

Exploring the Dynamic Radio Sky with the Allen Telescope Array

By

Peter Kelsey George Williams

Unofficial version (February 3, 2014)

The official version is [ProQuest document 1322973228](#)

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Astrophysics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Geoffrey C. Bower, Chair

Professor Steven E. Boggs

Professor Carl E. Heiles

Fall 2012

Exploring the Dynamic Radio Sky with the Allen Telescope Array

Unofficial version (February 3, 2014)

The official version is [ProQuest document 1322973228](#)

Copyright 2012

by

Peter Kelsey George Williams

Abstract

Exploring the Dynamic Radio Sky with the Allen Telescope Array

by

Peter Kelsey George Williams

Doctor of Philosophy in Astrophysics

University of California, Berkeley

Professor Geoffrey C. Bower, Chair

The revolution in digital technology that has had so many obvious effects in recent decades has not spared the field of astronomy. It has led to an enormous improvement in astronomers' ability to study the "time domain," the expected and unexpected ways in which celestial objects change on timescales ranging from milliseconds to centuries. In the field of radio astronomy a variety of advances have led to a new breed of observatories that are orders of magnitude more efficient at surveying the sky than previous facilities. These new observatories produce data at prodigious rates, however, and require sophisticated analysis to take full advantage of their capabilities. With several major facilities coming online in the next few years, there is an urgent need to prove that terabytes of data can be reliably turned into genuine astrophysical results.

This dissertation develops tools and techniques for coping with this challenge and applies them to data obtained with the Allen Telescope Array (ATA), a pioneering next-generation radio observatory located in Northern California. The ATA was built from the ground up to be a fast survey instrument, incorporating a suite of the new technologies that figure prominently in the new telescopes.

I develop and describe `miriad-python`, a framework for the rapid development of interferometric analysis software that is used in a variety of ways in my subsequent research. I also present a robust software system for executing multiple observing campaigns cooperatively ("commensally") at the ATA. Data from the ATA are difficult to analyze due to nontraditional features such as a large instantaneous field of view; continuous coverage of a large, interference-prone frequency range; and broadband, movable feeds; I describe and implement several methods for coping with these challenges.

This technical work is driven by the needs of a variety of astrophysical applications. I use broadband spectra of starforming galaxies to investigate the "calorimeter" interpretation of their cosmic ray energetics. The data are consistent with a recent hypothesis that the magnetic fields in these galaxies are stronger than traditionally thought. I use the survey

capabilities of the ATA to conduct ASGARD, a large survey of the Galactic component of the dynamic radio sky, which has remained poorly-explored due to the limitations of previous observatories and the technical challenges involved. I discuss in detail the methods used to analyze the data and provide catalogs, maps, completeness functions, and variability statistics. I map extended radio structures in Galactic fields and show how they can be subtracted from the data to simplify the search for transient Galactic sources. I limit the density of transient sources brighter than 10 mJy to be $<0.5 \text{ deg}^{-2}$ at 95% confidence. One of the areas of emphasis in this survey was the fascinating system Cygnus X-3, which shows prominent flares across the electromagnetic spectrum. Observations from 2010 May show a bright (1 Jy) radio flare followed by a 4.3σ γ -ray flare ($E > 100 \text{ MeV}$) ~ 1.5 days later. This timing is inconsistent with standard inverse-Compton models, suggesting that multiple mechanisms may be responsible for the system's high-energy emission.

Contents

Acknowledgments	v
1 Introduction	1
1.1 Synoptic Astronomy	2
1.2 Designing Radio Survey Telescopes	3
1.3 The Promise of the Dynamic Radio Sky	6
1.4 The 42-Element Allen Telescope Array	8
1.5 Dissertation Outline	9
2 Rapid Interferometric Software Development Using MIRIAD and Python	11
2.1 Introduction	12
2.2 Design Considerations in <code>miriad-python</code>	14
2.3 Comparisons to Related Packages	15
2.3.1 Bindings	16
2.3.2 Hybrid Packages	17
2.4 Implementation	18
2.5 Example Code	20
2.6 Applications	22
2.6.1 Algorithms for Millisecond Transients	22
2.6.2 Retroactive SEFD Calibration	23
2.6.3 Interactive u - v Data Visualization	26
2.6.4 An ATA Reduction Pipeline	27
2.7 <code>miriad-python</code> as a Software Project	28
2.7.1 Intellectual Property Issues	28
2.7.2 Availability	29
2.7.3 Development Model	29
2.7.4 Documentation and Examples	29
2.7.5 Quality Assurance	30
2.8 Performance Considerations	30
2.8.1 Serial Tests	30
2.8.2 Parallelization	31

2.9	Summary	33
3	The ATA Commensal Observing System	35
3.1	Introduction	35
3.2	Survey and System Design	36
3.3	System Implementation	37
3.3.1	Scheduling	38
3.3.2	Launching and Managing Observations	40
3.3.3	Synchronization Protocol	41
3.3.4	Correlator Observer	42
3.3.5	Observation Database	44
3.4	Recommendations	46
4	Evaluating the Calorimeter Model with Broadband, Continuous Spectra of Starburst Galaxies	53
4.1	Introduction	54
4.2	Observations	56
4.2.1	Details of observing methods	57
4.3	Data Reduction	59
4.3.1	Flagging	61
4.3.2	Calibration	64
4.3.3	u - v modeling	69
4.3.4	Uncertainties	75
4.3.5	Spectra	77
4.3.6	Self-consistency Checks	79
4.4	Spectral Modeling	86
4.4.1	Isolating the thermal emission	87
4.4.2	Modeling results	89
4.5	Discussion	92
4.6	Conclusions	97
5	ASGARD, A Large Survey for Slow Galactic Radio Transients: Overview and First Results	99
5.1	Introduction	100
5.2	Observations	101
5.3	Analysis	104
5.3.1	Subset of Data Presented in This Work	105
5.3.2	Calibration & Flagging	106
5.3.3	Imaging and Source Extraction	107
5.3.4	Detection of Uncataloged Sources	108
5.3.5	Primary Beam Modeling	109

5.3.6	Multi-Epoch Photometry	110
5.3.7	Timescales	111
5.4	First Results	114
5.4.1	Images	114
5.4.2	Source Catalog	123
5.4.3	Detection Limits and Completeness	124
5.4.4	Search for Transients	126
5.4.5	Variability Analysis	129
5.5	Discussion & Conclusions	132
5.5.1	Galactic Radio Transient Areal Densities	132
5.5.2	The Nature of the Most-Variable Sources	143
5.5.3	Outlook	144
6	The 2010 May Flaring Episode of Cygnus X-3 in Radio, X-Rays, and γ-Rays	147
6.1	Introduction	147
6.2	Observations & Data Analysis	149
6.2.1	Radio	149
6.2.2	X-Ray	150
6.2.3	γ -Ray	150
6.3	Results	151
6.4	Discussion	152
6.5	Summary & Conclusions	155
7	Conclusion	157
7.1	Principal Results	157
7.2	Looking Ahead	158
A	Confused, Wide-Field Interferometric Imaging and Photometry	161
A.1	The Wide-Field Challenge	161
A.2	Large-Scale Structure	163
A.3	Photometry	169
B	Probability of Constancy From a Set of Measurements	171
	Bibliography	173

Acknowledgments

This dissertation isn't mine. At least, this dissertation isn't mine alone. It would never have been finished — or started — without the help, love, and support of many, many people. I'd like to mention a few of them, with the recognition that many more names could be put here.

But first, a few of the formalities. Virtually all of the research presented herein was aided by NASA's Astrophysics Data System; the SIMBAD database, operated at CDS, Strasbourg, France; and the NASA/IPAC Extragalactic Database (NED) which is operated by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration. These acknowledgments are so routine that I suspect many people cease to notice them, but the prospect of doing research without these pieces of infrastructure is scary indeed.

Virtually all of my research is likewise done sitting in front of a computer, and as a scientist it's extremely important to me that the tools I use be open for anyone in the world to inspect and use. This wouldn't be possible without the work of innumerable people who have helped build the foundation of open-source and Free Software tools that I use every day.

I acknowledge the generous support of the Paul G. Allen Family Foundation, which has provided major support for design, construction, and operations of the Allen Telescope Array (ATA). Contributions from Nathan Myhrvold, Xilinx Corporation, Sun Microsystems, and other private donors have been instrumental in supporting the ATA. The ATA has been supported by contributions from the US Naval Observatory in addition to National Science Foundation grants AST-050690, AST-0838268, and AST-0909245.

I personally have been supported by a National Science Foundation Graduate Research Fellowship and a Space Sciences Laboratory Summer Fellowship.

My research has intimately involved the ATA and its home, the Hat Creek Radio Observatory. I've been consistently inspired and impressed by the dedication of the staff that I got to meet during my time at Berkeley: Samantha Blair, Rick Forster, Gary Gimblin, Colby Gutiérrez-Kraybill, Les Jordan, Susie Jorgensen, Jeff Kaufman, Garrett Keating, Marc Masters, and Carol Shumaker. It's all too easy for their contributions to be overlooked, despite the fact that without them none of my work would have been possible. On the other end of the spectrum, my dissertation would look very different if not for Paul Allen's

support of our little undertaking. Peter Backus, Tim Kilsdonk, and Jon Richards at the SETI Institute were responsible for several *essential* parts of the ASGAR commensal observing campaign, including the execution of most of the observations; their sleepless nights prevented many of my own.

Here in Berkeley, for the past six years the Radio Astronomy Lab has been not just an intellectual home but also a great group of colleagues and friends. It's been a privilege to work with such a talented and likeable group of scientists and engineers. I have trouble thinking of anyone who exemplifies the positive aspects of the RAL spirit more than the late Don Backer. I only knew him for a relatively short time, but he's still a role model for me: if not by his presence, then by the obvious impact he's had on so many of the people I know.

Within the excellent RAL community, of course, are my adviser Geoff Bower and the rest of his group, especially Steve Croft and Casey Law. They've taught me so many of the little and big things that go into being a good scientist, and a whole lot more besides. Geoff's hard work, enthusiasm, and resilience have been a consistent example for me.

Although I've become great friends with a bunch of the BADgrads, the rest of MONSTER deserves special mention. It's hard to believe that it's been six years and that the last of us are finally moving on out. It must be some bizarre statistical fluke for my classmates to all be so awesome, but I won't complain.

It's been quite a ride experiencing the West Coast way of life. In my time at Berkeley I've gotten the chance to meet some very special people outside of the astronomy department (!), and sometimes even outside of the university (!!). I'm a better, happier person for having been here and I'll miss all of you. In the words of the Governor, though . . . I'll be back.

Then there are the people who helped me get to Berkeley in the first place. That includes you, CompSquad, and now you're enshrined in the academic literature. You guys are the best and I love you.

I've had the good fortune to be able to learn from some fantastic teachers, but the two who stand out are George Swift and Christopher "Doc" Jones. Thanks so much for everything you do. David Newberg played a less formal role but was a brilliant coach in every sense.

Finally, though my entire family has been so important to me, I have to end with just Mom and Dad. You've given me so much more than any one person could reasonably ask for, and you've asked for so little in return. You've led by example every day of my life. Without a doubt, I'm proud to be handing in this dissertation, but I'd like to think that I can and should be much more proud of the kind of person I am, and that comes entirely from you two. I love you so much.

Chapter 1

Introduction

One hundred years ago, astronomers observed the universe using only a tiny slice of the electromagnetic (EM) spectrum. Since that time, more and more windows on the universe have been opened, now spanning not only more than twenty orders of magnitude of EM wavelength, but also observing, or striving to observe, other carriers of information such as neutrinos, cosmic rays, the solar wind, and gravitational waves. Over the course of the past century of technical development, important and frequently unanticipated discoveries have been made virtually every time a new way of observing the universe has become available ([Harwit 1981](#)). Sometimes these new windows have indeed been previously unexplored wavelength ranges; other times they have involved significant increases in sensitivity, field of view, or time resolution. The quest to open new observational windows continues unabated, as expressed in plans for ever-more sophisticated facilities such as the Atacama Large Millimeter-Submillimeter Array (ALMA; [Brown et al. 2004](#)), the Large Synoptic Survey Telescope (LSST; [Tyson 2002](#)), the James Webb Space Telescope ([Gardner et al. 2006](#)), the Advanced Laser Interferometer Gravitational-Wave Observatory ([Harry & LIGO Scientific Collaboration 2010](#)), and the Square Kilometer Array (SKA; [Carilli & Rawlings 2004](#)).

The birth of radio astronomy, as heralded by [Jansky \(1932\)](#) and [Reber \(1940\)](#), was the opening of one of these new windows, and indeed it led to some of the fundamental astrophysical results of the 20th century. Key breakthroughs include the development of multi-element interferometry ([Ryle & Vonberg 1946](#)), the detection of the hydrogen hyperfine transition ([Ewen & Purcell 1951](#)), the discovery of the cosmic microwave background radiation ([Penzias & Wilson 1965](#)), the discovery of pulsars ([Hewish et al. 1968](#)), the discovery of molecules in the interstellar medium ([Weinreb et al. 1963](#); [Weaver et al. 1965](#); [Cheung et al. 1969](#); [Wilson et al. 1970](#)), and the measurement of superluminal expansion in distant quasars ([Cohen 1975](#); [Cohen et al. 1977](#)). The discoveries described in [Ryle & Vonberg \(1946\)](#) and [Hewish et al. \(1968\)](#) resulted in the awarding of the 1974 Nobel Prize for Physics to Ryle and Hewish. (Infamously, Hewish's female postgraduate student S. Jocelyn Bell did not receive a share of the prize despite being widely considered to have played a key role in the pulsar discovery.)

The explosion of development in the fields of digital computing and telecommunications over the past several decades has led to a constellation of technologies that promise to revolutionize observational radio astronomy, and if the past is any indicator, this revolution will likely lead to another round of astrophysical breakthroughs.

1.1 Synoptic Astronomy

As I will describe in more detail below, the revolution in radio astronomy is one towards faster surveying of the sky: compared to old telescopes, new ones effectively see much more of the sky at once. The most obvious impact of this change is to increase astronomers' ability to search for astronomical variability: to discover and monitor the appearance, evolution, and disappearance of astronomical objects. The gains are substantial enough to constitute a truly new capability to explore the *time domain* of the sky.

Humans have of course recognized various forms of variability in the night sky since well before the dawn of scientific astronomy, but it has taken information technology to make it possible to systematically search for it at the scales commonly seen today. Although there is no bright line differentiating modern time-domain techniques and more traditional ones, the accumulation of improvements in these techniques can credibly be said to have led to a qualitative shift in the thinking of many astronomers. This shift is exemplified by the Large Synoptic Survey Telescope, a planned major optical facility (Tyson 2002). The term *synoptic* here is analogous to *synopsis*, that is, “looking at all aspects of something,” and is meant to capture this change in thinking: the model goes from one in which telescopes are used to take snapshots of small parts of the sky to one in which they continually output an ongoing, full-color movie of nearly the entire heavens.

Although synoptic astronomy is still more of a goal than a reality, the current state of the art represents a substantial change from the past. Current optical wide-field survey projects such as the Palomar Transient Factory (Law et al. 2009), PanSTARRS (Kaiser et al. 2002), the Catalina Real-time Transient Survey (Drake et al. 2009), and SkyMapper (Keller et al. 2007) observe and process enormous amounts of data every night. To give a sense of recent advances, eleven supernovae were discovered by astronomers in 1981; in 2011, that number was 290. In the field of radio astronomy, there has been a similar burst of improvements, often involving wide-field, survey capabilities. ALMA; the Karl G. Jansky Very Large Array (JVLA, formerly EVLA; Perley et al. 2011); the Westerbork Synthesis Radio Telescope focal plane array, Apertif (Verheijen et al. 2008); the Arecibo L-Band Feed Array (ALFA; Giovanelli et al. 2005); the Murchison Wide-field Array (Lonsdale et al. 2009); the Precision Array for Probing the Epoch of Reionization (Parsons et al. 2010); the Long-Wavelength Array (Ellingson et al. 2009); the Low Frequency Array (Kassim et al. 2004); the Allen Telescope Array (ATA; Welch et al. 2009); the Australian SKA Pathfinder (ASKAP; DeBoer et al. 2009); the South African SKA pathfinder, MeerKAT (Jonas 2009); and the planned SKA itself are just a few examples.

The exploration of the astronomical time domain has already proved to be as scientifically fruitful as the opening of previous observational domains. Most strikingly, the systematic supernova surveys mentioned above provided the data necessary to first demonstrate the accelerating expansion of the universe (Riess et al. 1998; Perlmutter et al. 1999). Similar wide-area searches led to the discovery of the trans-Neptunian objects (Jewitt et al. 1992). Microlensing surveys were important for constraining the amount of baryonic dark matter in compact objects (Alcock et al. 1993; Aubourg et al. 1993). Although it would be a stretch to call the discovery of the first exoplanet around a main-sequence star (Mayor & Queloz 1995) a specially “time-domain” accomplishment, the related wide-area monitoring techniques being used by the *Kepler* mission (Koch et al. 2010) and the OGLE microlensing planet search (e.g., Bond et al. 2004) fit the mold. Monitoring of the sky in X-rays and γ -rays has yielded the detection of the tidal disruption of a star by a supermassive black hole (Bloom et al. 2011) and unexpected new populations of millisecond pulsars (Abdo et al. 2009b).

It is notable that there are no examples from radio astronomy on the above list. Though there have been some tantalizing early results (Hyman et al. 2005; McLaughlin et al. 2006; Lorimer et al. 2007; Lovell et al. 2008; Brunthaler et al. 2009), a bread-and-butter application of wide-field time-domain radio astronomy has yet to emerge. The challenge is twofold: the radio sky appears to be relatively quiet, without supernova-like sources that brighten by factors of billions; and the techniques needed to survey the radio sky are perhaps more challenging, and less well-understood, than those at other wavelengths: most of the new radio facilities mentioned above are just coming online now. The flip side of the coin is that the variable radio sky is comparatively poorly-explored, and the opportunity for genuine discovery is correspondingly larger (Cordes et al. 2004; Fender & Bell 2011; Frail et al. 2012). Although there is excitement for the possibility of truly unexpected breakthroughs, there are also many domains in which there are well-motivated expectations for new discoveries, as I shall describe. The pursuit of these phenomena is mainly limited by current radio observational tools, and it is these tools that are in the midst of a revolution.

1.2 Designing Radio Survey Telescopes

There is no one single technology that is leading the way toward synoptic radio astronomy. Some examples of important developments are wide-bandwidth feeds and low-noise amplifiers, inexpensive electronic cryogenics systems, focal plane array systems, composite reflector materials, and ever-more-sophisticated microwave hardware operating up to terahertz frequencies. For radio interferometers, the development of fast, commodity digital computing is by far the most significant advance, since computing power is a key bottleneck limiting how much information can be extracted from the incoming EM radiation. Advances in specialized hardware, to perform the low-level interferometric correlations, and general-purpose computers, to implement essential, sophisticated analysis algorithms, have already had an enormous impact in interferometric observation, and promise to continue to do so.

The advances in the radio facilities can be understood by considering a “figure of merit” (FOM) for the survey speed of a telescope. For radio telescopes, a useful formulation is

$$\text{FOM} \propto \Omega \left(\frac{A_{\text{eff}}}{T_{\text{sys}}} \right)^2, \quad (1.1)$$

where Ω is the instantaneous field of view (FOV) of the telescope and the ratio $A_{\text{eff}}/T_{\text{sys}}$, the system’s “effective collecting area” divided by its “system temperature,” specifies its sensitivity (Lazio et al. 2009). That is, a fast survey telescope is one that can see a large portion of the sky at once and that is sensitive to faint objects. In Figure 1.1, I plot an equivalent of this FOM for a variety of existing and new radio facilities. In the terms of this metric, new facilities, including ones that are online today, are in some cases two orders of magnitude more capable than older ones.

Although a portion of this improvement is due to increases in the $A_{\text{eff}}/T_{\text{sys}}$ sensitivity term, the larger share is usually due to growth in the FOV, Ω . (The notable exception to this generalization is the JVLA.) The FOV of a parabolic radio antenna generically scales as

$$\Omega \propto \left(\frac{\lambda}{D} \right)^2, \quad (1.2)$$

where λ is the observing wavelength and D is the characteristic size of the antenna (dish diameter, *etc.*). Treating λ as fixed, increases in the FOV therefore come at the expense of the collecting area of the antenna, and hence its sensitivity. Single antennas become faster survey tools as they become larger, but one runs into practical engineering limitations at a diameter of ~ 100 m, which not coincidentally is the size of the world’s largest radio dishes: Arecibo, the Green Bank Telescope, and the Effelsberg 100 m dish.

