

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.

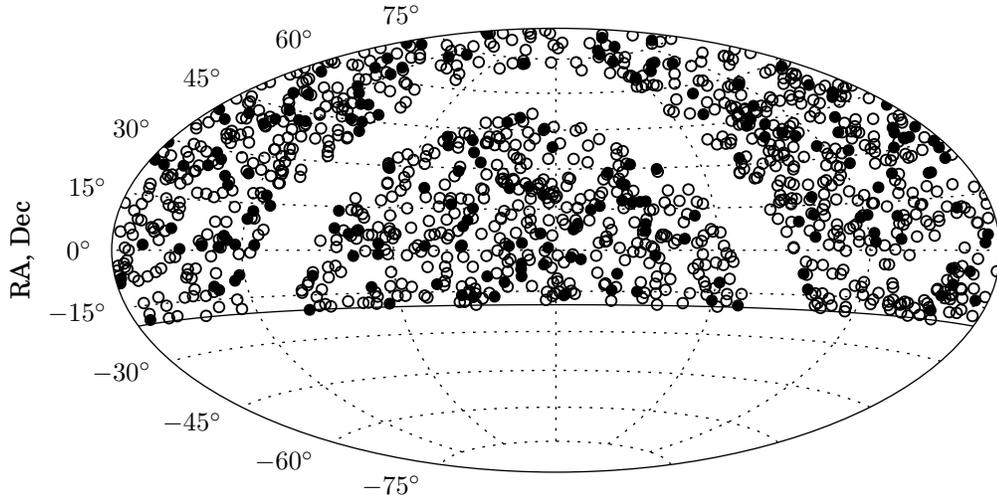


Figure 4.1. Positions of the CGRaBS sources in our program in equatorial coordinates. Filled circles indicate CGRaBS that are also in the 1LAC sample. The solid line marks the -20° declination limit of our program.

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

Sample	Total		CGRaBS Subset	
	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.

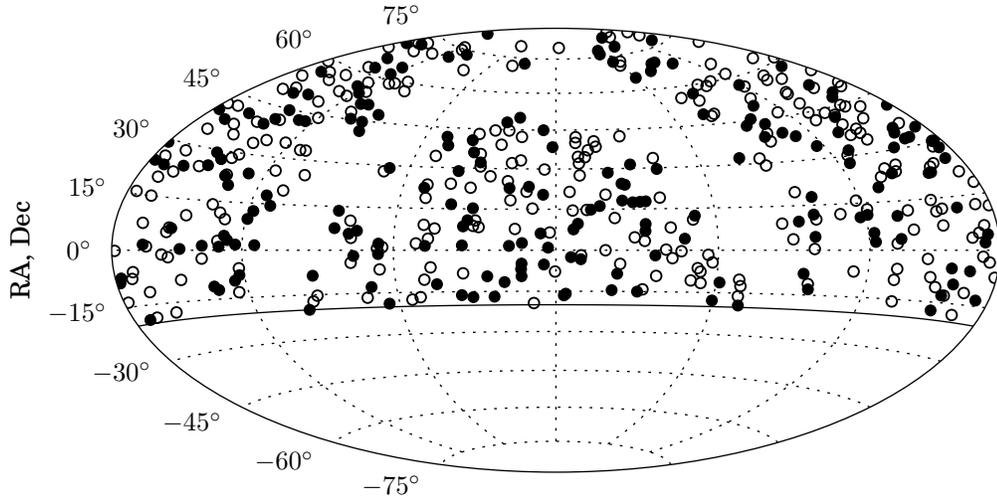


Figure 4.2. Positions of the 1LAC sources in our program in equatorial coordinates. Filled circles indicate 1LAC sources that are also in the CGRaBS sample. The solid line marks the -20° declination limit of our program.

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

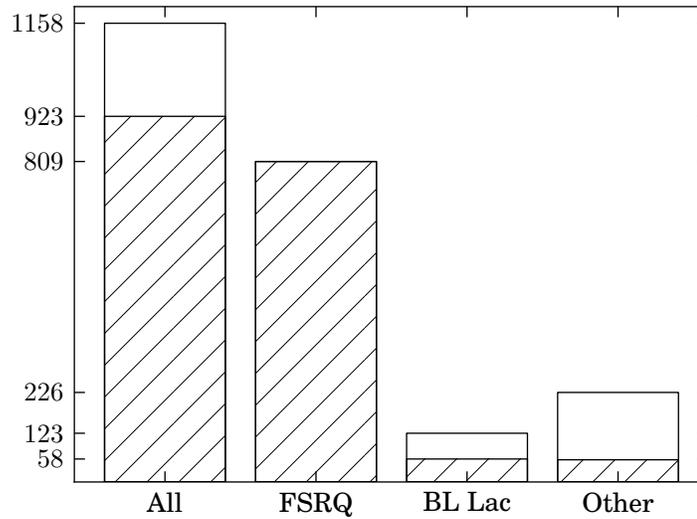


Figure 4.3. Optical classifications for the CGRaBS sample. Shaded regions indicate the number of sources in each class with measured redshifts. The “Other” category has 56 sources with known redshifts.

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).

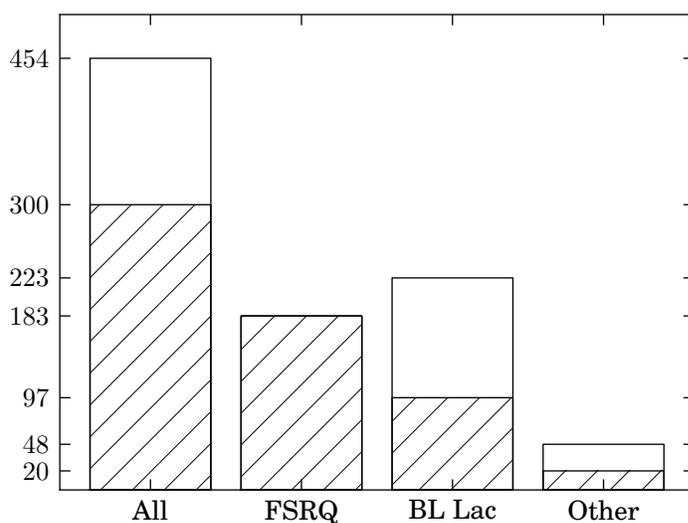


Figure 4.4. Optical classifications for the 1LAC sample. Shaded regions indicate the number of sources in each class with measured redshifts.

