlated noise or spurious interference of some sort.

3.2 Knee Frequency Note

We use the 1 Hz gain fluctuation level rather than the knee frequency as a parameter because it is independent of the radiometer bandwidth. The conversion between these parameters is straightforward, but depends on the type of radiometer in use. Following Wollack and Pospieszalski (1998), for our test set, the total fluctuations at frequency f are given by

$$\frac{\delta V}{V_{dc}} = \sqrt{\frac{2}{\beta} + \delta g^2(f)} \ . \tag{7}$$

Here, $\delta g^2(f) = \delta g_0^2 \times (1\ \mathrm{Hz}/f)$ is the $1/f$ gain fluctuation in 1/Hz and β is the RF bandwidth. The first term gives the white noise floor. Solving this to find the knee frequency yields

$$f_{knee} = \frac{\beta \cdot \delta g_0^2}{2} \ . \tag{8}$$

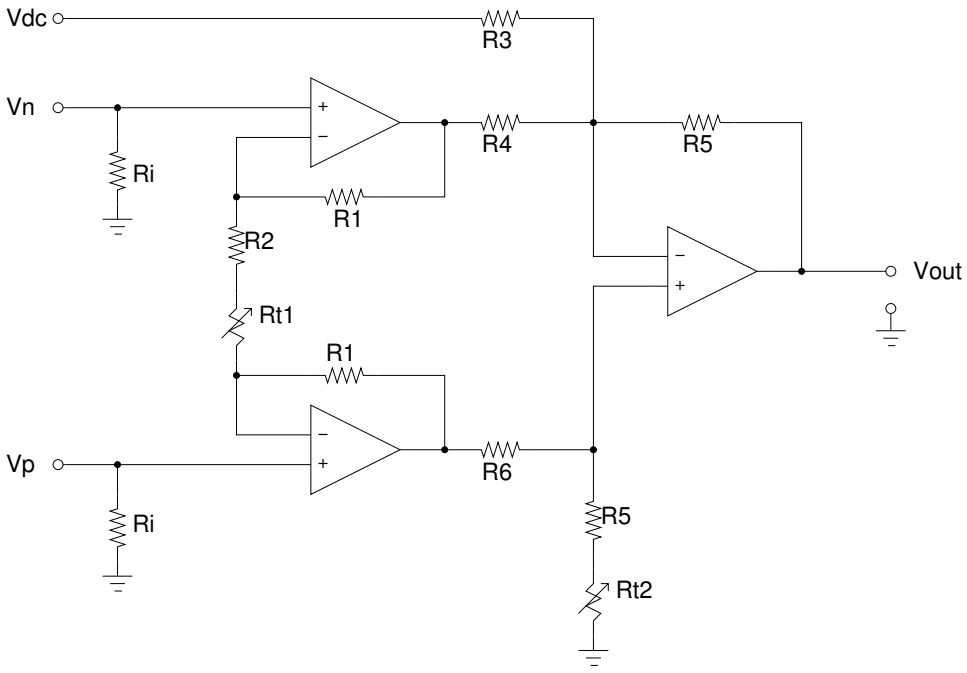

Figure D.66. Full preamp schematic. Note that resistor value R6 should be $R_4 + \frac{R_4 R_5}{R_3}$. When this is the case, the output of the amplifier will be $V_{out} = \frac{R_5}{R_4}\left(1 + 2\frac{R_1}{R_2}\right)(V_p - V_n) - \frac{R_5}{R_3}V_{DC}$. The trim pots Rt1 and Rt2 tune the gain and common-mode rejection ratio (CMRR), respectively.

D.5 Low-Noise Preamplifier Design

In this section, my design for a low-noise, high-input-impedance differential preamplifier with an independent subtraction input is presented. Results from testing of a prototype of the circuit are briefly described as well. A four-channel version was constructed on a custom printed circuit board and used extensively for radiometer characterization. This circuit was designed as an option for use in the QUIET back-end, and a scheme for connecting this amplifier to a biased RF detector diode is presented and analyzed. Ultimately a different design was selected for deployment, although this work influenced the final design.

The circuit shown in figure D.66 is a standard three-op amp instrumentation amplifier with the addition of a subtraction input in the second stage. It can be configured to have a wide bandwidth (>1 MHz), low input-referenced (RTI) noise ($\sim 6\ \mathrm{nV}/\sqrt{\mathrm{Hz}}$), and an input impedance limited only by the input impedance of the op amps ($>1\ \mathrm{M}\Omega$).

D.5.1 Analysis

The standard instrumentation amplifier is built from two pieces: an input stage consisting of a pair of noninverting amplifiers connected to a differential amplifier output stage. The design described here includes an additional subtraction node on the differential amplifier to allow a DC offset voltage to be removed from the input in order to avoid dynamic range problems due to detector diode biasing.

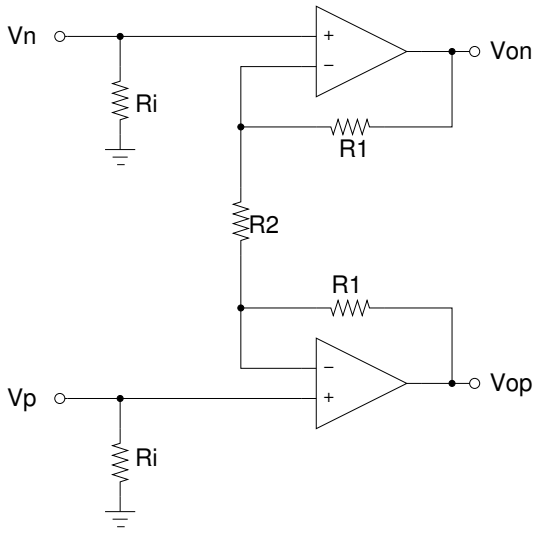

Figure D.67. Schematic of input amplifier stage. Note that the Rt1 trim pot has been included in R2 for analysis.

D.5.1.1 Input Stage

Figure D.67 shows the input stage of the preamp. The input impedance of the amplifier is $2R_i$. These shunt resistors are required to provide DC bias paths for the op amp inputs. Because these resistors are in parallel with the source, they will not contribute to the noise floor of the amplifier as long as the source impedance $R_s < R_i$. Furthermore, as long as $R_s \ll R_i$, the impedance of the source will not affect the effective gain of the amplifier.

The outputs of this stage, V_{op} and V_{on}, are related to the inputs as follows.

$$V_{on} = \left(1 + \frac{R_1}{R_2}\right) V_n - \frac{R_1}{R_2} V_p, \tag{D.1}$$

$$V_{op} = \left(1 + \frac{R_1}{R_2}\right) V_p - \frac{R_1}{R_2} V_n. \tag{D.2}$$

This is equivalent to a differential gain of

$$V_{op} - V_{on} \equiv G_{d1} \left(V_p - V_n\right) = \left(1 + 2\frac{R_1}{R_2}\right) \left(V_p - V_n\right). \tag{D.3}$$

D.5.1.2 Differential Stage

Figure D.68 shows the differential / DC offset subtraction stage of the preamp. First, we will assume V_{on} and V_{DC} are grounded, and calculate the output response due to V_{op} alone. Since the current into the op amp

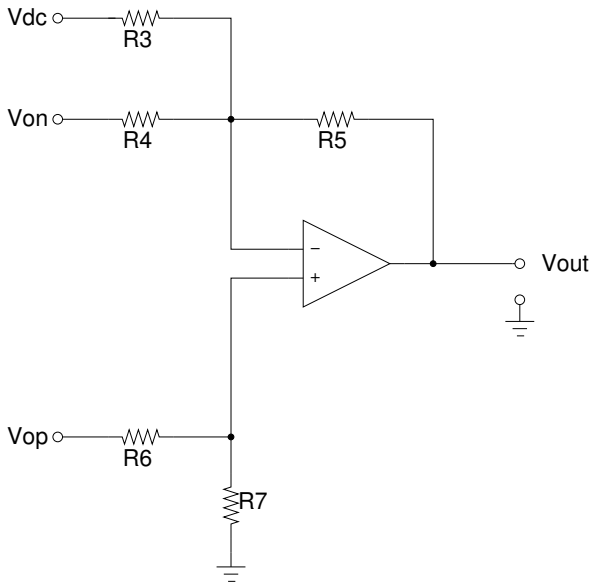

Figure D.68. Schematic of differential amplifier stage. Note that the Rt2 trim pot has been included in R7 for analysis. R7 was labeled R5 in figure D.66. In our analysis, we will show that R7 should have the same value as R5 to achieve balanced gain.

inputs is negligible, the voltage at the noninverting op amp input is

$$V_+ = V_{op} \times \frac{R_7}{R_6 + R_7}, \tag{D.4}$$

where R_7 in this equation equals the trimmed value $R_7 + R_{t7}$ on the schematic. Assuming that the feedback is working, this will be equal to the voltage at the inverting input, which is

$$V_- = V_{out} \times \frac{R_3 \parallel R_4}{(R_3 \parallel R_4) + R_5}. \tag{D.5}$$

Solving for the output voltage, this yields

$$V_{out} = \left(\frac{(R_3 \parallel R_4) + R_5}{R_6 + R_7} \right) \left(\frac{R_7}{R_3 \parallel R_4} \right) V_{op}. \tag{D.6}$$

Next, we will calculate the response due to V_{on} alone. The noninverting input is now at ground, so the inverting input will be a virtual ground. Using Kirchhoff's current law at the inverting input node, we have

$$\frac{V_{on}}{R_4} + \frac{V_{out}}{R_5} = 0, \tag{D.7}$$

so

$$V_{out} = -\frac{R_5}{R_4} V_{on}. \tag{D.8}$$

Likewise, the response due to V_{DC} is

$$V_{out} = -\frac{R_5}{R_3} V_{DC}. \tag{D.9}$$

The combined output is then

$$V_{out} = \left(\frac{(R_3 \parallel R_4) + R_5}{R_6 + R_7} \right) \left(\frac{R_7}{R_3 \parallel R_4} \right) V_{op} - \frac{R_5}{R_4} V_{on} - \frac{R_5}{R_3} V_{DC}. \tag{D.10}$$

If we set $R_6 = R_4 + \frac{R_4 R_5}{R_3}$ and $R_7 = R_5$, then this becomes

$$V_{out} = \frac{R_5}{R_4} (V_{op} - V_{on}) - \frac{R_5}{R_3} V_{DC} \equiv G_{d2} (V_{op} - V_{on}) - G_{DC} V_{DC}. \tag{D.11}$$

D.5.1.3 Full Amplifier

Combining the results of the preceding sections, we find that the total response of the preamplifier is as follows.

$$V_{out} = G_{d2} G_{d1} (V_p - V_n) - G_{DC} V_{DC}, \tag{D.12}$$

$$V_{out} = \frac{R_5}{R_4} \left(1 + 2 \frac{R_1}{R_2} \right) (V_p - V_n) - \frac{R_5}{R_3} V_{DC}. \tag{D.13}$$

D.5.2 Prototype Circuit

A prototype circuit has been assembled using Analog Devices OP37 op amps and the following component values. All resistors have 1% tolerance.

Component	Value
Ri	100 kΩ
R1	4.87 kΩ
R2	1.00 kΩ
R3	1.00 kΩ
R4	1.00 kΩ
R5	10.0 kΩ
R6	11.0 kΩ
Rt1	200 Ω single-turn potentiometer
Rt2	200 Ω single-turn potentiometer

This prototype was tuned using Rt1 to have a total DC gain of 100, divided equally between the two stages. The DC offset gain, G_{DC} is 10 in this prototype. Rt2 can be used to tune the common mode rejection ratio of the preamp, although this has not been tested.

The bandwidth of the differential inputs exceeds 1 MHz. The gain appears to be flat from DC to about 400 kHz. Starting at 400 kHz, there is a rise in the gain, peaking at a gain of about 150 at 1 MHz, then

dropping back to a gain of 100 at about 1.5 MHz and decreasing steadily above 1.5 MHz. It should be possible to eliminate this gain peak by adding some capacitors to the circuit.

A version of the prototype without the DC offset input was measured to have a noise floor of about $6\,\mathrm{nV}/\sqrt{\mathrm{Hz}}$. The inclusion of the DC offset input should not affect the noise floor significantly. The $1/f$ knee frequency appears to occur at a few hertz on the prototype circuit.

D.5.3 Detector Diode Biasing

A scheme using a single digital-to-analog converter (DAC) to provide both the diode bias current and a DC level for removal is shown in figure D.69. In this circuit, a DAC is operated with an output range $10\times$ the desired bias voltage range for V_{bias}, then divided down and low-pass filtered, and buffered to minimize the DAC noise contribution. Although a simple RC low-pass filter is shown in the schematic, a more aggressive filter could be integrated in the buffer amp if necessary.

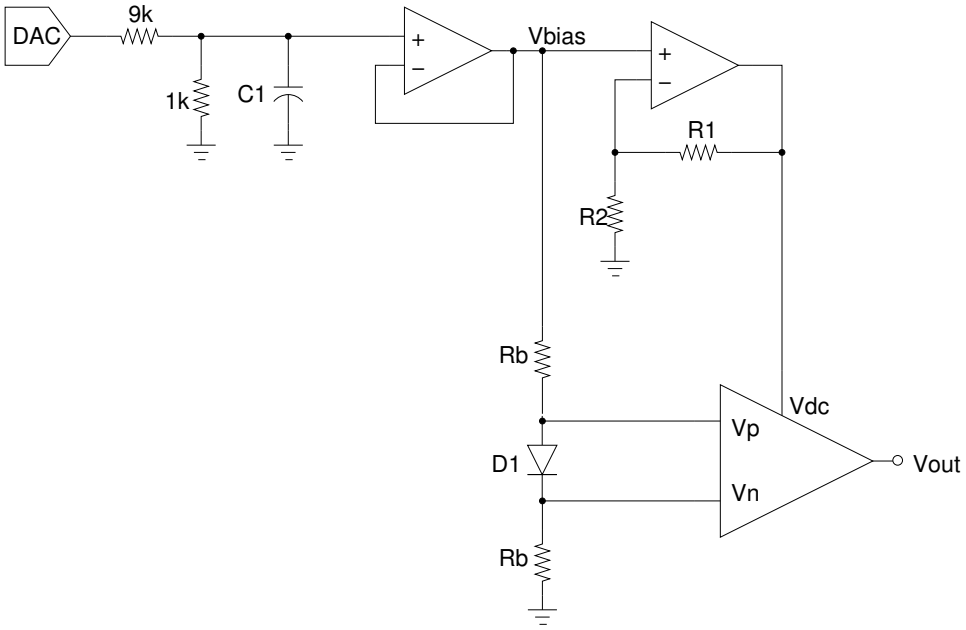

Figure D.69. One possible scheme for biasing the detector diode and generating a DC offset signal using a single DAC channel. The amplifier block in the lower-right represents the preamp described in this note.

This buffered voltage, V_{bias}, biases the detector diode D1 through the Rb resistors. These resistors must be chosen to allow sufficient bias current with the available V_{bias} range. Additionally, these resistors will determine the effective input impedance of the preamp circuit, so making these as large as possible will maximize transfer of the diode signal into the preamp.

The V_{bias} signal is also used to generate the DC offset subtraction signal for the preamp. In the circuit shown, it is amplified by a noninverting amplifier with gain $G = 1 + \frac{R1}{R2}$. These resistors should be set such

that

$$G = \frac{G_{d1}G_{d2}}{G_{DC}} \times \frac{R_{D1}}{R_{D1} + 2R_b}, \tag{D.14}$$

where G_{DC} is the DC offset gain of the preamp, G_{d1} and G_{d2} are the differential input gains in the first and second stages of the preamp, and R_{D1} is the DC resistance of the detector diode at its bias point. Error in the gain matching, caused by, e.g., an incorrect estimate of R_{D1}, will leave a residual DC signal contribution in V_{out}.

Because the DC offset gain of the preamp can be independently configured, it is possible to eliminate the noninverting amplifier and inject V_{bias} directly into the V_{DC} input of the preamp. In this case, G_{DC} can be set to compensate for the voltage drops in the bias resistors. This would allow the entire bias circuit and preamp to be constructed from a single quad op amp package.

D.5.4 Conclusions

This circuit is a promising candidate for a preamplifier. It exhibits a large input impedance, eliminating concerns that variations in detector diode impedance will affect the gain of the preamp. The noise floor, while slightly higher than the previous circuit (6 $\mathrm{nV}/\sqrt{\mathrm{Hz}}$ versus about 5 $\mathrm{nV}/\sqrt{\mathrm{Hz}}$), should be acceptably low. Finally, the inclusion of a DC offset input can eliminate the need for a separate unity-gain DC path while still allowing measurement of the DC level of the diode. Although only an OP37-based circuit has been tested so far, it should be possible to build an equivalent circuit using Linear Technologies LT1125 or LT1127 quad op amps. This would allow the entire preamp to be constructed using a single integrated circuit and a few resistors.

D.6 QUIET 43 GHz Paper

This section contains the preprint manuscript of the QUIET first-results paper, which describes the operation of and results from operation of the Q-band instrument on the Chajnantor plateau in the Atacama Desert in Chile (QUIET Collaboration et al. 2011). The present author's contributions to this work consisted of development, testing, and characterization of the polarimeter modules and support electronics, including the work described in the previous sections of this chapter.