The progress in radio surveying comes from the development of various strategies for working around this basic limitation; these strategies are only being deployed today because they are uniformly complex and technically demanding. In the case of single antennas, focal plane array technologies, especially phased array feeds as being developed for ASKAP and Apertif, are a focus of active research. I will not discuss these further, however, and instead will focus on an approach that is specific to radio interferometers, telescopes in which the signals from multiple antennas are combined. Each antenna in an interferometer collects radio waves, so one with N dishes of diameter D has $A_{\text{eff}} \propto ND^2$. But the interferometer FOV is the same as one of its component dishes, so that

$$\text{FOM} \propto (D^{-2})(ND^2)^2 = N^2 D^2 = N A_{\text{eff}}. \quad (1.3)$$

If the antenna costs, which roughly scale as A_{eff} , dominate other costs, you get the most bang for the buck by building as many small dishes as possible. This approach is generally referred to as a “large number of small dishes,” “large N small D ,” or simply “LNSD,” design. Studies indicate that the cost of a single radio reflector in fact scales as $D^{2.7}$, making the LNSD

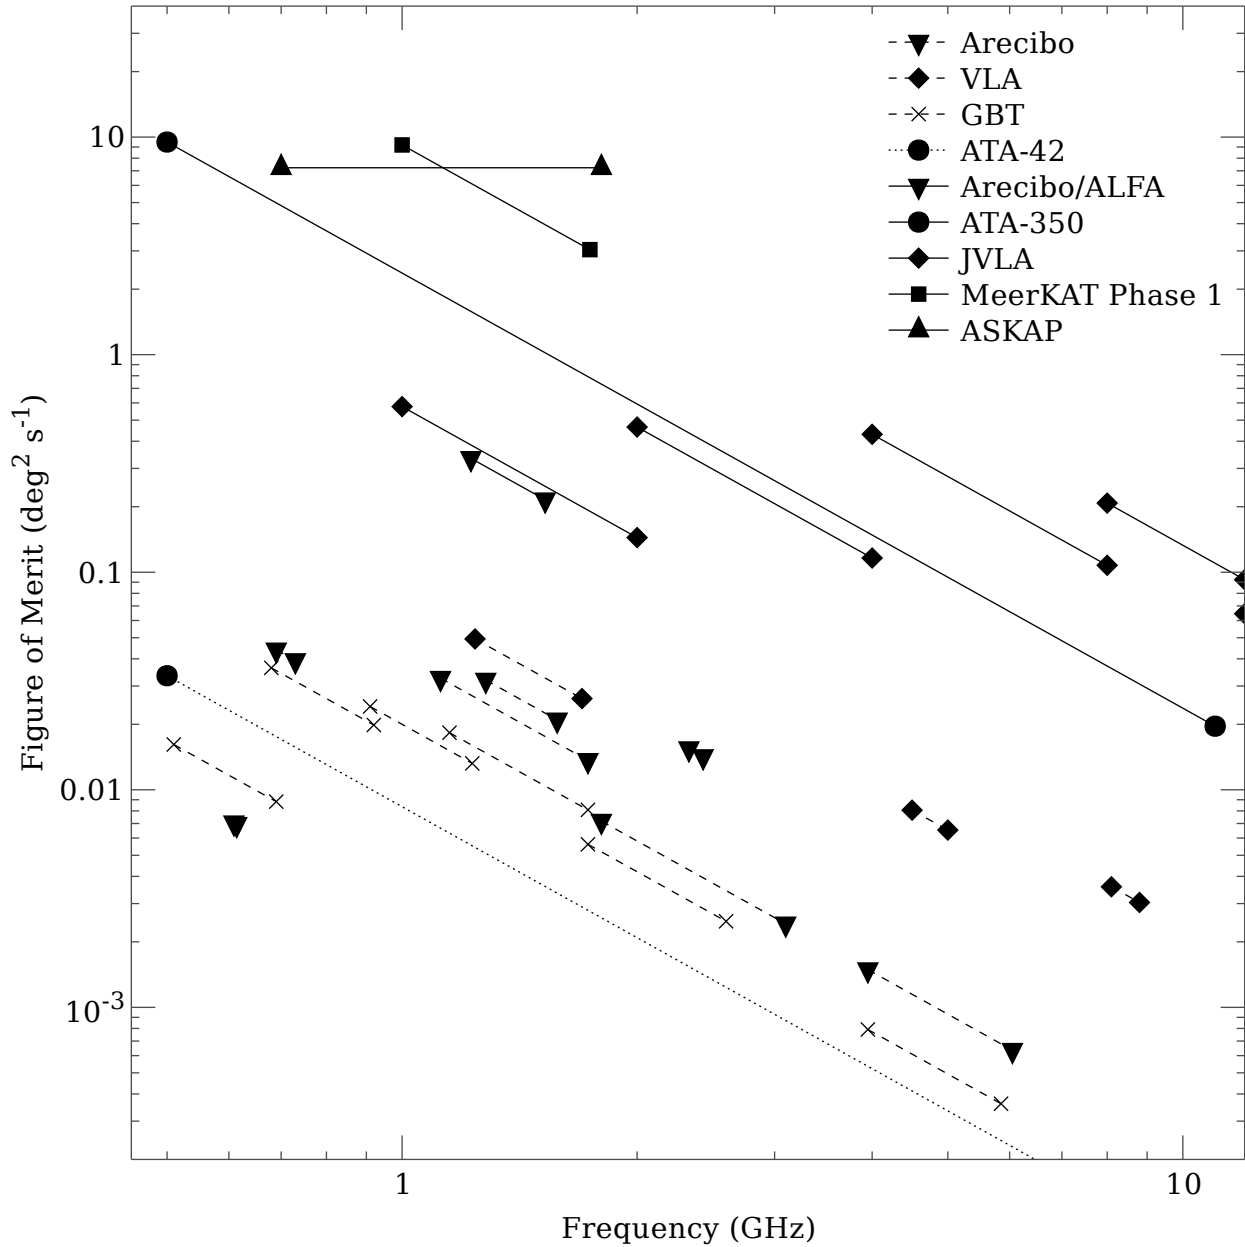


Figure 1.1: Survey speed figures of merit (FOMs) for various cm-wavelength facilities. The FOM is rate of sky coverage to a point-source sensitivity of 1 mJy. *Dashed lines* denote previous-generation facilities and *solid lines* denote next-generation or upgraded facilities. The *dotted line* denotes the ATA-42, which uses next-generation technologies but unlike all the other observatories shown is not a national-scale facility. System parameters for ASKAP, MeerKAT, and Arecibo/ALFA are from Johnston et al. (2009), Booth & Jonas (2012), and Giovanelli et al. (2005), respectively.

approach even more compelling (Bunton et al. 2002, Appendix A and references therein) in these terms.

The difficulty with radio interferometers is that the signals from the N dishes must be combined (“correlated”), and this is a computationally-demanding problem that scales as N^2 ; that is, if the number of antennas doubles, the size of the required computer increases by a factor of four. Historically, N has been limited to a few dozen. But digital technology has enabled the construction of far more powerful correlators; modern ones are highly specialized supercomputers that in some cases can handle hundreds of antennas. (The JVLA correlator, WIDAR, runs at 10 Pflop s^{-1} [Perley et al. 2011], about equal to the world’s current top supercomputer, Japan’s “K,” though the comparison is unfair because the latter is a general-purpose system.) Because correlators remain complicated machines, the optimal balance of telescope costs does not necessarily point toward maximizing N above all else (Bunton & Hay 2004), but without digital computing LNSD designs would simply be impossible to realize.

Other approaches to increasing the survey speed of radio telescopes involve different challenges, but in general, as Figure 1.1 shows, digital technologies are enabling a new era of synoptic radio astronomy.

1.3 The Promise of the Dynamic Radio Sky

What might we see once this next generation of radio telescopes is up and running? Broadly speaking, variable radio emission comes from energetic phenomena: explosions, shocks, and jets. There are many potential sources of time-variable radio emission, and several schemes for classifying them.

A pervasive yet frustrating classification in time-domain astronomy is whether a source can be considered “transient;” sources that change significantly with time but are not transients are usually termed “variables.” The general notion is that a transient is a source that has suddenly appeared for some reason, but this depends on the observer’s state of knowledge and is thus an awkward definition. (If you already knew something was there, can it be a transient?) One way of thinking about a transient is that it is a source that brightens by a large factor ($\gtrsim 10$, say), so that routine observations may fail to detect it for a while, then start detecting it robustly. Alternatively, a transient may be a source that wasn’t known to exhibit variability, but suddenly began to do so. In some cases I prefer to refer to “highly variable” sources, but the broader point is that this particular distinction is not something to get hung up on, unless the terms have been defined more precisely in a given context.

Another distinction is commonly made between “fast” and “slow” radio transients. The former are usually defined as pulses lasting $\lesssim 1 \text{ s}$, and the latter are everything else. Slow transients are usually more easily found by conventional imaging techniques (*e.g.* Murphy et al. 2012), while fast transients are more commonly found by analyzing time-stream data using thresholding techniques (*e.g.* Wayth et al. 2011; Law & Bower 2012). This line also roughly divides between emission mechanisms: short radio pulses generally come from coherent

processes, while longer-lived emission is usually the result of incoherent ones. The most famous short radio transients (or are they variables?) are the pulsars. Recently a larger variety of related phenomena has been discovered, such as the rotating radio transients (RRATs; [McLaughlin et al. 2006](#)), magnetar bursts ([Camilo et al. 2006](#)), and giant pulses ([Lundgren et al. 1995](#)). The Solar System planets, cool dwarfs, and perhaps even exoplanets ([Berger et al. 2001](#); [Berger 2002](#); [Hallinan et al. 2007](#); [Griessmeier et al. 2007](#)) also emit radio pulses due to maser-like processes. Finally, there have been suggested detections of giant extragalactic radio pulses ([Lorimer et al. 2007](#); [Keane et al. 2011](#)) although they may be explicable as terrestrial interference ([Burke-Spolaor et al. 2011](#)). In this work I concern myself only with slow transients.

“Intrinsic” and “extrinsic” variables may also be distinguished. Extrinsic variables are those in which the source itself is not variable, but its apparent brightness changes due to propagation effects between it and the Earth. Compact radio sources at GHz frequencies are subject to interstellar scintillation — twinkling — that can be a significant source of extrinsic variability ([Spangler et al. 1989](#); [Rickett 1990](#); [Narayan 1992](#)). Small-scale structure in the ionized component of the interstellar medium (ISM) gives rise to the class of intraday variable sources ([Dennett-Thorpe & de Bruyn 2002](#)). In “extreme scattering events” ([Fiedler et al. 1987](#); [Walker 2006](#)), which are still poorly understood, strong decreases in source intensity lasting months are seen. Although extrinsic variability can be considered a nuisance in some cases, it can also be used to perform exquisite measurements of the microstructure of the ISM ([Dennett-Thorpe & de Bruyn 2002](#)).

There is also the distinction between Galactic and extragalactic radio variables. To a certain extent, these do not represent separate classes of sources, but different ends of a luminosity function: extragalactic sources are rare and bright whereas Galactic sources are comparatively common and faint. Certain events, however, such as the γ -ray bursts described below, are sufficiently rare that it makes little sense to interpret them as a potential Galactic population. Others, such as flare stars, are sufficiently low-luminosity that there is virtually no hope of being able to detect them outside of the Milky Way.

Some extragalactic radio variability originates with supermassive black holes at the centers of galaxies: either from long-lived active galactic nucleus (AGN) jets ([Lister et al. 2011](#)), or from short-lived bursts due to a spike in the matter accretion rate, as happens in the “tidal disruption events” in which a star is sheared apart during a close encounter with such a black hole ([Rees 1988](#); [Bloom et al. 2011](#); [Berger et al. 2012](#)). Other extragalactic variability originates with cataclysmic stellar events, as in radio-loud supernovae ([Weiler et al. 2002](#); [Brunthaler et al. 2009](#); [Muxlow et al. 2010](#)) or phenomena related to γ -ray bursts (GRBs). Of particular interest are “orphan GRB afterglows” ([Frail et al. 2001](#); [Levinson et al. 2002](#)). As their name would suggest, GRBs are detected as powerful jetted bursts of γ rays. What is not known, however, is just how powerful those jets are, because it is unclear how tightly collimated they are. After a GRB explosion, an “afterglow” is seen, and the radio components of these afterglows emit isotropically, *i.e.*, they are not jetted like the γ rays themselves. Therefore there should exist radio afterglows that can be seen even if the initial burst of γ

rays was pointed away from the Earth. If these can be detected, their numbers can be used to measure the typical collimation of GRB jets and hence their total energy output.

Galactic radio variables include many of the fast transient sources mentioned above (pulsars, RRATs, magnetars, cool dwarfs, exoplanets), X-ray binaries (Waltman et al. 1995; Marscher & Brown 1975), active stellar binaries (Hall 1976; Bopp & Fekel 1977; Eker et al. 2008), and flare stars (Güdel 2002; Jackson et al. 1989; Richards et al. 2003). The “microquasar” subclass of X-ray binaries (Mirabel & Rodríguez 1998, 1999; Gallo 2010) show relativistic jets like those of their full-scale extragalactic brethren, but are valuable laboratories for the study of accretion on to black holes because they evolve on much faster timescales.

Unlike the extragalactic population, the Galactic population is expected to have a strong spatial dependence, with the most sources being found towards the Galactic Center (GC), of course. There is indeed a history of mysterious transient radio sources being detected in this direction (Davies et al. 1976; Zhao et al. 1992; Bower et al. 2005; Hyman et al. 2005, 2009), usually disappearing before follow-up observations can conclusively identify them: followup efforts are complicated by Galactic extinction, which makes identification at other wavelengths more difficult if not impossible, depending on the sightline and wavelength in question. Along with the higher density of sources in general towards the GC, there is also a higher density of ionized plasma, leading to an expectation of more extrinsic variability induced on background sources due to scintillation effects (Gregory & Taylor 1986; Becker et al. 2010).

1.4 The 42-Element Allen Telescope Array

Many aspects of this dissertation are intimately connected to the Allen Telescope Array (ATA), an LNSD radio interferometer located in Hat Creek, California at the Hat Creek Radio Observatory. A portion of it is pictured in Figure 1.2. From its construction until mid-2011 the ATA was operated jointly by the UC Berkeley Radio Astronomy Laboratory and the SETI¹ Institute.

The ATA is a pioneer of the new class of wide-field radio astronomy facilities. At the time of the work described in this dissertation, the ATA consisted for forty-two 6.1-m offset Gregorian dishes equipped with log-periodic, dual-linear-polarized feeds. The vision for the full ATA involves an expansion to 350 dishes. As shown in Figure 1.1, such a facility (the “ATA-350”) would be competitive with state-of-the-art radio survey telescopes. Technical innovations of the ATA include continuous frequency coverage from 1 to 11 GHz, thanks to the broadband feeds and amplifiers; the offset Gregorian dish design, which provides imaging advantages thanks to the unblocked aperture; high survey speed, thanks to the LNSD design; multiple interchangeable data-processing backends, made possible by the transmission of the full incoming voltage stream from each antenna to the signal processing systems; and innovative low-cost antenna and support system designs, necessary for the affordable

¹Search for Extraterrestrial Intelligence



Figure 1.2: Elevated view of a portion of the Allen Telescope Array, looking approximately east. The large white structure is a tent for antenna construction, and the smaller yellow structure is the control building, housing the signal processing room and observer facilities.

construction of hundreds of high-quality dishes. More details about the ATA are presented in [Welch et al. \(2009\)](#).

1.5 Dissertation Outline

Although the promise of synoptic studies of the radio sky is great, the techniques and facilities needed to perform them are still young. This dissertation is concerned with pushing forward this exciting new field.

I begin by describing some of the new tools and techniques that are needed to perform these studies. Chapter 2 describes `miriad-python`, a software framework used in the rest of my research to rapidly develop the tools needed to analyze the kinds of data that come out of modern telescopes such as the ATA. Chapter 3 describes the commensal observing system deployed at the ATA to allow the observatory resources to be shared between multiple cooperative observing projects. In Chapter 4, I explore the ATA’s continuous broadband

frequency coverage, using it to test the “calorimeter” interpretation of cosmic ray energetics in several well-known star-forming galaxies.

In Chapter 5, I describe how this infrastructure was used to execute a large survey for slow Galactic radio transients known as ASGAR, the ATA Survey of Galactic Radio Dynamism. The ASGAR dataset probes a hitherto unexplored parameter space for Galactic radio transients and has applications to other astrophysical questions as well. In Chapter 6 I examine the jet physics of microquasar Cyg X-3, a fascinating object that displays massive flares across the electromagnetic spectrum, using data from ASGAR as well as from X-ray and γ -ray satellites. Finally, in Chapter 7 I summarize and conclude. Appendix A provides a more detailed technical discussion of the challenges of searching for variability and transient events in wide-field, confused interferometric images.

Chapter 2

Rapid Interferometric Software Development Using MIRIAD and Python

State-of-the-art radio interferometers are complex systems that unleash torrents of data. If current and planned instruments are to routinely meet their performance goals, standard analysis techniques must be significantly improved, becoming simultaneously more sophisticated, more automatic, and more scalable. While there is no shortage of ideas for next-generation algorithms, there is a shortage of development resources, so it is vital that programming environments for interferometric software allow for rapid, flexible development. We present an open-source software package, `miriad-python`, that provides access to the MIRIAD interferometric reduction system in the Python programming language. The modular design of MIRIAD and the high productivity and accessibility of Python provide an excellent foundation for rapid development of interferometric software. Several other projects with similar goals exist and we describe them and compare `miriad-python` to them in detail. Along with an overview of the package design, we present sample code and applications, including the detection of millisecond astrophysical transients, determination and application of nonstandard calibration parameters, interactive data visualization, and a reduction pipeline using a directed acyclic graph dependency model analogous to that of the traditional Unix tool `make`. The key aspects of the `miriad-python` software project are documented. We find that `miriad-python` provides an extremely effective environment for prototyping new interferometric software, though certain existing packages provide far more infrastructure for some applications. While equivalent software written in compiled languages can be much faster than Python, there are many situations in which execution time is profitably exchanged for speed of development, code readability, accessibility to nonexpert programmers, quick interlinking with foreign software packages, and other virtues of the Python language.¹

¹This chapter has been previously published as Williams *et al.*, [2012 PASP 124 624](#), and is reproduced with the permission of all coauthors and the copyright holder. Copyright 2012 Astronomical Society of the

2.1 Introduction

Advances in the fields of digital computing and commercial telecommunications have fueled an explosion of innovation in radio interferometry. New and upgraded facilities such as the Allen Telescope Array (ATA; [Welch et al. 2009](#)), the Low-Frequency Array ([Kassim et al. 2004](#)), the Precision Array for Probing the Epoch of Reionization (PAPER; [Parsons et al. 2010](#)), the Murchison Wide-Field Array (MWA; [Lonsdale et al. 2009](#)), the Karl G. Jansky Very Large Array (formerly EVLA; [Perley et al. 2011](#)), the Westerbork Synthesis Radio Telescope (WSRT) Apertif project ([Verheijen et al. 2008](#)), the Australian Square Kilometer Array Pathfinder ([DeBoer et al. 2009](#)), and MeerKAT ([Jonas 2009](#)) have sophisticated designs including large-number-of-small-dishes architectures, multipixel or phased-array feeds, wide bandwidths, and phased array substations. Several of these facilities are pathfinders for the proposed Square Kilometer Array (SKA; [Carilli & Rawlings 2004](#)). They all aim to push the limits of interferometric techniques to image large fields of view with unprecedented spatial, spectral, temporal, and polarimetric fidelity, at unprecedented data rates.

Achieving these goals will require substantial amounts of new software to process the data coming out of these facilities. The data rates present a challenge in and of themselves: SKA-class facilities will require exaflop-scale computing ([Cornwell & Humphreys 2010](#)). Another challenge is the development of the necessary new algorithms, which is certain to require extensive experimentation. Techniques already under investigation include w -projection ([Cornwell et al. 2008](#)), A -projection ([Bhatnagar et al. 2008](#)), multi-scale multi-frequency CLEAN ([Rau & Cornwell 2011](#)), delay/delay-rate filtering ([Parsons & Backer 2009](#)), space-alternating generalized expectation (SAGE) maximizing calibration ([Kazemi et al. 2011](#)), generalized measurement-equation-based instrumental modeling ([Noordam & Smirnov 2010](#)), scale-invariant rank detection of radiofrequency interference (RFI; [Offringa et al. 2012](#)), subspace-tracking RFI mitigation ([Ellingson & Hampson 2002](#)), visibility stacking ([Hancock et al. 2011](#)), bispectral pulse detection ([Law & Bower 2012](#)) and improved source finding tools ([Whiting 2012](#)), to name a few. Until the next generation of algorithms is thoroughly explored, the quality of many results will be set not by the capabilities of observatory hardware but by the sophistication of the reduction pipeline that can be brought to bear before publication: a condition of *software-limited* science.

Although improvements in software development efficiency are always desirable, they're particularly salient now as the next generation of radio interferometers comes online. We discuss the efficiency of a programming environment in terms of *productivity*, which we take to measure the amount of useful functionality that a programmer can implement per unit time, and *accessibility*, which measures how much effort it takes for non-experts to begin successfully working within the environment without hand-holding. Although many studies attempt to quantify these metrics (e.g., [Petersen 2011](#), and references therein), our discussion will remain qualitative. These attributes are relevant in situations in which developer resources are constrained: a more productive environment allows individual developers to accomplish more,

Pacific.

while a more accessible environment has a lower barrier to entry from informal contributors.

Interpreted, dynamic programming languages can provide such an environment. While several of these exist, the language Python² in particular has seen a broad uptake in the astronomical community over the past decade (*e.g.* Barrett & Bridgman 1999; Greenfield & White 2000; Blakeslee et al. 2003; Kettenis et al. 2006; Magee et al. 2007). It is intended to be easy to learn and offers conveniences including object-oriented programming, lambda expressions, exception handling, and an enormous software ecosystem, including interactive interpreters (*e.g.*, IPython; Pérez & Granger 2007), visualization tools (*e.g.*, matplotlib; Hunter 2007), file format interfaces (*e.g.*, pyFITS; Barrett & Bridgman 1999), database interfaces (*e.g.*, SQLite³), statistics routines, web service support, and so on. The existence of this ecosystem and the rapid rise in the popularity of Python in the astronomical community in particular speak to its productivity and accessibility.

The programming environments of “classic” astronomical software packages tend to be intimidating and frustrating in comparison. Two of the major traditional radio interferometric reduction packages, MIRIAD (Multichannel Image Reconstruction, Analysis, and Display; Sault et al. 1995) and AIPS (Astronomical Image Processing System; Greisen 2002), are largely implemented in FORTRAN-77, and in both cases the process to go from source code to usable installed executable code is complicated and fragile. In light of the clear need for improved interferometric algorithms, it should be no surprise that a variety of projects have sought to build on or replace these systems with more modern ones. Perhaps the most prominent such undertaking is the Common Astronomy Software Applications (CASA; McMullin et al. 2007), the successor to AIPS. Other packages include Obit (Cotton 2008b), AIPY (Astronomical Interferometry in Python⁴), MeqTrees (Noordam & Smirnov 2010), and several more narrowly-targeted binding layers (*e.g.*, pyramid; Mehringer & Plante 2004).

To this list, we add `miriad-python`, a creatively-named package exposing MIRIAD tasks and subroutines in Python. Its design and philosophy are discussed (§2.2) and compared those of related projects (§2.3), including the ones mentioned above. We then describe the implementation of `miriad-python` (§2.4) and provide some very brief examples of its use (§2.5). Some applications in which it has been used are presented, including the detection of millisecond astrophysical transients, determination and application of nonstandard calibration parameters, interactive data visualization, and pipeline processing (§2.6). We document the nature of `miriad-python` as a software project (§2.7), discuss some of its performance characteristics (§2.8), and finally summarize (§2.9).

²<http://python.org/>

³<http://sqlite.org/>

⁴<http://purl.org/net/pkgwpub/aipy>

2.2 Design Considerations in `miriad-python`

The `miriad-python` project aims to encourage the rapid development of interferometric software by providing convenient access to MIRIAD tasks, datasets, and subroutines in the Python language. An important aspect of the `miriad-python` design is that it provides unified interfaces for accessing both high- and low-level MIRIAD functionality. While there is traditionally a rigid dichotomy between “users” (who stereotypically run shell scripts to drive existing tasks) and “gurus” (who compile FORTRAN-77 programs to provide new tasks), we intend `miriad-python` to permit the population of a broad middle ground of empowered user-developers, from those who need just a bit more sophistication than the Unix shell can provide to those who are implementing substantial new algorithms.