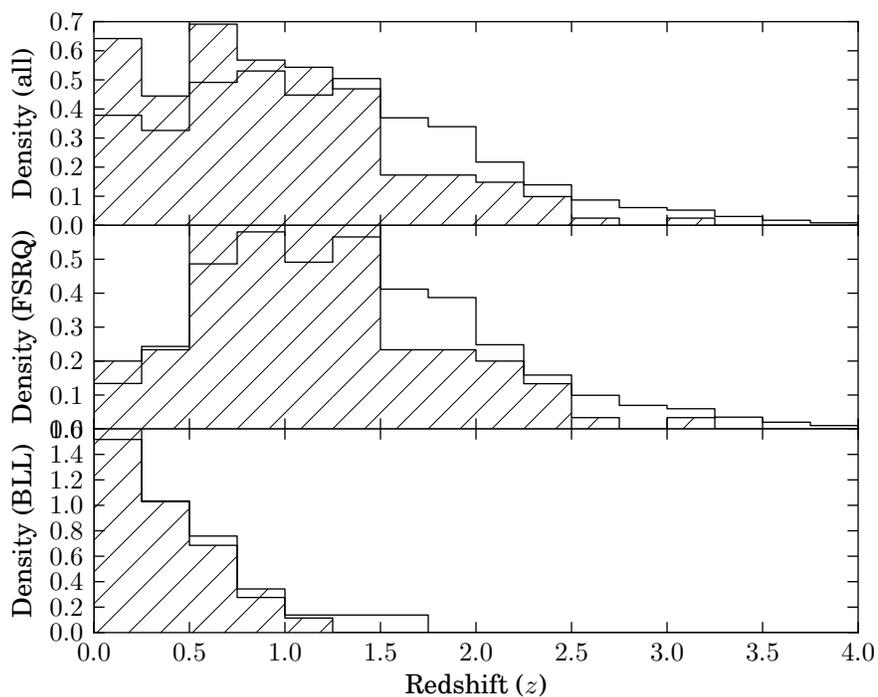


Figure 4.5. Histograms of redshifts for the CGRaBS sample (top), the FSRQ subset (middle), and the BL Lac subset (bottom). In each plot, the subset of sources that are also in the 1LAC sample is shown by the shaded region. Each histogram is normalized to integrate to unity.

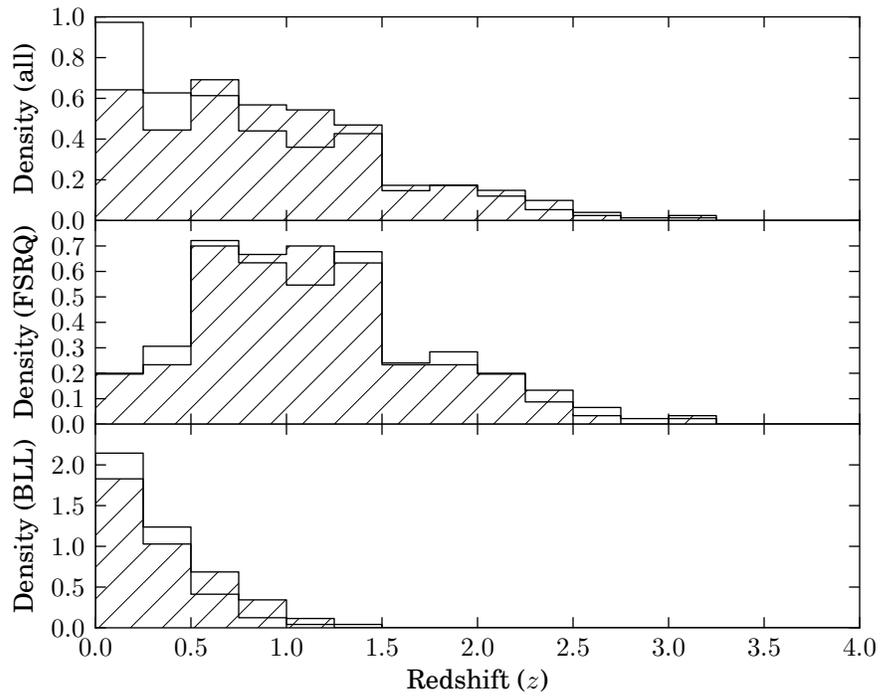


Figure 4.6. Histograms of redshifts for the 1LAC sample (top), the FSRQ subset (middle), and the BL Lac subset (bottom). In each plot, the subset of sources that are also in the CGRaBS sample is shown by the shaded region. Each histogram is normalized to integrate to unity.

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) test¹ rejects the hypothesis that the CGRaBS and the 1LAC subset are drawn from the same redshift population for all sources ($p = 5.8 \times 10^{-5}$) 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.