Submitted to ApJ—Draft version December 30, 2010
Preprint typeset using LaTeX style emulateapj v. 08/13/06

FIRST SEASON QUIET OBSERVATIONS: MEASUREMENTS OF CMB POLARIZATION POWER SPECTRA AT 43 GHZ IN THE MULTIPOLE RANGE $25 \leq \ell \leq 475$

QUIET Collaboration—C. Bischoff[1,22], A. Brizius[1,2], I. Buder[1], Y. Chinone[3,4], K. Cleary[5], R. N. Dumoulin[6], A. Kusaka[1], R. Monsalve[7], S. K. Næss[8], L. B. Newburgh[6,23], R. Reeves[5], K. M. Smith[1,23], I. K. Wehus[9], J. A. Zuntz[10,11,12], J. T. L. Zwart[6], L. Bronfman[13], R. Bustos[7,13,14], S. E. Church[15], C. Dickinson[16], H. K. Eriksen[8,17], P. G. Ferreira[10], T. Gaier[18], J. O. Gundersen[7], M. Hasegawa[3], M. Hazumi[3], K. M. Huffenberger[7], M. E. Jones[10], P. Kangaslahti[18], D. J. Kapner[1,24], C. R. Lawrence[18], M. Limon[6], J. May[13], J. J. McMahon[19], A. D. Miller[6], H. Nguyen[20], G. W. Nixon[21], T. J. Pearson[5], L. Piccirillo[16], S. J. E. Radford[5], A. C. S. Readhead[5], J. L. Richards[5], D. Samtleben[2,25], M. Seiffert[18], M. C. Shepherd[5], S. T. Staggs[21], O. Tajima[1,3], K. L. Thompson[15], K. Vanderlinde[1,26], R. Williamson[6,27], B. Winstein[1]

Submitted to ApJ—This paper should be cited as "QUIET (2010)"

ABSTRACT

The Q/U Imaging ExperimenT (QUIET) employs coherent receivers at 43 GHz and 95 GHz, operating on the Chajnantor plateau in the Atacama Desert in Chile, to measure the anisotropy in the polarization of the CMB. QUIET primarily targets the B modes from primordial gravitational waves. The combination of these frequencies gives sensitivity to foreground contributions from diffuse Galactic synchrotron radiation. Between 2008 October and 2010 December, over 10,000 hours of data were collected, first with the 19-element 43-GHz array (3458 hours) and then with the 90-element 95-GHz array. Each array observes the same four fields, selected for low foregrounds, together covering ≈ 1000 square degrees. This paper reports initial results from the 43-GHz receiver which has an array sensitivity to CMB fluctuations of $69\,\mu\mathrm{K}\sqrt{\mathrm{s}}$. The data were extensively studied with a large suite of null tests before the power spectra, determined with two independent pipelines, were examined. Analysis choices, including data selection, were modified until the null tests passed. Cross correlating maps with different telescope pointings is used to eliminate a bias. This paper reports the EE, BB, and EB power spectra in the multipole range $\ell = 25$–475. With the exception of the lowest multipole bin for one of the fields, where a polarized foreground, consistent with Galactic synchrotron radiation, is detected with 3-σ significance, the E-mode spectrum is consistent with the ΛCDM model, confirming the only previous detection of the first acoustic peak. The B-mode spectrum is consistent with zero, leading to a measurement of the tensor-to-scalar ratio of $r = 0.35^{+1.06}_{-0.87}$. The combination of a new time-stream "double-demodulation" technique, Mizuguchi–Dragone optics, natural sky rotation, and frequent boresight rotation leads to the lowest level of systematic contamination in the B-mode power so far reported, below the level of $r = 0.1$.

Subject headings: cosmic background radiation—Cosmology: observations—Gravitational waves—Inflation—Polarization

[1] Kavli Institute for Cosmological Physics, Department of Physics, Enrico Fermi Institute, The University of Chicago, Chicago, IL 60637, USA; send correspondence to A. Kusaka, **akito@kicp.uchicago.edu**
[2] Max-Planck-Institut für Radioastronomie, Auf dem Hügel 69, 53121 Bonn, Germany
[3] High Energy Accelerator Research Organization (KEK), 1-1 Oho, Tsukuba, Ibaraki 305-0801, Japan
[4] Astronomical Institute, Graduate School of Science, Tohoku University, Aramaki, Aoba, Sendai 980-8578, Japan
[5] Cahill Center for Astronomy and Astrophysics, California Institute of Technology, 1200 E. California Blvd M/C 249-17, Pasadena, CA 91125, USA
[6] Department of Physics and Columbia Astrophysics Laboratory, Columbia University, New York, NY 10027, USA
[7] Department of Physics, University of Miami, 1320 Campo Sano Drive, Coral Gables, FL 33146, USA
[8] Institute of Theoretical Astrophysics, University of Oslo, P.O. Box 1029 Blindern, N-0315 Oslo, Norway
[9] Department of Physics, University of Oslo, P.O. Box 1048 Blindern, N-0316 Oslo, Norway
[10] Department of Astrophysics, University of Oxford, Keble Road, Oxford OX1 3RH, UK
[11] Oxford Martin School, 34 Broad Street, Oxford OX1 3BD, UK
[12] Department of Physics and Astronomy, University College London, Gower Street, London WC1E, UK
[13] Departamento de Astronomía, Universidad de Chile, Casilla 36-D, Santiago, Chile
[14] Departamento de Astronomía, Universidad de Concepción, Casilla 160-C, Concepción, Chile
[15] Kavli Institute for Particle Astrophysics and Cosmology and Department of Physics, Stanford University, Varian Physics Building, 382 Via Pueblo Mall, Stanford, CA 94305, USA
[16] Jodrell Bank Centre for Astrophysics, Alan Turing Building, School of Physics and Astronomy, The University of Manchester, Oxford Road, Manchester M13 9PL, UK
[17] Centre of Mathematics for Applications, University of Oslo, P.O. Box 1053 Blindern, N-0316 Oslo, Norway
[18] Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA, USA 91109
[19] Department of Physics, University of Michigan, 450 Church Street, Ann Arbor, MI 48109, USA
[20] Fermi National Accelerator Laboratory, Batavia, IL 60510, USA
[21] Joseph Henry Laboratories of Physics, Jadwin Hall, Princeton University, Princeton, NJ 08544, USA
[22] Current address: Harvard-Smithsonian Center for Astrophysics, 60 Garden Street MS 43, Cambridge, MA 02138, USA
[23] Current address: Joseph Henry Laboratories of Physics, Jadwin Hall, Princeton University, Princeton, NJ 08544, USA
[24] Current address: Micro Encoder Inc., Kirkland, WA 98034, USA

1. INTRODUCTION

The inflationary paradigm resolves several outstanding issues in cosmology, including the flatness, horizon, and monopole problems, and it provides a compelling explanation for the origin of structure in the Universe (e.g. Liddle & Lyth 2000, and references therein). So far all cosmological data, including measurements of Cosmic Microwave Background (CMB) anisotropies, support this paradigm; still the underlying fundamental physics responsible for inflation is unknown. Inflation produces a stochastic background of gravity waves that induce odd-parity tensor "B modes" at large angular scales in the CMB polarization. If these primordial B modes, parametrized by the tensor-to-scalar ratio r, are detected, one can learn about the energy scale of inflation. In many attractive slow-roll models, this scale is given approximately by $r^{1/4} \times 10^{16}\,\mathrm{GeV}$. For large-field models, the energy scale is near the Grand Unification Scale in particle physics, so that $r \gtrsim 0.01$. A new generation of experiments aims for good sensitivity in this range of r. Establishing the existence of primordial B modes would both verify an important prediction of inflation and provide access to physics at an incredibly high energy scale.

The most stringent limit to date is $r < 0.20$ at the 95% confidence level (Komatsu et al. 2010) set by a combination of CMB–temperature-anisotropy measurements, baryon acoustic oscillations, and supernova observations, but cosmic variance prohibits improvements using only these measurements.

E-mode polarization has now been detected by many experiments (e.g., Kovac et al. 2002; Leitch et al. 2005; Montroy et al. 2006; Sievers et al. 2007; Wu et al. 2007; Bischoff et al. 2008; Larson et al. 2010). These measurements are consistent with predictions from CMB–temperature-anisotropy measurements, and they provide new information on the epoch of reionization. Only BICEP has accurately measured E-mode polarization in the region of the first acoustic peak (Chiang et al. 2010); that paper also reports the best limit on r coming from cosmological B modes: $r < 0.72$ at the 95% confidence level.

Experiments measuring B-mode polarization in the CMB should yield the best information on r, but this technique is still in its infancy. B modes are expected to be at least an order of magnitude smaller than the E modes so control of systematic errors and foregrounds will be particularly critical. Below $\approx 90\,\mathrm{GHz}$, the dominant foreground comes from Galactic synchrotron emission, while at higher frequencies, emission from thermal dust dominates. Most planned or operating CMB polarization experiments employ bolometric detectors observing most comfortably at frequencies $\gtrsim 90\,\mathrm{GHz}$, so they cannot estimate synchrotron contamination from their own data.

[25] Current address: Nikhef, Science Park, Amsterdam, The Netherlands
[26] Current address: Department of Physics, McGill University, 3600 Rue University, Montreal, Quebec H3A 2T8, Canada
[27] Current address: Kavli Institute for Cosmological Physics, Enrico Fermi Institute, The University of Chicago, Chicago, IL 60637, USA

The Q/U Imaging ExperimenT (QUIET) is one of two CMB polarization experiments to observe at frequencies suitable for addressing synchrotron contamination, making observations at 43 GHz (Q band) and 95 GHz (W band) and with sufficient sensitivity to begin to probe primordial B modes. The other is *Planck* (Tauber et al. 2010).

QUIET uses compact polarization-sensitive modules based upon High–Electron-Mobility Transistor (HEMT) amplifiers, combined with a new time-stream "double-demodulation" technique, Mizuguchi–Dragone (MD) optics (for the first time in a CMB polarization experiment), natural sky rotation, and frequent rotation about the optical axis to achieve a very low level of contamination in the multipole range where a primordial–B-mode signal is expected.

Between 2008 October and 2010 December, QUIET collected over 10,000 hours of data, split between the Q-band and W-band receivers. Here we report first results from the first season of 3458 hours of Q-band observation. After describing the instrument, observations, and detector calibrations (Sections 2, 3, and 4), we discuss our analysis techniques and consistency checks (5 and 6). CMB power spectra are then presented together with a foreground detection (7). We evaluate our systematic errors (8) and then conclude (9).

2. THE INSTRUMENT

The QUIET instrument comprises an array of correlation polarimeters cooled to 20 K and coupled to a dual-reflector telescope, installed on a three-axis mount inside a comoving ground screen. The instrument is illustrated in Figure 1. Further details are given below and in Newburgh et al. (2010), Kusaka et al. (2010), and Buder (2010).

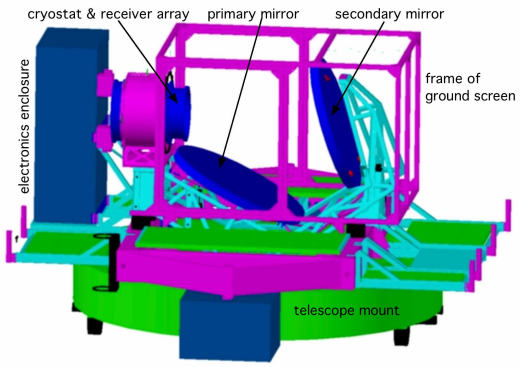

FIG. 1.— Overview of the QUIET instrument. The cryostat and 1.4-m telescope mirrors are enclosed in a rectangular comoving absorbing ground screen; in this figure its walls are transparent. The telescope, cryostat and electronics are mounted on a single platform attached to the deck bearing, which allows rotations around the instrument's optical axis.

The Q-band QUIET receiver is a 19-element array containing 17 low-noise correlation polarimeters, each simultaneously measuring the Stokes Q, U, and I parameters, and two CMB differential-temperature monitors.

The first element in the QUIET optical chain is a 1.4-m crossed Mizuguchi–Dragone dual-reflective telescope

(Mizugutch et al. 1976; Dragone 1978). The crossed MD configuration is very compact, with low cross polarization and a large diffraction-limited field of view. The telescope is described in detail in Imbriale et al. (2010). Light incident on the mirrors is focused into an array of corrugated circular feed horns (Gundersen & Wollack 2009), yielding a full-width half-maximum (FWHM) beam size of $27\farcm3$ and a roughly circular field of view of 7° diameter. Radiation from each feed horn enters a septum polarizer (Bornemann & Labay 1995) which separates left and right circularly-polarized components (L and R) into two waveguide ports which mate to a QUIET correlation module, detailed below.

The module array and feed horns are cooled to 20 K in a cryostat to reduce instrumental noise. An electronics enclosure mounted next to the cryostat houses the electronics necessary for biasing the modules and recording their data. The cryostat, electronics, and telescope are installed on the former CBI mount (Padin et al. 2002). This mount provides three-axis motion: azimuth, elevation, and rotation about the optical axis. This last is called "deck" rotation.

The cryostat and telescope are enclosed by an absorbing comoving ground screen. The ground screen was designed to have two parts, but the upper section (not shown in Fig. 1) was not installed until after the Q-band instrument was removed. Its absence was correctly anticipated to result in two far sidelobes, which were mapped with a high-power source by the QUIET W-band instrument in the field and measured to be $\lesssim -60$ dB with the QUIET Q-band instrument when the Sun passed through them. The effects of these sidelobes are mitigated through filtering and data selection (Sections 5.1.3 and 5.2). Section 8.4 shows that any residual contamination is small.

Each QUIET Q-band correlation module, in a footprint of only $5.1 \times 5.1\,\mathrm{cm}^2$, receives the circular polarization modes of the celestial radiation and outputs Stokes Q, U and I as follows. Each input is independently amplified and passed through a phase switch. One phase switch alternates the sign of the signal voltage at 4 kHz, while the other switches at 50 Hz. The two signals are combined in a 180° hybrid coupler, with outputs proportional to the sum and difference of the inputs. Since the module inputs are proportional to $(L, R) = (E_x \pm iE_y)/\sqrt{2}$, where E_x and E_y are orthogonal components of the incident electric field, the coupler outputs are amplified versions of E_x and iE_y, with the phase switch reversing their roles. Half of each output is bandpass filtered and rectified by a pair of detector diodes, while the other half passes into a 90° hybrid coupler. A second pair of bandpass filters and detector diodes measures the power from this coupler's outputs (Kangaslahti et al. 2006).

Synchronous demodulation of the 4-kHz phase switching yields measurements of Stokes $+Q$ and $-Q$ on the first two diodes and Stokes $+U$ and $-U$ on the remaining two. This high-frequency differencing suppresses low-frequency atmospheric fluctuations as well as $1/f$ noise from the amplifiers, detector diodes, bias electronics, and data-acquisition electronics. Subsequent demodulation of the 50-Hz phase switching removes spurious instrumental polarization generated by unequal transmission coefficients in the phase-switch circuits. The resulting four "double-demodulated" time streams are the polarization channels.

Averaging the output of each diode rather than demodulating it results in a measurement of Stokes I, hereafter called total power, denoted "TP." The TP time streams are useful for monitoring the weather and the stability of the detector responsivities, but suffer too much contamination from $1/f$ noise to constrain the CMB temperature anisotropy. Therefore, the Q-band instrument includes two correlation modules that are coupled to a pair of neighboring feed horns to measure the temperature difference between them, in a scheme similar to the *WMAP* differencing assemblies (Jarosik et al. 2003). These differential-temperature modules provide calibration data for the telescope pointing, beams, and sidelobes, as well as CMB data. Their feed horns are in the outer ring of the close-packed hexagonal array, $\approx 3°$ from the center.

Here we summarize several array-wide characteristics of the polarimeters. Bandpass measurements in the lab and at the start of the observing season find that the average center frequency is 43.1 ± 0.4 GHz, and the average bandwidth is 7.6 ± 0.5 GHz. We calculate the noise power spectra of the double-demodulated polarimeter time streams from each 40–90-minute observation to assess their $1/f$ knee frequencies and white-noise levels (see Section 5.1). The median $1/f$ knee frequency is 5.5 mHz, well below the telescope scan frequencies of 45–100 mHz.

From the white-noise levels and responsivities (Section 4.1) we find an array sensitivity[28] to CMB fluctuations of $69\,\mu\mathrm{K}\sqrt{\mathrm{s}}$, such that the mean polarized sensitivity per module is $280\,\mu\mathrm{K}\sqrt{\mathrm{s}}$.

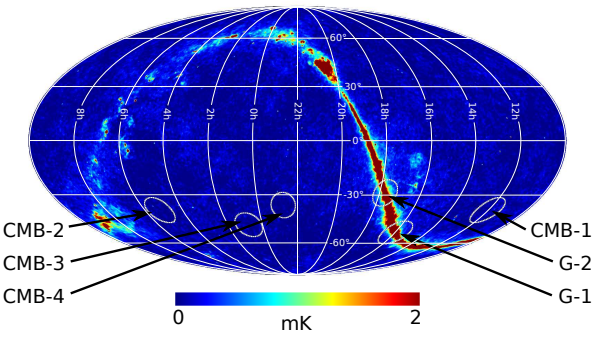

FIG. 2.— The CMB and Galactic patches, in equatorial coordinates, superimposed on a Q-band all-sky *WMAP* 7-year temperature map (Jarosik et al. 2010). Note that the Galactic-plane temperature signal saturates the color scale. Patch G-2 is the Galactic center.

3. OBSERVATIONS

QUIET is located on the Chajnantor plateau in the Atacama Desert of northern Chile (67°45′42″ W, 23°01′42″ S). A combination of high altitude (5080 m) and extreme dryness results in excellent observing conditions for most of the year. During the eight months of QUIET Q-band observations, the median precipitable

[28] This is the sensitivity for 62 polarization channels. Six of 68 polarization channels are non-functional—an array yield of 92%.

water vapor (PWV) measured at the nearby APEX site (Güsten et al. 2006) was 1.2 mm.

We began observations with the Q-band receiver on 2008 October 24, and took 3458 hours of data until 2009 June 13 (when the receiver was replaced on the telescope by the 90-element W-band receiver). Of these data, 77% are for CMB, with 12% of the observing time used for Galactic fields, 7% for calibration sources, and 4% cut due to obvious instrumental problems such as lack of telescope motion. We observe 24 hours a day, except when interrupted. Our full-season operating efficiency is 63%; causes of downtime include occasional snow, power outages, and mechanical failures.