Although `miriad-python` aims to ease the rapid development of new interferometric algorithms, it does not provide any nontrivial algorithms itself. Instead it focuses on providing the best possible interfaces for access to MIRIAD infrastructure and is as flexible as possible regarding what is done with it. To use terminology dating back to at least the design of the X Windows system, it provides *mechanism* but not *policy* (Scheifler & Gettys 1986). We believe that the wide range of applications described in §2.6 is evidence of the power of this approach.

Work on `miriad-python` was initially inspired by a practical desire for more efficient scripting of MIRIAD reductions. The MIRIAD package, however, is a good fit to the larger goals of the `miriad-python` project. MIRIAD offers its wide variety of tools and algorithms in a modular task architecture, and its data formats are simple and efficient. In particular, it is fairly straightforward to (ab)use the streaming MIRIAD *u-v* data format for novel applications (see, e.g., AIPY, §2.3.2). Other applications are built upon MIRIAD for similar reasons (e.g. Teuben 2011; Pound & Teuben 2012). MIRIAD is routinely used to process data from facilities including the ATA, the Berkeley-Illinois-Maryland Association array (BIMA; Welch et al. 1996), the Combined Array for Research in Millimeter/submillimeter Astronomy (CARMA; Woody et al. 2004), the Australia Telescope Compact Array (ATCA), and the Submillimeter Array (SMA; Ho et al. 2004). There are multiple forks of MIRIAD: at least two nontrivially divergent MIRIAD codebases are maintained for use with the ATCA and CARMA, with some sharing of modifications between them. `miriad-python` is referenced to the CARMA MIRIAD codebase. Recent work on MIRIAD includes complete support for 64-bit file offsets and pointers, allowing imaging of arbitrarily large datasets; integration of the `wcslib` library (Calabretta 2011) for more comprehensive support of coordinate manipulations⁵; and of course bug fixes, new features, and documentation improvements.

When creating a package such as `miriad-python` that interfaces with a lower-level one, one must decide how closely to hew to the APIs (application programming interfaces) of the original package. In `miriad-python`, the APIs have been significantly reworked: they are object-oriented and aim to take full advantage of builtin language features. When they meet

⁵While `wcslib` is targeted at FITS images, the MIRIAD image format stores coordinate information in a virtually identical manner to FITS.

the goal of providing mechanism rather than policy, substantial new features are implemented in the Python layer. One example is approximate cryptographic hashing of u - v datasets, which can be used to quickly and fairly robustly check for modifications to the data regardless of file modification times. (This is useful in dependency-tracking pipelines, cf. §2.6.4.) To write useful programs, `miriad-python` authors must be familiar with the semantics of MIRIAD datasets or tasks, and good knowledge of Python is helpful, but they do not need to understand the design of the MIRIAD subroutine library.

Developers of packages that build on lower-level software must also decide whether to link to a separate installation of the lower-level package, or to compile and install their own version of it, bundling a copy of its source code with their own. We have chosen to go the former route with `miriad-python`. Although `miriad-python` offers very different APIs and new features, we conceive of it as fundamentally a layer above MIRIAD and not a standalone package: if the user has customized their MIRIAD installation in some way, for instance, it's appropriate for `miriad-python` to reflect that customization, and not override it with its own copy of the subroutine library. Another reason is that MIRIAD is still evolving, and tracking changes from an upstream codebase to a forked copy can be tedious and error-prone. (This has been the authors' experience with the existing forks of MIRIAD.) Another is that `miriad-python` is a general-purpose interface to not just the MIRIAD subroutine library but its task collection, so the entire, substantial, MIRIAD codebase would need to be bundled. Finally, we expect that `miriad-python` users will be existing MIRIAD users who are likely to already have installed MIRIAD on their systems, so the convenience of bundled source code is lessened.

To touch on a broader philosophical issue, it may be argued that accessibility is not a desirable feature when it comes to interferometric software. Interferometric analysis tends to be subtle, and many scientists are inexperienced programmers. On the other hand, scientific research is a fundamentally creative and exploratory process in which new or improved techniques are often required, and this is clearly the case in the domain of interferometry at present. The tension here is a variation on the well-explored “cathedral versus bazaar” theme first described by [Raymond \(1999\)](#). Considering the amount of algorithmic development necessary to fully exploit next-generation radio interferometers, we take the stance that the empowerment of interferometric software users is a good thing.

2.3 Comparisons to Related Packages

As alluded to in §2.1, there are several other efforts underway to provide developer-friendly environments for producing interferometric software. In this section we attempt to situate `miriad-python` among them. We consider “bindings”, which do not provide substantial new algorithms, and “hybrid” packages, which do.

2.3.1 Bindings

Several other packages expose MIRIAD functionality in the Python language, but we found them unsatisfactory. The `pyramid` package (Mehring & Plante 2004), providing a module named `Miriad`, includes similar functionality to that of `mirexec` in a somewhat less structured manner. The package `mirpy`⁶ also provides similar features, with an aim towards data processing inside interactive Python sessions. The modules `mirlib` (via the WSRT) and `miriadwrap` both bind the MIRIAD subroutine library. Unlike `miriad-python`, they provide a fairly direct mapping to the MIRIAD subroutines and do not extensively “Pythonify” the API. These packages are less mature than `miriad-python` and do not appear to be under active development. None of them includes substantial documentation, high-level convenience features, or an integrated system for invoking both MIRIAD tasks and subroutines.

There also exist several other packages that appear to share the same motivations as `miriad-python` but involve different technologies. Other dynamic languages can be used. For instance, MIRIAD functionality can be accessed in Ruby⁷ with the `MIRIAD-Ruby`⁸ or `mirdl`⁹ packages. The two packages, written by the same author, cover the link-versus-bundle tradeoff discussed above: `MIRIAD-Ruby` includes a portion of the MIRIAD source code, while `mirdl` links to a separate installation of the full MIRIAD libraries. Although it is beyond the scope of this chapter to compare the merits of the languages, we note that Ruby has seen less uptake in the astronomical community than Python but is starting to make inroads (*e.g.* Lammers et al. 2002; Fuentes et al. 2007; Gutierrez-Kraybill et al. 2010).

Other packages build on top of different interferometry frameworks. `ParselTongue` (Kettenis et al. 2006) provides a Python interface to AIPS, building on top of parts of the `Obit` package (Cotton 2008b), described in the next subsection. The `pyrap` project¹⁰ provides a Python interface to the core C++ support libraries of CASA.

In many cases, the choice of the appropriate binding will be constrained by external limitations in either the high-level language or the underlying interferometry package. Our reasons for preferring Python and MIRIAD, respectively, are sketched in §2.1 and §2.2. We emphasize the simplicity and efficiency of MIRIAD’s data formats as being important factors in facilitating algorithmic experimentation. Another distinction between `miriad-python` and alternative bindings is its rate of development. Even in otherwise stable and mature packages, documentation improvements and subtle bugfixes tend to require fairly frequent modifications. In the calendar year of 2011, there were 5 commits to the `mirdl` version control system (VCS), 10 commits to the `pyrap` VCS, and 110 commits to that of `miriad-python`. The last known update to `ParselTongue` was in 2010 and the last known update to `pyramid` was in 2008. Finally, we believe that `miriad-python` acquits itself well when it comes to the subjective factors of API cleanliness, code readability, and code quality.

⁶<http://pypi.python.org/pypi/mirpy/>

⁷<http://www.ruby-lang.org/>

⁸<http://purl.org/net/pkgwpub/miriad-ruby>

⁹<http://purl.org/net/pkgwpub/mirdl>

¹⁰<http://purl.org/net/pkgwpub/pyrap>

2.3.2 Hybrid Packages

The projects described in this subsection are more ambitious than `miriad-python` and include significant functionality along with Python-based environments: they provide policy, as well as mechanism (cf. §2.2). While the packages described in the previous subsection are intended to provide suitable bases for development for a wide range of applications, the ones described below may not be appropriate for certain applications, depending on both technical and architectural factors. All of these packages are actively developed.

CASA (McMullin et al. 2007) is a substantial package including high-level Python interfaces to lower-level routines implemented in C++ and FORTRAN. The first iteration of CASA, known as AIPS++ (McMullin et al. 2004), included flexibility, ease of programming, and modifiability in its mission statement (AIPS++ Steering Committee 1993). We find that these are difficult to achieve in practice because most CASA functionality is implemented in fairly large C++ modules, with the built-in Python interfaces providing mostly coarse-grained access to data and algorithms; the combination of the core CASA libraries with `pyrap` is more friendly to Python development. CASA bundles its own Python interpreter, rather than linking to one provided by the operating system, which avoids several thorny installation issues but can make it more complicated for Python developers to integrate CASA and other site-specific packages (cf. §2.2). CASA’s imager is quite flexible and can combine multiple advanced algorithms, for instance simultaneous multi-scale, multi-frequency deconvolution (Rau & Cornwell 2011).

Obit is billed as “a development environment for astronomical algorithms” (Cotton 2008b). Although many of its fundamental features are general-purpose, Obit’s main emphasis is radio astronomy, and in particular it includes several algorithms important to low-frequency radio interferometry. The package includes a set of base libraries implemented in C, a set of tasks built on top of them, and a Python layer for interacting with both of these. This layer is based on ParselTongue, although the two projects remain separate (W. D. Cotton, priv. comm.). Two aspects of the Obit design are of particular note. Firstly, Obit uses AIPS and FITS file formats for data storage, so it interoperates seamlessly with certain existing packages and developer effort need not be spent reimplementing existing tools. Secondly, the Obit libraries are architected to allow multithreaded execution, and this capability has been explored during Obit’s development. In one example problem, the time taken to generate a complex image was reduced by a factor of 6.5 when the number of CPU cores used for computations was increased from one to twelve (Cotton & Perley 2010).

The AIPY system represents a sharper break from tradition than those previously mentioned. It is a standalone package implementing many advanced synthesis and calibration techniques in Python to meet the needs of the low-frequency dipole array PAPER. Its design emphasizes flexibility and is Python-centric: the primary interfaces are the Python ones, and code from preexisting packages such as HEALPix (Górski et al. 2005) and XEphem (Downey 2011) is liberally borrowed to provide fast, well-tested implementations of many features. AIPY can use the MIRIAD format for u - v data storage, and hence includes an internal

Python binding of the relevant portions of the MIRIAD subroutine library. A particularly interesting strength of AIPY is its infrastructure for creating maps of the whole sky at once.

The MeqTrees project (Noordam & Smirnov 2010) also uses Python to provide a productive environment for exploring new interferometric algorithms using the measurement equation formalism (Hamaker et al. 1996; Smirnov 2011a). Its input/output (I/O) and astronomical infrastructure are built on the core CASA libraries and `pyrap`, with certain numerical routines implemented in C++. MeqTrees provides a substantial infrastructure for developing and applying sophisticated calibration methodologies with numerous visualization features. As an example of its power, MeqTrees has been used to generate a noise-limited radio astronomical image of the source 3C 147 with a dynamic range of 1.6×10^6 (Noordam & Smirnov 2010).

Of these packages, only AIPY uses the MIRIAD u - v data format. (Packages based on the CASA core libraries interoperate with MIRIAD images, which is a useful feature. We do not believe that similar support for MIRIAD u - v data could feasibly be implemented, because the streaming MIRIAD format and table-based CASA and AIPS formats require different online data processing models, even though the formats can be interconverted offline.) AIPY, meanwhile, is aimed at a very specific problem space and does not aspire to be a general-purpose reduction package; there are no plans to increase its MIRIAD integration beyond its use of the MIRIAD format for u - v data storage. As such, the existing hybrid packages are not well-matched to the needs of many applications that intend to take advantage of the MIRIAD u - v data format.

2.4 Implementation

`miriad-python` provides three Python modules: `mirtask`, `mirxec`, and `miriad`. The first of these uses C and FORTRAN-77 based extension modules to interface with the MIRIAD subroutine library, providing most of the low-level functionality that makes `miriad-python` useful. Numerical functionality is provided by NumPy¹¹ and linking with FORTRAN-77 is accomplished using `f2py` (Peterson 2009). The second module, `mirxec`, is Python-only and provides a uniform framework for executing MIRIAD tasks inside the Python language. The third, `miriad`, is Python-only and provides a simple abstraction for referencing and manipulating MIRIAD datasets, with hooks into `mirtask` and `mirxec` to provide convenient access to certain common operations. The three modules are loosely-coupled such that they do not actually depend on each other, although most nontrivial tasks end up using functionality from all three modules. The relationships of the APIs are diagrammed in Figure 2.1.

As mentioned above, the MIRIAD APIs are significantly reworked and expanded. In most cases, Python’s runtime type information is used to automatically choose data types when performing lowlevel I/O. Because these types are important for task interoperability, they must be specified explicitly when writing new data. MIRIAD’s internal errors are mapped into Python exceptions using a mechanism based on the C library functions `setjmp` and

¹¹<http://numpy.scipy.org/>

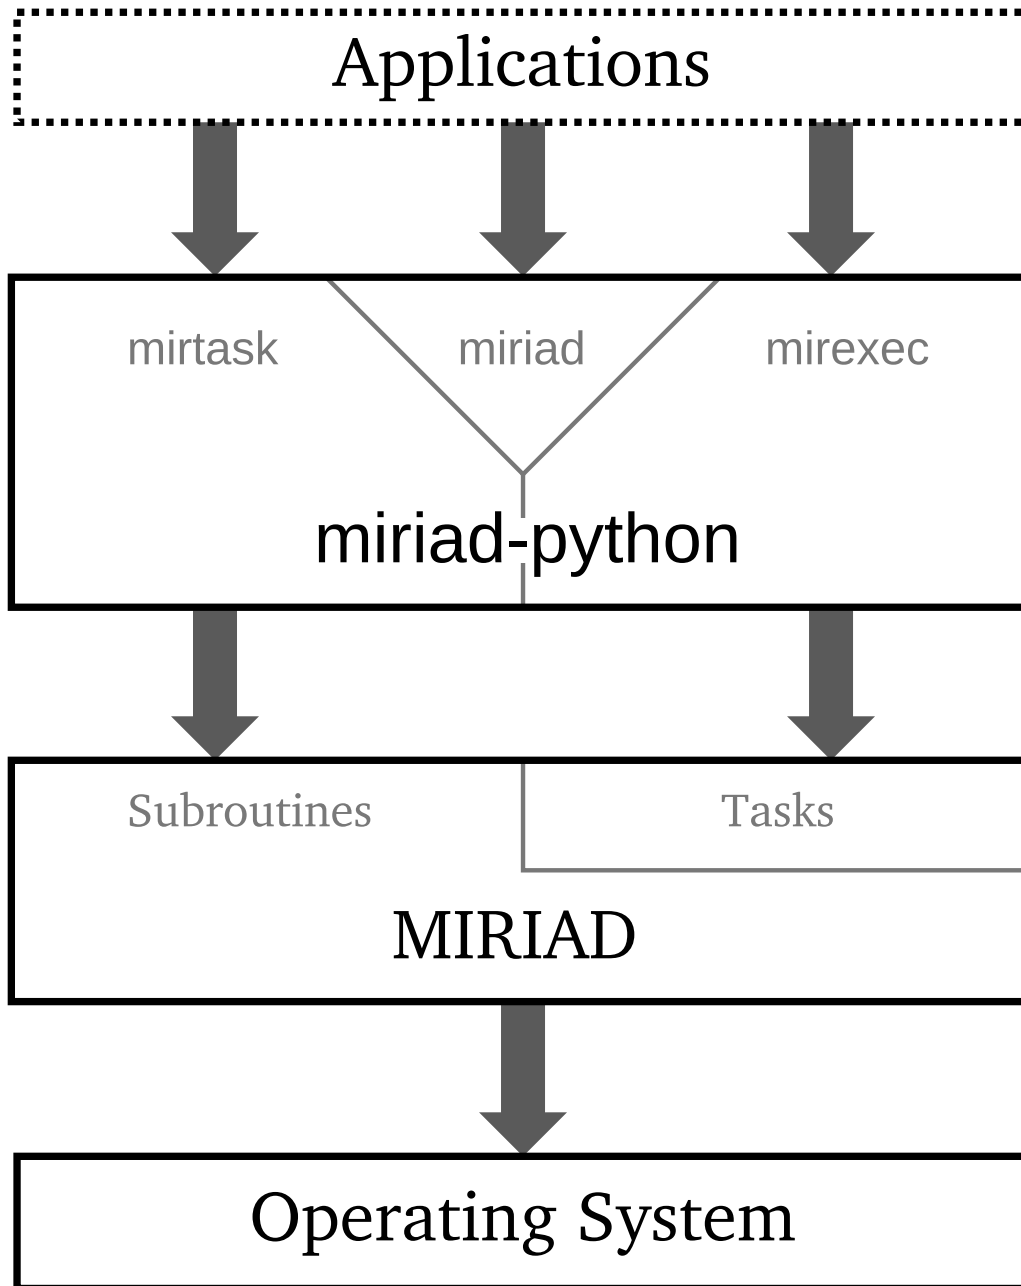


Figure 2.1: Diagram of the relationships between different components of the `miriad-python` system. Applications use the three `miriad-python` modules, which in turn access the tasks and subroutines in MIRIAD, which interact with the base operating system. Both applications and `miriad-python` use the NumPy module (not depicted) for numerical array operations.

longjmp.

`miriad-python` also provides a program that can serve as a help viewer for either existing MIRIAD tasks or programs written in Python, in the latter case using standard Python “documentation strings” written in the standard MIRIAD documentation format. Tasks written in Python can thus seamlessly merge with traditional compiled MIRIAD tasks from the user standpoint.

2.5 Example Code

Although it is impractical to provide extensive code examples in this chapter, in this section we give a few brief ones. These are necessarily simplistic and do not demonstrate the most interesting capabilities of `miriad-python`. As discussed below (§2.7.4), the `miriad-python` source distribution contains longer samples¹². We assume a familiarity with typical MIRIAD usage and the Python language. Many valuable coding practices are ignored below for the sake of concision.

MIRIAD tasks can be executed with the `mirexec` module. The following example demonstrates the parallel execution of multiple tasks: the `launch` method starts a task process, returning a handle to it (a subclass of Python’s `subprocess.Popen`), and the handle’s `wait` method pauses the caller until the process completes. (For I/O-intensive operations as shown here, it may actually be faster to execute the tasks serially to avoid disk thrashing.)

```
import sys
from mirexec import TaskUVAver

prochandle = [ ]

for p in sys.argv[1:]:
    taskdesc = TaskUVAver (vis=p, out=p+".avg", interval=10)
    prochandle = taskdesc.launch ()
    prochandle.append (prochandle)

for prochandle in prochandle:
    prochandle.wait ()
```

The above example is intentionally verbose. To stimulate the reader’s imagination, we provide the following variation that uses Python’s “list comprehension” syntax and avoids the use of several local variables:

¹²Currently browseable online at <https://github.com/pkgw/miriad-python/tree/master/examples>.

```

import sys
from mirexec import TaskUVAver

for ph in [TaskUVAver (vis=p, out=p+".avg", interval=10).launch () for p in sys.argv[1:]]:
    ph.wait ()

```

The following example demonstrates how the “header variables” of MIRIAD datasets may be accessed using `miriad-python`. The variable `vispath` is some string giving the filesystem path of a MIRIAD *u-v* dataset. As mentioned in §2.4, storage types of variables are detected automatically on read, but must be specified explicitly on write to enforce consistency with preexisting MIRIAD tasks.

```

import miriad, numpy as np

vispath = ...
ref = miriad.VisData (vispath)
handle = ref.open ("rw")
ncorr = handle.getScalarItem ("ncorr")
handle.setScalarItem ("mine", np.double, np.pi * ncorr)

```

The following example replaces the contents of an image with the magnitudes of its fast Fourier transform (FFT). This can be used, for instance, to recover the weights used in the imaging process from the dirty beam image. The variable `impath` is analogous to `vispath`. The value of the variable `whichplane` indicates that the first image plane should be selected.

```

import miriad
from numpy import abs, float32
from numpy.fft import fft2, fftshift, ifftshift

impath = ...
whichplane = [ ]
handle = miriad.ImData (impath).open ("rw")
d = handle.readPlane (whichplane).squeeze ()
d = abs (ifftshift (fft2 (fftshift (d))))
handle.writePlane (d.astype (float32), whichplane)
handle.close ()

```

Finally, this example inserts information about the synthesized beam shape of an image into a preexisting relational database.

```
import miriad, sqlite3

impath = ...
handle = miriad.lmData (impath).open ("rw")
dbpath = ...
db = sqlite3.connect (dbpath)

info = [impath]
for item in "bmaj bmin bpa".split ():
    info.append (handle.getScalarItem (item))

db.execute ("INSERT INTO data VALUES (?, ?, ?, ?)", info)
db.commit ()
db.close ()
```

Such functionality would be difficult to achieve using FORTRAN. We emphasize that in the interests of simplicity and clarity, we have avoided some of the most important features of the Python language such as object orientation, first-class functions, and exception handling.

2.6 Applications

The fast Python development cycle makes it relatively easy to produce new algorithms and tools for astronomical data analysis. In this section we present some example applications of `miriad-python`.

2.6.1 Algorithms for Millisecond Transients

We have investigated with new algorithms to study millisecond transients in the visibility domain using data from the ATA and VLA (Law et al. 2011b; Law & Bower 2012). Since interferometers have not traditionally operated at this time scale, many interferometric packages lack features needed to work in this regime, e.g., the ability to dedisperse visibilities and high-resolution timestamps. We used `miriad-python` to develop new ways to visualize and search these unusual data streams for transient radio sources.

Figure 2.2 demonstrates imaging of millisecond pulses from the Crab pulsar (B0531+21) using a Python-based toolchain. We observed the Crab pulsar at a time resolution of 1.2 ms and frequencies between 720 to 800 MHz using the ATA (Law et al. 2011b). At this time resolution, the dispersion introduced by the interstellar medium (ISM) delays the pulse arrival time by a few tens of milliseconds across the band. The upper panel of Figure 2.2 shows this effect in a spectrogram formed by summing visibilities, effectively forming a synthesized beam toward the Crab pulsar. The spectrogram shows that the dispersive delay follows a quadratic shape that has a slope consistent with that expected from the Crab pulsar. After identifying the time and dispersion of the pulse, we can image it to see if it is consistent with a point

source in the location of the Crab pulsar. The u - v input/output facilities of `miriad-python` are used to write a new dedispersed visibility dataset. The lower panel of Figure 2.2 shows an image made from visibilities during the pulse. The dispersive delay is large enough that it must be corrected in order to detect the source.