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

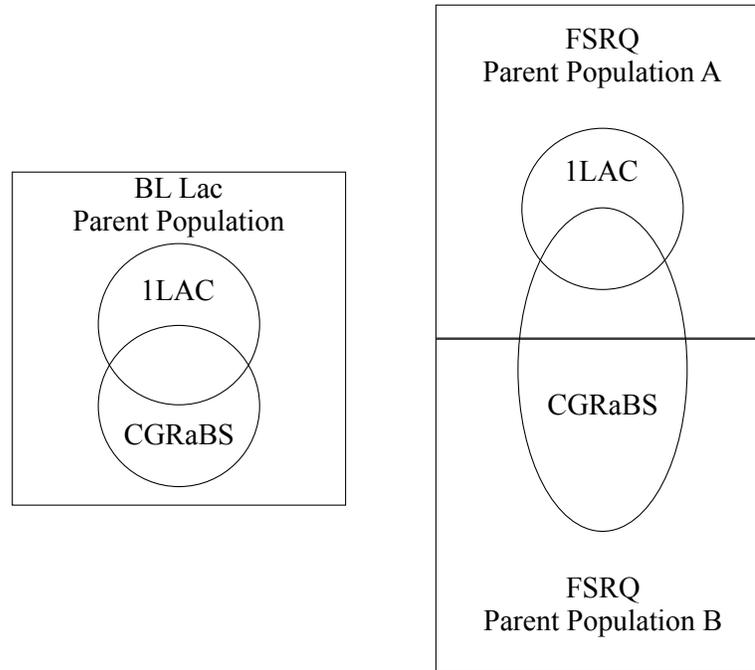


Figure 4.7. Venn diagram showing the relationship between the CGRaBS and 1LAC samples suggested by comparison of optical classification and redshift distributions.

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-Moerbeek. The new algorithm used from 2009 until the present was developed and written by Walter Max-Moerbeek. 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 $\lesssim 20^\circ$, using the HEALPix² 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/>

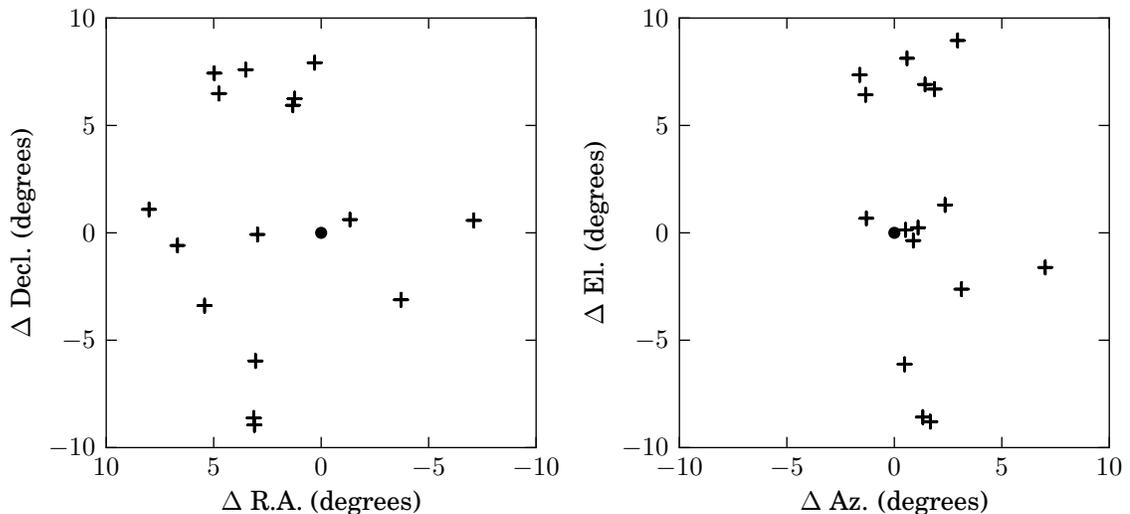


Figure 4.8. Example of an observation region, Region 83 in the 11 April 2011 schedule revision. *Left:* Source positions in equatorial coordinates offset from the pointing calibrator, indicated by the filled circle. *Right:* Region as it was observed on 15 April 2011, plotted in telescope azimuth and 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_{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.