TABLE 1
Patch Locations and Integration Times

Patch	RA (J2000)	Dec. (J2000)	Integration Hours
CMB-1	12^h04^m	$-39°00'$	905
CMB-2	05^h12^m	$-39°00'$	703
CMB-3	00^h48^m	$-48°00'$	837
CMB-4	22^h44^m	$-36°00'$	223
G-1	16^h00^m	$-53°00'$	311
G-2	17^h46^m	$-28°56'$	92

Note. — The central equatorial coordinates and integration times for each observing patch. G-1 and G-2 are Galactic patches.

3.1. *Field Selection*

We observe four CMB fields, referred to henceforth as "patches." Table 1 lists their center positions and total integration times, while Figure 2 indicates their positions on the sky. The number of patches is determined by the requirement to always have one patch above the lower elevation limit of the mount (43°). The specific positions of each patch were chosen to minimize foreground emission using *WMAP* 3-year data. The area of each patch is $\approx 250\,\mathrm{deg}^2$. In addition to the four CMB patches, we observe two Galactic patches. These allow us to constrain the spectral properties of the polarized low-frequency foregrounds with a high signal-to-noise ratio. The results from the Galactic observations will be presented in a future publication.

3.2. *Observing Strategy*

Scanning the telescope modulates the signal from the sky, converting CMB angular scales into frequencies in the polarimeter time streams. Since QUIET targets large angular scales, fast scanning ($\approx 5°\,\mathrm{s}^{-1}$ in azimuth) is critical to ensuring that the polarization modes of interest appear at higher frequencies than the atmospheric and instrumental $1/f$ knee frequencies.

So that each module sees a roughly-constant atmospheric signal, each QUIET scan is a constant-elevation scan (CES): periodic motion solely in azimuth with both the elevation and deck-rotation axes fixed. Each CES has an amplitude of 7.5° on the sky, with period 10–22 s. Typical CESes last 40–90 minutes. We repoint the telescope when the patch center has moved by 15° in order to build up data over an area of $\approx 15° \times 15°$ for each patch. Note that a central region $\simeq 8°$ across is observed by all polarimeters since the instrument's field of view has a diameter of $\simeq 7°$. Diurnal sky rotation and weekly deck rotations provide uniform parallactic-angle coverage of the patch, and ensure that its peripheral regions are also observed by multiple polarimeters.

TABLE 2
Regular Calibration Observations

Source	Schedule	Duration (min.)
sky dips	every 1.5 hours	3
Tau A	every 1–2 days	20
Moon	weekly	60
Jupiter	weekly	20
Venus	weekly	20
RCW38	weekly	20

4. CALIBRATION

Four quantities are required to convert polarimeter time streams into polarization power spectra: detector responsivities, a pointing model, detector polarization angles, and beam profiles. To this end, a suite of calibration observations is performed throughout the season using astronomical sources (Taurus A–hereafter Tau A, Jupiter, Venus, RCW38, and the Moon); atmospheric measurements ("sky dips," which typically consist of three elevation nods of ±3°); and instrumental sources (a rotating sparse wire grid and a polarized broadband noise source). From these we also measure instrumental polarization, as described below. QUIET's regular calibration observations are summarized in Table 2.

We typically use two or more methods to determine a calibration constant, taking the spread among the methods as an indication of the uncertainty. We show in Section 8 that aside from the case of absolute responsivity, all calibration uncertainties lead to estimates of systematic effects on the power spectra well below statistical errors. This immunity comes from having a large number of detectors and highly-crosslinked polarization maps.

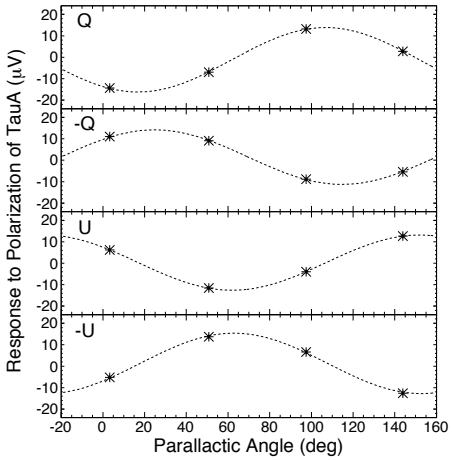

Fig. 3.— Polarimeter responses from the central feed horn to the polarization of Tau A at four parallactic angles. These data were collected with one correlation module in about 20 minutes. The errors are smaller than the points. From top to bottom, responses are shown for the detector diodes sensitive to the Stokes parameters +Q, −Q, +U, and −U, respectively. For each, the fitted model is plotted as a dashed line.

4.1. *Responsivity*

The polarized flux from Tau A provides a 5 mK signal which we observe at four parallactic angles. The sinusoidal modulation of the signal induced by the changing parallactic angles is fitted to yield responsivity coefficients for each detector. Figure 3 shows the response of the four polarization channels from the central feed horn to Tau A. A typical responsivity is 2.3 mV K^{-1}, with a precision from a single set of observations of 6%. The absolute responsivity from Tau A was measured most frequently for the central feed horn. We choose its +Q diode detector to provide the fiducial absolute responsivity.

The responsivities of other detectors relative to the fiducial detector are determined with the sky dips as described below. We have three independent means of assessing the relative responsivities among polarimeters: from nearly-simultaneous measurements of the Moon, from simultaneous measurements of responses to the rotating sparse wire grid in post-season tests, and from Tau A measurements. The errors from these methods are 4%, 2%, and 6% respectively, while the error from the sky-dip method is 4%. All the methods agree within errors.

Sky dips generate temperature signals of several 100 mK and thus permit measurement of the TP responsivities. The signals vary slightly with PWV. We estimate the slope from the data as 4% mm^{-1} and correct for it. This slope is consistent with the atmospheric model of Pardo et al. (2001). Because the ratios of the responsivities for the TP and polarized signals from each detector diode are stable quantities within a few percent of unity, we use sky dips performed at the beginning of each CES to correct short-term variations in the polarimeter responsivities. The responsivities vary by $\lesssim 10\%$ over the course of a day, due to changing thermal conditions for the bias electronics. Further post-season tests provide a physical model: the relevant temperatures are varied intentionally while the responsivities are measured with sky dips. We confirm the results with the polarized broadband source.

We bound the uncertainty in the absolute responsivity of the polarimeter array at 6%. The largest contributions to this estimate are uncertainties in (1) the beam solid angle (4%, see below), (2) the response difference between polarized and TP signals for each diode detector (3%), and (3) the Tau A flux (3%, Weiland et al. 2010). The first enters in converting the flux of Tau A into μK, while the second enters because although one fiducial diode detector is calibrated directly from Tau A, for the rest we find relative responsivities from sky dips and normalize by the fiducial diode's responsivity.

For the differential-temperature modules, all detectors observe the signal from Jupiter simultaneously, providing the absolute responsivity for all channels upon comparison with the Jupiter flux from Weiland et al. (2010). Observations of Venus (Hafez et al. 2008) and RCW38 agree with the Jupiter measurements within errors, and sky dips track short-term variations. We calibrate the absolute responsivity with 5% accuracy.

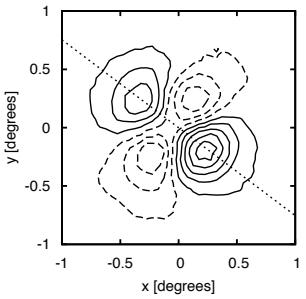

FIG. 4.— Map of the polarization of the Moon from one detector diode. The amplitude of the quadrupole polarization visible here is $\approx$ 400 mK. Similar maps are produced for all 17 polarization modules in the array with a single $\approx$ hour-long observation. The dotted line indicates the polarization orientation of the detector. Contours are spaced at intervals of 100 mK, with negative contours indicated by dashed lines.

4.2. *Pointing*

The global pointing solution derives from a physical model of the 3-axis mount and telescope tied to observations of the Moon with the central feed horn in the array, as well as Jupiter and Venus with the differential-temperature feed horns. Optical observations are taken regularly with a co-aligned star camera and used to monitor the time evolution of the pointing model.

During the first two months in the season, a mechanical problem with the deck-angle encoder results in pointing shifts. The problem was subsequently repaired. Based on pointing observations of the Moon and other astronomical sources, we verify that these encoder shifts are less than 2°. Systematic uncertainties induced by this problem are discussed in Section 8.1.

After the deck-angle problem is fixed, no significant evolution of the pointing model is found. The difference in the mean pointing solution between the start and the end of the season is smaller than 1′. Observations of the Moon and Jupiter also provide the relative pointing among the feed horns. The root mean square (RMS) pointing error in the maps is 3.′5.

4.3. *Detector Polarization Angles*

Our primary measurement of the polarization angle for each detector comes from observing the radial polarization of the Moon, as illustrated in Figure 4. The polarization angles are stable, changing by $< 0.^\circ 2$, except during the period with the deck-angle–encoder problem mentioned above.

Two other less precise methods also give estimates of the detector angles: fits to the Tau A data, and determination of the phases of the sinusoidal responses of all the detectors to rotation of the sparse wire grid. In each case, the differences between the detector angles determined by the secondary method and the Moon are described by a standard deviation of $\approx 3°$. However, we find a mean shift between the Tau A-derived and Moon-derived angles of $1.^\circ 7$. To estimate the errors in the angles in light of this shift, we use an empirical approach: in Section 8.2 we estimate the impact on the power spectra from using the Tau A results instead of the Moon results, and find it to be small.

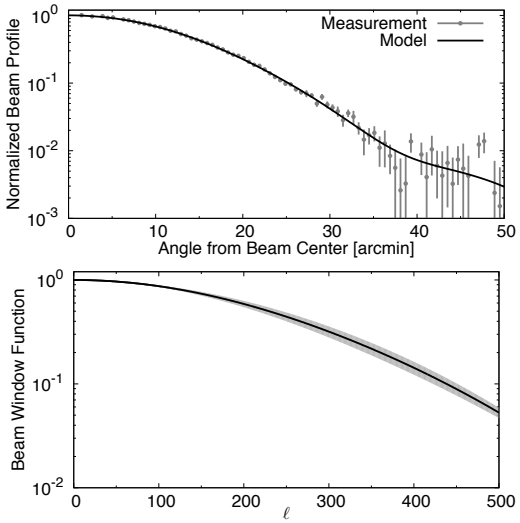

FIG. 5.— Top panel: Polarization beam profile from Tau A observations with the central feed horn. The data are overplotted with the expansion in Gauss-Hermite polynomials described in the text. Bottom panel: Beam window function with errors shown by the gray band.

4.4. *Beam Profile and Window Function*

The polarization and differential-temperature beams are obtained from maps created using the full data sets of Tau A and Jupiter observations respectively, with square pixels of 1.′8 on a side. For polarization, this process produces the main and leakage beam maps simultaneously, with the latter describing the instrumental polarization. The average FWHM for the beams across the array is 27.′3, measured with 0.′1 precision for the central feed horn and for the differential-temperature feed horns at the edge of the focal plane. The non-central–polarization-horn FWHMs are measured less frequently and thus are less precisely known, with an uncertainty of 1.′5. The beam elongation is typically small (1%), and its effect is further reduced by the diurnal sky rotation and weekly deck rotations which result in a symmetrized effective beam in the CMB maps. We compute 1-dimensional symmetrized beam profiles, with a resolution of 0.′6. These profiles are modeled as a sum of six even Gauss-Hermite terms (Monsalve 2010). The main-beam solid angles are computed by integrating these models out to 54′ (roughly -28 dB), yielding $78.0 \pm 0.4\,\mu$sr for the differential-temperature horns and $74.3 \pm 0.7\,\mu$sr for the central horn. An average gives $76\,\mu$sr for all horns in the array. We also examine alternative estimates such as integrating the raw beam map instead of the analytical fit. We assign a systematic uncertainty of 4% based on the differences among these different estimates. The systematic error includes possible contributions from sidelobes, which we constrain to $0.7 \pm 0.4\,\mu$sr with antenna range measurements carried out before the observation season.

The window functions, encoding the effect of the finite resolution of the instrument on the power spectra, are computed from the central-horn and the temperature-horn–profile models. The central-horn beam profile and window function are shown in Figure 5. The uncertainty accounts for statistical error and differences between polarization and differential-temperature beams, as described in Section 8.1.

4.5. *Instrumental Polarization*

Instrumental imperfections can lead to a spurious polarization signal proportional to the unpolarized CMB temperature anisotropy. We call this the I to Q (or U) leakage term. In our instrument, a fraction of the power input on one port of the correlation module is reflected because of a bandpass mismatch to the septum polarizer, and a fraction of the reflected power re-enters the other port. The dominant monopole term comes from this effect. We measure the monopole term from the polarimeter responses to temperature changes, using sky dips; Moon, Tau A, and Galactic signals; as well as variations from the weather. The average magnitude is 1.0% (0.2%) for the Q (U) diodes. Note that the discrepancy in the Q and U averages was predicted from measurements of the properties of the septum polarizers and confirmed in the field.

5. DATA ANALYSIS PROCEDURE

QUIET employs two independent analysis pipelines to derive CMB power spectra. We present the methods used for analysis in each pipeline, including data selection, filtering, map making, and power-spectra estimation.

Pipeline A is based on the pseudo-C_ℓ analysis framework, first described by Hivon et al. (2002), which is used by numerous experiments (Netterfield et al. 2002; Brown et al. 2009; Chiang et al. 2010; Larson et al. 2010; Lueker et al. 2010). This pipeline made all analysis choices in accordance with a strict (blind) analysis validation policy described in Section 6. An advantage of the pseudo-C_ℓ framework is computational efficiency, which is critical for completing the more than 30 iterations of the null-test suite. For the same reason, this pipeline is used for the systematic-error evaluations found in Section 8. Pseudo-C_ℓ analysis also enables us to perform cross correlation, making the resultant power spectra immune to possible misestimation of noise bias.

Pipeline B implements a maximum-likelihood framework (e.g., Tegmark 1997; Bond et al. 1998), which has a long history of use by CMB experiments (e.g., Mauskopf et al. 2000; Page et al. 2007; Wu et al. 2007; Bischoff et al. 2008). This framework yields minimum-variance estimates of the power spectra, naturally accounts for E/B mixing, and directly provides the exact CMB likelihood required for estimation of cosmological parameters, without the use of analytical approximations. In addition to power spectra, it produces unbiased maps with full noise-covariance matrices, useful for comparisons with other experiments. On the other hand, this approach is also computationally more expensive than the pseudo-C_ℓ framework, and a reduced set of null tests is therefore used to evaluate data consistency.

The processing of the time-ordered data (TOD) and the methodology used for data selection are treated in Sections 5.1 and 5.2, respectively. Brief descriptions of the pseudo-C_ℓ and maximum-likelihood techniques are found in Section 5.3. TOD processing, data selection, and analysis for temperature-sensitive modules are discussed in Section 5.4.

5.1. *Time-Ordered–Data Processing*

To prepare the TOD for map making, we execute three steps: pre-processing, noise modeling, and filtering. Of these steps, only the filtering is significantly different between the two pipelines.

5.1.1. *Pre-processing*

The first data-processing step is to correct for a small non-linearity that was discovered in the analog-to-digital converter (ADC) system. The non-linearities occur every 1024 bits; roughly 14% of the data are affected. Systematic uncertainty from this effect is estimated in Section 8.5. Next, the receiver data are synchronized with the telescope pointing. The double-demodulation step, described in Section 2, is applied, reducing the sample rate from 100 Hz to 50 Hz. A model of the detectors' polarized responsivities converts the data from ADC counts into thermodynamic temperature. The two pipelines use different responsivity models. Pipeline A applies a constant responsivity throughout each CES, addressing possible variability within a CES as part of the systematic error (Section 8); pipeline B updates responsivities on 2-minute timescales (Dumoulin 2010).

5.1.2. *Noise Model*

After pre-processing, the time streams for each detector diode in each CES are Fourier-transformed and their noise power spectra are fit to a model[29] with three parameters: the amplitude of white noise, the $1/f$ knee frequency, and the power-law slope of the $1/f$ noise. We also compute the white-noise correlations among detector diodes in the same module: the most important are between the two Q or the two U detector diodes (with an average coefficient of 0.22). A small fraction of the noise spectra contain features not accounted for in the noise model: beam sidelobes (see Section 2) scanning across features on the ground create a narrow spike at the scan frequency; slowly-changing weather patterns during a CES create a broader peak also at the scan frequency; and there are some narrow spikes at high ($\gtrsim 6$ Hz) frequencies. To prevent these features from biasing the noise model, the fit excludes a region around the scan frequency as well as frequencies above 4.6 Hz. In addition to the noise-model parameters, several statistics quantifying the agreement between the data and noise model are also used for data selection as described in Section 5.2.

5.1.3. *Filtering*

In pipeline A, three filters are applied. These were chosen from the results of many runs of the null-test suite (see Section 6). First, to remove the high-frequency narrow spikes, we apply a low-pass filter that cuts signals off sharply above 4.6 Hz[30]. Second, to suppress contamination from atmospheric fluctuations and detector $1/f$ noise, we subtract a linear function from each telescope half scan (left-going or right-going) removing modes below twice the scan frequency[31]. The third filter, designed to eliminate signal from ground emission, removes any azimuthal structure that remains after summing over all half scans in the CES.

In pipeline B, an apodized bandpass filter is used that accepts modes from 2.5 times the scan frequency to 4.5 Hz; the highpass component of this filter is designed to suppress scan-synchronous contamination. Further, a time-independent ground-emission model is subtracted. The model of ground emission is generated by building low-resolution and high–signal-to-noise maps in horizon coordinates from the full-season data for each deck angle and module, using large ($55'$) pixels. Only features that are stable in time, azimuth, elevation, and deck angle contribute to this model. The amplitude of the ground correction is $\lesssim 1\,\mu$K.

5.2. *Data Selection*

The fundamental unit of data used for analysis is the double-demodulated output of one detector diode for a single CES, referred to as a "CES-diode." Selecting only those CES-diodes that correspond to good detector performance and observing conditions is a critical aspect of the data analysis. The data-selection criteria began with a nominal set of cuts and evolved into several distinct configurations, as many as 33 in the case of pipeline A. For each configuration, analysis validation (see Section 6) was performed yielding statistics quantifying the lack of contamination in the data set. The final data set was chosen when these statistics showed negligible contamination and were little affected by changes to the cuts.