Low-level access to visibility data in `miriad-python` also makes it possible to experiment with new algorithms. We’ve used this capability to develop a technique to detect millisecond transients based on an interferometric closure quantity called the bispectrum (Cornwell 1987; Law & Bower 2012). The bispectrum is formed by multiplying three visibilities from baselines that form a closed loop. This product is sensitive to transients anywhere in the field of view, which makes it powerful for surveys. However, most software packages only use the bispectrum for calibration, so none have low-level access to functions for prototyping algorithms. Figure 2.3 shows millisecond light curves toward pulsar B0329+54 made with `miriad-python`. With direct access to dedispersed visibilities, we were able to compare traditional beamforming with the bispectrum technique. In this case, we can detect pulses for four of the five rotations of the pulsar during this observation. Being able to compare the techniques on the same pulses allowed us to show that the bispectrum responds more strongly than beamforming, confirming an unusual theoretical property of the bispectrum (Law & Bower 2012).

2.6.2 Retroactive SEFD Calibration

The ATA lacks online measurement of the system temperature (T_{sys}) or system equivalent flux density (SEFD), information that is important for assessing data quality and thermal noise limits. We have written a task, `calnoise`, to assess per-antenna SEFDs after-the-fact using observations of a bandpass calibrator and the assumption that the variance across the channels of each spectral window is entirely thermal.

Each ATA antenna has two orthogonally linearly polarized feeds. The signal from each of these is identified as originating from a given antenna-polarization pairing (“antpol”). The SEFDs of the two feeds on a single antenna may differ appreciably, and all of the SEFDs may be time-variable. Rather than directly computing SEFDs, `calnoise` determines calibration coefficients relating SEFDs to the RMS value across the spectral window of each antpol’s uncalibrated autocorrelation amplitude, a value we refer to as the RARA (raw autocorrelation RMS amplitude). While the value of this coefficient is assumed to be static over the course of an observing session, the derived SEFD value will vary proportionally with the RARA. Although the RARA is a function of both system noise and a gain factor, the time variation in the latter is small enough so that the derived SEFDs will be sufficiently accurate.

`calnoise` takes as an input a dataset with a gains (and optionally bandpass) table. The dataset is scanned through once, without applying the gains table, to accumulate RARA values for each timestamp and antpol. The dataset is then reread, with the gains being

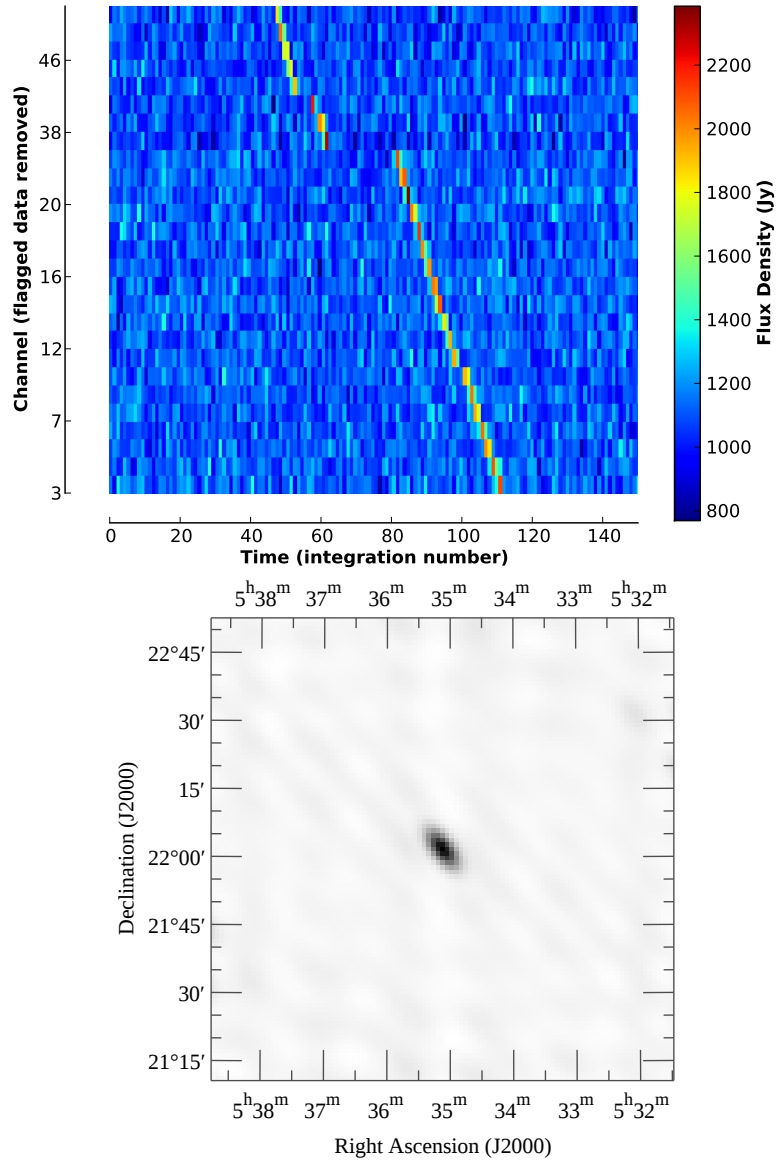


Figure 2.2: *Upper panel:* Spectrogram of a pulse from the Crab pulsar, summing visibilities coherently using the known location of the emitter. Dispersion due to the ISM causes pulse arrival times to vary with frequency (channel number). Apparent discontinuities in the pulse are due to flagged channels not being rendered (note the gaps in the left axis values). *Lower panel:* An image of the pulsar combining all visibilities in the observation, applying a correction for the dispersive time delay. The grayscale is linear from -42 Jy (white) to 1189 Jy (black). See §2.6.1 and [Law et al. \(2011b\)](#).

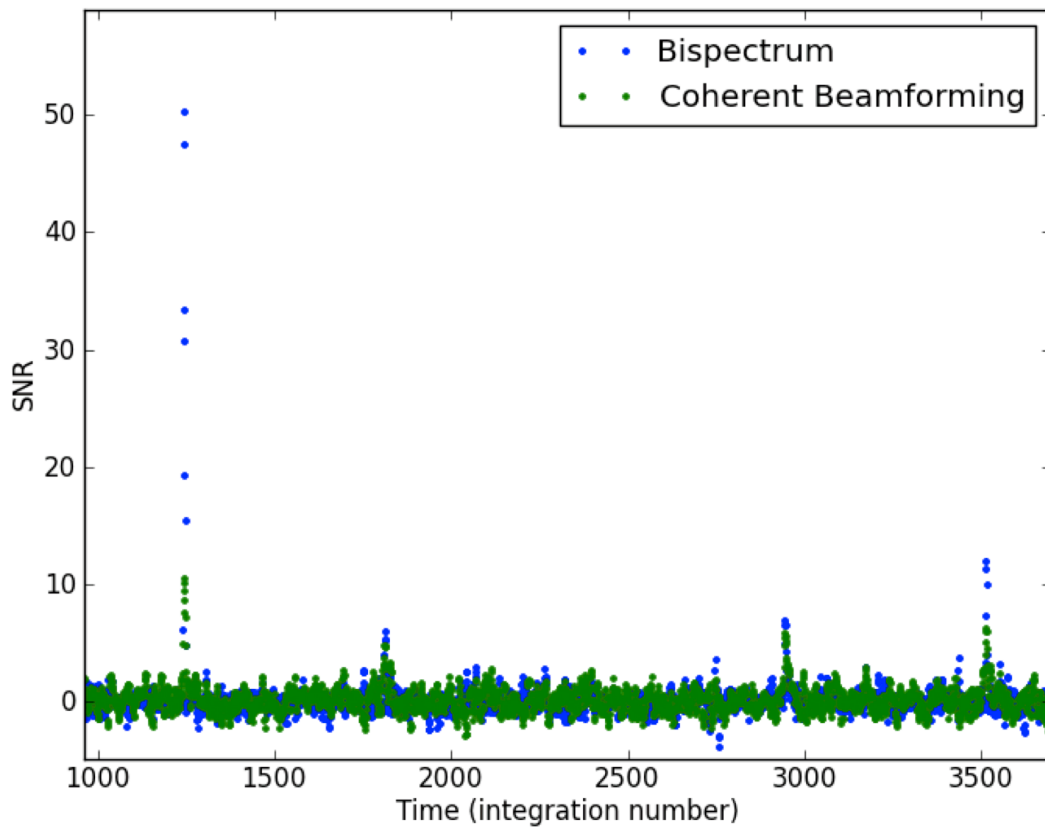


Figure 2.3: Comparison of two methods for detecting millisecond astrophysical pulses, one using the bispectrum and one using traditional beamforming, both implemented in Python. Each method is applied to a visibility data stream containing five pulses from the bright pulsar B0329+54 and the signal-to-noise ratio (SNR) of each detection is plotted. The bispectrum-based method can be more effective (higher SNR detections for the same data) than coherent beamforming and is sensitive to pulses coming from any direction. See §2.6.1 and [Law & Bower \(2012\)](#).

applied, to compute per-baseline SEFDs:

$$\text{SEFD} = \frac{1}{2} (\sigma_r + \sigma_i) \eta \sqrt{2\Delta\nu\tau},$$

where σ_r and σ_i are the standard deviation in the real and imaginary parts, respectively, of the calibrated visibilities across the spectral window (thus in Jansky units), η is a tabulated efficiency of the correlation (depending, e.g., on the number of bits used in digitization), $\Delta\nu$ is the channel width in the spectral window, and τ is the integration time of the record. Per-antenna calibration coefficients are then calculated using a least-squares fit of the baseline-based values assuming a simple geometric dependence:

$$\text{SEFD}_{i,j,t} = \sqrt{c_i \text{RARA}_{i,t} \cdot c_j \text{RARA}_{j,t}},$$

where the c_i are the desired calibration parameters. Outlier values (possibly due to RFI or hardware failures) are identified by user-specified clipping limits and iteratively removed from the fits. MIRIAD's on-the-fly calibration system does not support SEFD information, so the c_i are then recorded in a textual table on disk; a companion task reads the table along with an input dataset and creates a new dataset with the SEFD information inserted. (For technical reasons, a constant T_{sys} is assumed and varying values of the MIRIAD variable `jyperk` are inserted into the dataset, where $\text{SEFD} = T_{\text{sys}} \cdot \text{jyperk}$ and thus `jyperk` is related to the effective area of each receiving element.)

`calnoise` is written in Python and uses `miriad-python` to read the visibility data and for miscellaneous astronomical routines. It uses NumPy for its numerics and other Python libraries for optional plotting of diagnostic information, least-squares solving, storage of numerical metadata, and integration into the ATA pipeline described below (§2.6.4). The main implementation comprises ~600 statement lines of code (i.e., excluding whitespace and comments). The chief benefits of using Python and `miriad-python` were quick creation of the task skeleton (thanks mainly to high-level data structures), rapid turnaround during refinement (thanks mainly to the lack of a compilation step), and easy investigation of algorithmic tweaks (thanks mainly to easy access to other libraries, e.g., only one line of code needed to interactively plot variables).

2.6.3 Interactive u - v Data Visualization

Data visualization is an important part of the development of any processing pipeline. This is especially true for radio interferometers, in which the data undergo complex transformations, are high-dimensioned, and closed-loop instrumental modeling is often a key aspect of the reduction. We have used `miriad-python` to implement a powerful u - v data visualizer, first described in Williams (2010). As shown in Fig 2.4, the main display is a dynamic spectrum of the visibilities on a baseline as a function of frequency and time. The user can quickly switch between different displays (real, imaginary, amplitude, phase), apply various processing steps on-the-fly (e.g., average, rephase), and navigate through the dataset. There is also

substantial support for visualizing and creating data flags when necessary, although it is widely recognized that manual flagging of data is rapidly becoming an impossible task as data rates increase (Keating et al. 2010).

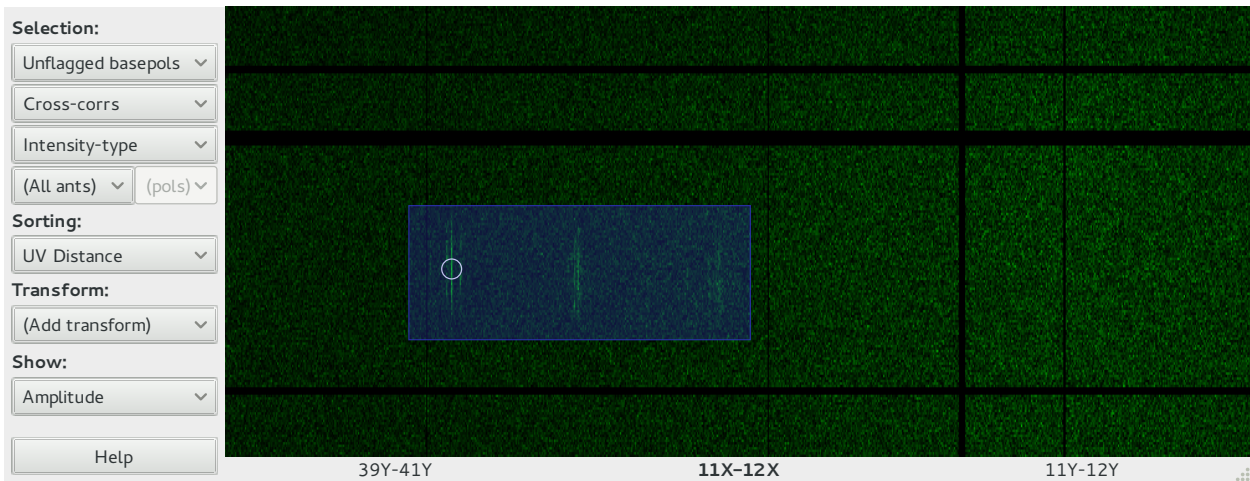


Figure 2.4: Sample screenshot of the u - v data visualizer described in §2.6.3. The main graphical display is a dynamic spectrum of amplitude on one baseline as a function of frequency (horizontal axis) and time (vertical axis). The left-hand panel provides controls for filtering and transforming the data. Python bindings to well-established graphical toolkits allow the data to be displayed in an attractive and responsive user interface.

The use of Python is a key aspect to the u - v visualizer because it allows easy combination of the existing MIRIAD libraries with very different software, in this case the modern graphical toolkits Cairo¹³ and GTK+¹⁴. Cairo provides routines for fast rendering of the gridded visibility data, while GTK+ provides a higher-level widget library that allows new user interactions to be implemented quickly. Achieving this functionality without the use of these (or similar) libraries would be a massive undertaking: Cairo and GTK+ comprise 240000 and 560000 lines, respectively, of well-tested, efficient C code. The main visualizer codebase, on the other hand, comprises only 2900 statement lines of Python.

2.6.4 An ATA Reduction Pipeline

Modern and next-generation radio observatories are expected to produce data at rates significantly higher than those of older facilities (*e.g.* Cornwell & Humphreys 2010). Under these conditions, rapid automated processing of data is changing from a luxury to a necessity (Keating et al. 2010). Python is often identified as a good system in which to construct pipelines to perform such processing, thanks to its amenability to implementing high-level

¹³<http://cairographics.org/>

¹⁴<http://gtk.org/>

application logic, revisability, and ability to glue together existing tools and libraries (Sanner 1999; Myers et al. 2007; Pérez et al. 2011). We have used `miriad-python` to implement a commensal observing system (Chapter 3) for the ATA Survey of Galactic Radio Dynamism (ASGARD; Chapter 5) and just such a pipeline for processing the data. In the latter case, the `mirexec` module is used extensively to invoke existing MIRIAD tasks and `mirtask` is used for low-level examination and manipulation of the data, often to accomplish tasks specific to ATA data reduction.

Of particular note is that the pipeline internally manages data flow through a directed acyclic graph model analogous to that of the traditional Unix tool `make` (Feldman 1979). While traditional shell scripting languages do not support the data structures necessary to conveniently implement such a design, doing so is easy in Python. The `drPACS` package (Teuben 2011) takes a similar approach but actually uses `make` for its underlying management. We find this approach to be unsatisfactory for two reasons. Firstly, `make` can only track dependencies between files on disk, and can only invoke tools via the Unix shell. This model maps very inefficiently onto certain aspects of typical data reduction workflows. Secondly, and more fundamentally, `make` checks whether steps need to be rerun by comparing file modification times rather than actual contents. This suffers from a variety of minor issues and the major issue that if a pipeline product is regenerated, all products downstream of it must also be regenerated, even if the rebuilt product did not actually change. In a data reduction pipeline the costs of this inefficiency can be severe. The tool described here avoids this problem by detecting modifications with cryptographic hashes as alluded to in §2.2.

2.7 `miriad-python` as a Software Project

In this section we describe `miriad-python` as a software project. We also hope that this section may provide future authors a useful reference of issues that, in our opinion, are important to discuss when documenting any other software project.

2.7.1 Intellectual Property Issues

`miriad-python` is an open source project, almost entirely licensed under the GNU General Public License, version 3 or later. Small portions of it relating to the compilation process are licensed under different permissive licenses. As such, `miriad-python` is available for inspection and use by anyone for any purpose. The source code copyright is owned by the authors. `miriad-python` is not known to be subject to any software patents.

2.7.2 Availability

The source code is available for free download from the project website¹⁵ either in the form of file archives or via the `git`¹⁶ version control system. The `git` repository is also available on the website GitHub¹⁷. Thanks to the distributed nature of `git`, clones of either repository stand alone and contain the complete, cryptographically-verified revision history of the entire project¹⁸. Most installations of `miriad-python` track the `git` repository and so official releases are rare. When made, they are documented and linked to on the project website. The most recent release is version 1.2.2, made on 2012 June 28¹⁹.

Installation instructions are beyond the scope of this chapter and are provided on the website. We note, however, that `miriad-python` must be compiled against MIRIAD libraries from the CARMA codebase built with the `autoconf`-based build system.

2.7.3 Development Model

Development of `miriad-python` occurs in the `git` repository. The authors aim to conduct development in an open, welcoming fashion. Contributions from the community are encouraged via (e.g.) a web-based issue tracker (accessible via the project website), email, or “pull requests” on the GitHub site. Based on the experiences of other community-based software projects, copyright assignment is not required for external contributions (see, for instance, the discussion in O’Mahony 2003).

The scale of the `miriad-python` project is such that other pieces of project coordination infrastructure such as mailing lists are not currently deemed necessary. These will arise as the `miriad-python` community grows, and will be linked to on the project website.

2.7.4 Documentation and Examples

A moderately-complete manual to `miriad-python` is available on the website, providing both an API reference and somewhat higher-level guidance as to the intended usage of the package. The primary documentation format is HTML served over the web, but the documentation is generated using Sphinx²⁰, the official documentation system of the Python language, and

¹⁵<http://purl.org/net/pkgwpub/miriad-python>. This link is a “permanent URL” providing a durable, reroutable link suitable for inclusion in the academic literature. The authors encourage the use of this URL and not its current, possibly-ephemeral destination when linking to the `miriad-python` website and related pages. See <http://purl.org/> for more information.

¹⁶<http://git-scm.com/>

¹⁷Currently <https://github.com/pkgw/miriad-python/>, but the canonical link may be found via the project website.

¹⁸The SHA1 checksum of the version of `miriad-python` described in this chapter is `c22527691e0bb608-fba7e41912a3d856c5fe6195`. The secure architecture of `git` ensures that a valid commit of this checksum is guaranteed to correspond to the exact source code described in this work.

¹⁹The corresponding `git` commit is `4b706bada8d47ae5a55ec948db34382add7b85ab`.

²⁰<http://sphinx.pocoo.org>

in principle the documentation can be generated in several other output formats, including printable PDF files.

Besides the brief examples in §2.5, the `miriad-python` source distribution contains a small set of longer samples implementing features as both standalone scripts and importable modules. Exercised features include reading and writing of u - v data, executing MIRIAD tasks with `mirtask`, manipulating dataset header items, reading gains tables, and integrating into the MIRIAD documentation system.

2.7.5 Quality Assurance

`miriad-python` is in line with the vast majority of scientific software in that, unfortunately, it has virtually no quality-assurance systems in place beyond the fact that it is exercised in day-to-day use. In particular, it has no automated suite of validation or regression tests. Code changes in diff format are reviewed before being committed to the repository, hopefully preventing unintended changes from being integrated into the source tree. An automatic test framework is planned, using short unit tests of the `miriad-python` APIs as well as more involved functional tests using sample MIRIAD datasets either generated on-the-fly (with `uvgen`) or optionally downloaded from the project website.

Given time constraints and the potential payoffs it is certainly not hard to understand why scientific software so often lacks systematic testing infrastructure. As a point of reference, testing can approach half the cost of a piece of software in the corporate context (Machado et al. 2010). We emphasize, however, that software defects are extremely common — about one for every *twenty* statement lines of code for undisciplined coding (Shull et al. 2002) — and that even rudimentary testing can capture many of these defects (*e.g.* Binder 1996). In particular, regression tests for previously-discovered bugs can be valuable and lend themselves to consistent implementation even in informally managed projects. Shull et al. (2002) present an excellent summary of the prevalence and impact of defects in typical software systems.

2.8 Performance Considerations

Python, being a dynamic interpreted language, is not optimal for achieving raw numerical throughput. Indeed, the *raison d'être* of a tool such as `miriad-python` is that in many cases numerical throughput can profitably be traded for other desirables such as code readability, programmer time savings, and ease of integration with existing codebases. In our experience, data-processing algorithms written in Python are almost always “fast enough,” even for fairly complex algorithms.

2.8.1 Serial Tests

To give a very rough quantification of the overhead of using Python compared to FORTRAN-77, we wrote two versions of a simple task that iterates through a MIRIAD

u - v dataset and computes the RMS value of the unflagged visibilities in each record. One version was written in Python and one in FORTRAN-77, the latter being compiled with GNU `gfortran` version 4.4.1. For a 0.3 GB visibility dataset containing 2,812,072 u - v records of 16 channels each, the Python version ran about 21 times slower than the FORTRAN-77 version (63 s vs 3 s) on our test system, a machine running a quad-core AMD Opteron 1385 CPU at 2.7 GHz. (Tests were run repeatedly with no effort to flush filesystem caches, so performance was CPU-limited.) For a 1.6 GB dataset containing 406,260 u - v records of 1,024 channels each, the Python version ran about 4 times slower (26 s vs 7 s). These results suggest that there is indeed a nonnegligible overhead to the Python interpreter, although its significance can be highly dependent on the structure of the input data. The performance of the Python version in the 1,024-channel case is relatively good because relatively more work is done inside a few vectorized NumPy function calls; manually iterating over each channel in Python increases the runtime by a factor of ~ 90 . For cases where most of the work is performed inside a few vectorized NumPy functions, the overhead to using Python can be fairly small.