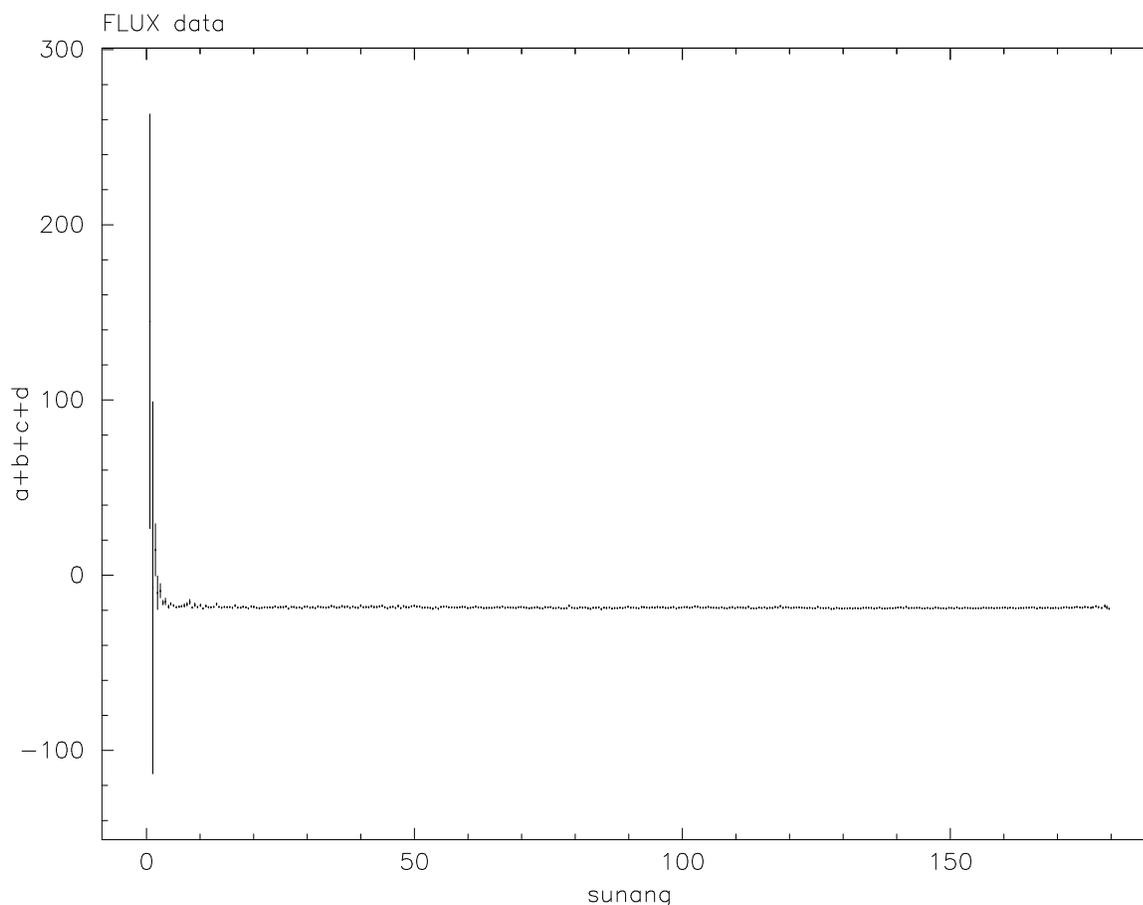


Figure 4.9. Source-nulled flux (S_{null}) (arbitrary units) as a function of solar elongation (degrees). Plotted points and error bars are the means and standard deviations of 2009 data in $\sim 0.5^\circ$ bins.

Figure 4.9 shows S_{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.

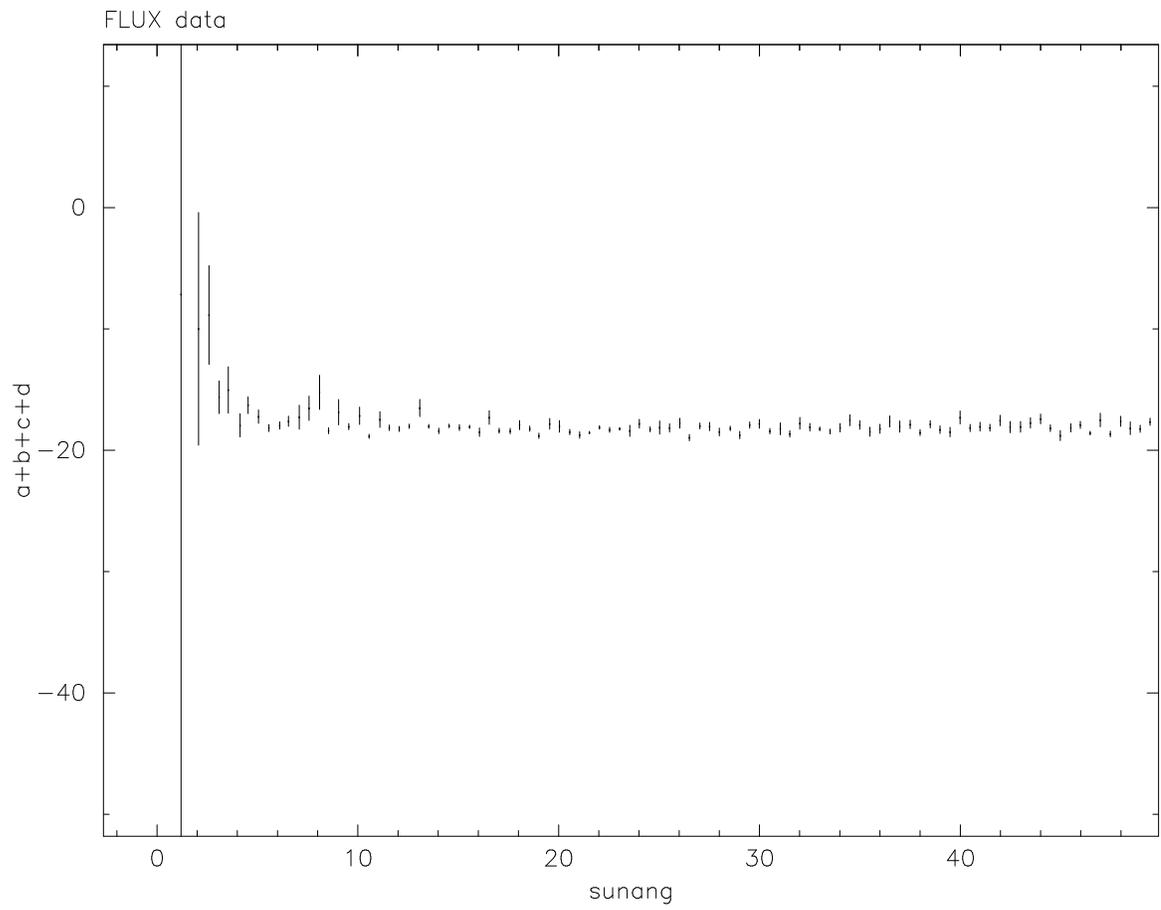


Figure 4.10. Detail of figure 4.9, showing features at small solar elongation.