Cut efficiencies, defined as the fractions of CES-diodes accepted for the analysis, are given for both pipelines in Table 3. While each pipeline applies its own cuts uniformly to all four patches, the efficiencies among patches are non-uniform because of differences in weather quality. Over the course of the eight month observing season, patch CMB-1 is primarily visible at night, when the atmosphere tends to be more stable; patch CMB-3 is mostly observed during the day.

The first step of the data selection is simply to remove known bad data: data from six non-functional detector diodes, data during periods of mount malfunctions, and CESes lasting less than 1000 s. Further, we cut individual CES-diodes that show deviation from the expected linear relationship between the demodulated and TP signals. This cut removes data with poor thermal regulation of the electronics or cryostat, or residual ADC non-linearity.

The beam sidelobes, described in Section 2, introduce contamination to the data if the telescope scanning motion causes them to pass over the ground or the Sun. Ground pickup is dealt with by filtering as described in Section 5.1.3. The less frequent cases of Sun contamination are handled by cutting those CES-diodes for which the Sun's position overlaps with the measured sidelobe regions for each diode.

Additional cuts are specific to each pipeline. Pipeline A removes data taken during bad weather using a statistic calculated from fluctuations of the TP data during 10-s periods, averaged across the array. This cut removes entire CESes. Several more cuts remove individual CES-diodes. While these additional cuts are derived from the noise modeling statistics, they also target residual bad weather. During such marginal weather conditions only some channels need to be cut, since the sensitivity for

[29] At the level of a single CES, the TOD of each detector diode are dominated by noise; the contribution of the CMB is negligible.

[30] For QUIET's beam size and scanning speed a low-pass filter of 4.5–4.6 Hz results in a minimal loss of sensitivity to the CMB.

[31] Typical scan frequencies range from 45 mHz to 100 mHz.

a given detector diode to atmospheric fluctuations depends on its level of instrumental polarization. Next, we reject CES-diodes with poor agreement between the filtered data and the noise model in three frequency ranges: a narrow range (only 40 Fourier modes) about the scan frequency, from twice the scan frequency to 1 Hz, and from 1 Hz to 4.6 Hz. We also cut CES-diodes that have higher than usual $1/f$ knee frequencies, or large variations during the CES in the azimuthal slopes of the double-demodulated time streams; both these cuts help eliminate bad weather periods. Finally, we also remove any CES-diodes with an outlier greater than $6\,\sigma$ in the time domain on three timescales (20 ms, 100 ms, and 1 s).

For pipeline B, the weather cut rejects CESes based on a statistic computed from fluctuations of the double-demodulated signals from the polarization modules on 10-s and 30-s timescales. Three cuts are applied to remove individual CES-diodes. The first is a cut on the $1/f$ knee frequency, similar to that of pipeline A. Second, a cut is made on the noise model χ^2 in the frequency range passed by the filter, and third, we reject CES-diodes having a large χ^2 in the azimuth-binned TOD. This cut rejects data with possible time variation in the ground signal. Finally, an entire CES is removed if more than 40% of its detectors have already been rejected.

5.3. *Map Making and Power-Spectra Estimation*

After filtering, the TOD for all diodes are combined to produce Q and U maps for each of the QUIET patches. The maps use a HEALPix $N_{\rm side} = 256$ pixelization (Gorski et al. 2005). This section describes the map making and power-spectra estimation from the maps for each of the pipelines.

5.3.1. *Pipeline-A Map Making*

Polarization maps (Q and U) are made by summing samples into each pixel weighted by their inverse variance, calculated from the white-noise amplitudes. The full covariance matrix is not calculated. Two polarized sources, Centaurus A and Pictor A, are visible in the maps and are removed using circular top-hat masks with radii of 2° and 1°, respectively.

Separate maps are made for each range of telescope azimuth and deck-angle orientations. The coordinates are binned such that there are 10 divisions in azimuth[32] and six distinct ranges of deck-angle orientation. Making separate maps for different telescope pointings enables the cross correlation described in the next section.

TABLE 3
Total Hours Observed and Data-Selection Efficiencies

Patch	Total Hours	A %	B %	Common %
CMB-1	905	81.7	84.3	76.7
CMB-2	703	67.3	70.0	61.2
CMB-3	837	56.0	61.4	51.4
CMB-4	223	70.6	74.2	65.9
All Patches	2668	69.4	72.9	64.2

Note. — Selection efficiencies for each pipeline. "Common" gives the efficiencies if both sets of cuts were applied.

[32] The azimuth divisions are the same for all patches, which means that not all divisions are populated for patches CMB-3 and CMB-4.

5.3.2. *Power-Spectra Estimation in Pipeline A*

The MASTER (Monte Carlo Apodized Spherical Transform Estimator) method is used in pipeline A (Hivon et al. 2002; Hansen & Gorski 2003); it is based on a pseudo-C_ℓ technique and takes account of effects induced by the data processing using Monte Carlo (MC) simulations. The pseudo-C_ℓ method allows estimation of the underlying C_ℓ using spherical-harmonics transformations when the observations do not cover the full sky uniformly (Wandelt et al. 2001). The pseudo-C_ℓ spectrum, designated by $\tilde{C}_\ell$, is related to the true spectrum C_ℓ by:

$$\langle \tilde{C}_\ell \rangle = \sum_{\ell'} M_{\ell\ell'} F_{\ell'} B^2_{\ell'} \langle C_{\ell'} \rangle. \quad (1)$$

There is no term corresponding to noise bias, which would arise if we did not employ a cross-correlation technique. Here B_ℓ is the beam window function, described in Section 4.4, and $M_{\ell\ell'}$ is a mode-mode–coupling kernel describing the effect of observing only a small fraction of the sky with non-uniform coverage. It is calculable from the pixel weights, which are chosen to maximize the signal-to-noise ratio (Feldman et al. 1994). We bin in ℓ and recover C_ℓ in nine band powers, C_b, and F_ℓ is the transfer function (displayed in Section 7) due to filtering of the data; its binned estimate, F_b, is found by processing noiseless CMB simulations through pipeline A and used to obtain C_b. For the polarization power spectra, equation (1) is generalized for the case where $\tilde{C}_\ell$ contains both $\tilde{C}^{EE}_\ell$ and $\tilde{C}^{BB}_\ell$.

In the power-spectra estimates, we include only the cross correlations among pointing-division maps, excluding the auto correlations. Because the noise is uncorrelated for different pointing divisions, the cross-correlation technique allows us to eliminate the noise-bias term and thus the possible residual bias due to its misestimate. Cross correlation between different pointing divisions also suppresses possible effects of ground contamination and/or time-varying effects. Dropping the auto correlations creates only a small increase in the statistical errors ($\approx 3\%$) on the power spectra.

The errors estimated for the pipeline-A power spectra are frequentist two-sided 68% confidence intervals. A likelihood function used to compute the confidence intervals is modeled following Hamimeche & Lewis (2008) and calibrated using the MC simulation ensemble of more than 2000 realizations with and without CMB signal. We also use the likelihood function to put constraints on r and calculate the consistency to ΛCDM.

The partial sky coverage of QUIET generates a small amount of E/B mixing (Challinor & Chon 2005), which contributes an additional variance to the BB power spectrum. We incorporate it as part of the statistical error. This mixing can be corrected (Smith & Zaldarriaga 2007) in future experiments where the effect is not negligible compared to instrumental noise.

5.3.3. *Pipeline-B Map Making*

In pipeline B, the pixel-space sky map $\hat{\mathbf{m}}$ ($N_{\rm side} = 256$) is given by

$$\hat{\mathbf{m}} = \left(\mathbf{P}^T\mathbf{N}^{-1}\mathbf{F}\mathbf{P}\right)^{-1}\mathbf{P}^T\mathbf{N}^{-1}\mathbf{F}\mathbf{d}, \quad (2)$$

where $\mathbf{P}$ is the pointing matrix, $\mathbf{N}$ is the TOD–noise-covariance matrix, $\mathbf{F}$ corresponds to the apodized band-pass filter discussed in Section 5.1.3, and $\mathbf{d}$ denotes the TOD. This map is unbiased, and for the case $\mathbf{F} = \mathbf{1}$ it is additionally the maximum-likelihood map, maximizing

$$\mathcal{L}(\mathbf{m}|\mathbf{d}) = \mathrm{e}^{-\frac{1}{2}(\mathbf{d}-\mathbf{P}\mathbf{m})^T\mathbf{N}^{-1}(\mathbf{d}-\mathbf{P}\mathbf{m})}. \quad (3)$$

The corresponding map–noise-covariance matrix (e.g., Tegmark 1997; Keskitalo et al. 2010) is

$$\mathbf{N}_{\hat{\mathbf{m}}} = \left(\mathbf{P}^T\mathbf{N}^{-1}\mathbf{F}\mathbf{P}\right)^{-1}\left(\mathbf{P}^T\mathbf{F}^T\mathbf{N}^{-1}\mathbf{F}\mathbf{P}\right)\left(\mathbf{P}^T\mathbf{N}^{-1}\mathbf{F}\mathbf{P}\right)^{-1}. \quad (4)$$

Note that one often encounters the simplified expression $\mathbf{N}_{\hat{\mathbf{m}}} = \left(\mathbf{P}^T\mathbf{N}^{-1}\mathbf{F}\mathbf{P}\right)^{-1}$ in the literature. This corresponds effectively to assuming that $\mathbf{F} = \mathbf{F}^2$ in the Fourier domain, and is strictly valid for top-hat–filter functions only. For our filters, we find that the simplified expression biases the map-domain $\chi^2 (\equiv \hat{\mathbf{n}}^T\mathbf{N}_{\hat{\mathbf{m}}}^{-1}\hat{\mathbf{n}}$, where $\hat{\mathbf{n}}$ is a noise-only map) by $\approx 3\,\sigma$, and we therefore use the full expression, which does lead to an unbiased χ^2.

Equations (2–4) apply to both polarization and temperature analysis. The only significant difference lies in the definition of the pointing matrix, $\mathbf{P}$. For polarization, $\mathbf{P}$ encodes the detector orientation, while for temperature it contains two entries per time sample, $+1$ and -1, corresponding to the two horns in the differential-temperature assembly.

After map making, the maps are post-processed by removing unwanted pixels (i.e., compact sources and low–signal-to-noise edge pixels). All 54 compact sources in the 7-year *WMAP* point source catalog (Gold et al. 2010) present in our four patches are masked out, for a total of 4% of the observed area. We also marginalize over large-scale and unobserved modes by projecting out all modes with $\ell \leq 5$ ($\ell \leq 25$ for temperature) from the noise-covariance matrix using the Woodbury formula, assigning infinite variance to these modes.

5.3.4. *Power-Spectra Estimation in Pipeline B*

Given the unbiased map estimate, $\hat{\mathbf{m}}$, and its noise-covariance matrix, $\mathbf{N}_{\hat{\mathbf{m}}}$, we estimate the binned CMB power spectra, C_b, using the Newton–Raphson optimization algorithm described by Bond et al. (1998), generalized to include polarization. In this algorithm one iterates towards the maximum-likelihood spectra by means of a local quadratic approximation to the full likelihood. The iteration scheme in its simplest form is

$$\delta C_b = \frac{1}{2}\sum_{b'}\mathcal{F}^{-1}_{bb'}\mathrm{Tr}\left[(\hat{\mathbf{m}}\hat{\mathbf{m}}^T - \mathbf{C})(\mathbf{C}^{-1}\mathbf{C}_{,b'}\mathbf{C}^{-1})\right], \quad (5)$$

where b denotes a multipole bin, $\mathbf{C}$ is the signal-plus-noise pixel-space covariance matrix, and $\mathbf{C}_{,b}$ is the derivative of $\mathbf{C}$ with respect to C_b. The signal component of $\mathbf{C}$ is computed from the binned power spectra, C_b, and the noise component is based on the noise model described in Section 5.1.2, including diode-diode correlations. Finally,

$$\mathcal{F}_{bb'} = \frac{1}{2}\mathrm{Tr}(\mathbf{C}^{-1}\mathbf{C}_{,b}\mathbf{C}^{-1}\mathbf{C}_{,b'}) \quad (6)$$

is the Fisher matrix. Additionally, we introduce a step length multiplier, α, such that the actual step taken at iteration i is $\alpha\,\delta C_b$, where $0 < \alpha \leq 1$ guarantees that $\mathbf{C}$ is positive definite. We adopt the diagonal elements of the Fisher matrix as the uncertainties on the band powers.

We start the Newton–Raphson search at $C_\ell = 0$, and iterate until the change in the likelihood value is lower than 0.01 times the number of free parameters, corresponding roughly to a 0.01-σ uncertainty in the position of the multivariate peak. Typically we find that 3 to 10 iterations are required for convergence.

Estimation of cosmological parameters, θ, is done by brute-force grid evaluation of the pixel-space likelihood,

$$\mathcal{L}(\theta) \propto \frac{-\frac{1}{2}\mathbf{d}^T\mathbf{C}^{-1}(\theta)\mathbf{d}}{\sqrt{|\mathbf{C}(\theta)|}}. \quad (7)$$

Here $\mathbf{C}(\theta)$ is the covariance matrix evaluated with a smooth spectrum, C_ℓ, parametrized by θ. In this paper, we only consider 1-dimensional likelihoods with a parametrized spectrum of the form $C_\ell = a\,C_\ell^{\mathrm{fid}}$, a being a scale factor and C_ℓ^{fid} a reference spectrum; the computational expense is therefore not a limiting factor. Two different cases are considered, with a being either the tensor-to-scalar ratio, r, or the amplitude of the EE spectrum, q, relative to the ΛCDM model.

5.4. *Temperature Data Selection and Analysis*

As described in Section 2, we dedicate one pair of modules to differential-temperature measurements. While these modules are useful for calibration purposes, when combined with our polarization data they also enable us to make self-contained measurements of the TE and TB power spectra.

For temperature, both pipelines adopt the pipeline-A data-selection criteria used for polarization analysis (see Section 5.2). The temperature-sensitive modules, however, are far more susceptible to atmospheric contamination than the polarization modules. Thus, these cuts result in reduced efficiencies: 12.4%, 6.9%, and 6.8% for patches CMB-1, CMB-2, and CMB-3, respectively[33]. More tailoring of the cuts for these modules would improve efficiencies.

In pipeline A, the analysis proceeds as described in Sections 5.1.3, 5.3.1, and 5.3.2 except for two aspects. First, in the TOD processing a second-order polynomial is fit and removed from each telescope half scan instead of a linear function. This suppresses the increased contamination from atmospheric fluctuations in the temperature data. Second, we employ an iterative map maker based on the algorithm described by Wright et al. (1996). Map making for differential receivers requires that each pixel is measured at multiple array pointings or crosslinked. In order to improve crosslinking we divide the temperature data into only four maps by azimuth and deck angle, rather than the 60 divisions used for polarization analysis. To calculate TE and TB power spectra, polarization maps are made for these four divisions, plus one additional map that contains all polarization data with pointings not represented in the temperature data.

For pipeline B the algorithms for making temperature maps and estimating power spectra are identical to the polarization case, as described in Sections 5.3.3 and 5.3.4.

[33] Patch CMB-4 is excluded due to low data-selection efficiency and a lack of sufficient crosslinking.

6. ANALYSIS VALIDATION

The QUIET data analysis follows a policy of not looking at the power spectra until the analysis is validated using a set of predefined tests for possible systematic effects[34]. The validation tests consist of a suite of null tests, comparisons across multiple analysis configurations, and consistency checks among power spectra from different CMB patches. Data-selection criteria, filtering methods, and the division of data into maps for cross correlation in pipeline A are all evaluated based on the test results. We finalize all aspects of the data analysis including calibration and evaluation of the systematic error before unveiling the power spectra (blind analysis). The risk of experimenter bias is thereby eliminated.

Details of tests found in this section describe pipeline A. While the pipeline B analysis follows a similar program of null tests to verify the result, the increased computational requirements of the maximum-likelihood framework limit the number of tests that could be performed and require those tests to be run using lower-resolution maps than for the non-null analysis. The bulk of this section treats validation of the polarization analysis; at the end, we briefly describe the temperature analysis validation.

In a null test, the data are split into two subsets. Maps, m_1 and m_2, are made from each subset. The power spectra of the difference map, $m_{\rm diff} \equiv (m_1 - m_2)/2$, are analyzed for consistency with the hypothesis of zero signal. The null suite consists of 42 tests[35], each targeting a possible source of signal contamination or miscalibration. These are highly independent tests; the data divisions for different null tests are correlated at only 8.8% on average. Nine tests divide the data by detector diode based on susceptibility to instrumental effects, such as instrumental polarization. Ten tests target effects that depend on the telescope pointing such as data taken at high or low elevation. Five tests divide based on the proximity of the main or sidelobe beams to known sources such as the Sun and Moon. Eight tests target residual contamination in the TOD using statistics mentioned in Section 5.2. Ten tests divide the data by environmental conditions such as ambient temperature or humidity.

Each null test yields EE and BB power spectra in nine ℓ bins, calculated separately for each CMB patch. Figure 6 shows the power spectra from one null test. Although the EB spectra are also calculated for each null test, they are assigned lesser significance since sources of spurious EB power will also result in the failure of EE and BB null tests. Combining all EE and BB points for all patches and null tests in the null suite yields a total of 3006 null-spectrum points. For each power-spectrum bin b, we calculate the statistic $\chi_{\rm null} \equiv C_b^{\rm null}/\sigma_b$, where $C_b^{\rm null}$ is the null power and σ_b is the standard deviation of $C_b^{\rm null}$ in MC simulations. We evaluate both $\chi_{\rm null}$ and its square; $\chi_{\rm null}$ is sensitive to systematic biases in the null spectra while $\chi^2_{\rm null}$ is more responsive to outliers. We run MC simulations of the full null suite to take into account the small correlation among the null tests and the slight non-Gaussianity of the $\chi_{\rm null}$ distribution. Non-Gaussianity is caused by the small number of modes at low ℓ.