In most software, the majority of the execution time is spent in only a small portion of the code. For those cases where serial throughput is a limitation, large improvements can often be achieved for low effort by identifying the most-executed codepaths and porting their particular implementations to a compiled language. This strategy is adopted by several of the projects discussed in §2.3. In some instances, it may be possible and profitable to generate speed-critical code on-the-fly using a compiler library such as the Low-Level Virtual Machine (Lattner & Adve 2004), avoiding the use of customized precompiled modules altogether. In other cases, it may make sense to prototype an algorithm in Python and then develop a production implementation in a compiled language. Of course, there will always exist problems for which Python is an inappropriate tool, but in our experience the combination of Python and a few basic numerical libraries is sufficient to address a wide range of challenges.

2.8.2 Parallelization

It is clear that parallelized processing will play a fundamental role in the reduction of future interferometric datasets. There is a special opportunity for next-generation interferometric processing packages to establish themselves in this space because the existing packages were generally built purely for serial processing.

There are several classes of parallelization that may be considered. For some applications, multiple independent reduction processes may be executed simultaneously on a single machine or on a cluster with no intra-process communication or synchronization. In this “embarrassingly parallel” case, a language such as Python may ease the development of both the reduction processes and the system to manage the overall processing job. Indeed, there are several preexisting Python frameworks for constructing parallel applications, e.g. Parallel Python²¹.

²¹<http://www.parallelpython.com/>

Other parallel applications can be implemented via multiple independent processes that communicate and synchronize. Such applications are often implemented with the standard Message Passing Interface (MPI; Walker 1994), which is accessible in Python via several different toolkits (e.g., `mpi4py`; Dalcin et al. 2008). Somewhat surprisingly, MPI-based parallel Python programs with core numerical routines written in a compiled language can perform as well as parallel programs written purely in a compiled language (Cai et al. 2005). The comparative ease of writing the driving logic of such a program in Python versus compiled languages makes this an intriguing model for parallel algorithmic development.

For applications that require more intensive data-sharing, one may wish to parallelize MIRIAD operations within a single process by using multithreading. Unfortunately, Python’s “global interpreter lock” prevents multiple threads from actually executing concurrently. (In many cases, naïvely multithreading a Python program makes it much slower!) In the particular case of `miriad-python`, extreme care would also be needed to synchronize access to the MIRIAD subroutine libraries, which are not threadsafe and maintain substantial shared state. Cotton (2008a) describes some of the challenges and successes encountered in tackling this kind of problem in *Obit*. The recommended system for threaded-style computation in Python is the `multiprocessing` module, which provides an infrastructure for launching concurrent Python subprocesses and communicating objects between parent and child. In the context of this work, `multiprocessing` is a more limited, Python-specific variant of the standard MPI approach. For certain I/O-limited applications, multithreaded Python code can obtain a performance gain above single-threaded code, but the difficulties of serializing access to the MIRIAD libraries remain.

Finally, certain classes of problems are well-suited to the highly-parallel capabilities of GPUs (graphics processing units). Current uses of GPUs in interferometric applications include pulse dedispersion (e.g. Magro et al. 2011; Straten & Bailes 2011), RFI mitigation (Ait-Allal et al. 2012), cross-correlation (Clark et al. 2012), and visualization (Hassan et al. 2011). Python code can take advantage of GPUs using frameworks such as PyCUDA (Klöckner et al. 2012). The nature of these examples suggest that the algorithms most suited to GPU processing would not likely need to build on the MIRIAD/`miriad-python` infrastructure, except perhaps to use MIRIAD data formats for I/O.

Regardless of the kind of parallelization in question, many interferometric algorithms are I/O-intensive. The options for improving the I/O bandwidth of MIRIAD/`miriad-python` are somewhat limited due to the lack of built-in support for parallel I/O (e.g. Thakur et al. 1999) in the MIRIAD I/O subroutines. With investment in the appropriate hardware, speedups can be achieved using transparently parallel filesystems (e.g., Lustre²² or PVFS; Carns et al. 2000). Significant gains in effective I/O bandwidth may alternatively be achieved with commodity hardware by spreading data between many independent processing nodes if the application permits partitioning of the data (e.g., by spectral channel or pointing direction).

²²<http://lustre.org/>

2.9 Summary

The new generation of interferometric hardware demands a new generation of interferometric software implementing a new generation of algorithms. Fortunately, there's a wide variety of projects aiming to provide the infrastructure needed to develop these tools as well as going ahead and implementing them.

`miriad-python` is one of these projects. It is “broad” rather than “tall”: it provides a wide range of APIs for accessing MIRIAD tasks and data, but does not provide its own algorithms built on top of those APIs. While certain applications will be best matched to the facilities provided by the higher-level packages `Obit`, `AIPY`, or `MeqTrees`, `miriad-python` is a good foundation on which to build those that are not, especially if one wishes to take advantage of the simple and efficient MIRIAD u - v data format. We actively encourage contributions to `miriad-python` development and are interested in supporting new `miriad-python` applications.

It seems likely that SKA-scale interferometric software will be only tangentially, if at all, related to any of the packages that exist today, and it is unclear if humans will “reduce” SKA data in any meaningful way given the rates involved. If the SKA is to meet both its goals and its schedule, however, the techniques needed to conquer exascale data rates must be discovered and prototyped, piece by piece, using today's software packages. Experience suggests that the innovations needed to solve these kinds of massive challenges do not come out of one organization; instead, wide-ranging and vibrant experimentation by a broad community of contributors will be essential.

Acknowledgments

I thank W. D. Cotton for useful discussions and the anonymous referee of the published form of this chapter for helpful comments. `miriad-python` is built upon the open-source Python language as well as an extensive foundation of other open-source and Free Software packages. It could not exist without the work of the innumerable people whose contributions have helped build this infrastructure.

Chapter 3

The ATA Commensal Observing System

This chapter describes the system used to conduct commensal correlator and beamformer observations at the ATA. This system was deployed for ~ 2 years until the ATA hibernation in 2011 and was responsible for collecting >5 TB of data during thousands of hours of observations. The general system design is presented (§3.2) and the implementation is discussed in detail (§3.3). I emphasize the rationale for various design decisions and attempt to document a few aspects of ATA operations that might not be obvious to non-insiders. I close with some recommendations (§3.4) from my experience developing the software infrastructure and managing the correlator observations. These include: reuse existing systems; solve, don't avoid, tensions between projects, and share infrastructure; plan to make standalone observations to complement the commensal ones; and be considerate of observatory staff when deploying new and unusual observing modes. The structure of the software codebase is documented (§3.3, Tables 3.3 and 3.4).¹

3.1 Introduction

One of the major design goals of the Allen Telescope Array (ATA; Welch et al. 2009) was the ability to share observatory resources to conduct multiple independent observing programs simultaneously — a feat generally referred to as commensal observing.² This chapter describes the system used to conduct the largest commensal observing campaigns before the ATA's hibernation in 2011. The primary campaign was a joint survey of regions in the Galactic plane: a traditional radio-astronomical survey looking for broadband Galactic radio transients (*à la* Hyman et al. 2005) and a SETI (Search for Extraterrestrial Intelligence) search for narrowband sources. In both cases, the Galactic population is clearly best probed by looking

¹This chapter is also available as Williams, 2012, Allen Telescope Array Memo Series #89 or [arxiv:1201.5413](#).

²In biology, commensalism is a symbiotic arrangement between mutualism and parasitism: it is the coexistence of two organisms in which one benefits and the other neither benefits nor suffers. It derives from the Latin *cum mensa*, “sharing a table,” originally alluding to the sharing of food scraps in particular.

in the plane rather than searching the sky uniformly. Various components of the traditional search have been referred to as the GCS (Galactic Center Survey), the AGCTS (ATA Galactic Center Transient Survey), AGILITE (ATA Galactic Transient and Lightcurve Experiment), gal90 (Galactic plane, $l = 90^\circ$ region), or the Kepler or Cygnus X-3 surveys. These are now all grouped together as ASGARD, the ATA Survey of Galactic Radio Dynamism (Chapter 5). Because of the low overhead to doing so, commensal correlator observations were also made during the SETI “exoplanets” and “waterhole” surveys, but there are currently no specific plans to use the correlator data from these undertakings.

The original vision that motivated the goal of substantial commensal observing on the ATA was one in which scheduling was based on traditional radio astronomy applications but SETI searches ran continuously in the background as well. The traditional observing program would drive the telescope, controlling the pointing and most hardware systems, and use the ATA’s correlator backends to take data. Meanwhile SETI observations would be performed using the ATA’s beamformer backends more-or-less passively, choosing any promising observation targets lying within whatever field of view (FOV) was chosen by the traditional observing program. It’s worth noting that this is not the only commensal observing scheme that might feasibly be implemented at the ATA. For instance, with a sufficiently large number of dishes and two relatively undemanding traditional radio astronomical observing projects, one could partition the array into two subarrays and run the two projects simultaneously, each using separate antennas and correlators. (One could argue that this is in fact not a commensal mode since so few resources would be shared between the two projects.)

The earliest commensal observations at the ATA were performed in August 2008 by Karto Keating and Tom Kilsdonk, but these and a few other efforts never became routine observing modes. As such, I refer to the system under discussion as “the ATA commensal observing system,” without further qualification. A design for an observatory-wide commensal observing system, and a deeper discussion of the observatory-wide software systems, is presented in [Gutierrez-Kraybill et al. \(2010\)](#).

3.2 Survey and System Design

Observations for the commensal campaigns are scheduled in blocks as per the standard system in use at the ATA. In contrast to the long-term vision of ATA commensal observations, it is SETI, not the traditional radio observing program, that’s “in the driver’s seat” for the observations: SETI software takes responsibility for all of the telescope hardware, most importantly the pointing. This arrangement came about because SETI already had a well-established observing system called Prelude, which had been adapted from its Arecibo roots to work at the ATA as well. Given this existing codebase, the project was approached with the plan of minimizing the amount of changes required to Prelude, while adding a separate “commensal observer” component that would take care of everything related to the commensal correlator campaigns.

With Prelude in charge of pointing the antennas and taking care of SETI’s data-taking, the responsibilities of the commensal component of the observing campaigns are *extremely* constrained: essentially, all it can do, and all it needs to do, is turn the correlators on and off. One could envision a much more complex system in which the commensal observer dynamically notifies Prelude of various needs (“please visit a phase calibrator”), but the system is vastly simplified if all such decision-making is centralized in Prelude. This simplification was made possible by pre-negotiating such decisions as pointing centers, dwell times, calibrator intervals, and focus settings.

In principle, the commensal observer could be completely ignorant of Prelude and its workings; by monitoring the current pointing directions of the ATA dishes, it could decide when a source was being observed and when the array was slewing, and in the former case it could do the necessary work to take data. It turns out, however, that “where is the ATA pointing?” and “are we tracking or slewing?” are questions that are more difficult to answer than one might think: obtaining ATA antenna telemetry requires fairly complex code (cf. the implementation of the `atastatus` command), and glitches in the data stream make it difficult to interpret the data robustly. Prelude knows exactly what its intentions are, however, so it was decided that the commensal observer would monitor a telemetry stream from Prelude to obtain pointing information. This stream is described more fully below (§3.3.3).

An important thread running through the design of the commensal system is automation: with thousands of hours of commensal observations scheduled, it’s desirable to execute them using as little human intervention as possible. The existing Prelude system fortunately dealt with the difficult task of choosing targets and planning a night’s observing.

3.3 System Implementation

The ATA commensal observing system is composed of a group of software tools and practices for using them. As mentioned above, some aspects of the commensal system were implemented in Prelude and its successor, SonATA (“SETI on the ATA”). This code is internal to SETI and is not discussed here.

Although the rest of the commensal software runs on a diverse set of hosts in a diverse range of situations (*e.g.*, at HCRO during observing; at UCB during data reduction), the whole codebase is maintained in a single Git (<http://git-scm.com/>) repository by the author. This chapter describes Git commit 88263be44c3d724e700e46b156e9b1dcfd0b1089, made on 2011 November 17. (Due to the nature of the Git version control system, this one value identifies an exact snapshot of the commensal source tree as well as its entire revision history.) A public version of the commensal repository is currently available at <http://astro.berkeley.edu/~pkwill/repos/commensal.git>. This URL will likely go stale as I am soon leaving Berkeley. A public copy of the repository may be established at his GitHub account, <https://github.com/pkgw/>, and at a minimum the repository will be available upon request. (Once again, thanks to the design of Git, each repository stands

alone and contains every version of every file as well as the complete revision history of the source tree.)

The vast majority of the source code is written in Python, with some tools written in shell script. The majority of the latter are Bourne shell scripts (most likely only compatible with `bash`), but a few are `tcsh` scripts, since the latter was the language used for most ATA observatory infrastructure. The repository also contains scheduling and analysis metadata. The Python scripts reference some support modules for controlling the ATA that are distributed in the `mmm`³ repository at <https://svn.hcero.org/mmm/pwilliams/pyata/>. A few secondary packages use NumPy, `miriad-python` (Chapter 2) and/or my plotting package, `Omegaplot` (<https://github.com/pkgw/omegaplot/>). The system design is shown in schematic form in Figure 3.1.

3.3.1 Scheduling

As mentioned above, commensal observations are allocated blocks in the ATA schedule as per standard practice.

At the time of the observations, the observatory-wide system for enacting this schedule was fairly *ad hoc*. The main HCRO computing environment is a distributed Unix-type system, with user logins and home directories shared over the network to a variety of workstations and servers in the signal processing room (SPR). One user account, `obs`, has a shared password and by convention is used to conduct the vast majority of observatory operations. SETI software systems, however, are segregated from this shared system, and generally run on separate machines with distinct logins. I’m not familiar with their configuration.

Correlator observations are scheduled by constructing a large shell script in the `obs` account, `obs.run`, that serially invokes other shell scripts to perform standard observing tasks. Before the 2011 hibernation, most observatory time was devoted to correlator observations so this could be considered the standard ATA scheduling system. For purely SETI observations, the `obs.run` script idles for a fixed amount of time while the array is controlled from SETI computers.

Because the `obs.run` system proved to be unreliable and the commensal campaigns involved many identical observations, the commensal observer was designed to be scheduled and launched separately in a more automatic fashion. Commensal observing blocks are entered into a separate schedule file in the repository (`sched/current.tab`) using the commands `plan-schedsess` or, later, `plan-importtxt`. The observing processes are launched via `cron` (see below).

Meanwhile, the SETI observations are launched via a `cron` system as well in the SETI part of the network. This means that there are in effect at least *four* relevant versions of the ATA schedule: the textual/graphical version distributed to users, that schedule as incarnated in `obs.run`, the schedule used by the commensal observer, and the schedule used by Prelude. The large PiGSS (Pi GHz Sky Survey; Bower et al. 2010, 2011) transient survey

³The name derives from our weekly data analysis meeting, “MIRIAD Masterclass with Melvyn” (Wright).

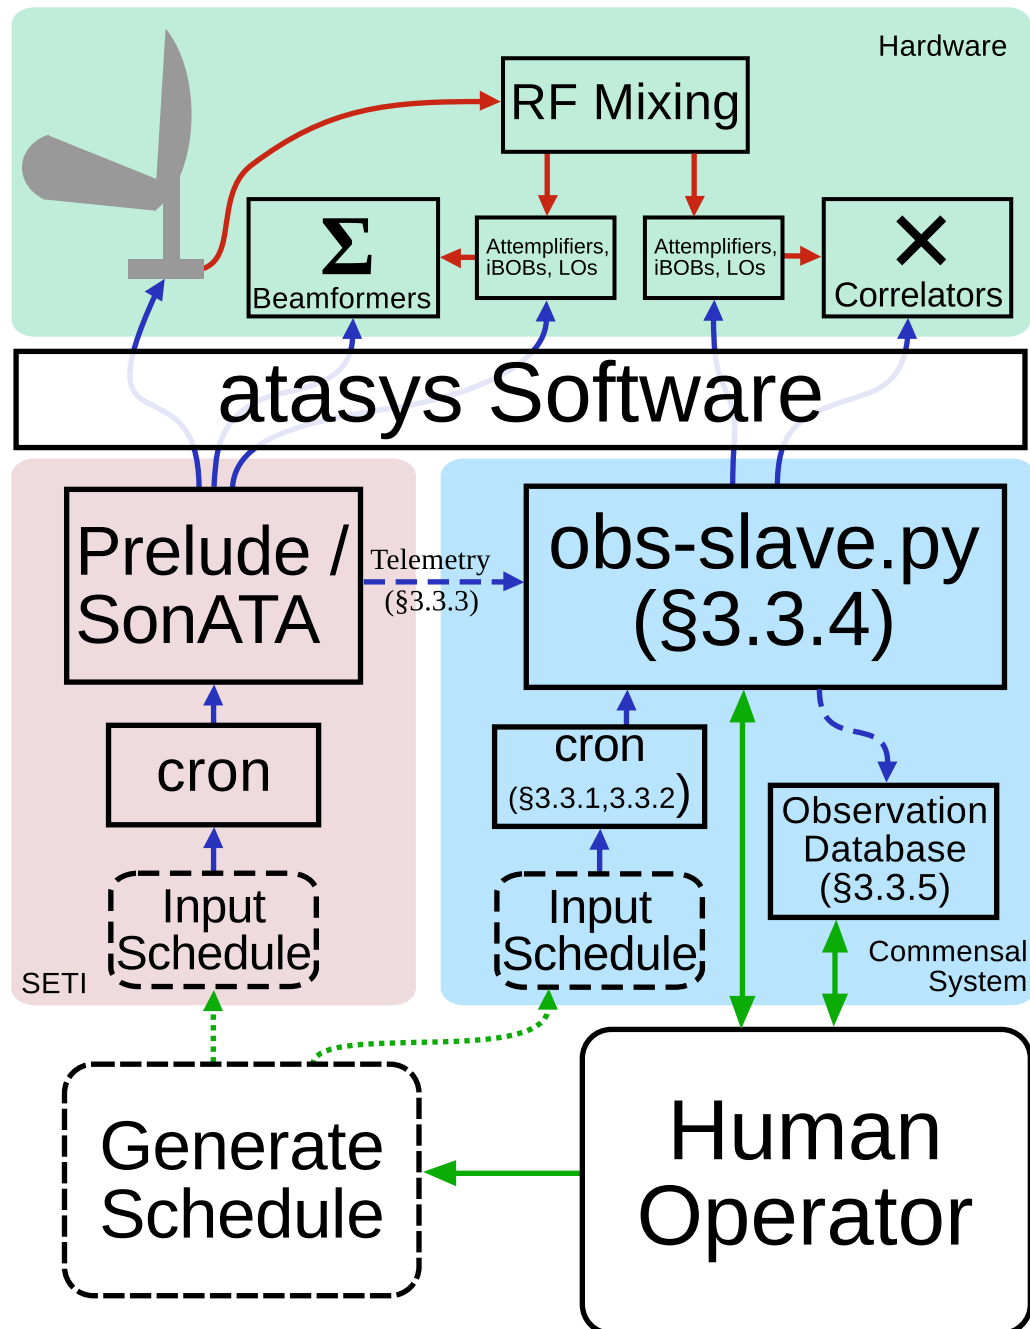


Figure 3.1: Schematic of the commensal observing system. See §3.3.

also used its own scheduling system. The suboptimality of this situation should be obvious. Regarding the particular case of the commensal campaigns, however, things weren't too bad. While it was sometimes necessary to coordinate schedule changes with SETI staff over email, and mistakes could result in lost observing opportunities, most observations went smoothly. Rarely, miscommunications would result in `obs.run` and the commensal launcher attempting to observe simultaneously; due to poor lockouts in the observatory infrastructure, it's actually possible for both systems to operate at the same time, with expectedly poor results.

Each observing session scheduled in the commensal system was assigned a random 32-bit hexadecimal identifier, *e.g.* `8592a0ce`. This identifier was used for tracking data and metadata throughout the rest of the observing and analysis chain. (The original scheme, used in the 2009 GCS survey, used universally unique identifiers [UUIDs], which are a particular kind of 128-bit identifier. They are 36 characters long and were judged to be overkill.) Note that identifying a session by the day on which it occurs, for instance, is insufficient if there are multiple sessions in one day. In theory, it doesn't even work to identify a session by its start time, since the array could be partitioned into two subarrays with different sessions being executed during the same block of time.

3.3.2 Launching and Managing Observations

The main challenges in automatically launching commensal observations are robustness and control. The loose nature of the ATA online system makes it possible for the commensal observations to start up with the array in odd states, and there are inevitably network outages, computer crashes, hardware problems, *etc.* Meanwhile, the array schedule isn't written in stone, and the operator needs to be able to see what's going on and potentially alter array operations at will.

Commensal correlator observations are launched automatically under the author's user account, `pkwill`. A launcher program, `obs-hourly`, runs hourly via `cron` (see `resources/obs.crontab`) on the host `obsctl` and consults the schedule file in a checkout of the repository found in `pkwill`'s home directory. If an observation is upcoming in the next hour, the launcher stays resident and `execs` the appropriate observing script at the right time. Because most observations are scheduled for at least a few hours, the hourly invocation of the launcher was intended to provide some insurance if the observing script crashes; this feature has been helpful only a handful of times.

Observations can also be "kickstarted" manually if the automatic invocation fails. This usually happens when an unusual hardware condition needs to be resolved before correlator observations can start. This is done with the `obs-kickstart` command, which uses GNU `screen` to launch the observing script in a detached, persistent virtual terminal. Both launchers have interlocks to attempt to prevent multiple observations from running simultaneously and to prevent observations from being launched outside of a scheduled commensal observing block.

Due to the distributed nature of the HCRO computing environment, observations can in theory be run from a variety of hosts. This capability was not usually exploited in ATA

Name	Code	Description
idle	0	Observing is not active, and the coordinates are undefined
slewing	1	The telescopes are slewing, and the coordinates specify their destination
pointed	2	The telescopes are pointed up on the specified coordinates

Table 3.1: Array states enumerated in the synchronization protocol.

operations. Commensal observations were designed to be run from the `cron` host `obsctl` (AKA `auxctl`) or, in early days, `tumulus`.

The `obs` user can monitor and control commensal observations with the command `commensal-mc`, which was set up to be easily accessible to the `obs` user and was intended to be straightforward to use. The command can perform diagnostics such as checking whether a commensal observing session is currently active (according to the commensal schedule), whether any commensal observing processes are currently active, killing or kickstarting observer processes, *etc.* Communication between processes and users is accomplished with special files in the shared ATA archive disk, `/ataarchive`, which is accessible to all HCRO computers and is writeable by most user accounts (notably, `obs` and `pkwill`). Commensal observing processes should be run as the user `pkwill`, which presents a difficulty if the user `obs` is attempting to kickstart a new process. This is dealt with by doing the kickstarting through a special passwordless SSH key with a preset command.