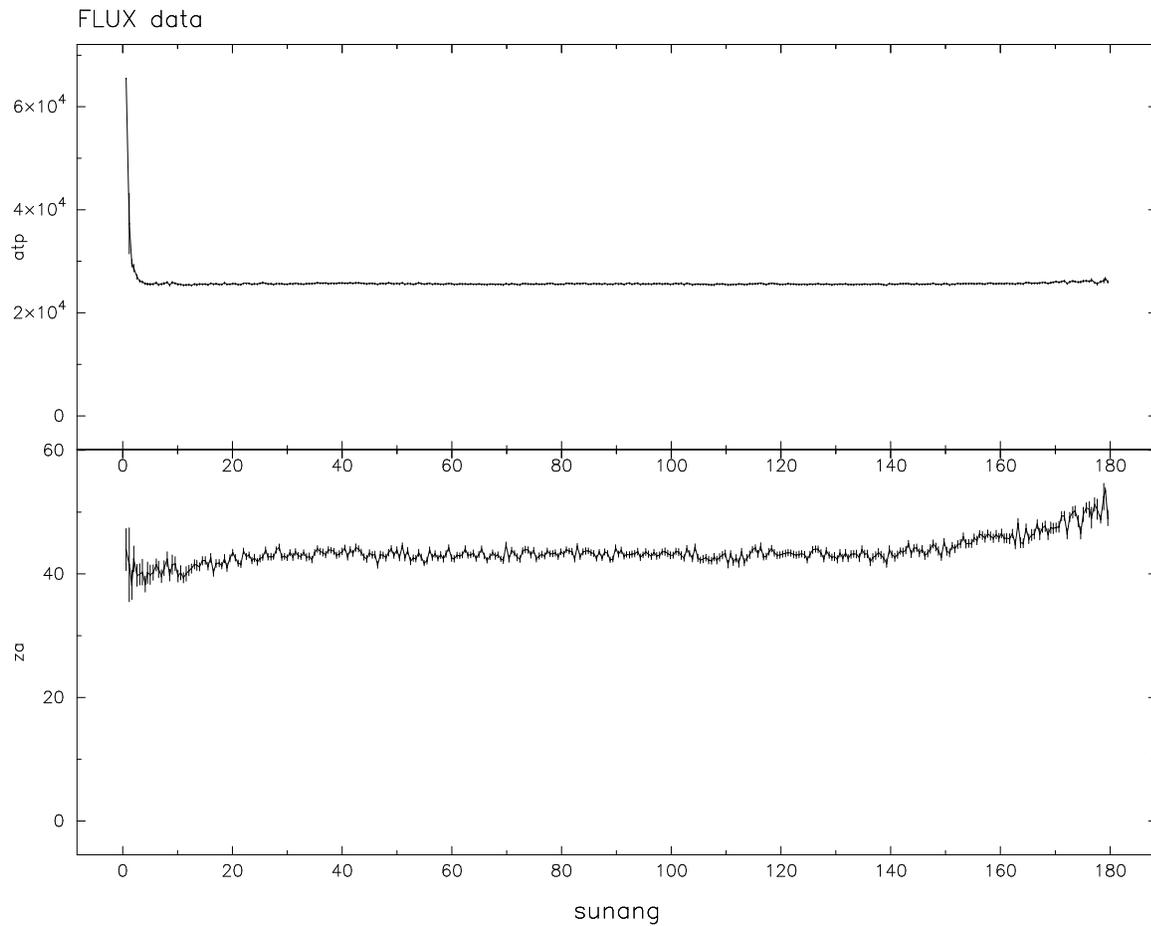


Figure 4.11. Top panel: Total power FLUX A segment data plotted against solar elongation. Data points (errors) are the means (standard deviations) within 0.5° bins. Bottom panel: Zenith angle of the observations in the upper panel, using the same binning. The absence of low-zenith angle observations 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.

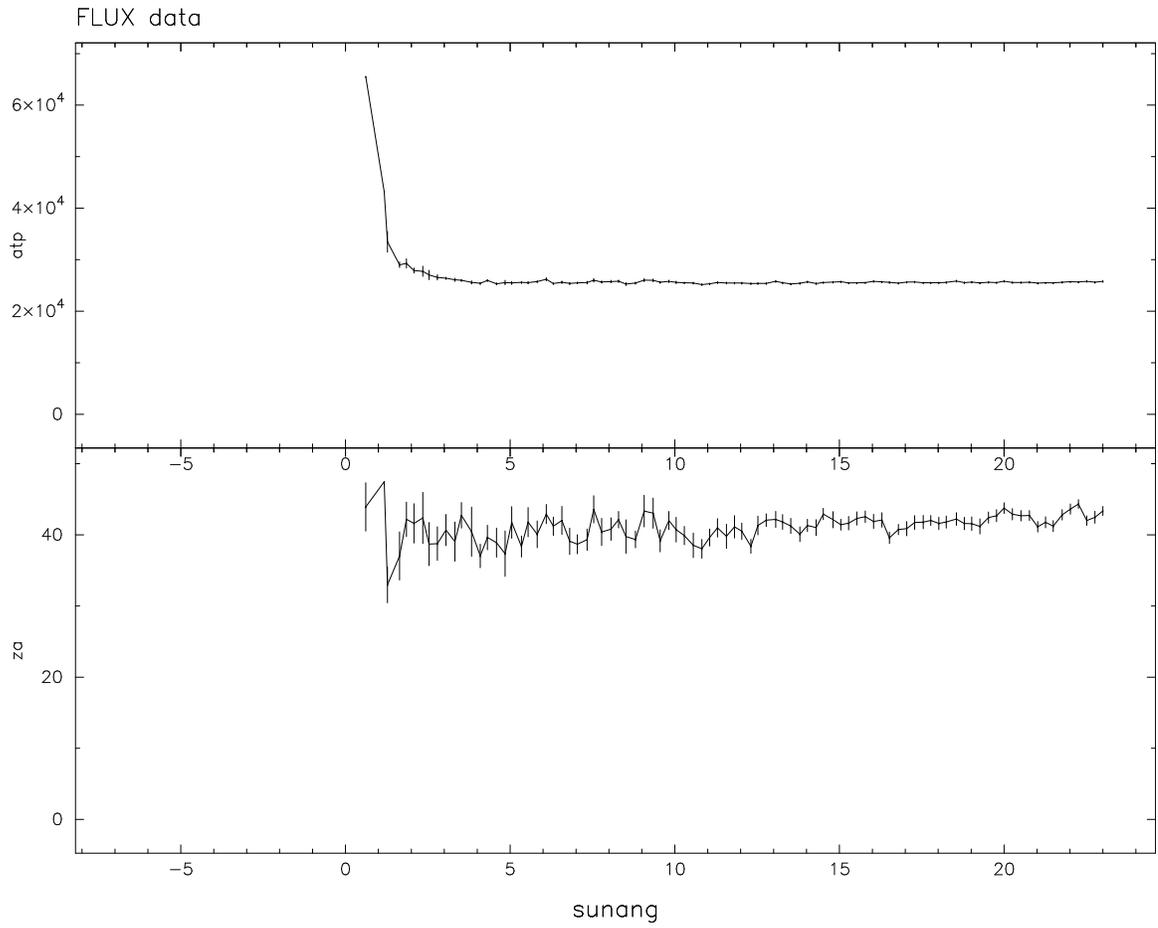


Figure 4.12. Detail of figure 4.11 with bins of $\sim 0.25^\circ$.