[34] Some systematic effects, such as a uniform responsivity-calibration error, cannot be detected by these techniques, and are addressed in Section 8.

[35] Only 41 null tests are performed for patch CMB-4; one test is dropped because there are no data in one of the subsets.

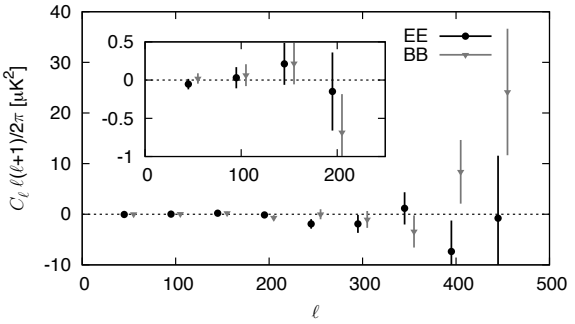

FIG. 6.— EE and BB power spectra for the patch CMB-1 null test between Q and U detector diodes. The inset shows the low-ℓ region in detail.

As we refine the data-selection criteria based on the results of the null suite, we use a second test to monitor changes in the non-null power spectra. Using a blind analysis framework, we compute the difference of the power spectra between any two iterations of the data selection without revealing the non-null spectra. Further, we randomize the sign of the difference to hide the direction of the change; knowledge of the direction could allow experimenter bias (e.g. a preference for low BB power). Figure 7 shows the differences in the power spectra between the final configuration and several intermediate iterations of the data selection, starting with data sets that showed significant failures for the null-test suite. Statistically significant differences indicate a change in the level of contamination in the selected data set. Our data-selection criteria are finalized when further iterations only result in statistically expected fluctuations. The sensitivity of this test is demonstrated by the fact that the expected fluctuations are much less than the statistical error of the final result.

Finally, the non-null power spectra are compared among the four CMB patches. A χ^2 statistic is computed from the deviation of each patch's non-null power spectra from the weighted average over all patches. The total χ^2 is compared to MC simulations to compute probabilities to exceed (PTE).

When all aspects of the analysis are finalized, the last round of null tests and CMB patch comparisons validates the non-null–power-spectra results. Figure 8 shows the distributions of the $\chi_{\rm null}$ statistic and of the PTEs corresponding to all $\chi^2_{\rm null}$ values from the full null suite. In pipeline A, the distribution of $\chi_{\rm null}$ is consistent with the expectation from MC simulations. The mean of the $\chi_{\rm null}$ distribution is 0.02 ± 0.02; the mean of the MC-ensemble $\chi_{\rm null}$ distribution is also consistent with zero. The distribution of the $\chi^2_{\rm null}$ PTEs is uniform as expected. Table 4 lists the PTEs for the sums of the $\chi^2_{\rm null}$ statistic over all bins in each patch. Examinations of various subsets of the null suite, such as EE or BB only, do not reveal any anomalies. The EB null spectra do not indicate any failure either. Patch comparison PTEs are 0.16, 0.93, and 0.40 for EE, BB, and EB, respectively, demonstrating no statistically significant difference among the patches.

A similar, but smaller, null suite is run by pipeline B.

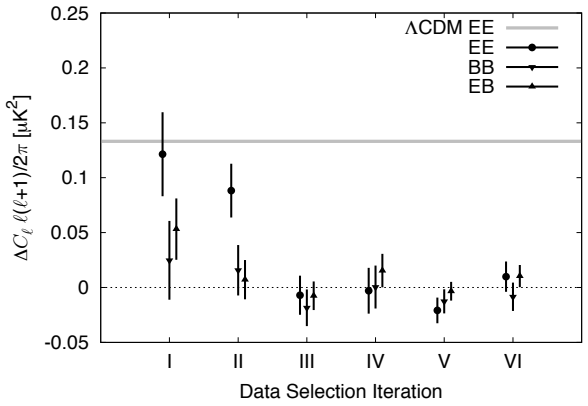

FIG. 7.— Power-spectra differences between the final data selection and six of the 32 earlier data-selection iterations, ordered by date. The lowest-ℓ bin of patch CMB-1 is shown. The error bars correspond to the expected fluctuations due to the differences in data selected, which are much smaller than the final statistical errors in this bin ($\approx 0.10\,\mu$K^2 for BB). Iterations that are closer to the final data selection have smaller errors. The expected EE power in this bin from the ΛCDM model is also shown for comparison.

Specifically, 21 null tests are made at a HEALPix resolution of $N_{\rm side} = 128$. The results obtained in these calculations are summarized in the bottom panel of Figure 8, and total PTEs for each patch are listed in Table 4. As in pipeline A, no anomalous values are found.

Finally, we make a comment on the usefulness of the $\chi_{\rm null}$ distribution (as opposed to the $\chi^2_{\rm null}$ distribution) for identifying and quantifying potential contaminants. During the blind stage of the analysis, a positive bias in the $\chi_{\rm null}$ distribution of 0.21 (0.19) was identified using pipeline A (B) (corresponding to 21% (19%) of the statistical errors). The number from pipeline A was obtained when including auto correlations in its power-spectra estimator. When excluding auto correlations, and cross-correlating maps made from data divided by time (day by day), the bias decreased to 0.10. Further detailed studies lead to the division of data into maps based on the telescope pointing, as described in Section 5.3; the result is an elimination of the observed bias.

The maximum-likelihood technique employed by pipeline B intrinsically uses auto correlations, and a corresponding shift in the $\chi_{\rm null}$ distribution is seen in Figure 8. However, as will be seen in Section 7, the power spectra from the two pipelines are in excellent agreement, thereby confirming that any systematic bias coming from including auto correlations is well below the level of the statistical errors. We close this section by mentioning that we know of no other CMB experiment reporting an examination of the $\chi_{\rm null}$ distribution, which is sensitive to problems not detected by examining the $\chi^2_{\rm null}$ distribution only.

6.1. *Validation of the Temperature Analysis*

A smaller number of null tests is used for the temperature analysis. Several are not applicable and others are discarded due to lack of data with sufficient crosslinking. Even so, we are able to run suites of 29, 27, and 23 TT null tests on patches CMB-1, CMB-2, and CMB-3, respectively. We calculate the sums of $\chi^2_{\rm null}$ statistics, yielding PTEs of 0.26 and 0.11 for patches CMB-1 and CMB-2, respectively. No significant outliers are found for these patches. However, a 5-σ outlier in a single test[36] is found in patch CMB-3, implying contamination in its temperature map. CMB-3 is therefore excluded from further analysis. We confirm consistency between the patches CMB-1 and CMB-2 with a PTE of 0.26.

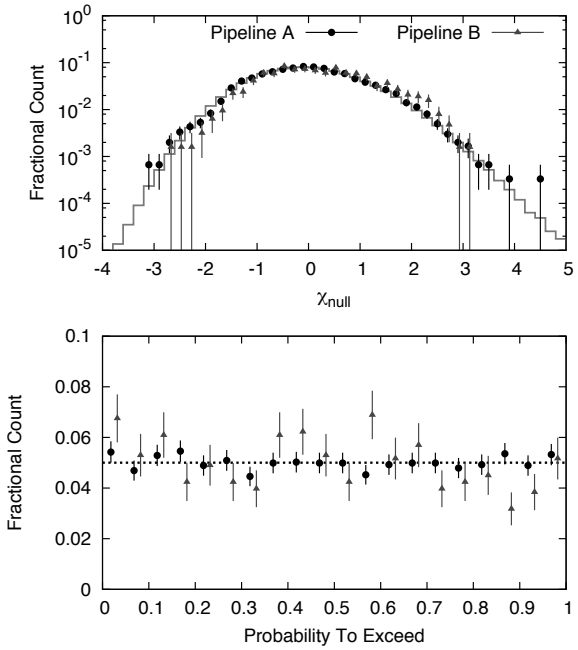

FIG. 8.— Null-Suite Statistics. The upper panel shows a histogram of the $\chi_{\rm null}$ values for the pipeline-A null suite (circles), pipeline-B null suite (triangles), and the average of 1024 MC realizations of the pipeline-A null suite (gray histogram). Both data and MC distributions show similar non-Gaussianity in the $\chi_{\rm null}$ statistic. The shift in $\chi_{\rm null}$ seen for pipeline B, also seen in earlier iterations of pipeline A, is discussed in the text. The lower panel shows a histogram of PTEs calculated from the $\chi^2_{\rm null}$ statistic (outliers from either side of the upper distribution manifest as low PTEs).

TABLE 4
NULL SUITE PROBABILITY TO EXCEED BY PATCH

Patch	Pipeline A %	Pipeline B %
CMB-1	44	7
CMB-2	19	43
CMB-3	16	23
CMB-4	68	28

NOTE. — PTEs calculated from the sums of the $\chi^2_{\rm null}$ statistics, for EE and BB spectra points, over the null tests for each patch.

With no significant contamination in TT, EE, or BB spectra, one may be confident that the TE and TB spectra are similarly clean. For confirmation, we calculate TE and TB null spectra for the five null tests that are common to the temperature and polarization analyses. These yield PTEs of 0.61 and 0.82 for TE, and 0.16 and 0.55 for TB, for patches CMB-1 and CMB-2, respectively, with no significant outliers. Patch consistency checks give PTEs of 0.48 for TE and 0.26 for TB. Thus, the TE and TB power spectra, as well as the TT, pass all validation tests that are performed.

[36] This null test divides the data based on array pointing.

7. RESULTS

We report results from the first season of QUIET Q-band observations: CMB power spectra, derived foreground estimates, and constraints on the tensor-to-scalar ratio, r.

7.1. *Polarization Power Spectra*

The CMB power spectra are reported in nine equally-spaced bands with $\Delta\ell = 50$, beginning at $\ell_{\rm min} = 25$. Given the patch size, modes with $\ell < \ell_{\rm min}$ cannot be measured reliably. The correlation between neighboring bins is typically -0.1; it becomes negligible for bins further apart.

The EE, BB, and EB polarization power spectra estimated by both pipelines are shown in Figure 9. The agreement between the results obtained by the two pipelines is excellent, and both are consistent with the ΛCDM concordance cosmology. Our findings and conclusions are thus fully supported by both pipelines. Only the statistical uncertainties are shown here; we treat systematic errors in Section 8. Because the systematic error analysis was only done for pipeline A, we adopt its power-spectra results (tabulated in Table 5) as the official QUIET results.

The bottom sub-panels in Figure 9 show the window and transfer functions for each bin computed by pipeline A. Figure 10 shows the maps for patch CMB-1 computed by pipeline B, and Figure 11 shows the QUIET power spectra in comparison with the most relevant experiments in our multipole range. Additional plots and data files are online[37].

Fitting only a free amplitude, q, to the EE spectrum[38] relative to the 7-year best-fit *WMAP* ΛCDM spectrum (Larson et al. 2010), we find $q = 0.87 \pm 0.10$ for pipeline A and $q = 0.94 \pm 0.09$ for pipeline B. Taking into account the full non-Gaussian shapes of the likelihood functions, both results correspond to more than a 10-σ detection of EE power. In particular, in the region of the first peak, $76 \leq \ell \leq 175$, we detect EE polarization with more than 6-σ significance, confirming the only other detection of this peak made by BICEP at higher frequencies. The χ^2 relative to the ΛCDM model, with $C_\ell^{EB} = C_\ell^{BB} = 0$, is 31.6 (24.3) with 24 degrees of freedom, corresponding to a PTE of 14% (45%) for pipeline A (B).

7.2. *Foreground Analysis*

In order to minimize possible foreground contamination, QUIET's four CMB patches were chosen to be far from the Galactic plane and known Galactic synchrotron spurs. In these regions, contributions from thermal dust emission are negligible in Q band. Spinning dust is expected to be polarized at no more than a few percent in Q band (Battistelli et al. 2006; Lopez-Caraballo et al. 2010), so we expect the contribution to polarized foreground emission in our patches to be small. We therefore consider only two dominant sources of possible foreground contamination, namely compact radio sources and Galactic diffuse synchrotron emission.

[37] http://quiet.uchicago.edu/results/index.html

[38] Only $\ell \geq 76$ are used in the EE fit and the χ^2 calculation relative to ΛCDM because the first EE bin has a significant foreground contribution; see Section 7.2.

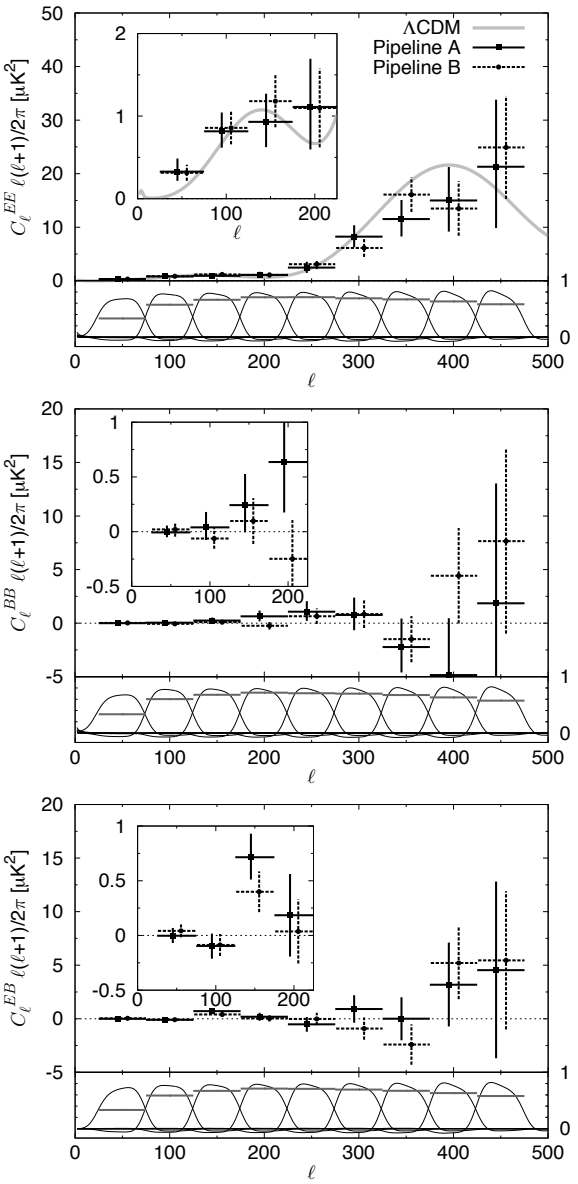

FIG. 9.— EE, BB, and EB power spectra from each QUIET pipeline, all four patches combined. The insets show the low-ℓ region in detail. Window and transfer functions for each ℓ bin are shown below the corresponding power spectra in black and gray, respectively. The window function combines the mode-mode–coupling kernel $M_{\ell\ell'}$ with the beam (B_ℓ) and represents, in combination with the transfer function (F_ℓ), the response in each band to the true C_ℓ spectrum. The EE point in the lowest-ℓ bin includes foreground contamination from patch CMB-1. For this display, pipeline A shows frequentist 68% confidence intervals while pipeline B uses the diagonal elements of the Fisher matrix; the difference is most pronounced in the lowest-ℓ bin where the likelihood is the most non-Gaussian.

To limit the effect of compact radio sources, we apply a compact-source mask to our maps before computing the power spectra, as described in Section 5. We also evaluate the CMB spectra both with and without the full *WMAP* temperature compact-source mask (Gold et al. 2010), and find no statistically significant changes. The possible contribution from compact radio sources with

TABLE 5
CMB-Spectra Band Powers from QUIET Q-Band Data

ℓ bin	EE	BB	EB
25-75	[a]$0.33^{+0.16}_{-0.11}$	$-0.01^{+0.06}_{-0.04}$	$0.00^{+0.07}_{-0.07}$
76-125	$0.82^{+0.23}_{-0.20}$	$0.04^{+0.14}_{-0.12}$	$-0.10^{+0.11}_{-0.12}$
126-175	$0.93^{+0.34}_{-0.31}$	$0.24^{+0.28}_{-0.25}$	$0.71^{+0.22}_{-0.20}$
176-225	$1.11^{+0.58}_{-0.52}$	$0.64^{+0.53}_{-0.46}$	$0.18^{+0.38}_{-0.38}$
226-275	$2.46^{+1.10}_{-0.99}$	$1.07^{+0.98}_{-0.86}$	$-0.52^{+0.68}_{-0.69}$
276-325	$8.2^{+2.1}_{-1.9}$	$0.8^{+1.6}_{-1.4}$	$0.9^{+1.3}_{-1.3}$
326-375	$11.5^{+3.6}_{-3.3}$	$-2.2^{+2.7}_{-2.4}$	$0.0^{+2.0}_{-2.0}$
376-425	$15.0^{+6.2}_{-5.8}$	$-4.9^{+5.3}_{-4.9}$	$3.2^{+3.9}_{-3.9}$
426-475	21^{+13}_{-11}	2^{+11}_{-10}	$4.5^{+8.3}_{-8.2}$

Note. — Units are thermodynamic temperatures, μK^2, scaled as $C_\ell \ell(\ell+1)/2\pi$.

[a]Patch CMB-1 has significant foreground contamination in the first EE bin.