A more detailed look at the data coming out of a commensal observation can be obtained with `misc-parseplog`, which parses log output from the commensal observer and does some simple diagnostics on the datasets being written by the data catchers. This program has proven to be quite helpful for checking target selection and catching occasional correlator failures which were both serious and undetected by the usual diagnostics.

3.3.3 Synchronization Protocol

As alluded to above, the commensal observing software monitors a telemetry stream from Prelude. This stream was designed to convey the minimal amount of information necessary for the commensal observer to easily drive the correlators. The stream is broken into messages, each of which transmits a timestamp, an array state, and observing coordinates. The states are listed in Table 3.1.

This architecture lends itself to a state-machine implementation in the commensal observer. The lack of state in the telemetry stream means that the observer and/or Prelude can die at arbitrary times and the observer will recover sensibly.

The telemetry stream is implemented with broadcast UDP packets on port 24243. (This port was chosen arbitrarily.) Each packet consists of 32 bytes: a 64-bit double-precision floating point Unix timestamp, a 32-bit unsigned integer state code, an unused 32-bit padding field, a double-precision right ascension coordinate measured in hours, and a double-precision

declination coordinate measured in degrees. The values are big-endian and can be decoded with the Python `struct.unpack` specifier `>dlldd`. The numerical codes corresponding to the states are also listed in Table 3.1. The specified transmission rate is one packet per second, though there’s nothing special setting this rate. Protocol decoding is implemented in `pylib/protocol.py`.

Clearly, it is not difficult to implement this protocol. The program `obs-listener.py` can monitor the telemetry stream and print out low-level packet contents and check for unexpected state transitions. For example, it was used to diagnose a situation in which two Prelude processes were simultaneously running, one broadcasting a series of idle packets, the other attempting to run observations. The apparent oscillation between the idle and active states unsurprisingly confused the observer.

The program `obs-master.py` can be used to perform commensal-style observations when Prelude is inactive. Like Prelude, it performs target selection, drives the telescopes, and emits telemetry packets. Thanks to the clear division of responsibilities and the simplicity of the commensal observing protocol, the commensal observer can operate completely normally even though the software “in the driver’s seat” is completely different. This was also the experience when Prelude was replaced with SonATA, SETI’s next-generation observing system. Besides a change in the active port number (to allow simultaneous testing of Prelude and SonATA), the replacement of several racks of computer equipment and the entire software codebase was completely transparent to the commensal observer.

Compared to most elements of the commensal software package, the `obs-master.py` program is relatively complicated. It can either loop through a list of targets in order, with calibrator scans interspersed, or it can always observe the lowest-declination visible source. A few other parameters, such as the antenna focus position and interval between calibrator scans, are also configurable. It was not used extensively so its target-selection algorithms are relatively unsophisticated.

The current code to monitor broadcast UDP packets doesn’t work in the case that the emitter and listener are on the same host, which corresponds to situations when `obs-master.py` is being used. I suspect that this could be fixed with a better understanding of the proper low-level socket setup, but since this problem arises during manually-arranged observations, during those times I just tweak the socket code to skip actual network communication (by using the loopback device). This “solution” did lead to problems once or twice when I forgot to undo the tweak once observations were over.

3.3.4 Correlator Observer

The commensal correlator observer, `obs-slave.py`, is the most complex element of the commensal observing system. It is a single program that is launched by the hourly cronjob, monitors the telemetry stream, and operates the ATA correlators to take data for transient searches.

On a low level, the commensal observer is responsible for driving a specific set of subsys-

tems:

- the ATA correlators
- the amplifiers used by those correlators
- the local oscillators (LOs) used by those correlators
- the fringe-rotator iBOBs

The observer must take appropriate action based on state transitions in the telemetry stream or other external inputs (*e.g.*, abort commands).

The observer must also pay attention to some secondary systems to support its observations:

- It must have access to a catalog (mapping from coordinates to source names) so that correct target names can be shown to the operator and embedded in output datasets.
- It must generate ATA “ephemeris files” to be used by the fringe rotators and correlator data-catchers.
- It must check for abort signals, which is accomplished by monitoring the `/ataarchive` filesystem for the appearance of special files.
- It must generate detailed logs, since one wants to be able to debug subtle problems, check that hardware is being driven correctly, search for efficiency losses, and monitor the array.
- It must save appropriate state so that if it crashes and is restarted, the array hardware is reset *or not reset* so that datasets from before and after the crash may be calibrated consistently.

In order to accomplish all this, the observer is a multithreaded program with a main coordinator thread and various subthreads responsible for the subsystems. Because it is important that the observer be resilient to crashes, there’s also complex code to deal with Python exceptions as robustly as possible. The final program is still less than 700 SLOC (statement lines of code), a nice demonstration of the concision of Python. Looking at the source will confirm, however, that the code is quite dense.

The commensal observer stores each individual scan in a separate MIRIAD dataset simply named with a serial number. This approach makes a directory listing relatively opaque, but was hoped to be more robust than the standard ATA practice of naming datasets with some combination of source name, observing frequency, *etc.*, and appending to these datasets when multiple scans occur. Most MIRIAD tasks can stream through multiple datasets, but none can truly analyze only a portion of a dataset, so it should be more efficient to create many datasets and sometimes deal with them in bulk, rather than to create large datasets and

sometimes subdivide them. There have also been instances where a large dataset has been rendered invalid due to some problem at the very end of an observation, and subdividing datasets helps minimize the damage incurred in these cases.

The datasets generated by the commensal observer are augmented with a few extra data items. A file within the dataset named `c_instr` records the identifier of the ATA correlator used to generate the data; there is perhaps an oversight in the ATA data catcher software that this information is not recorded in datasets otherwise. (This information is useful because failures in the digital electronics have correlator-specific effects.) Another file named `c_tstop` records the timestamp at which Prelude reported leaving the current pointing for another target — there will be a delay between Prelude issuing the slew command and the correlator data catchers shutting down and closing their datasets, so there may be some bad data at the end of a scan taken while the antennas are slewing.

One particular challenge faced by the commensal observer is that the standard ATA data catcher program, `atafx`, was designed to be run with a predefined integration time. By the nature of the commensal observing system, however, the commensal observer does not know how long each integration will last. (If the expected integration time were transmitted from Prelude to the observer, one would still have to check that it was honored, and dealing with unexpected circumstances would require all of the flexibility needed by the current implementation.) Given the current software, the best solution is actually to SSH to `strato`, the data-taking host, and kill the `atafx` processes. A project was started to change the data takers to be always-running servers so that integration could be stopped and started quickly and on-demand, but that code never reached deployment.

3.3.5 Observation Database

The commensal observing system includes not only software for making observations but also a set of tools for understanding the observations that have been made. These are built on a database of observing sessions and individual scans.

These post-observation tools could plausibly have been separated into a different software package, and that might arguably have been a better design decision. Different observing programs may have different post-observation analysis needs and thus could benefit from multiple post-observation toolsets. On the other hand, any post-observation analysis toolset requires some knowledge of available datasets and the context in which they were obtained, so there’s an advantage to grouping this code with the observing code. It seemed helpful to not splinter the relevant software into too many pieces, so the two components were kept together.

That being said, an important aspect of the design of the post-observation analysis toolkit was a reliance on only the datasets stored in the ATA archive and no additional metadata. The reasoning was that while one might have plenty of expectations about the data on disk from the observing plan or even observing logs, the data on disk are the “ground truth,” and there are always unexpected ways for the logged metadata and recorded visibilities to

Name	Type	Units	Description
uid	string		UID of the scan's session
fncode	int		identifies the scan filename within the session data directory via reference to a separate table
ccode	int		identifies the equatorial coordinates of the scan pointing via reference to a separate table
freq	int	MHz	the observing frequency of the scan
focus	int	MHz	the predominant focus setting of the antennas during the scan
tst	int	seconds	the start time of the scan (as a Unix time-since-Epoch)
dur	int	seconds	the duration of the scan
lst0	float	hours	the LST at the beginning of the scan
failtype	int		failure flag if scan is unusable for some reason

Table 3.2: Metadata recorded in the scan database.

disagree. There were indeed many cases in practice in which the metadata and the actual datasets disagreed.

For each correlator observing campaign, a table of observing sessions and individual scans is maintained. The session table is populated from the schedule with the program **pdb-stubsess**, with new sessions marked as “unscanned”. After being observed, each session is eventually scanned with **pdb-scansess** and marked as “scanned”. The **pdb-scansess** program creates entries for each scan in the session and in fact reads all of the visibility files completely to search for any potential problems in the data. It also records useful per-scan metadata, the fields of which are listed in Table 3.2. Of particular note are the **lst0** field, which allows quick analysis of the hour-angle coverage of observations of a source, and the **failtype** field, which records any issues that make the scan unusable. To paraphrase *Anna Karenina*, “Successful observations are all alike; every unsuccessful observation is unsuccessful in its own way.” Thus there’s one value of **failtype**, zero, which indicates success, but a variety of nonzero values indicate possible failure modes, as in the familiar **errno** values returned in Unix error codes.

The observation databases are recorded in flat, line-oriented text files using a simple database layer implemented in **pylib/flatdb.py**. While it’s probably foolish to implement one’s own database layer, the **flatdb** system is simple, fairly efficient, and was not a major sink of programmer time. The motivation for creating it was to take advantage of the fact that the Git repository would effectively provide change tracking and data backup. While a few existing text-based Python database modules were found, they were generally not well-engineered.

Additional utilities were created to populate the observation databases more fully as certain new needs came to light. For instance, **pdb-filldatasizes** computes and records the total data size of each session, to allow reporting of the total survey data volume. **pdb-fillfocus** determines and inserts focus setting information because the importance of this information

was not initially obvious.

Several schedule blocks were used to observe ASGARD sources outside of the commensal observing system. Several tools were written to retroactively populate the database with information from these sessions so that all of the relevant information would be centralized. Not all sessions run outside of the commensal system map onto the schema of the observation database, but many do. The program `misc-scantree` performs a similar task to `pdb-scansess`, printing out processed scan information, but it does not rely on the existence of metadata from `obs-slave.py`. The program `pdb-retrosess` does the job of inserting this information into the database. The two can be linked together in a shell pipeline.

Various other utilities query the databases to analyze, *e.g.*, hour angle coverage of a source, total observing time, or data quality. The most widely-used of these tools is `qdb-dumpsess`, which can summarize the status of an entire observing campaign or one of its component sessions. The simultaneous use of two correlators complicates some queries, since separate scan records are created for each correlator’s data stream. For instance, if one naïvely adds up the integration time of a group of scans on a particular source, the total will be about twice the actual integration time, because two correlators were active. A different database schema could certainly trade off this particular issue for other ones.

A final group of tools integrate the information in the observation database to ARF, the ATA Reduction Framework, the system currently used to analyze the commensal correlator observations. While a discussion of ARF is outside the scope of this chapter, I’ll mention that the programs `rx-create` and `rx-recreate` stub out ARF “reduction specifications” and link ARF work directories with the appropriate raw datasets.

3.4 Recommendations

I conclude with some recommendations to be kept in mind when designing and implementing future commensal observing campaigns. These are naturally targeted toward projects similar to the one described here and won’t apply to every campaign that could be described as “commensal.”

- *KISS: Keep It Simple, Stupid.* Perhaps the only universal engineering maxim, and it’s as relevant in the commensal context as it is everywhere else. It’s almost always better to get something small working and build out. We certainly had many ideas for the campaign described in this chapter that, in retrospect, would have been complete wastes of time to implement.

The following group of recommendations is actually more specific *large* projects than *commensal* projects. For a more authoritative perspective, [Kleinman et al. \(2008\)](#) present some lessons learned from the operations of the Sloan Digital Sky Survey.

- *It’s all about the data.* You’ll be able to write new software, but you won’t be able to take new data, so get the data right. The highest priority, of course, is getting whatever

allows you to accomplish your science goals. Beyond that, the more uniform your data are, the easier processing will be — so not only is it important to get the data right, but it makes life a lot easier to think about these things hard before the campaign even starts. Large quantities of data are extremely slow to move around, so production-scale processing needs to require as little copying as possible.

- *Get an end-to-end processing pipeline online as soon as possible.* You don't want to wait until after half your observing is done to realize that you need a new kind of calibration observation, or you're missing some very useful metadata, or something's wrong in the archives. Once a pipeline is in place, you can also start inserting sanity checks to discover bad data just days, not months, after they start coming in. Start out by stubbing out as much as possible (cf. KISS) and fill in the details as you can.
- *Define and monitor observing efficiency metrics.* You want to know if you're on track to reach your science goals, and your observations will almost surely be less efficient in practice than in theory. Choose a just few key metrics to avoid information overload. As with the processing pipeline, the earlier these can be put into regular use, the better.
- *Schedule a time to step back and review.* If observations are running near-continuously, it becomes difficult to take the time to review how the campaign has progressed and ponder how efficiency might be improved. After a large project has gotten rolling, it's probably worthwhile to fall behind on data processing in order to spend a week or so conducting such a review.

These recommendations are more specific to commensal campaigns:

- *“Electronics before concrete.”* This is a slogan promoting cost-effective design adopted by Swiss railroad planners ([Schwager 2008](#)). The idea is that it's much cheaper and faster to retrofit existing systems (new signaling systems on existing lines, in the Swiss context) than it is to build new ones from scratch. This is certainly also true when building the infrastructure to run a commensal campaign: you should take advantage of existing infrastructure for non-commensal observations. It only took a few weeks to bolt a small telemetry module onto the existing SETI Prelude system; meanwhile, it took several years to get SonATA, the from-scratch Prelude replacement, on the air.
- *“Organization before electronics before concrete.”* This variation is used by some German planners ([Baumgartner et al. 2011](#)). Their point is large projects often involve multiple unaffiliated actors (rail transit agencies) whose turf wars and internal preferences can lead to plans that are *much* costlier than what would be arrived at by an apolitical efficiency-focused team; thus, when planning a project, one of the most useful things you can do is constructively tackle points of conflict, even though it's always tempting to avoid them. In a commensal campaign, of course, there are also multiple actors with diverging interests. It was possible to design the telemetry system (§3.3.3) in such a

reliable way only because decisions about observing strategy were negotiated in person and not left to be implicitly made in software.

- *Share as much infrastructure as possible.* This is related to the previous item. In a commensal context, multiple systems doing the same job will find a way to get out of sync and cause problems. The scheduling system of this chapter is an example of this. It isn't bad considered on its own, but its interactions with the SETI and HCRO systems are problematic. Much time has been spent emailing back and forth to negotiate and confirm schedule changes. There's no good reason for the correlator and beamformer observations to be scheduled separately. The *bad* reason for this is that the two sets of observations are run on separate computer systems and so it is difficult for them to share information. I have no doubt, however, that some arrangement could have been arrived at, avoiding not only the tedious emailing but also making observing more robust thanks to the elimination of *an entire class* of potential failures.
- *Make standalone pilot observations, and expect to perform standalone "patch-up" observations.* It will help the campaign design process if you've performed a few test observations without the commensal infrastructure to check out the mechanics and hopefully learn about any potential pitfalls. There will almost definitely be *some* kind of observation that would be good to get that won't be made (or will be botched) during regular operations, so plan to make occasional standalone observations to cover these gaps. (Don't forget to propose for the time to make these observations, if necessary!)
- *Be considerate of observatory staff.* This should go without saying. In the particular case of commensal campaigns, hacks to the observatory infrastructure will likely be necessary, and it's vital that these occur in a fashion acceptable and comprehensible to the staff who will have to deal with them.

Finally, these recommendations may be relevant to the software implementation of a commensal campaign:

- *Avoid bidirectional communications.* There are many more ways to implement a one-way communications channel than a two-way channel. The one-way telemetry stream described in this chapter (§3.3.3) was straightforward to implement and robust, as demonstrated by the smooth replacement of Prelude with SonATA and the easy creation of simple tools such as `obs-master.py` or `obs-listener.py`. The success of the synchronization protocol can be contrasted with that of JSDA, the messaging system used to interlink the ATA subsystems (Gutierrez-Kraybill et al. 2010). JSDA uses a two-way "remote procedure call" model for communications. While it undoubtedly offers important functionality, the JSDA system is complicated, and obscure lockups or communications failures have not been uncommon. The heaviness of the protocol also makes it time-consuming to integrate new software into the messaging network.

- *Use stateless protocols.* Software and hardware will inevitably fail, and testing of your observing systems is likely to be inexhaustive. Using a stateless architecture makes it so you don't need to even try to handle a variety of tricky problems.

Speaking as the person who drove the commensal correlator system and who has to reduce all of the data it generated, I feel that it performed well, and that this was largely due to some good design decisions early on in the commensal campaign. It's not likely that much of the particular implementation will be portable to future facilities, but I hope that some of these recommendations will be.

Table 3.3: Commensal Command Summaries

Name	Description
commensal-mc	Monitor & control of commensal obs by obs user
commensal-mc-helper	Helper for the above
misc-archdir	Print the <code>/ataarchive</code> directory for a session's data
misc-crsync	Copy session data from strato to cosmic
misc-decodeunixtime	Print a Unix time in human-friendly format
misc-diffunixtime	Print the difference between two Unix times in a human-friendly format
misc-dsetnames	Print the names of datasets including a particular source
misc-latestdircmd	Print an eval'able command to change to the directory containing the latest observations
misc-makeunixtime	Convert calendar information into a Unix time
misc-parseplog	Summarize an ongoing commensal observation from its pointing logfile
misc-scantree	Scan an arbitrary tree of visibility data
misc-sessalias	Given a session UID, print its alias
misc-sessuid	Given a session alias, print its UID
misc-whatsnear	Given coordinates, find the nearest observed pointings
obs-hourly	The hourly observing launcher
obs-kickstart	Launch observations right now
obs-launcher	Backend program to launch observations
obs-launcher-helper	Helper for the above
obs-listener.py	Debug the commensal synchronization protocol
obs-master.py	Drive the array simplistically
obs-slave.py	Perform commensal correlator observations
pdb-datealias	Set session aliases from their observing dates
pdb-filldatasizes	Fill in data size information for sessions missing it
pdb-fillfocus	Fill in focus information for sessions missing it
pdb-filllists	Fill in LST information for sessions missing it
pdb-fillobsfcgs	Fill in observing frequency/focus configuration information for sessions missing it
pdb-importagcts	Import database information from the first-generation AGCTS database
pdb-importproj	Import information from one commensal observing project (campaign) into another

Table 3.3 (cont'd): Commensal Command Summaries

Name	Description
<code>pdb-retrosess</code>	Retroactively import the scans and metadata for a session
<code>pdb-scansess</code>	Import scan information for a given session
<code>pdb-stubsess</code>	Stub out session information in the database from the current schedule
<code>plan-archdir</code>	Print out the expected archive directory of a scheduled observing session
<code>plan-importtxt</code>	Import the ATA textual schedule into the commensal schedule
<code>plan-schedsess</code>	Manually schedule an observing block
<code>plan-showsched</code>	Print the current observing schedule
<code>qdb-dumpsess</code>	Print information about an observing campaign or session
<code>qdb-sesssrchacov</code>	Show how well each session covers a given source in a given hour angle range
<code>qdb-srccov</code>	Summarize the coverage of a particular source in a project on a session-by-session basis
<code>qdb-srchacov</code>	Show the hour angle coverage of a given source over the course of the campaign
<code>qdb-summarize</code>	Summarize the overall observing statistics of a campaign
<code>qdb-toscan</code>	Print out a list of sessions that probably need to be scanned
<code>rx-create</code>	Create an ARF reduction specification and workspace for a given session
<code>rx-recreate</code>	Re-realize an ARF reduction workspace for a given session
<code>rx-status</code>	Set the reduction status of a session
<code>rx-suggest</code>	Suggest a useful session to reduce

Acknowledgments

Special thanks to Peter Backus, Tom Kilsdonk, and Jon Richards for their work to make the commensal observing campaign happen; and to Geoff Bower for carefully reading drafts of this chapter. Finally, the MMM group was an invaluable resource for discussing all things observational.

Table 3.4: Commensal Python Module Summaries

Name	Description
<code>catalog</code>	Loading and using source catalogs
<code>files</code>	Computing paths for accessing file resources
<code>flatdb</code>	Simple line-oriented textual database
<code>projdb</code>	Session and scan database implementation
<code>protocol</code>	Synchronization protocol implementation
<code>sched</code>	Loading and using observing schedules
<code>util</code>	Miscellaneous utilities

Chapter 4

Evaluating the Calorimeter Model with Broadband, Continuous Spectra of Starburst Galaxies

Although the relationship between the far-infrared and cm-wave radio luminosities of normal galaxies is one of the most striking correlations in astronomy, a solid understanding of its physical basis is lacking. In one interpretation, the “calorimeter model,” rapid synchrotron cooling of cosmic ray electrons is essential in reproducing the observed linear relationship. Observed radio spectra, however, are shallower than what is expected of cooled synchrotron emission. In [Thompson et al. \(2006, hereafter T+06\)](#), a simple parameterized model is presented to explain how relatively shallow observed spectra might arise even in the presence of rapid synchrotron cooling by accounting for ionization losses and other cooling mechanisms. During the commissioning of the 42-element Allen Telescope Array (ATA), we observed the starburst galaxies M82, NGC 253, and Arp 220 at frequencies ranging from 1 to 7 GHz, obtaining unprecedented broadband continuous radio spectra of these sources. We combine our observations with high-frequency data from the literature to separate the spectra into thermal and nonthermal components. The nonthermal components all steepen in the cm-wave regime and cannot be well-modeled as simple power laws. The model of [T+06](#) is consistent with our M82 results when plausible parameters are chosen, and our results in fact significantly shrink the space of allowed model parameters. The model is only marginally consistent with our NGC 253 data. Assuming the [T+06](#) model, a steep electron energy injection index of $p = -2.5$ is ruled out in M82 and NGC 253 to >99% confidence. We describe in detail the observing procedures, calibration methods, analysis, and consistency checks used for broadband spectral observations at the ATA.¹

¹This chapter has been previously published as Williams *et al.*, [2010 ApJ 710 1462](#), and is reproduced with the permission of all coauthors and the copyright holder. Copyright 2010 American Astronomical Society.

4.1 Introduction

The correlation between the far-infrared (FIR) and centimeter-wave radio emission of normal galaxies has been recognized for almost four decades (van der Kruit 1971). It holds over more than four orders of magnitude in L_{FIR} (Yun et al. 2001), is constant out to at least $z = 1$ (Appleton et al. 2004), and is observed in a large variety of galaxies (Condon 1992). The near-linearity (Yun et al. 2001) and universality of the FIR-radio correlation strongly suggest that it is intimately linked to the fundamental astrophysics of galaxies.