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

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

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 $n_0 = 51 \text{ Jy}^{-1} \text{ sr}^{-1}$ 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

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

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 $g_A(\theta, \phi) = \exp(-\theta^2/\theta_0^2)$ with $\theta_0 = 4.57 \times 10^{-4}$ rad. The region to integrate is bounded by $\theta_{max} = \theta_0 [\ln(S/S_c)]^{1/2}$, so

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

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

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

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

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

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

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

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 $n_0 = 51 \text{ Jy}^{-1} \text{ sr}^{-1}$, $\beta = -2.15$, $\theta_0 = 4.57 \times 10^{-4} \text{ rad}$, 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 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

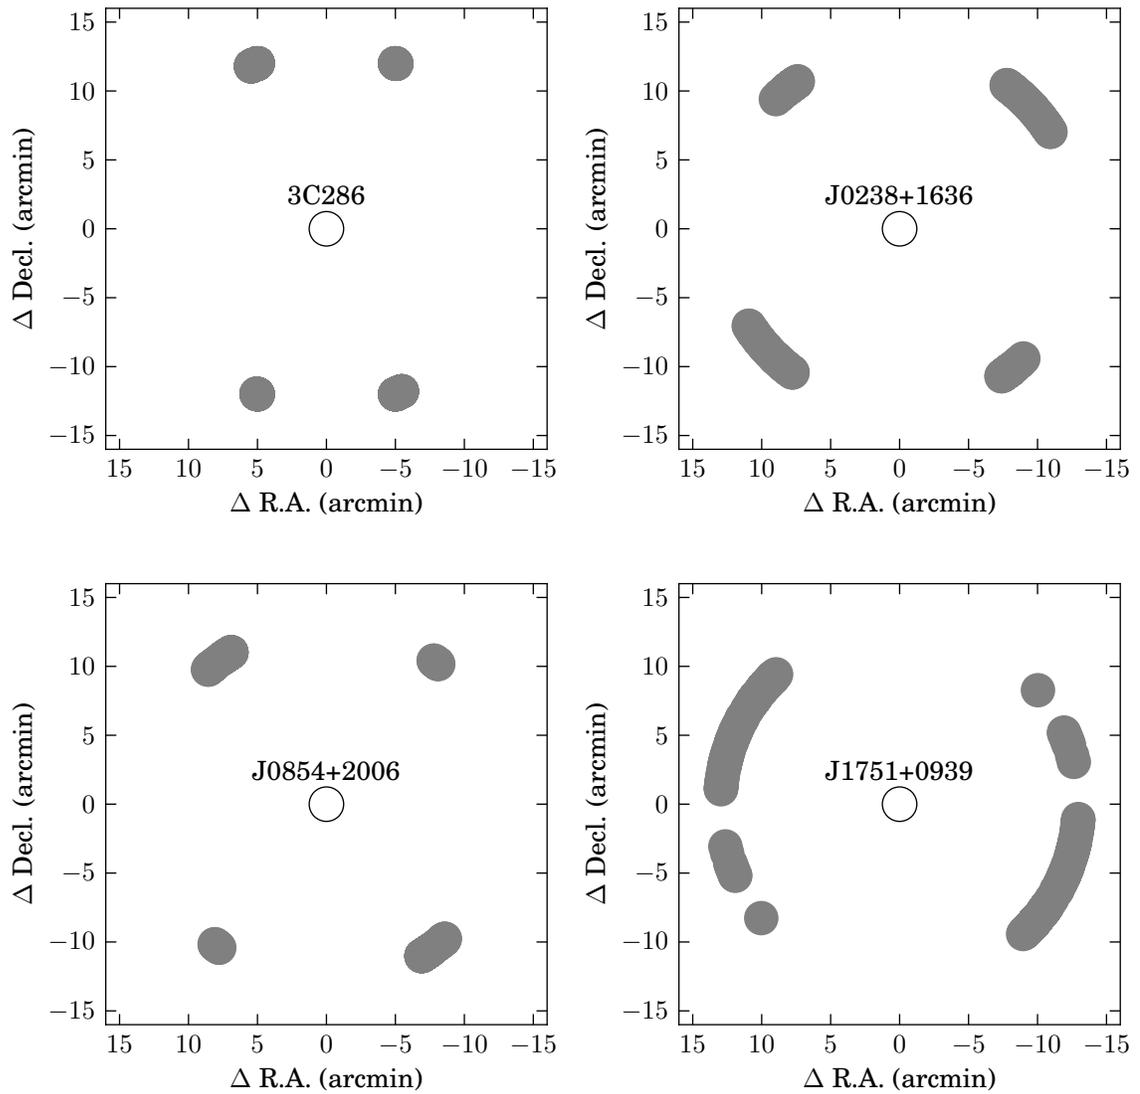


Figure 4.13. Reference field coverage during the full data set for four sources. The size (FWHM) of a single field is illustrated by the open circle at the source field location. J1751+0939 is a particularly extreme case, the other sources are more typical.