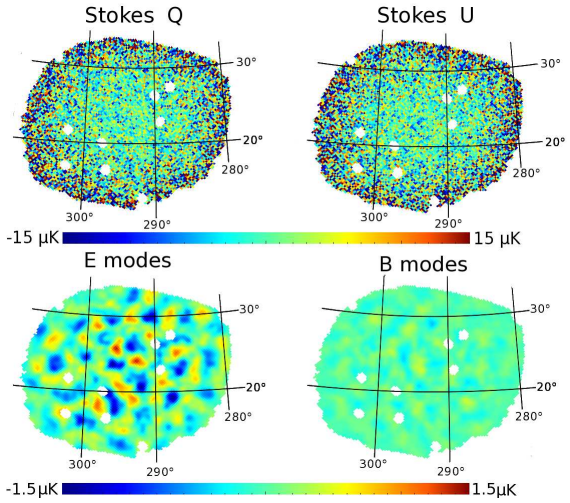

Fig. 10.— Maps of patch CMB-1 in Galactic coordinates. The top row shows our polarization maps with compact sources masked (white disks). The bottom row shows E and B modes decomposed using a generalized Wiener filter technique, implemented through Gibbs sampling (Eriksen et al. 2004; Larson et al. 2007), including only modes for $\ell \geq 76$ and smoothed to 1° FWHM; lower multipoles are removed due to a significant foreground contribution. Note the clear difference in amplitude: the E modes show a high–signal-to-noise cosmological signal while the B modes are consistent with noise.

fluxes below the *WMAP* detection level (1 Jy) is small: $0.003\,\mu$K^2 at $\ell = 50$ and $0.01\,\mu$K^2 at $\ell = 100$ (Battye et al. 2010). We therefore conclude that our results are robust with respect to contamination from compact radio sources and that the dominant foreground contribution comes from diffuse synchrotron emission.

In Figure 12 we show the power spectra measured from each patch. The CMB-1 EE band power for the first bin is $0.55 \pm 0.14\,\mu$K^2, a 3-σ outlier relative to the expected ΛCDM band power of $0.13\,\mu$K^2; while not significant enough to spoil the overall agreement among the patches as shown in Section 6, this is a candidate for a bin with foreground contamination.

To estimate the Q-band polarized synchrotron contamination in our CMB patches, we process the *WMAP*7 K-band (23-GHz) map through pipeline A and estimate its band power, $\hat{C}_b^{\rm KK}$, as well as the cross spectra with the QUIET Q-band data, $\hat{C}_b^{\rm QK}$. These results are shown for the first bin ($25 \leq \ell \leq 75$; $b = 1$) in Table 6, together with the corresponding QUIET band powers, $\hat{C}_b^{\rm QQ}$. Since foregrounds do not contribute to the sample variance, the uncertainties for $\hat{C}_{b=1}^{\rm KK}$ and $\hat{C}_{b=1}^{\rm QK}$ are given by instrumental noise only, including contributions from both *WMAP* and QUIET. For $\hat{C}_{b=1}^{\rm QQ}$, sample variance as predicted by the ΛCDM model is also included.

There is significant EE power in patch CMB-1 as measured by $\hat{C}_{b=1}^{\rm KK}$. We also find a correspondingly significant cross correlation between the *WMAP* K band and the QUIET Q band, confirming that this excess power is not due to systematic effects in either experiment and is very likely a foreground. No significant power is found in any other case. The non-detection of foreground power at $\ell > 75$ is consistent with the expected foreground dependence: $\propto \ell^{-2.5}$ (Carretti et al. 2010), and the low power found in $\hat{C}_{b=1}^{\rm KK}$.

The excess power observed in the first EE bin of CMB-1 is fully consistent with a typical synchrotron frequency spectrum. To see this, we extrapolate $\hat{C}_{b=1}^{\rm KK}$ from K band to Q band, assuming a spectral index of $\beta = -3.1$ (Dunkley et al. 2009), and calculate the expected power in $C_{b=1}^{\rm QK}$ and $C_{b=1}^{\rm QQ}$,

$$C_{b=1}^{\rm QK} = \frac{1.05}{1.01}\left(\frac{43.1}{23}\right)^{\beta} \hat{C}_{b=1}^{\rm KK} = 2.57 \pm 0.69\,\mu{\rm K}^2\,, \quad (8)$$

$$C_{b=1}^{\rm QQ} = \left[\frac{1.05}{1.01}\left(\frac{43.1}{23}\right)^{\beta}\right]^2 \hat{C}_{b=1}^{\rm KK} = 0.38 \pm 0.10\,\mu{\rm K}^2\,, \quad (9)$$

where the prefactor accounts for the fact that β is defined in units of antenna temperature, and the uncertainties are scaled from that of $\hat{C}_{b=1}^{\rm KK}$. These predictions are fully consistent with the observed values of $\hat{C}_{b=1}^{\rm QK}$ and $\hat{C}_{b=1}^{\rm QQ}$, when combined with the ΛCDM-expected power. We conclude that the excess power is indeed due to synchrotron emission.

7.3. *Constraints on Primordial B modes*

We constrain the tensor-to-scalar ratio, r, using the QUIET measurement of the BB power spectrum at low multipoles ($25 \leq \ell \leq 175$). Here r is defined as the ratio of the primordial–gravitational-wave amplitude to the curvature-perturbation amplitude at a scale $k_0 = 0.002\,{\rm Mpc}^{-1}$. We then fit our measurement to a BB-spectrum template computed from the ΛCDM concordance parameters with r allowed to vary. For simplicity, we fix the tensor spectral index at $n_t = 0$ in computing the template[39]. This choice makes the BB–power-spectrum amplitude directly proportional to r.

For pipeline A, we find $r = 0.35^{+1.06}_{-0.87}$, corresponding to $r < 2.2$ at 95% confidence. Pipeline B obtains $r = 0.52^{+0.97}_{-0.81}$. The results are consistent; the lower panel of Figure 11 shows our limits on BB power in comparison with those from BICEP, QUaD, and *WMAP*. QUIET lies between BICEP and *WMAP* in significantly limiting r

[39] Our definition of r agrees with Chiang et al. (2010)

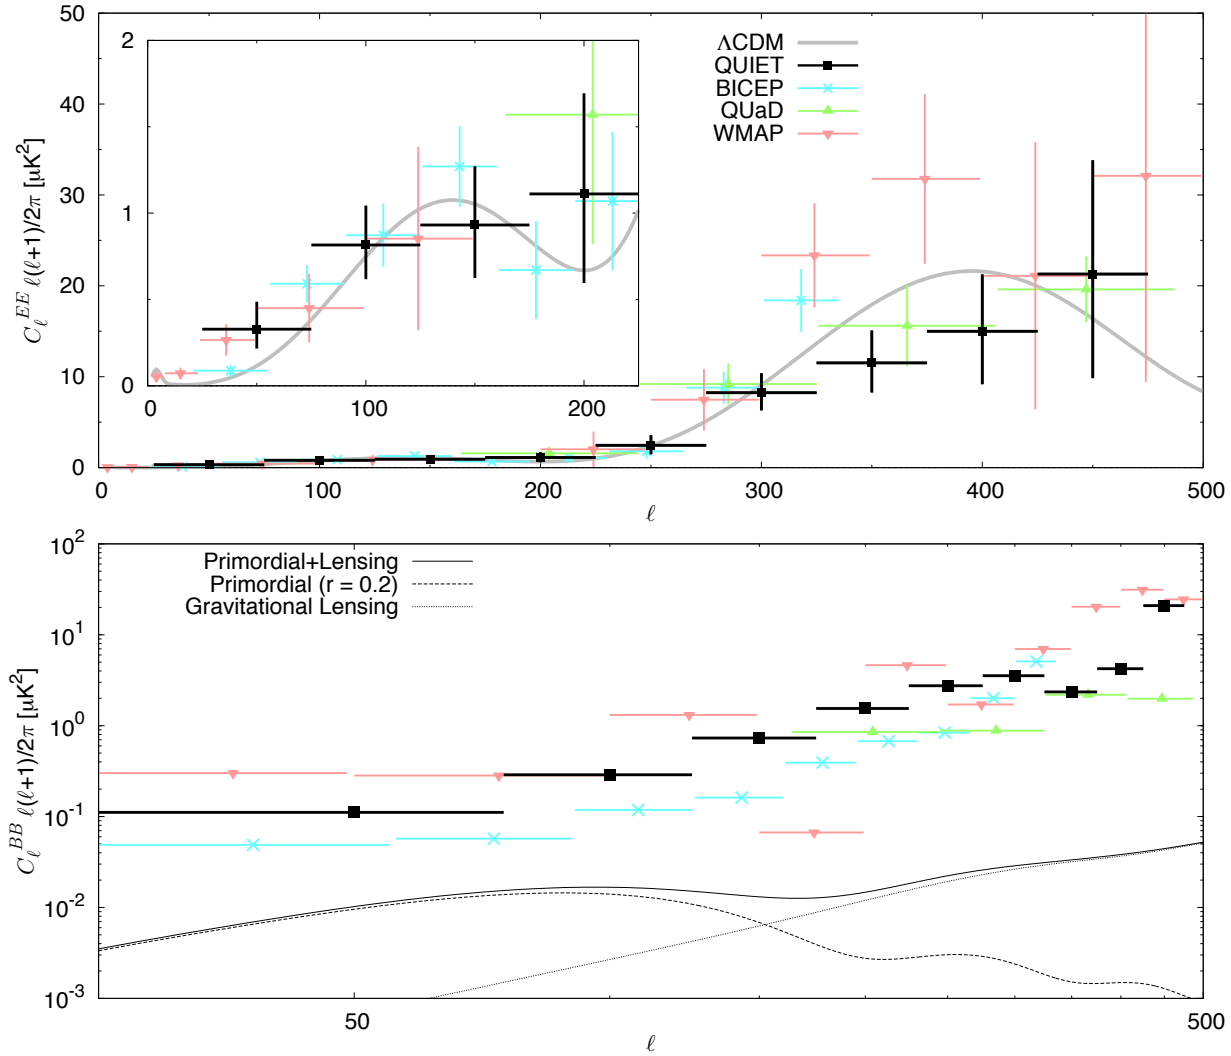

FIG. 11.— The top panel shows EE results with 68% C.L. error bars; the bottom panel shows BB 95% C.L. upper limits. For comparison, we also plot results from previous experiments (Brown et al. 2009; Chiang et al. 2010; Larson et al. 2010) and the ΛCDM model (the value $r = 0.2$ is currently the best 95% C.L. limit on tensor modes).

from measurements of CMB–B-mode power in our multipole range. Although we neither expected nor detected any BB foreground power, the detection of an EE foreground in patch CMB-1 suggests that BB foregrounds might be present at a smaller level. We emphasize that the upper limit we report is therefore conservative.

7.4. *Temperature Power Spectra*

Figure 13 compares the QUIET and *WMAP* Q-band temperature maps and TT, TE, and TB power spectra. Agreement with the ΛCDM model is good. This is a strong demonstration of the raw sensitivity of the QUIET detectors; the single QUIET differential-temperature assembly produces a high–signal-to-noise map using only 189 hours (after selection) of observations. The high sensitivity of these modules makes them very useful for calibration, pointing estimation, and consistency checks (see Section 4).

8. SYSTEMATIC ERRORS

The passing of the null suite itself limits systematic uncertainty, but to get well below the statistical errors, dedicated studies are needed. They are important in gaining confidence in the result and also in evaluating the potential of the methods and techniques we use for future efforts. We pay special attention to effects that can generate false B-mode signals. Our methodology is to simulate and then propagate calibration uncertainties (see Section 4) and other systematic effects through the entire pipeline. The systematic errors in the power spectra are shown in Figure 14. The possible contaminations are well below the statistical errors; in particular, the levels of spurious B modes are less than the signal of $r = 0.1$. This is the lowest level of BB contamination yet reported by any CMB experiment. This section describes how each effect in Figure 14 is determined and considers three additional possible sources of contamination.

An uncertainty not shown in Figure 14 is that aris-

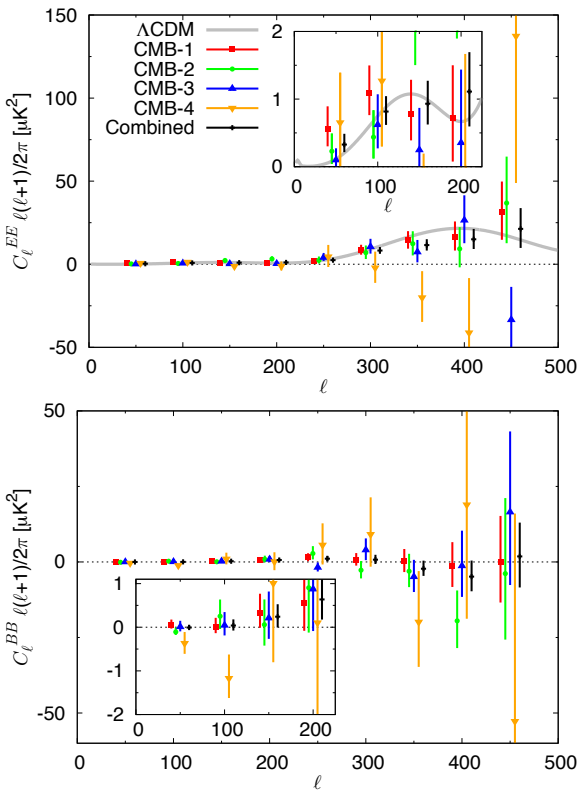

FIG. 12.— CMB power spectra are shown for each patch individually. The top and bottom panels show the EE and BB spectra, respectively. The different error bars for each patch mainly reflect the amounts of time each was observed.

TABLE 6
BAND AND CROSS POWERS FOR $\ell = 25$–75

Patch	Spectrum	$\hat{C}_{b=1}^{\rm KK}$	$\hat{C}_{b=1}^{\rm QK}$	$\hat{C}_{b=1}^{\rm QQ}$
CMB-1	EE	$\mathbf{17.4 \pm 4.7}$	$\mathbf{3.30 \pm 0.55}$	$\mathbf{0.55 \pm 0.14}$
	BB	4.8 ± 4.5	0.40 ± 0.41	0.06 ± 0.08
	EB	-6.2 ± 3.2	0.27 ± 0.38	0.10 ± 0.08
CMB-2	EE	5.5 ± 3.7	0.01 ± 0.56	0.23 ± 0.19
	BB	4.6 ± 3.4	0.18 ± 0.48	-0.11 ± 0.13
	EB	-5.5 ± 2.8	-0.39 ± 0.41	-0.20 ± 0.12
CMB-3	EE	0.2 ± 1.9	0.64 ± 0.43	0.10 ± 0.18
	BB	-0.3 ± 2.6	0.33 ± 0.35	0.01 ± 0.13
	EB	1.4 ± 1.7	-0.34 ± 0.30	$\mathbf{-0.27 \pm 0.11}$
CMB-4	EE	-5.2 ± 5.1	0.7 ± 1.2	0.65 ± 0.58
	BB	-2.6 ± 5.2	-0.1 ± 1.1	-0.37 ± 0.52
	EB	-1.0 ± 3.9	0.0 ± 0.9	-0.15 ± 0.47

NOTE. — Power-spectra estimates for the first multipole bin for each patch, computed from the *WMAP*7 K-band data and the QUIET Q-band data. The units are $\ell(\ell+1)C_\ell/2\pi$ (μK^2) in thermodynamic temperature. Uncertainties for $\hat{C}_{b=1}^{\rm KK}$ and $\hat{C}_{b=1}^{\rm QK}$ include noise only. For $\hat{C}_{b=1}^{\rm QQ}$ they additionally include CMB sample variance as predicted by ΛCDM. Values in bold are more than $2\,\sigma$ away from zero.

ing from the overall responsivity error estimate of 6% (12% in power-spectra units). After including the effect of possible time-dependent responsivity variations (4%, see below), the power-spectra uncertainty is 13%. It is multiplicative, affecting all power-spectra results independent of multipole.

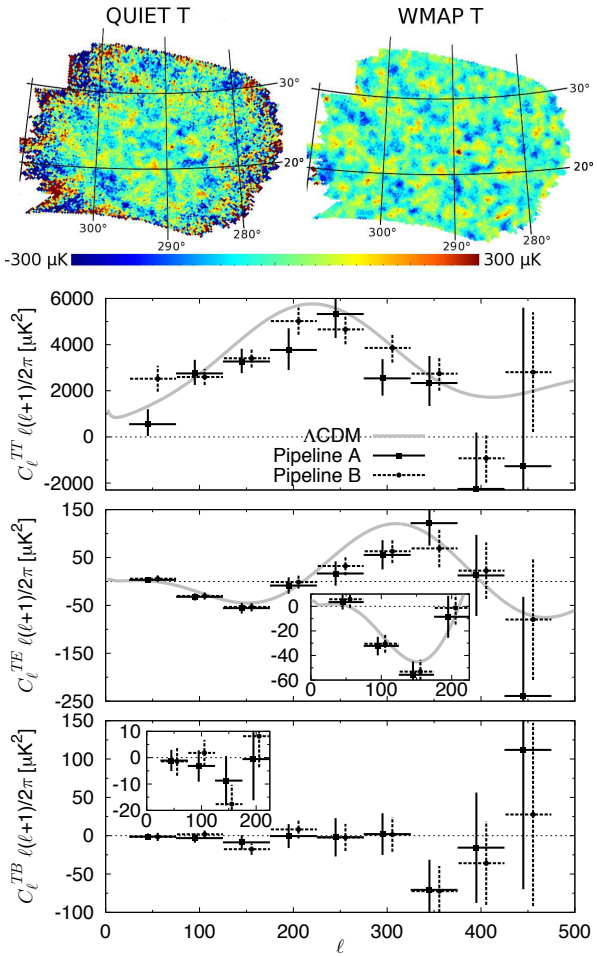

FIG. 13.— The top row compares our temperature map to the *WMAP* 7-year Q-band map (Jarosik et al. 2010) for patch CMB-1 in Galactic coordinates. Lower panels show the CMB temperature power spectra: TT, TE, and TB.

8.1. *Beam Window Function and Pointing*

The uncertainty in the beam window function is another multiplicative factor, one which increases with multipole. We estimate this uncertainty using the difference of the beam window functions measured for the central module and the modules of the differential-temperature assembly, which are at the edge of the array. The difference is statistically significant, coming from the different locations (with respect to the optics) in the focal plane; it is expected from the pre-season antenna range measurements.

Uncertainties in pointing lead to distortions in polarization maps. E power will be underestimated and spurious B power (if the distortions are non-linear) generated (Hu et al. 2003). We quantify these effects by using the differences in pointing solutions from two independent models: the fiducial model used for the analysis and an alternative model based on a different set of calibrating observations. We also modeled and included the effects of the deck-angle–encoder shift which occurred for a portion of the season (Section 4.3).