The physical underpinning of the correlation is broadly understood to be the formation of massive stars, which can believably be associated with the dominant contributions to both the FIR and continuum cm-wave radio emission of galaxies (Harwit & Pacini 1975; Wunderlich & Klein 1988). The FIR contribution is due to the hot dust envelopes surrounding these stars, which are energized by the copious amounts of Lyman continuum radiation that massive stars emit. The radio contribution includes two components. The first is synchrotron emission from relativistic cosmic ray (CR) electrons spiraling in the interstellar magnetic field. Such electrons can be accelerated in the supernova (SN) shocks that occur when massive stars die. The second, less significant, component is free-free emission, which predominantly comes from H II regions. These too are associated with massive stars.

This simple picture, however, neglects many important details. The timescales and dynamics of the radio and FIR emission are very different: CR electrons radiate over a timescale of 10^4 - 10^7 yr (see below), during which they diffuse through the galaxy, are possibly advected out of it via a galactic wind (Breitschwerdt et al. 1991), and cool via a variety of mechanisms, including inverse Compton (IC) scattering, ionization, free-free emission, and adiabatic expansion. Meanwhile, the FIR emission is prompt and comprised of comparable warm (~ 40 K) and cool (“cirrus,” ~ 20 K) dust components, the latter of which is primarily heated by the interstellar radiation field, not newly-formed stars (Walterbos & Schwing 1987; Boulanger & Perault 1988; Fitt et al. 1988).

Much theoretical and observational effort has been expended in the cause of understanding the FIR-radio correlation in detail, (*e.g.* de Jong et al. (1985); Helou et al. (1985); Cox et al. (1988); Hummel et al. (1988); Devereux & Eales (1989); Price & Duric (1992); Lisenfeld et al. (1996)). An influential model has been that of Völk (1989), which presents a theory for understanding the FIR-radio correlation in the case that CR electrons are accelerated only by SNe and cool completely within a galaxy. In this situation, the galaxy acts as a “calorimeter” with regard to the energy injected into CR electrons. For the FIR-radio correlation to hold, the synchrotron luminosity per unit energy injected into CR electrons must be a universal constant. Taking the primary cooling mechanisms to be IC scattering and synchrotron, it is found that $\bar{U}_B/(\bar{U}_{\text{ph}} + \bar{U}_B)$ must be nearly a constant value, where $\bar{U}_B = \langle B^2/8\pi \rangle$ is the mean energy density of the interstellar magnetic field and \bar{U}_{ph} is the mean energy density of the interstellar radiation field.

Because the synchrotron cooling time scales as $B^{-3/2}$ (T+06), the applicability of the calorimeter model depends on the strength of the interstellar magnetic field. In a stronger

field, the calorimeter model is more likely to hold as CR electrons will cool more rapidly and have less chance to escape the galactic disk before radiating all of their endowed energy. This is precisely the scenario suggested in T+06, where it is argued that starforming galaxies have much stronger magnetic fields than commonly assumed. The work assesses the applicability of the calorimeter model by considering the timescales for synchrotron cooling and escape from the galactic disk by CR electrons. Using the usual assumptions, the two are comparable; using the stronger magnetic field that T+06 argues for, the synchrotron cooling timescale is much shorter, favoring the assumptions of the calorimeter model.

T+06 specifically argues that in starburst galaxies the magnetic field is dynamically relevant: its energy density is comparable to the hydrostatic pressure of the ISM, or

$$B \approx B_{\text{eq}} = (8\pi^2 G)^{1/2} \Sigma_g \approx 2.3 \left(\frac{\Sigma_g}{1 \text{ g cm}^{-2}} \right) \text{ mG}, \quad (4.1)$$

where Σ_g is the surface density of the galactic gas disk. This is in contrast to the classic “minimum-energy” estimate (Burbidge 1956; Longair 1994),

$$B \approx B_{\text{min}} \approx 9.5 \eta^{2/7} \left(\frac{L_\nu}{10^{23} \text{ W Hz}^{-1}} \right)^{2/7} \left(\frac{V}{1 \text{ pc}^3} \right)^{-2/7} \left(\frac{\nu}{1 \text{ GHz}} \right)^{1/7} \text{ mG}, \quad (4.2)$$

where η is the ratio of total CR energy to energy in the CR electrons alone, L_ν is the synchrotron luminosity, V is the estimated volume of the emitting region, and ν is the observing frequency. For the parameter values commonly associated with starburst galaxies, $B_{\text{eq}} \gg B_{\text{min}}$ (T+06, fig. 1).

An important observational constraint on the argument of T+06 is the shape of the galactic nonthermal spectrum. It is typically argued that galactic spectra are insufficiently steep to be compatible with the calorimeter model: observations indicate $\alpha \approx -0.8$ ($S_\nu \propto \nu^\alpha$; Condon (1992); Niklas et al. (1997)) rather than $\alpha \approx -1$ as would be expected from the combination of SN shock physics (Blandford & Eichler 1987) and spectral steepening caused by strong cooling (Condon 1992). T+06 contends that more detailed analysis mitigates this issue. As CR electron energy decreases, the timescales of ionization losses and free-free emission are, respectively, shorter and unchanged, so both of these processes will tend to flatten observed spectra. In T+06 simple steady-state simulations were performed to quantify this effect (cf. its fig. 3 and §4.3). These simulations lead to the general prediction that the nonthermal radio spectra of starburst galaxies should steepen with increasing frequency, asymptoting to $\alpha \lesssim -1$.

Most spectral index measurements are based on only a handful of fluxes. To accurately measure the steepening of a nonthermal spectrum, many more data are needed, because both the steepening and the thermal contribution must be determined. The work of Duric et al. (1988) and Niklas et al. (1997) shows that determination of the thermal fraction is difficult and requires high-quality high-frequency data, even if a simple power-law nonthermal spectrum is assumed. As Condon (1992) points out, the thermal fraction and nonthermal

steepening can also be highly degenerate (cf. his Figure 6). One must sample a large number of points over a wide frequency range in uniform conditions to effectively constrain these parameters.

Though it is still being commissioned, the Allen Telescope Array (ATA; Welch et al. (2009)), with its continuous frequency coverage of 0.5 GHz to 10 GHz, is well-suited to testing the predictions of T+06 by observing high-resolution, continuous spectra of starburst galaxies. Section 4.2 describes observations of M82, NGC 253, and Arp 220 made with the ATA in the fall and winter of 2008. In §4.3 our data reduction methods are described along with results for standard calibrator sources. Section 4.4 describes our modeling of the observed spectra to extract thermal and nonthermal contributions. Finally, in §4.5 we discuss our results and in §4.6 we present our conclusions.

4.2 Observations

We observed the starforming galaxies M82, NGC 253, and Arp 220 with the Allen Telescope Array on seven nights over the course of 2008 September to 2009 June. The ATA is a large- N , small- D (LNSD) cm-wave array located in Hat Creek, California, and is a joint project of the SETI² Institute of Mountain View, CA and the Radio Astronomy Laboratory of the University of California, Berkeley. The scientific motivators and technical capabilities of the ATA are described in Welch et al. (2009). At the time of the observations the ATA was undergoing commissioning, with many aspects of the array being actively developed. Forty-two antennas and one correlator had been installed at the site. Two-thirds of the antennas were typically available for use during our observations.

We used the continuous frequency coverage of the ATA to observe the galaxies at almost a decade of sky frequency in a sequence of snapshot observations. The observing frequencies ranged between 1.0 GHz and 7.0 GHz with a spacing of 100 MHz, skipping some frequencies due to known radiofrequency interference (RFI) or low priority. The primary beam sizes of the ATA at 1.0 GHz and 7.0 GHz are $\sim 210'$ and $\sim 30'$, respectively, and the synthesized beam sizes are typically $\sim 250''$ and $\sim 35''$. The snapshots lasted between 60 and 130 seconds, with the correlator dump time always being 10 seconds. All observations used a bandwidth of 104.9 MHz and 1024 correlator channels and sampled both vertical and horizontal linear polarizations. Observations of the science targets were interleaved with observations of the calibrators 3C 48, 3C 147, and 3C 286. These calibrators are compact to the ATA and have up-to-date broadband flux density models available from the VLA.³ Because of the need to set the flux density scale at numerous frequencies and the desire to obtain a large amount of data for sanity-checking, a large portion of the observing time was dedicated to calibrator observations. The total integration time on the calibrators was ten hours, compared to nine hours spent on the target galaxies. An observing strategy that gathered fewer redundant

²Search for Extraterrestrial Intelligence

³<http://purl.oclc.org/net/pkgwpub/vla-analytic-cal-models>

data could spend approximately half as much time observing calibrators as observing science targets.

4.2.1 Details of observing methods

Because the observations presented herein use a novel technique on a novel instrument, we shall describe our methods in detail. Each night was divided into several segments during which a subset of the targets and calibrators would be observed as a group. The typical duration of a segment was a few hours. Over the course of each segment, an observing script would repeatedly slew to each of the selected sources in turn, observing each at several frequencies before moving on to the next, until the segment’s allotted time had elapsed. This scheme attempted to balance the desire to obtain calibrator observations temporally close to the science observations with the desire to not spend too much time moving from one target to the next. Most segments included observations at 1.4 GHz and 5.0 GHz to enable checks of repeatability and stability. The observations presented in this chapter are summarized in Table 4.1.

Observing over a large range of frequencies requires not just a large number of calibrator scans but also reconfiguration of the ATA hardware. In particular, the sensitivity of the log-periodic ATA feeds is a function of position along the axis of their pyramidal structure, with the point of peak sensitivity approaching the apex of the pyramid with increasing frequency. To allow for maximal performance, the ATA feeds can be focused by moving them in a piston motion along the normal to each antenna’s focal plane (see [Welch et al. 2009](#), fig. 4). The optimal focus positions have been determined theoretically and calibrated so that an ATA observer can specify the desired focus position in frequency space. The focusing operation, however, is slow. In our work, before each observation we set the feed focus frequency to the observing frequency rounded up to an integral number of GHz (*e.g.*, 4.0, 4.4, and 4.9 GHz all map to 5 GHz). This scheme was chosen because the sensitivity of the feed to frequencies above the nominal focus frequency drops off more rapidly than the sensitivity to frequencies below it. Within each segment we chose observing frequencies that lay within a single GHz span to avoid superfluous refocus operations. The focusing and other time-consuming operations are parallelized as much as possible to conserve observing time.

Because neither the feed sensitivity nor the noise power entering the ATA are spectrally flat, when the focus position of a feed changes, the total power entering the analog system also changes. Just before digitization, the downmixed and filtered analog signal is passed through an “attenuplier” (attenuator/amplifier) stage to adjust the incoming power for optimal quantization ([Backer 2007](#)). The amount of attenuation or amplification must be explicitly configured, but the best settings can be calibrated automatically during observations. Clearly, any two observations that are to be compared or used for flux calibration must be made with the same attenuplier settings. Furthermore, the attenuplier calibration step is also somewhat slow (~ 60 s). In these observations, snapshots made at similar frequencies shared attenuplier settings. Specifically, one set of attenuplier settings was used for each integral

Table 4.1: Observations

Date (UT)	Segment #	Sources	Frequencies (GHz)
20080912	1	3C 48, 3C 147, M82, NGC 253	1.4, 1.5, 2.0, 2.1, 3.1, 4.0 ^a , 4.1 ^a , 5.0, 5.1
20080913	2	3C 286, Arp 220	1.5, 1.6, 2.5, 2.6, 4.5, 4.6
	3	3C 147, NGC 253	1.1, 1.2, 1.3, 2.6, 2.7, 3.6, 3.7 ^a , 5.6, 5.7
	4	3C 48, 3C 147, M82	3.2, 3.3, 4.4, 4.5, 5.5, 5.6, 6.6, 6.7
20080914	5	3C 286, Arp 220	1.1, 1.2, 2.5, 2.6, 5.0, 5.7, 6.3
	6	3C 48, Arp 220	1.1, 1.2, 2.5, 2.6, 5.0, 5.7, 6.3
	7	3C 147, NGC 253	2.7, 3.1, 3.3, 4.5, 4.7, 6.1, 6.3
	8	3C 48, 3C 147, M82	1.0, 1.1, 2.6, 2.9, 4.8, 5.0
	9	3C 286, Arp 220	1.9, 2.9, 3.2, 3.3, 5.4
20080924	10	3C 286, Arp 220, M82	1.4, 1.7, 1.8, 2.0, 3.4, 3.5, 3.6, 5.0, 5.2, 6.0
20081015	11	3C 48, 3C 286, Arp 220	0.5, 0.7, 0.9, 1.4, 4.3, 4.4, 4.7, 5.0, 6.2, 6.7, 7.0
	12	3C 48, 3C 147, M82	1.4, 2.4, 2.5, 4.3, 4.7, 5.0
	13	3C 147, NGC 253	1.4, 1.7, 1.8, 1.9, 4.5, 4.6, 4.9, 5.0, 5.8, 5.9, 6.0, 6.4, 6.7, 7.0
	14	3C 48, 3C 147, M82	6.1, 6.2, 6.3, 6.4
20090307	15	3C 48, 3C 286	1.4, 3.1, 5.0, 6.0, 7.0
20090621	16	3C 147, 3C 286	1.0, 1.4, 1.5, 5.0, 5.1, 5.5, 6.1, 6.5, 7.0

^aDiscarded due to severe RFI.

GHz span (*e.g.*, $4 \leq \nu_{\text{GHz}} < 5$). The first time any frequency in a span is observed, the amplifier settings are calibrated for that frequency and saved; for subsequent observations during the same segment, the saved settings are reapplied.

As an example of all of these considerations, consider segment #12, as detailed in Table 4.2. This segment was scheduled for observation from 03:42 UT to 05:12 on 2008 Oct 15. (All dates and times are in UT.) The selected sources were 3C 48, 3C 147, and M82. The timing and (short) duration of this segment were set by the visibility of the science targets to the ATA: Arp 220 set just before this segment began, and NGC 253 rose just after it ended. Because M82 is circumpolar at the ATA’s latitude, the other two targets were prioritized whenever they were visible. The selected observing frequencies were 1.4, 2.4, 2.5, 4.3, 4.7, and 5.0 GHz. As described above, the observing script slewed to the sources in turn and observed each at several frequencies. The frequencies were chosen and grouped so as to reduce the amount of time spent adjusting the feed focus positions and amplifier settings. At 04:08, the script would have observed 3C 147 at the trio of lower frequencies if possible, but that source had not yet risen. Instead, the script skipped those observations and looped back to observe 3C 48 at the trio of higher frequencies.

Including time lost due to slewing, hardware initialization, software overhead, and so on, we made ~ 18 snapshots per hour while observing, integrating for 38% of the 53 hours spent using telescope resources during the observations presented in this chapter. The efficiency of future observations will be higher: we have made an unusually large number of calibrator observations and have made no effort to reduce slew time between sources. Furthermore, our observations were made before the installation of a second independently-tunable correlator at Hat Creek, which doubles the observational throughput.

4.3 Data Reduction

The data were reduced with a combination of the analysis package MIRIAD ([Sault et al. 1995](#)) and custom Python scripts based on the NumPy and IPython packages. Over the course of this project and other work on the ATA, significant improvements have been made to MIRIAD to ease its compilation, improve its ability to handle large datasets, and fix various bugs. A special module was written to allow easy management of and access to MIRIAD datasets in the Python language (Chapter 2).

The two orthogonally-polarized signals that are received by the ATA feed elements are amplified by separate low-noise amplifiers (LNAs) and subsequently treated as entirely separate signals. The ATA correlator at the time of the observations accepted 64 such antenna-polarization signals, or “antpols” for short, making it capable of correlating 32 antennas with full-polarization coverage, 64 antennas with single-polarization coverage, or, as encountered in practice, an intermediate case. Each ATA antenna is identified by a small integer and a single letter, *e.g.* 3g. (These specify the node and element identifier of the antenna within the ATA architecture.) An antpol is specified by suffixing the antenna name

Table 4.2: Snapshots made during segment #12

Time (UT)	Pointing ^a	Azimuth (deg)	Elevation (deg)	ν_{sky} (GHz)	Focus ^a (GHz)
03:50	3C 48	77	39	1.4	2.0
03:53		77	40	2.4	3.0
03:55		78	40	2.5	
03:59	M82	359	21	1.4	2.0
04:02		358	21	2.4	3.0
04:04		358	21	2.5	
04:08	3C 48 ^b	79	43	4.3	5.0
04:11		80	43	4.7	
04:13		80	44	5.0	
04:17	M82	359	21	4.3	
04:20		0	21	4.7	
04:22		0	20	5.0	
04:28	3C 48 ^b	82	47	1.4	2.0
04:31		82	47	2.4	3.0
04:34		83	48	2.5	
04:38	M82	1	21	1.4	2.0
04:41		2	21	2.4	3.0
04:43		2	21	2.5	
04:47	3C 147	40	19	1.4	2.0
04:50		40	19	2.4	3.0
04:52		40	20	2.5	
04:56	3C 48	86	52	4.3	5.0
04:58		86	52	4.7	
05:01		87	53	5.0	
05:04	M82	4	21	4.3	
05:07		4	21	4.7	
05:09		4	21	5.0	

^aA blank entry indicates no change from the previous value.^bHad it been above the 18° elevation limit of the ATA, 3C 147 would have been observed at these times.

with the letter X or Y, *e.g.* 2kY. We refer to a pair of antpols as a “basepol.”

Although the ATA correlator produces cross-hand polarization data, the calibration and analysis of these data have not yet been investigated and tested at the time of this work. In this chapter, cross-hand data were discarded, and all sources were assumed to be completely unpolarized. All images produced were of the Stokes I parameter. Because the linear polarization of 3C 286 is $\sim 10\%$, this treatment will lead to $\sim 1\%$ errors in some measurements. This effect contributes to, but does not dominate, the typical uncertainty in our measurements (see §4.3.4).

4.3.1 Flagging

The ATA does not yet have an online flagging system and our observations were performed outside of the protected bands. Fairly extensive manual editing of the data was performed. Several automatic criteria were applied. The 100 channels on both ends of the 1024-channel spectral window were flagged because the ATA digital filter aliases in large amounts of noise at the bandpass edges. Shadowed antennas were flagged, taking the effective antenna radius to be 7.5 m (somewhat more than the 6.1 m physical diameter of the primaries). Finally, visibilities with $\sqrt{u^2 + v^2} < 100\lambda$ were flagged as these short baselines were consistently sensitive to terrestrial and solar RFI, undesired large-scale structure, and possible cross-talk.

A fraction of the antpols had problems with their cryogenics, LNAs, or delay calibrations, resulting in data for the night that could easily be rejected after a cursory look. Occasional correlator hardware failures resulted in wildly erroneous amplitudes in a small subset of the spectral data, *e.g.* the first quarter of the spectrum of all visibilities involving antpols 1aX, 1cY, and 4kX. Such data were also flagged and the issue reported.

After an initial check for bad antpols or correlator hardware, RFI was flagged with a semi-automated, channel-based routine. *All* of the amplitude spectra at a given frequency observed over the course of a single segment were summed, regardless of time, baseline, polarization, or pointing. A simple peak-finding routine identified suspected RFI channels by searching for extremely high amplitudes. The suspect channels were presented to the user graphically and made available for interactive editing before a list of channels to flag was recorded. The listed channels for each frequency were flagged out for all observations at that frequency in the segment. Figure 4.1 shows an example of the summing and peaking-finding process. Figure 4.2 characterizes the RFI environment of Hat Creek via a schematic representation of the summed amplitudes across the entire observed frequency range.

This procedure is obviously suboptimal in the case of time-varying RFI, but it was found that most of the interference was virtually always present. Informal tests of more sensitive methodologies that attempted to better mitigate time-variable RFI failed to improve results and were much more time-intensive to the user.

A few bands had nearly omnipresent RFI that made astronomical observations nearly impossible. These were: 500-1000 MHz (terrestrial television, radar, pagers, aeronautical navigation, self-interference), portions of L band (1150-1280 MHz, aeronautical navigation

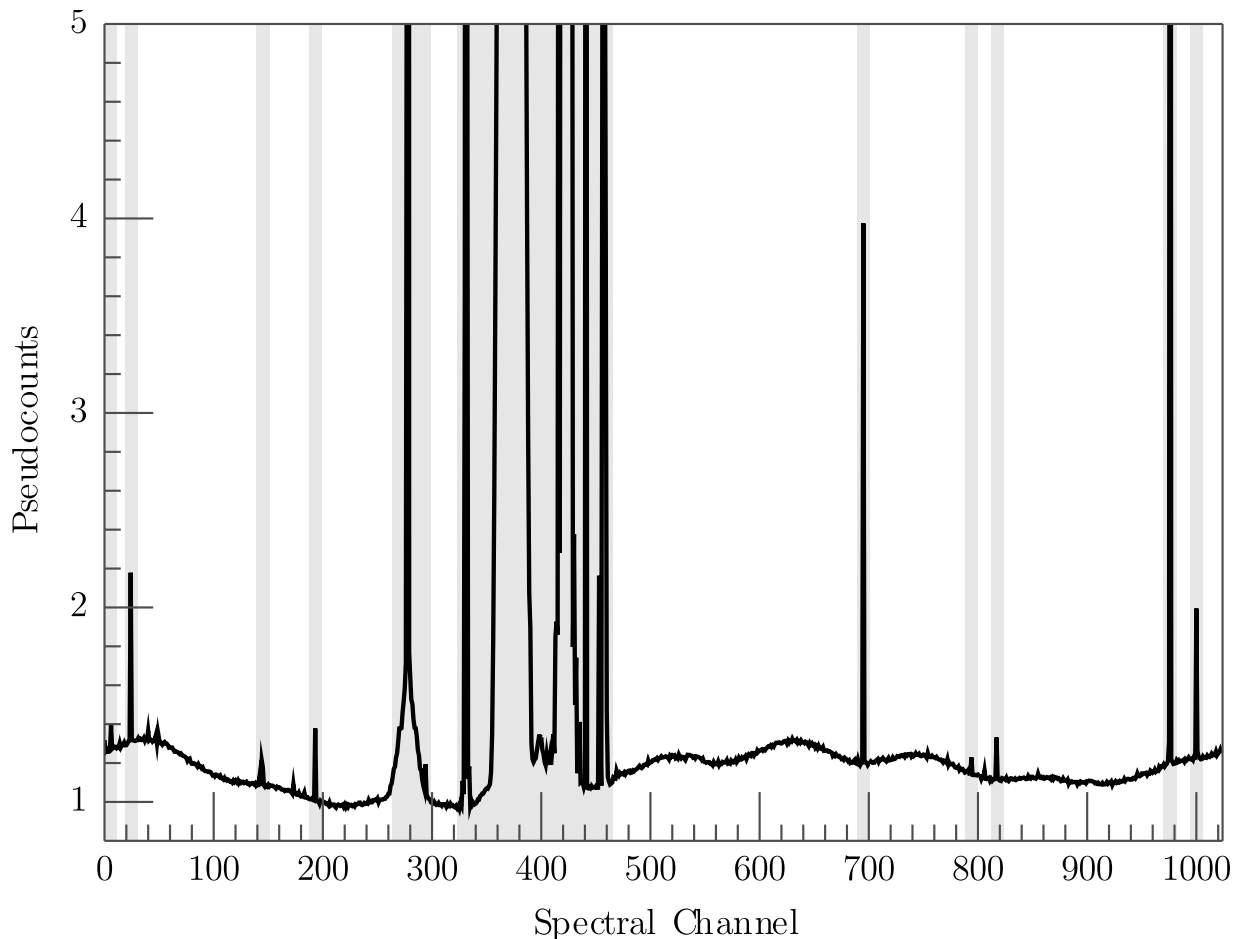


Figure 4.1: An assessment of the RFI present during a block of 1.7 GHz observations. This plot is typical of those produced during execution of the semi-automated RFI excision algorithm described in §4.3.1. *Solid lines* show summed amplitudes of the 1.7 GHz data obtained during segment #14 as a function of correlator channel, corresponding to frequencies ranging from 1.65 to 1.75 GHz. *Shaded regions* indicate RFI-afflicted regions detected by the peak-finding algorithm. Plots such as these were reviewed and sometimes edited to improve the flagging (*e.g.*, to widen the flagged region around the broad peak centered near channel 280). Channel 512 corresponds to a sky frequency of 1.700 GHz and the channel width is 102.4 kHz. The sinusoidal features across the bandpass are due to the ATA digital filter.