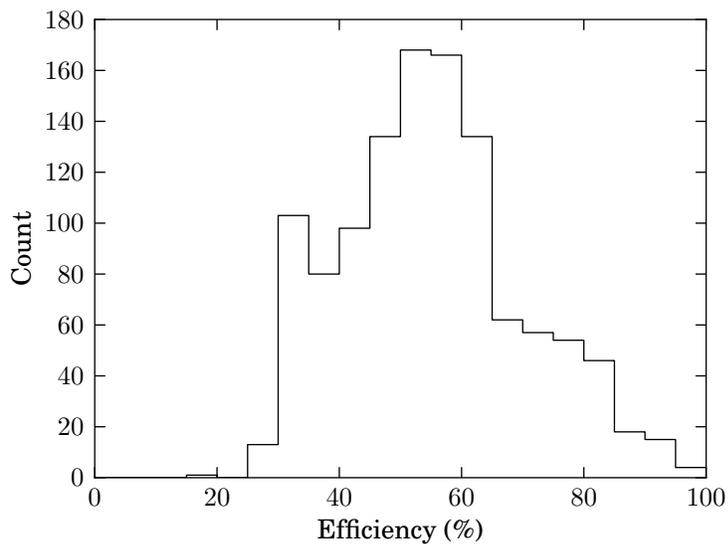


Figure 4.14. Histogram of per-source observing efficiency for CGRaBS sources. The expected number of observations per source was 365. Five sources exceeded 100% efficiency. The mean per-source efficiency was 55%.

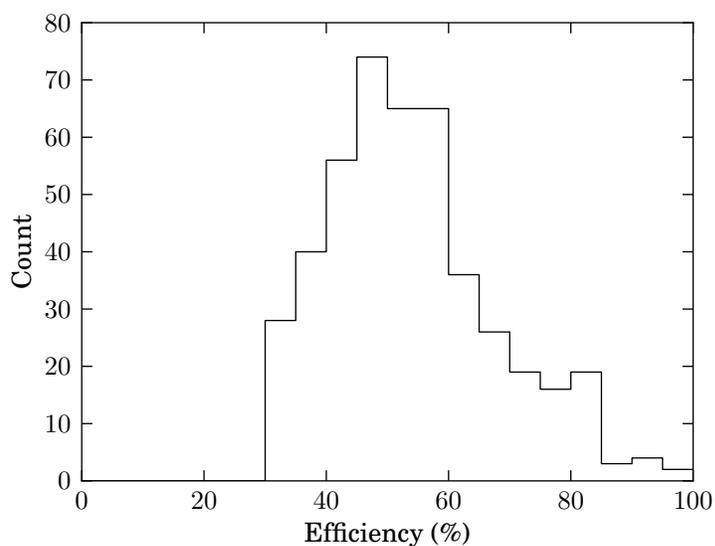


Figure 4.15. Histogram of per-source observing efficiency for ILAC sources. The expected number of observations per source was 365 for CGRaBS sources. For non-CGRaBS sources, the expected number was computed using the time between the first successful observation and the end of the 42-month interval,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

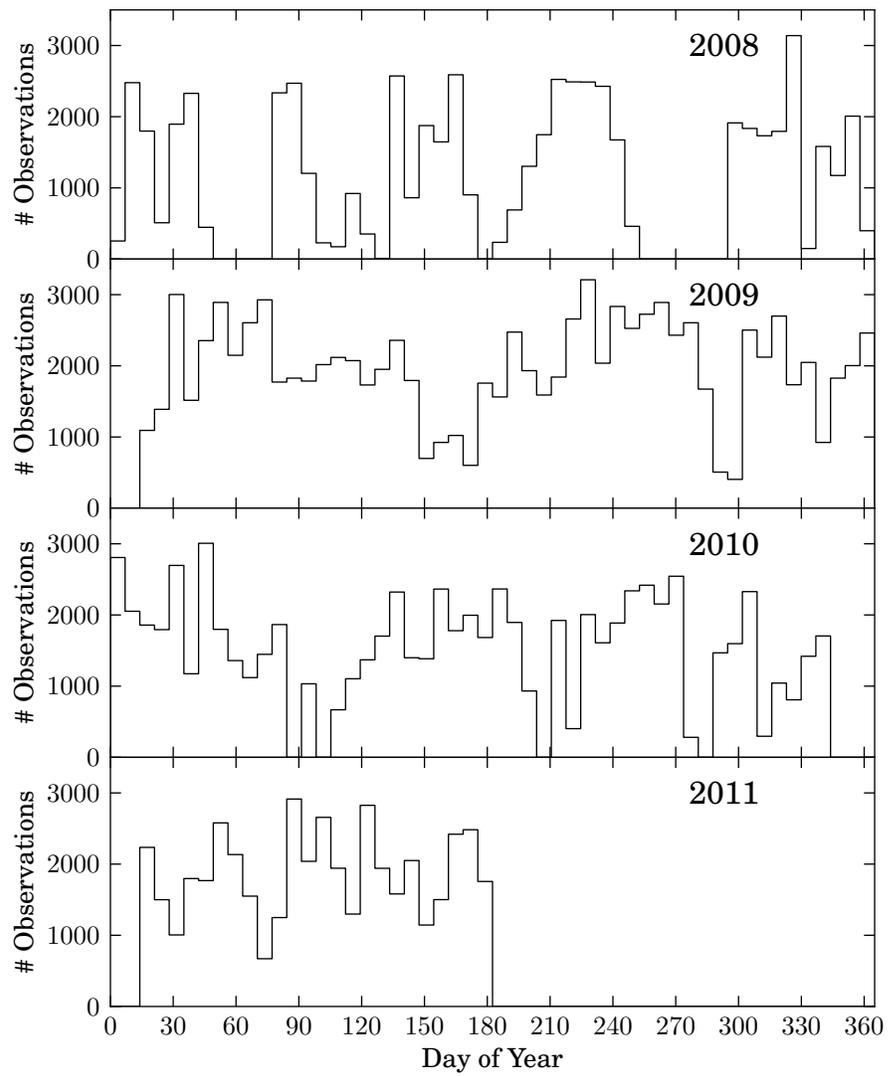


Figure 4.16. Weekly observation counts for each year of observations. Data plotted are the total numbers of flux density observations that survive to the end of the reduction pipeline.