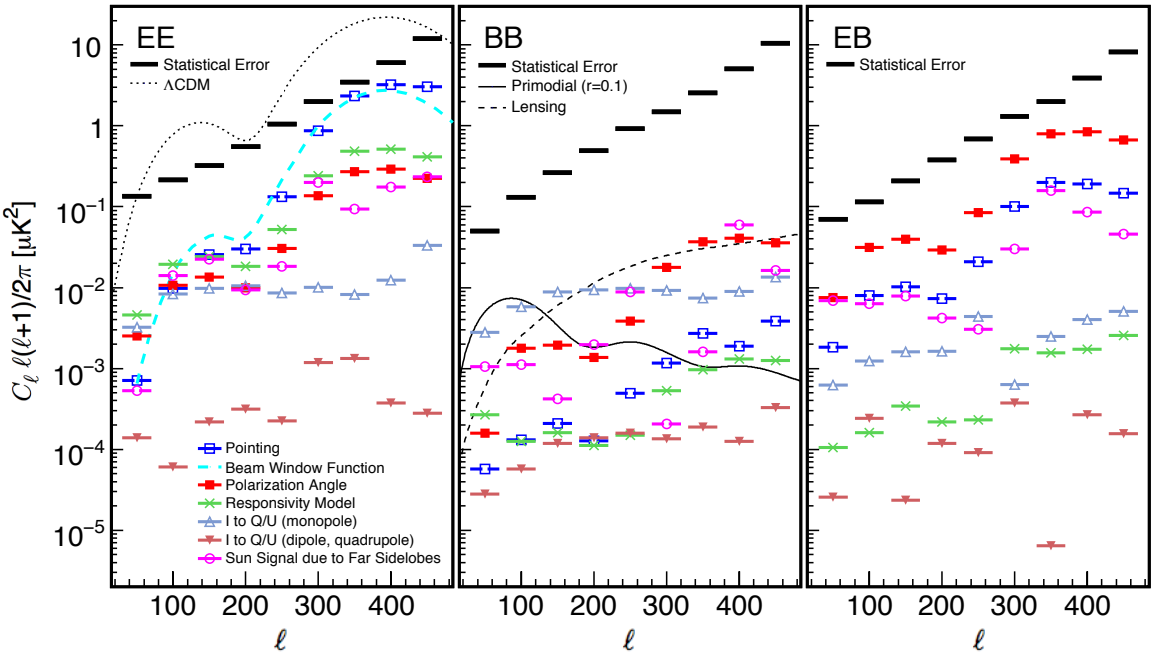

FIG. 14.— Systematic uncertainty estimates for EE, BB, and EB power spectra. Estimates for a variety of effects (see text) are shown for the three power spectra. In all cases, they are well below the statistical errors, which are also shown. In particular, the contaminations to the primordial–B-mode signal, at multipoles below 100, are below the level of $r = 0.1$, even though we do not make a correction for the largest contaminant, the monopole leakage.

8.2. *Responsivity and Polarization Angle*

Responsivity shifts, particularly within CESes, lead to distortions in the maps. Full-pipeline simulations quantify the shifts caused by variations in the cryostat or electronics temperatures. Similarly shifts from using responsivities determined from the Moon data, Tau A data, or from the sparse wire grid, rather than those from the sky dips, are determined. We also incorporate the uncertainty in the atmospheric-temperature model used in analyzing the sky-dip data. The largest possible effects on the power spectra are shown in Figure 14.

Uncertainties in the orientation of the polarization axes of the modules can lead to leakage between E and B modes. To quantify this leakage, we use the differences in power spectra where these angles are determined from Moon data, Tau A data, and the sparse–wire-grid data. As expected, the largest effects show up in EB power.

8.3. *Instrumental Polarization*

As described in Section 4.5, the I to Q (U) leakage coefficients for the QUIET detector diodes are small: 1% (0.2%). Except in the case of patch CMB-4, our scanning strategy significantly reduces this effect with the combination of sky and deck-angle rotation.

We estimate spurious Q and U in the maps for each CES-diode using the *WMAP* temperature map and our known leakages. Shown in Figure 14 are the estimates of spurious EE, BB, and EB powers from full-pipeline simulations, where for each realization the spurious Q and U are added to the Q and U from simulated ΛCDM E modes. While this method has an advantage of being able to use the real (not simulated) temperature map, it does not incorporate TE correlation, which only affects the spurious EE power. As a complement, we repeat the study, but using simulated ΛCDM maps for both temperature and polarization; this only changes the estimate of spurious EE power by 30% at most. Because the spurious power is as small as it is, we have treated it as a systematic rather than correcting for it. Doing so would give us a further order of magnitude suppression.

Differing beam ellipticities can also induce higher multipole polarization signals. We measure these leakages from Tau A and Jupiter observations and find that the higher-order multipoles are at most 0.1% of the main-beam peak amplitude. The corresponding effects on the power spectra, which are seen in Figure 14, are of little concern.

8.4. *Far Sidelobes Seeing the Sun*

While we make cuts to reduce the effects of far sidelobes seeing the Sun (Sections 2 and 5.1.3), small contaminations could remain. We make full-season maps for each diode in Sun-centered coordinates and then use these maps to add contamination to full-pipeline CMB simulations. The excess power found in the simulations is taken as the systematic uncertainty.

8.5. *Other Possible Sources of Systematic Uncertainty*

Here we discuss a few additional potential sources of systematic uncertainty, which are found to be subdominant.

Ground-Synchronous Signals. QUIET's far sidelobes do see the ground for some diodes at particular elevations and deck angles. Ground pickup that is constant throughout a CES is removed by our TOD filters; the net effect of this filtering in the full-season maps is a correction of $\approx 1\,\mu$K.

The only concern is ground pickup that changes over

the short span of a single CES. We find little evidence for changes even over the entire season, let alone over a single CES. We therefore conservatively place an upper limit on such changes using the statistical errors on the ground-synchronous signal. We start with the CES and module with the largest ground pickup. We then simulate one day's worth of data, inserting a ground-synchronous signal that changes by its statistical error. Given the distribution in the magnitude of the ground-synchronous signal and assuming that changes in this signal are proportional to the size of the signal itself, by considering that the signals from changing pickup add incoherently into the maps made from multiple CES-diodes at a variety of elevations and deck angles, we estimate an upper limit on residual B power from possible changing ground-pickup signals. The result is $\lesssim 10^{-4}\,\mu\mathrm{K}^2$ at multipoles below 100.

ADC Non-linearities. The possible residual after the correction for the non-linearity in the ADC system results in effects similar to the I to Q (or U) leakage and the variation of the responsivity during the CES. We estimate such effects based on the uncertainty in the correction parameters, confirming that there is at most a 3% additional effect for the leakage bias, and that the responsivity effect is also small, less than half of the systematic error shown for the responsivity in Figure 14.

Data-Selection Biases. Cuts can cause biases if they are, for example, too stringent. We expect none but to be sure we apply our selection criteria to 144 CMB + noise simulations. No bias is seen, and in particular we limit any possible spurious B modes from this source to $\lesssim 10^{-3}\,\mu\mathrm{K}^2$ at multipoles below 100.

9. CONCLUSIONS

QUIET detects polarization in the EE power spectrum at 43 GHz. We confirm with high significance the detection of polarization in the region of the first acoustic peak (Chiang et al. 2010) in the multipole region $\ell = 76$–175. We find no significant power in either BB or EB between $\ell = 25$ and $\ell = 475$. We measure the tensor-to-scalar ratio to be $r = 0.35^{+1.06}_{-0.87}$.

These results are supported by a very extensive suite of null tests in which 42 divisions of data were used for each of 33 different cut configurations. The selection criteria and systematic errors were determined before the power spectra themselves were examined. Biases were revealed during this process, the last of which was a contamination present in the null spectra at the level of about 20% of the statistical errors, but eliminated when cross-correlating maps with differing telescope pointings. The robustness of the final results is further supported by having two pipelines with results in excellent agreement, even though one uses only cross correlations while the other also uses auto correlations.

Several possible systematic effects are studied with full end-to-end simulations. The possible contaminations in the B-mode power are thereby limited to a level smaller than for any other published experiment: below the level of $r = 0.1$ for the primordial B modes; simply correcting for the known level of instrumental polarization would reduce this to $r < 0.03$. This very low level of systematic uncertainty comes from the combination of several important design features, including a new time-stream "double-demodulation" technique, Mizuguchi–Dragone optics, natural sky rotation, and frequent deck rotation.

The correlation modules we use have a polarization sensitivity (Q and U combined) of $280\,\mu\mathrm{K}\sqrt{\mathrm{s}}$, leading to an array sensitivity of $69\,\mu\mathrm{K}\sqrt{\mathrm{s}}$. Further, the $1/f$ noise observed in our detectors is small: the median knee frequency is just 5.5 mHz. One important outcome of this work, then, is the demonstration that our detectors, observing from a mid-latitude site, give excellent sensitivity and systematic immunity.

Because of our mid-latitude site, we are driven to collect data in four separate patches. While we lose some sensitivity (compared to going deeper on a single patch), there are a few advantages that we have exploited. The patches are scanned differently, in terms of time of day and the degree of crosslinking, and these differences allow some important systematic checks. Another advantage concerns foregrounds.

Foreground contamination is expected to be one of the main limiting factors in the search for primordial B modes. Indeed we report a 3-σ detection of synchrotron emission in one of our four CMB patches, originally chosen for their expected low foreground levels. Our detection is only in EE but assuming a similar BB level and extrapolating to the foreground minimum of about 95 GHz, we would have synchrotron contamination at the level of $r = 0.02$. Neither *WMAP* nor *Planck* will have enough sensitivity (Tauber et al. 2010) to sufficiently constrain the polarized synchrotron amplitude at this level. In fact, our Q-band polarization maps are already as deep or deeper than what *Planck* will achieve at the same frequency. Dedicated low-frequency observations are clearly needed to achieve such constraints. When foreground cleaning becomes important, consistency among separate patches will be an important handle on our understanding.

Further progress must be made through larger arrays and longer integration times. In hand we have data collected by the 90-element W-band array with similar sensitivity to our Q-band array and more than twice the number of observing hours. Results from the analysis of that data set will be reported in future publications. A W-band receiver with the sensitivity to reach below the level of $r = 0.01$ is under development.

Support for the QUIET instrument and operation comes through the NSF cooperative agreement AST-0506648. Support was also provided by NSF awards PHY-0355328, AST-0448909, AST-1010016, and PHY-0551142; KAKENHI 20244041, 20740158, and 21111002; PRODEX C90284; a KIPAC Enterprise grant; and by the Strategic Alliance for the Implementation of New Technologies (SAINT).

Some work was performed on the Joint Fermilab-KICP Supercomputing Cluster, supported by grants from Fermilab, the Kavli Institute for Cosmological Physics, and the University of Chicago. Some work was performed on the Titan Cluster, owned and maintained by the University of Oslo and NOTUR (the Norwegian High Performance Computing Consortium), and on the Central Computing System, owned and operated by the Computing Research Center at KEK. Portions of this work were performed at the Jet Propulsion Laboratory (JPL) and California Institute of Technology, operating under

a contract with the National Aeronautics and Space Administration. The Q-band polarimeter modules were developed using funding from the JPL R&TD program.

C.D. acknowledges an STFC Advanced Fellowship and an ERC IRG grant under FP7. P.G.F. and J.A.Z. gratefully acknowledge the support of the Beecroft Institute of Particle Astrophysics and Cosmology, the Oxford Martin School, and the Science and Technology Facilities Council. L.B., R.B., and J.M. acknowledge support from CONICYT project Basal PFB-06. A.D.M. acknowledges a Sloan foundation fellowship.

PWV measurements were provided by the Atacama Pathfinder Experiment (APEX). We thank CONICYT for granting permission to operate within the Chajnantor Scientific Preserve in Chile, and ALMA for providing site infrastructure support. We are particularly indebted to the engineers and technician who maintained and operated the telescope: José Cortés, Cristobal Jara, Freddy Muñoz, and Carlos Verdugo.

In addition, we would like to acknowledge the following people for their assistance in the instrument design, construction, commissioning, operation, and in data analysis: Augusto Gutierrez Aitken, Colin Baines, Phil Bannister, Hannah Barker, Matthew R. Becker, Alex Blein, Mircea Bogdan, April Campbell, Anushya Chandra, Sea Moon Cho, Emma Curry, Maire Daly, Richard Davis, Fritz Dejongh, Joy Didier, Greg Dooley, Hans Eide, Will Grainger, Jonathon Goh, Peter Hamlington, Takeo Higuchi, Seth Hillbrand, Christian Holler, Ben Hooberman, Kathryn D. Huff, William Imbriale, Koji Ishidoshiro, Eiichiro Komatsu, Jostein Kristiansen, Richard Lai, Erik Leitch, Kelly Lepo, Martha Malin, Mark McCulloch, Oliver Montes, David Moore, Makoto Nagai, Ian O'Dwyer, Stephen Osborne, Stephen Padin, Felipe Pedreros, Ashley Perko, Alan Robinson, Jacklyn Sanders, Dale Sanford, Mary Soria, Alex Sugarbaker, David Sutton, Matias Vidal, Liza Volkova, Edward Wollack, Stephanie Xenos, and Mark Zaskowski.

REFERENCES

Battistelli, E. S., Rebolo, R., Rubiño-Martín, J. A., Hildebrandt, S. R., Watson, R. A., Gutiérrez, C., & Hoyland, R. J. 2006, ApJ, 645, L141
Battye, R. A., Browne, I. W. A., Peel, M. W., Jackson, N. J., & Dickinson, C. 2010, arXiv:1003.5846v1 [astro-ph.CO]
Bischoff, C. et al. 2008, ApJ, 684, 771
Bond, J. R., Jaffe, A. H., & Knox, L. 1998, Phys. Rev. D, 57, 2117
Bornemann, J., & Labay, V. A. 1995, IEEE Trans. MTT, 43, 95
Brown, M. L., et al. 2009, ApJ, 705, 978
Buder, I. 2010, Proc. SPIE, 7741, 77411D
Carretti, E., et al. 2010, MNRAS, 405, 1670
Challinor, A., & Chon, G. 2005, MNRAS, 360, 509
Chiang, H. C. et al. 2010, ApJ, 711, 1123
Dragone, C. 1978, Bell Syst. Tech. J., 57, 2663
Dumoulin, R. N. 2010, Proc. SPIE, 7741, 77412N
Dunkley, J. et al. 2009, ApJ, 701, 1804
Eriksen, H. K. et al. 2004, ApJS, 155, 227
Feldman, H. A., Kaiser, N., & Peacock, J. A. 1994, ApJ, 426, 23
Gold, B. et al. 2010, arXiv:1001.4555v2 [astro-ph.GA]
Gorski, K. M., et al. 2005, ApJ, 622, 759, http://healpix.jpl.nasa.gov/
Gundersen, J., & Wollack, E. 2009, Journal of Physics: Conf. Ser., 155, 012005
Güsten, R., Nyman, L. Å., Schilke, P., Menten, K., Cesarsky, C., & Booth, R. 2006, A&A, 454, L13
Hafez, Y., et al. 2008, MNRAS, 388, 1775
Hamimeche, S., & Lewis, A. 2008, Phys. Rev. D, 77, 103013
Hansen, F. K., & Gorski, K. M. 2003, MNRAS, 343, 559
Hivon, E., Gorski, K. M., Netterfield, C. B., Crill, B. P., Prunet, S., & Hansen, F. 2002, ApJ, 567, 2
Hu, W., Hedman, M., & Zaldarriaga, M. 2003, Phys. Rev. D, 67, 043004
Imbriale, W. A., Gundersen, J., & Thompson, K. L. 2010, To appear in the Special Issue on Radio Astronomy in the IEEE Transactions on Antennas and Propagation
Jarosik, N. et al. 2003, ApJS, 145, 413
Jarosik, N., et al. 2010, arXiv:1001.4744v1 [astro-ph.CO]
Kangaslahti, P. et al. 2006, in Microwave Symposium Digest, 2006. IEEE MTT-S International, 89–92
Keskitalo, R. et al. 2010, A&A, 522, A94
Komatsu, E., et al. 2010, arXiv:1001.4538v3 [astro-ph.CO]
Kovac, J. M., Leitch, E. M., Pryke, C., Carlstrom, J. E., Halverson, N. W., & Holzapfel, W. L. 2002, Nature, 420, 772
Kusaka, A., et al. 2010, To appear in Proceedings of the 45th Rencontre de Moriond
Larson, D. et al. 2010, arXiv:1001.4635v2 [astro-ph.CO]
Larson, D. L., Eriksen, H. K., Wandelt, B. D., Górski, K. M., Huey, G., Jewell, J. B., & O'Dwyer, I. J. 2007, ApJ, 656, 653
Leitch, E. M., Kovac, J. M., Halverson, N. W., Carlstrom, J. E., Pryke, C., & Smith, M. W. E. 2005, ApJ, 624, 10
Liddle, A. R., & Lyth, D. H. 2000, Cosmological Inflation and Large-Scale Structure
Lopez-Caraballo, C. H., Rubino-Martin, J. A., Rebolo, R., & Genova-Santos, R. 2010, arXiv:1011.1242v1 [astro-ph.CO]
Lueker, M., et al. 2010, ApJ, 719, 1045
Mauskopf, P. D. et al. 2000, ApJ, 536, L59
Mizugutch, Y., Akagawa, M., & Yokoi, H. 1976, IEEE Antennas and Propagation Society International Symposium, 14, 2
Monsalve, R. A. 2010, Proc. SPIE, 7741, 77412M
Montroy, T. E. et al. 2006, ApJ, 647, 813
Netterfield, C. B., et al. 2002, ApJ, 571, 604
Newburgh, L., et al. 2010, To appear in Proceedings of the Twelfth Marcel Grossmann Meeting on General Relativity
Padin, S., et al. 2002, PASP, 114, 83
Page, L. et al. 2007, ApJS, 170, 335
Pardo, J., Cernicharo, J., & Serabyn, E. 2001, IEEE Transactions on Antennas and Propagation, 49, 1683
Sievers, J. L. et al. 2007, ApJ, 660, 976
Smith, K. M., & Zaldarriaga, M. 2007, Phys. Rev. D, 76, 043001
Tauber, J. A., et al. 2010, A&A, 520, A1
Tegmark, M. 1997, ApJ, 480, L87
Wandelt, B. D., Hivon, E., & Górski, K. M. 2001, Phys. Rev. D, 64, 083003
Weiland, J. L. et al. 2010, arXiv:1001.4731v1 [astro-ph.CO]
Wright, E. L., Hinshaw, G., & Bennett, C. L. 1996, ApJ, 458, L53
Wu, J. H. P. et al. 2007, ApJ, 665, 55

References

Abdo, A. A., et al. 2009a, *ApJ*, 700, 597.