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,

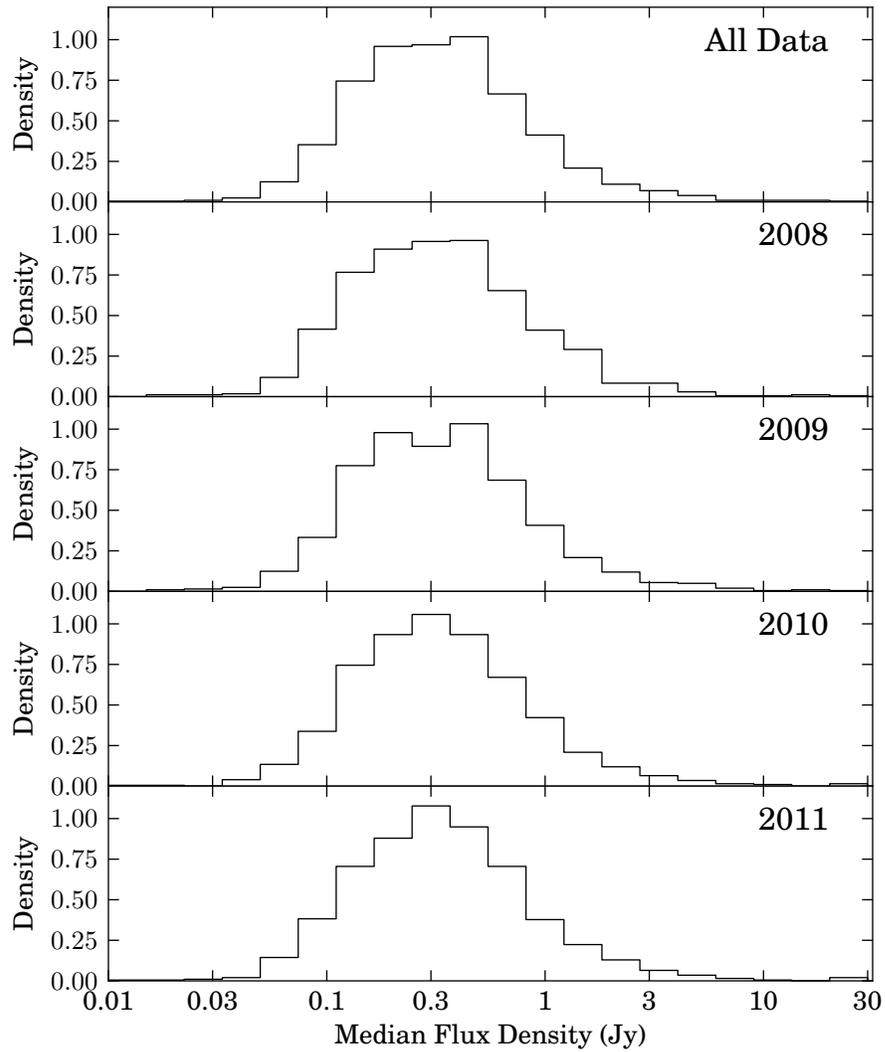


Figure 4.17. Histograms of the distributions of per-source median flux density in the CGRaBS sample for the entire 42-month data set and for individual years. Note that 2011 includes only 6 months of data. The visual similarity of the distributions is confirmed by K-S tests, which do not reject the hypotheses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

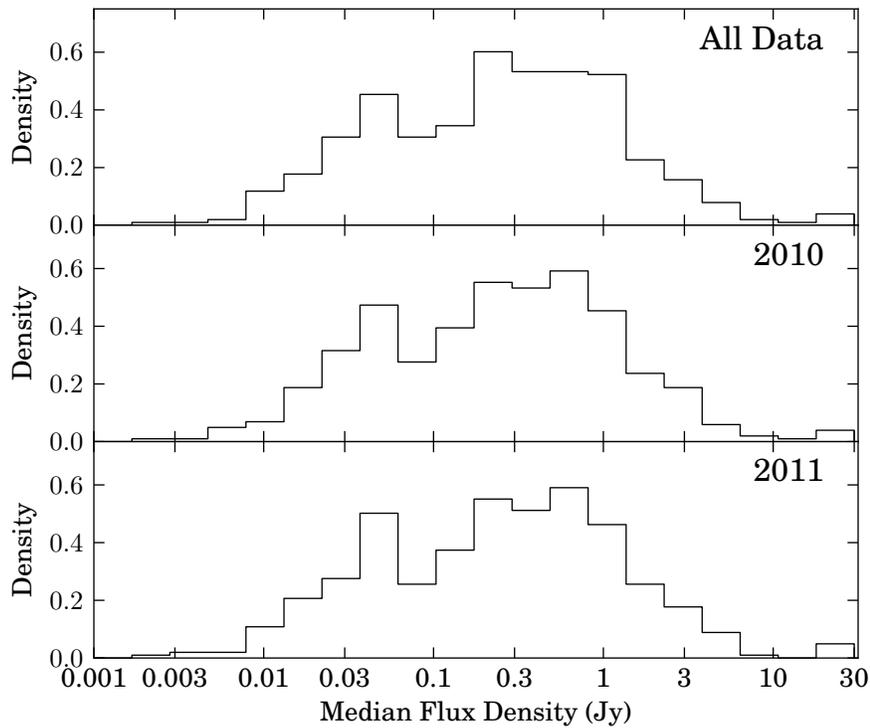


Figure 4.18. Histograms of the distributions of per-source median flux density in the 1LAC sample for the entire data set and for individual years. Note that 2011 includes only 6 months of data. K-S tests do not reject the hypotheses that these data are drawn from identical distributions ($p > 0.99$ 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.