—. 2009b, *Phys. Rev. Lett.*, 103, 251101.

—. 2009c, *ApJ*, 699, 976.

—. 2010a, *ApJS*, 188, 405.

—. 2010b, *ApJ*, 715, 429.

—. 2010c, *ApJ*, 716, 30.

Acciari, V. A., et al. 2009, *Science*, 325, 444.

—. 2010, *ApJ*, 709, L163.

Ackermann, M., et al. 2011, *ApJ*, 741, 30.

Aharonian, F., et al. 2007, *ApJ*, 664, L71.

—. 2009, *A&A*, 502, 749.

Aller, M. F., Aller, H. D., & Hughes, P. A. 1992, *ApJ*, 399, 16.

—. 2003, *ApJ*, 586, 33.

Aller, M. F., Aller, H. D., Hughes, P. A., & Latimer, G. E. 1999, *ApJ*, 512, 601.

Andrew, B. H., MacLeod, J. M., Harvey, G. A., & Medd, W. J. 1978, *AJ*, 83, 863.

Angelakis, E., Fuhrmann, L., Nestoras, I., Zensus, J. A., Marchili, N., Pavlidou, V., & Krichbaum, T. P. 2010, in *Proceedings of the Workshop "Fermi meets Jansky: AGN in Radio and Gamma-Rays"*, ed. T. Savolainen, E. Ros, R. W. Porcas, & J. A. Zensus, 81–84, `http://www.mpifr-bonn.mpg.de/div/vlbi/agn2010/PdfFiles/fmj2010_complete.pdf`.

Angelakis, E., Kraus, A., Readhead, A. C. S., Zensus, J. A., Bustos, R., Krichbaum, T. P., Witzel, A., & Pearson, T. J. 2009, *A&A*, 501, 801.

Antonucci, R., & Miller, J. S. 1985, *ApJ*, 297, 621.

Atwood et al. 2009, *ApJ*, 697, 1071.

Baade, W., & Minkowski, R. 1954a, *ApJ*, 119, 206.

—. 1954b, *ApJ*, 119, 215.

Baars, J. W. M., Genzel, R., Pauliny-Toth, I. I. K., & Witzel, A. 1977, *A&A*, 61, 99.

Barthel, P. D. 1989, *ApJ*, 336, 606.

Begelman, M. C., Blandford, R. D., & Rees, M. J. 1984, *Rev. Mod. Phys.*, 56, 255.

Bennett, A. S. 1962, *MmRAS*, 68, 163.

Blake, G. M. 1970, *Astrophys. Lett.*, 6, 201.

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

Blandford, R. D., & Levinson, A. 1995, *ApJ*, 441, 79.

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

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

—. 1978, in *Pittsburgh Conference on BL Lac Objects, Pittsburgh, Pa., April 24–26, 1978, Proceedings.*, ed. A. M. Wolfe, 328–341.

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

Błażejowski, M., Sikora, M., Moderski, R., & Madejski, G. M. 2000, *ApJ*, 545, 107.

Blondin, S., et al. 2008, *ApJ*, 682, 724.

Böttcher, M. 2007, *Ap&SS*, 309, 95.

Burbidge, G. R. 1959, *ApJ*, 129, 849.

Bustos, R. A. 2008, PhD thesis, Universidad de Concepción.

Cohen, M. H., Cannon, W., Purcell, G. H., Shaffer, D. B., Broderick, J. J., Kellermann, K. I., & Jauncey, D. L. 1971, *ApJ*, 170, 207.

Cohen, M. H., Lister, M. L., Homan, D. C., Kadler, M., Kellermann, K. I., Kovalev, Y. Y., & Vermeulen, R. C. 2007, *ApJ*, 658, 232.

Curtis, H. D. 1918, in Publications of the Lick Observatory, Vol. XIII, *Studies of the nebulae made at the Lick Observatory, University of California, at Mount Hamilton, California, and Santiago, Chile* (Berkeley: University of California Press), 9–42.

Dennett-Thorpe, J., & de Bruyn, A. G. 2000, *ApJ*, 529, L65.

Dent, W. A. 1965, *Science*, 148, 1458.

—. 1966, *ApJ*, 144, 843.

Dermer, C. D. 1995, *ApJ*, 446, L63.

Dermer, C. D., Schlickeiser, R., & Mastichiadis, A. 1992, *A&A*, 256, L27.

Dermer, C. D., Sturner, S. J., & Schlickeiser, R. 1997, *ApJS*, 109, 103.

Edelson, R., Turner, T. J., Pounds, K., Vaughan, S., Markowitz, A., Marshall, H., Dobbie, P., & Warwick, R. 2002, *ApJ*, 568, 610.

Edge, D. O., Shakeshaft, J. R., McAdam, W. B., Baldwin, J. E., & Archer, S. 1959, *MmRAS*, 68, 37.

Evans, G., & McLeish, C. W. 1977, *RF Radiometer Handbook* (Dedham, MA: Artech).

Fanaroff, B. L., & Riley, J. M. 1974, *MNRAS*, 167, 31P.

Fermi-LAT Collaboration et al. 2010, *Nature*, 463, 919.

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

Fragile, P. C. 2008, in *Proceedings of the VII Microquasar Workshop: "Microquasars and Beyond."*, 39, http://pos.sissa.it/cgi-bin/reader/conf.cgi?confid=62.

Fuhrmann, L., Zensus, J. A., Krichbaum, T. P., Angelakis, E., & Readhead, A. C. S. 2007, in American Institute of Physics Conference Series, Vol. 921, *The First GLAST Symposium*, ed. S. Ritz, P. Michelson, & C. A. Meegan, 249–251.

Ghirlanda, G., Ghisellini, G., Tavecchio, F., & Foschini, L. 2010, *MNRAS*, 407, 791.

Ghisellini, G., Celotti, A., Fossati, G., Maraschi, L., & Comastri, A. 1998, *MNRAS*, 301, 451.

Górski, K. M., Hivon, E., Banday, A. J., Wandelt, B. D., Hansen, F. K., Reinecke, M., & Bartelmann, M. 2005, *ApJ*, 622, 759.

Hargrave, P. J., & Ryle, M. 1974, *MNRAS*, 166, 305.

Hartman, R. C., et al. 1999, *ApJS*, 123, 79.

Healey, S. E., et al. 2008, *ApJ*, 175, 97.

Herbig, T. 1994, PhD thesis, California Institute of Technology.

Hovatta, T., Nieppola, E., Tornikoski, M., Valtaoja, E., Aller, M. F., & Aller, H. D. 2008, *A&A*, 485, 51.

Hovatta, T., Tornikoski, M., Lainela, M., Lehto, H. J., Valtaoja, E., Aller, M. F., & Aller, H. D. 2007, *A&A*, 469, 899.

Hughes, P. A., Aller, H. D., & Aller, M. F. 1992, *ApJ*, 396, 469.

Hunter, J. D. 2007, *Computing in Science & Engineering*, 9, 90.

Hunter, S. D., et al. 1997, *ApJ*, 481.

Ioannidis, J. P. A. 2005, *PLoS Med.*, 2, e124.

Jarosik, N. C. 1996, *IEEE Trans. MTT*, 44, 193.

Jauncey, D. L., Kedziora-Chudczer, L. L., Lovell, J. E. J., Nicolson, G. D., Perley, R. A., Reynolds, J. E., Tzioumis, A. K., & Wieringa, M. H. 2000, in *Astrophysical Phenomena Revealed by Space VLBI*, ed. H. Hirabayashi, P. G. Edwards, & D. W. Murphy, 147–150.

Jennison, R. C., & Das Gupta, M. K. 1953, *Nature*, 172, 996.

Jorstad, S. G., Marscher, A. P., Mattox, J. R., Aller, M. F., Aller, H. D., Wehrle, A. E., & Bloom, S. D. 2001a, *ApJ*, 556, 738.

Jorstad, S. G., Marscher, A. P., Mattox, J. R., Wehrle, A. E., Bloom, S. D., & Yurchenko, A. V. 2001b, *ApJS*, 134, 181.

Kedziora-Chudczer, L., Jauncey, D. L., Wieringa, M. H., Walker, M. A., Nicolson, G. D., Reynolds, J. E., & Tzioumis, A. K. 1997, *ApJ*, 490, L9.

Kellermann, K. I., Clark, B. G., Niell, A. E., & Shaffer, D. B. 1975, *ApJ*, 197, L113.

Kellermann, K. I., Sramek, R., Schmidt, M., Shaffer, D. B., & Green, R. 1989, *AJ*, 98, 1195.

Kellermann, K. I., et al. 2004, *ApJ*, 609, 539.

—. 2007, *Ap&SS*, 311, 231.

Konigl, A. 1981, *ApJ*, 243, 700.

Kovalev, Y. Y., et al. 2009, *ApJ*, 696, L17.

Kraus, A., et al. 2003, *A&A*, 401, 161.

Krolik, J. H. 1999, *Active Galactic Nuclei* (Princeton, NJ: Princeton University Press).

Ledlow, M. J., & Owen, F. N. 1996, *AJ*, 112, 9.

Leitch, E. M. 1998, PhD thesis, California Institute of Technology.

Lilliefors, H. W. 1967, *J. Am. Stat. Assoc.*, 62, 399.

Linford, J. D., Taylor, G. B., Romani, R. W., Helmboldt, J. F., Readhead, A. C. S., Reeves, R., & Richards, J. L. 2012, *ApJ*, 744, 177.

Linford, J. D., et al. 2011, *ApJ*, 726, 16.

Lister, M. L., & Homan, D. C. 2005, *AJ*, 130, 1389.

Lister, M. L., & Marscher, A. P. 1997, *ApJ*, 476, 572.

Lovell, J. E. J., et al. 2008, *ApJ*, 689, 108.

Lynden-Bell, D. 1969, *Nature*, 223, 690.

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

Macquart, J., & de Bruyn, A. G. 2007, *MNRAS*, 380, L20.

Mahony, E. K., Sadler, E. M., Murphy, T., Ekers, R. D., Edwards, P. G., & Massardi, M. 2010, *ApJ*, 718, 587.

Maltby, P., & Moffet, A. T. 1962, *ApJS*, 7, 141.

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

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

Markarian, B. E. 1967, *Astrofizika*, 3, 24.

Markaryan, B. E., Lipovetskii, V. A., & Stepanyan, D. A. 1981, *Astrophysics*, 17, 321.

Marscher, A. P., & Gear, W. K. 1985, *ApJ*, 298, 114.

Marscher, A. P., et al. 2008, *Nature*, 452, 966.

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

Mattox, J. R., et al. 1996, *ApJ*, 461, 396.

Max-Moerbeck, W., et al. 2010, in *Proceedings of the Workshop "Fermi meets Jansky: AGN in Radio and Gamma-Rays"*, ed. T. Savolainen, E. Ros, R. W. Porcas, & J. A. Zensus, 77–80, `http://www.mpifr-bonn.mpg.de/div/vlbi/agn2010/PdfFiles/fmj2010_complete.pdf`.

McKinney, J. C., & Blandford, R. D. 2009, *MNRAS*, 394, L126.

Meyer, E. T., Fossati, G., Georganopoulos, M., & Lister, M. L. 2011, *ApJ*, 740, 98.

Mills, B. Y., & Slee, O. B. 1957, *Aust. J. Phys.*, 10, 162.

Moffet, A. T. 1973, "Circular Polarization in Waveguide Feeds,", unpublished memorandum, 19 July, 1973.

Moffet, A. T., Gubbay, J., Robertson, D. S., & Legg, A. J. 1972, in IAU Symposium, Vol. 44, *External Galaxies and Quasi-Stellar Objects*, ed. D. S. Evans, D. Wills, & B. J. Wills, 228.

Mücke, A., et al. 1997, *A&A*, 320, 33.

Myers, S. T., Baker, J. E., Readhead, A. C. S., Leitch, E. M., & Herbig, T. 1997, *ApJ*, 485, 1.

NOAA. 2010, Average temperature profile of the atmosphere, `http://www.srh.noaa.gov/jetstream/atmos/atmprofile.htm`.

Ojha, R., et al. 2010, *A&A*, 519, A45.

Orr, M. J. L., & Browne, I. W. A. 1982, *MNRAS*, 200, 1067.

Ott, M., Witzel, A., Quirrenbach, A., Krichbaum, T. P., Standke, K. J., Schalinski, C. J., & Hummel, C. A. 1994, *A&A*, 284, 331.

Padovani, P., Ghisellini, G., Fabian, A. C., & Celotti, A. 1993, *MNRAS*, 260, L21.

Pavlidou, V., et al. 2012, *ApJ*, *in press*.

Pearson, T. J. 1999, Telescope Control System for the OVRO 40-m and 5.5-m antennas, `http://www.astro.caltech.edu/~tjp/ovroman/manual.html`.

Pearson, T. J., Unwin, S. C., Cohen, M. H., Linfield, R. P., Readhead, A. C. S., Seielstad, G. A., Simon, R. S., & Walker, R. C. 1981, *Nature*, 290, 365.

Penna, R. F., McKinney, J. C., Narayan, R., Tchekhovskoy, A., Shafee, R., & McClintock, J. E. 2010, *MNRAS*, 408, 752.

Punch, M., et al. 1992, *Nature*, 358, 477.

QUIET Collaboration et al. 2011, *ApJ*, 741, 111.

Readhead, A. C. S. 1994, *ApJ*, 426, 51.

Readhead, A. C. S., Cohen, M. H., & Blandford, R. D. 1978a, *Nature*, 272, 131.

Readhead, A. C. S., Cohen, M. H., Pearson, T. J., & Wilkinson, P. N. 1978b, *Nature*, 276, 768.

Readhead, A. C. S., Lawrence, C. R., Myers, S. T., Sargent, W. L. W., Hardebeck, H. E., & Moffet, A. T. 1989, *ApJ*, 346, 566.

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

Richards, J. L., et al. 2011, *ApJS*, 194, 29.

Rickett, B., Kedziora-Chudczer, L., & Jauncey, D. L. 2002, *PASA*, 19, 106.

Rickett, B. J., Lazio, T. J. W., & Ghigo, F. D. 2006, *ApJS*, 165, 439.

Rohlfs, K., & Wilson, T. L. 2000, *Tools of Radio Astronomy* (Berlin: Springer-Verlag).

Rybicki, G. B., & Lightman, A. P. 1979, *Radiative Processes in Astrophysics* (New York: John Wiley & Sons).

Salamon, M. H., & Stecker, F. W. 1994, *ApJ*, 430, L21.

Scargle, J. D. 2000, *J. Sci. Expl.*, 14, 91.

Schmid, H. 2007, *IEEE Circ. and Syst.*, 7, 32.

Schmidt, M. 1963, *Nature*, 197, 1040.

—. 1968, *ApJ*, 151, 393.

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

Schott, G. A. 1912, *Electromagnetic radiation and the mechanical reactions arising from it; being an Adams prize essay in the University of Cambridge* (Cambridge: Cambridge University Press).

Scott, P. F., & Ryle, M. 1961, *MNRAS*, 122, 389.

Seyfert, C. K. 1943, *ApJ*, 97, 28.

Shaw, M. S., et al. 2012a, *ApJ*, 748, 49.

—. 2012b, *in preparation*.

Shepherd, M. C. 1997, in Astronomical Society of the Pacific Conference Series, Vol. 125, *Astronomical Data Analysis Software and Systems VI*, ed. G. Hunt & H. Payne, 77.

—. 2011a, "The OVRO 40m telescope drive system," unpublished technical manual.

—. 2011b, The Telescope Scheduling Language, `http://www.astro.caltech.edu/~mcs/ovro/40m/help/`.

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

Slipher, V. M. 1917, *Lowell Obs. Bull.*, 3, 59.

Soldi, S., et al. 2008, *A&A*, 486, 411.

Stecker, F. W., Hunter, S. D., & Kniffen, D. A. 2008, *Astropart. Phys.*, 29, 25.

Stecker, F. W., Salamon, M. H., & Malkan, M. A. 1993, *ApJ*, 410, L71.

Stevens, J. A., Litchfield, S. J., Robson, E. I., Hughes, D. H., Gear, W. K., Terasranta, H., Valtaoja, E., & Tornikoski, M. 1994, *ApJ*, 437, 91.

Taylor, G. B., et al. 2007, *ApJ*, 671, 1355.

Teräsranta, H., et al. 2004, *A&A*, 427, 769.

Thompson, D. J., et al. 1993, *ApJS*, 86, 629.

Urry, C. M., & Padovani, P. 1995, *PASP*, 107, 803.

van der Ziel, A. 1988, *Proc. IEEE*, 76, 233.

Waldram, E. M., Pooley, G. G., Davies, M. L., Grainge, K. J. B., & Scott, P. F. 2010, *MNRAS*, 404, 1005.

Wall, J. V., & Jenkins, C. R. 2003, *Practical Statistics for Astronomers* (Cambridge: Cambridge University Press).

Weedman, D. W. 1977, *ARA&A*, 15, 69.

Weekes, T. C., et al. 2002, *Astropart. Phys.*, 17, 221.

Whitney, A. R., et al. 1971, *Science*, 173, 225.

Wollack, E. J., & Pospieszalski, M. W. 1998, in *IEEE MTT-S International Microwave Symposium Digest*, 669–672.