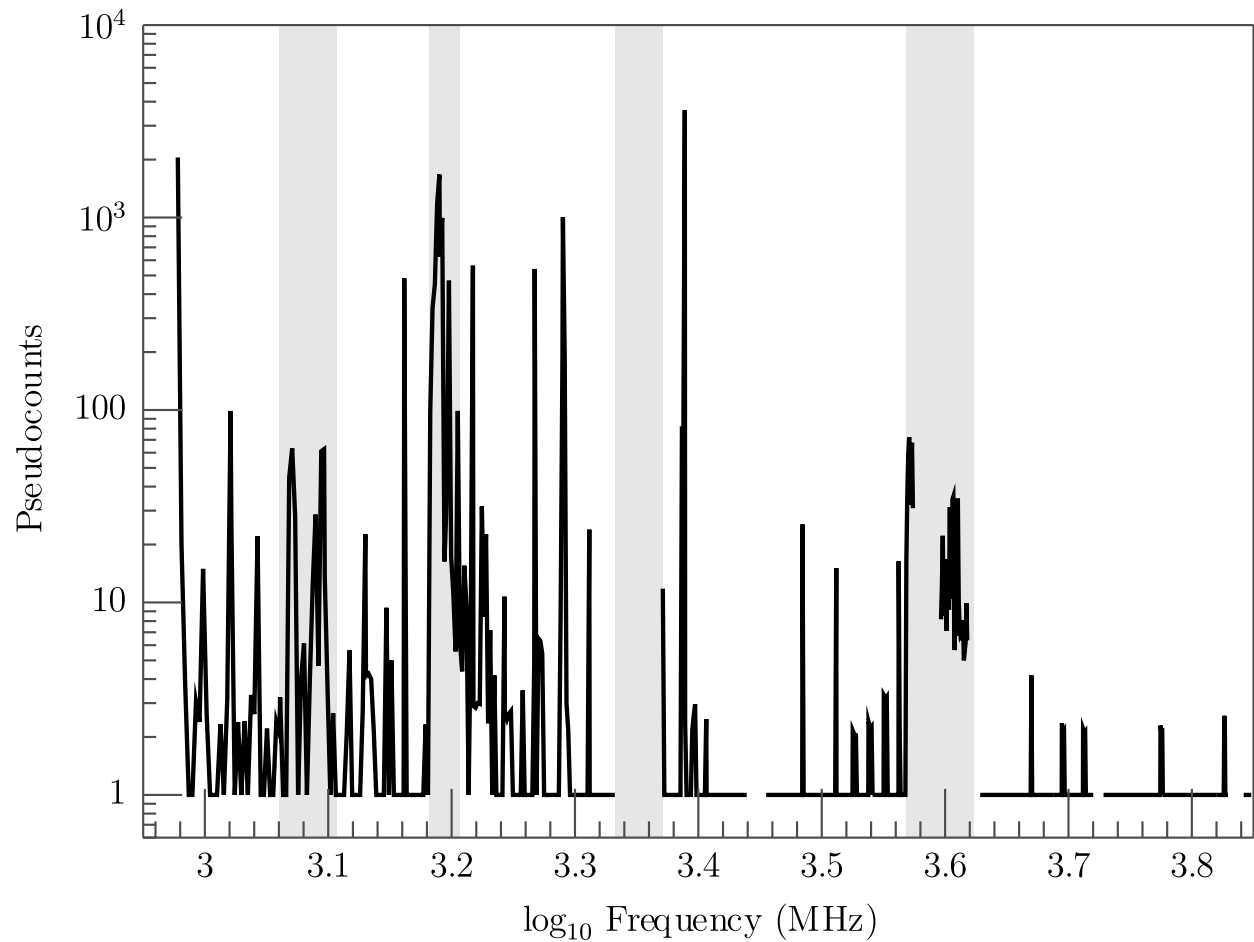


Figure 4.2: The RFI environment of the ATA, in schematic form. *Solid lines* show a processed version of the summed data amplitudes, as described in §4.3.1. The raw 100 kHz channelized amplitudes were processed to increase clarity while accurately portraying narrow RFI peaks. Gaps in the data are due to a lack of observations, usually because the RFI was known to dominate astronomical signals at certain frequencies. The *shaded regions* indicate bands afflicted with especially strong and persistent RFI.

and GPS L5; 1520-1610 MHz, which includes GPS L1 at 1575 MHz), the lower end of S band (satellite communications, 2150-2350 MHz, especially Sirius XM radio centered at 2339 MHz), and C band (satellite downlinks, 3700-4200 MHz).

After RFI-afflicted channels were flagged, various manual data-quality checks were run to check for remaining issues. These included phase-closure rankings of the calibrator observations, evaluation of the amplitude and phase variability of the baselines, and by-eye assessments of spectra. Many of these checks resulted in a few antpols or basepols being flagged for the duration of the segment, while occasionally a temporary system glitch was revealed and edited out.

Finally, after system temperatures were computed (see below), antpols with a T_{sys} above 300 K were flagged. At the end of the flagging process, $\sim 40\%$ of the antpols and $\sim 75\%$ of the visibility channels in a typical dataset were discarded. Figure 4.3 shows the average data retention rates as a function of frequency. The deleterious effect of RFI is apparent, as is the moderate degradation of system performance with increasing frequency. Approximately 30% fewer antpols and visibilities are retained at 7 GHz as compared to 1.4 GHz.

Ongoing retrofits to the array are increasing the amount of usable data generated during each observation. A large portion of this work comprises improved RFI mitigation techniques, including better shielding, especially of the signal processing room; improved techniques and algorithms for interference excision; and stricter on-site RFI protocols. Some data are lost because certain antpols have extremely poor system temperatures, for various reasons: failed cryogenics systems, damaged LNAs, or offset pointing due to imbalanced electrical connections. More robust replacement hardware will make these issues much rarer.

4.3.2 Calibration

The calibration of our data is challenging because of the modest u - v coverage, short integration times, and wide temporal and spatial separation between snapshots entailed by our observing mode. We will provide an overview of our calibration process with more detailed descriptions presented in the subsections below.

The snapshot observations of the calibrator sources were first self-calibrated to point-source models. All observations of the same calibrator and frequency were then imaged jointly to form a self-calibrated *reference image* with u - v coverage and sensitivity superior to those obtained in any one snapshot. The individual calibrator snapshots were then recalibrated to their corresponding reference images.

The snapshot observations of the science sources then had appropriate calibration solutions applied to them. Some snapshots (of both science targets and calibrators) were duplicated and had several different calibration solutions applied to allow for cross-checks of the calibrator flux models. Reference images were then generated from the calibrated observations and once again the individual snapshots were recalibrated to their corresponding reference images.

Given the nature of this work, accurate broadband calibrator flux density models are vital. The 1999.2 analytic models listed in the VLA Calibrator Manual were initially adopted

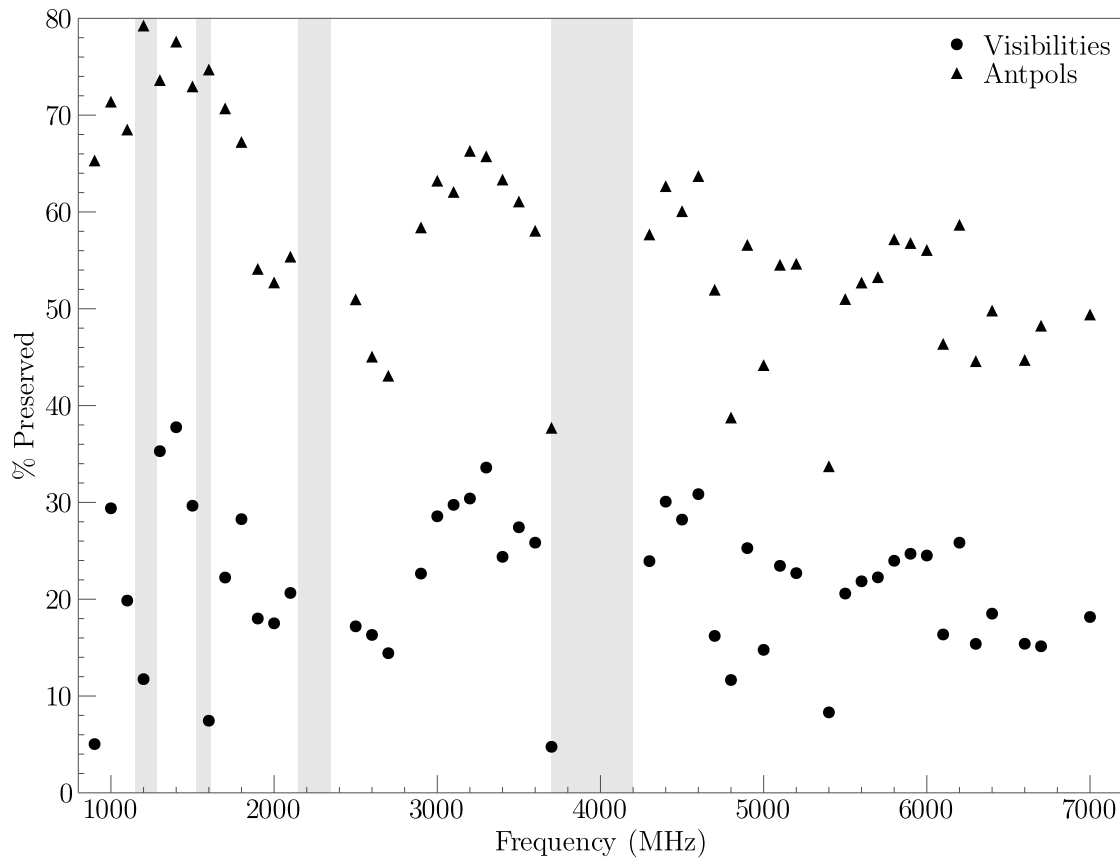


Figure 4.3: Average data retention rates over all observations, as a function of observing frequency. Entirely discarded datasets are not shown. *Triangles* denote the fraction of antpols used during imaging as compared to the number present in the raw data. *Circles* denote the same for individual visibilities. The *shaded regions* represent bands afflicted with strong, persistent RFI. The maximum possible visibility retention rate is $\sim 75\%$ due to the flagging of edge channels, shadowed antennas, and short baselines described in §4.3.1.

and found to yield good results. These models were, however, almost a decade old at the time of this work, and long-term VLA monitoring indicates that the fluxes of the calibrators evolve slightly on this timescale. Section 4.3.6 describes consistency checks used to verify and fine-tune the flux density scale used in this work.

Self-calibration

The calibration process started with self-calibration of the calibrator observations. Bandpasses and complex antenna gains in the individual snapshots were first determined via the MIRIAD task `smamfcal`, which assumes a point-source model. Spectral smoothing was used to increase the signal-to-noise ratio (SNR) in the high-frequency observations.

System temperature (T_{sys}) information was then computed and inserted into the datasets. This step was necessary because the ATA did not have online system temperature calibration at the time of this work. A Python script called `calctsys` was developed to estimate T_{sys} values from the variances in the real and imaginary parts of the visibilities across the band. (This was a precursor to the system described in §2.6.2.) The variances were used to compute per-basepol T_{sys} values assuming an antenna gain of 153 Jy/K. Per-antpol T_{sys} values were then computed by performing a least-squares minimization assuming that the T_{sys} of each basepol was the geometric mean of the T_{sys} of its two contributing antpols. The values obtained in this way had good reproducibility over the course of a segment and agreed well with lunar measurements made by ATA staff (R. Forster, private communication). The system temperatures could vary strongly as a function of frequency, with most, but not all, antpols demonstrating superior performance at lower frequencies.

Reference images were then created from the group of calibrator snapshot observations. All of the snapshots of each source and frequency were merged into a temporary dataset that was imaged using natural weighting. Multifrequency synthesis (Conway et al. 1990) was used to avoid bandwidth smearing and provide better u - v sampling. The images were deconvolved with the MIRIAD implementation of the CLEAN algorithm (Högbom 1974), which incorporates the algorithmic adjustments suggested by Clark (1980) and Steer et al. (1984). After an initial round of imaging, the merged visibilities were self-calibrated to the generated CLEAN component model and imaged once more. The second self-calibration was phase-only if the SNR of the data was less than 0.3. Testing showed that further iterations of calibration and imaging did not lead to improvements in the image quality. Figure 4.4 shows the reference image of NGC 253 at 1.3 GHz and an image generated from a single snapshot observation of the same source at the same frequency. Figure 4.5 shows reference images of M82 at four frequencies. The individual calibrator snapshots were then recalibrated using the CLEAN component model of the reference images.

Secondary calibration

The complex antenna gains and bandpass solutions were copied from each calibrator snapshot to any observations made at the same frequency within 40 minutes of the calibrator

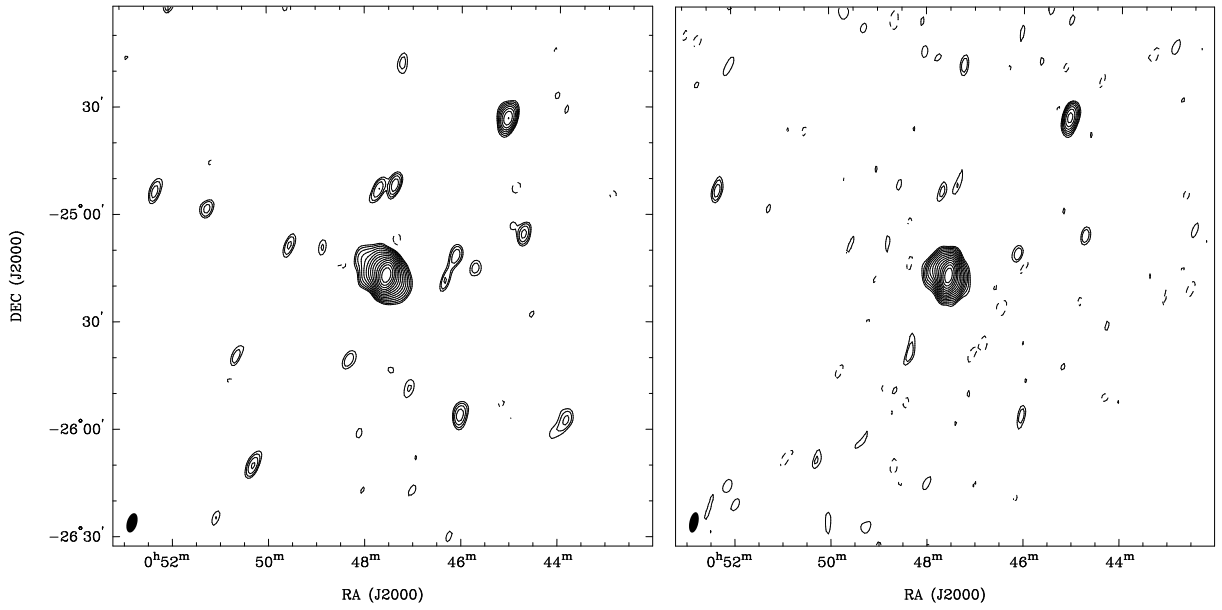


Figure 4.4: ATA maps of NGC 253 at 1.3 GHz. The *left panel* shows the reference image for this source and frequency while the *right panel* shows an image generated from a 70 s integration made during Segment #3. The total integration time for the reference image is 250 s. Contour levels in both panels are $30 \text{ mJy} \times 2^{n/2}$ for $n = 0, 1, 2, \dots, 13$. The RMS noise levels of the reference image and the snapshot are 9 and 11 mJy, respectively. Although the longer integration time and more complete u - v coverage of the reference image lead to a noticeable increase in its quality, the total fluxes of NGC 253 measured in the two images agree to 2%.

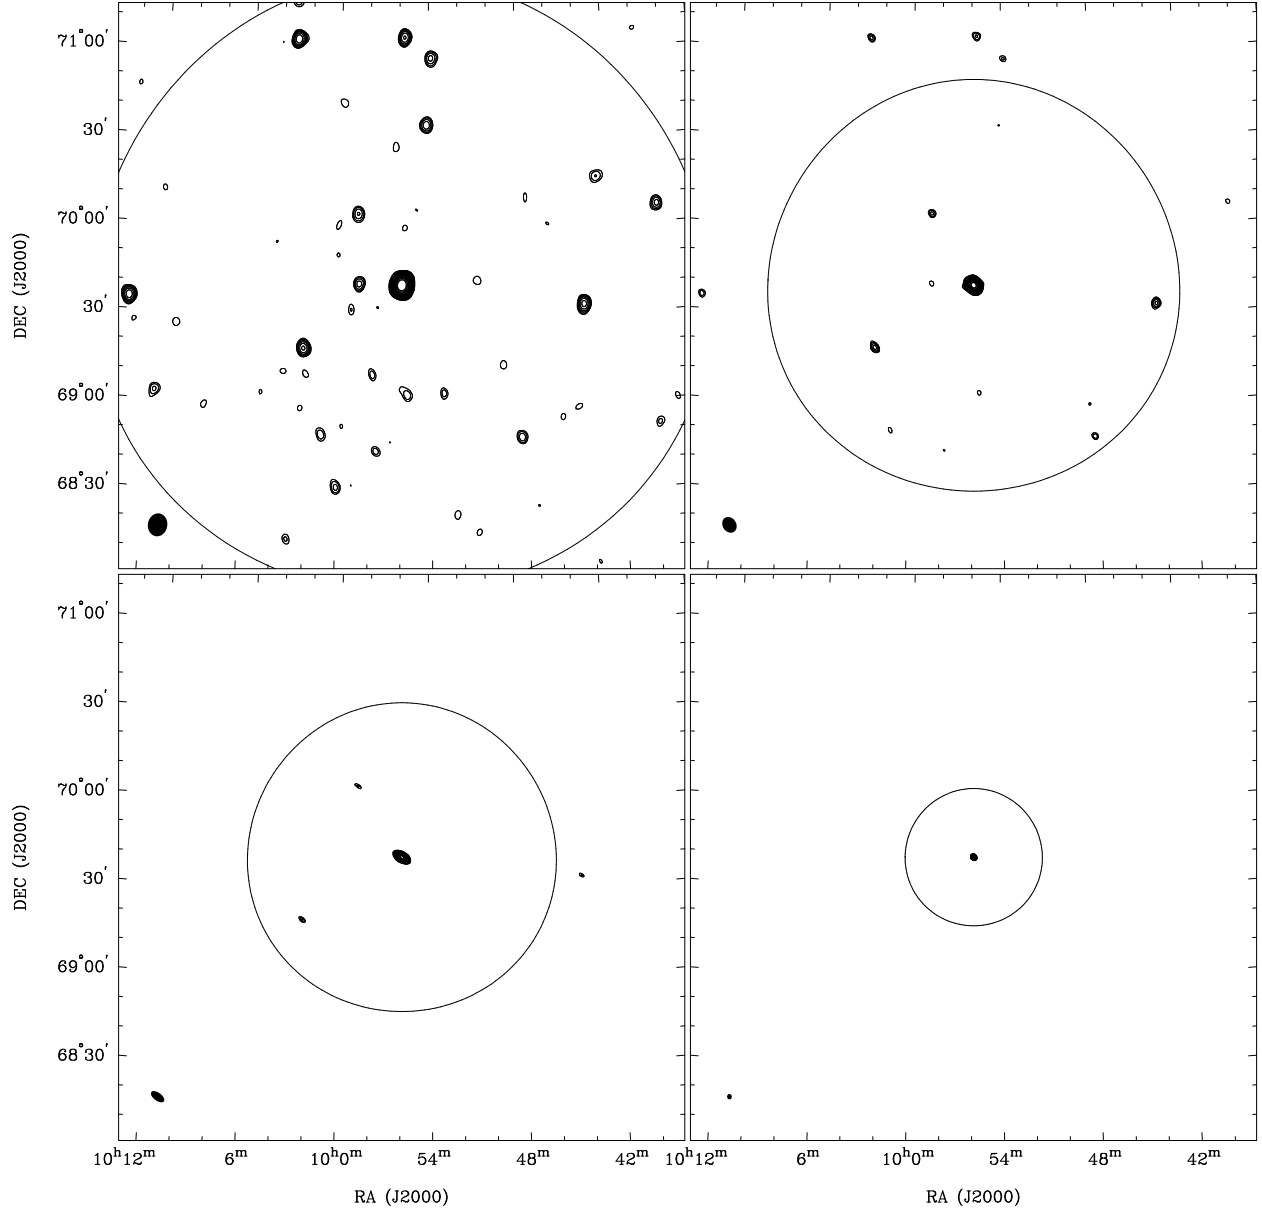


Figure 4.5: ATA maps of M82 at 1.0, 1.5, 2.0, and 4.5 GHz (*top left, top right, bottom left, and bottom right* panels, respectively). The depicted maps are the reference images at each frequency reimaged onto a common coordinate grid. *Circles*: half-power points of the ATA beam. *Filled ellipses*: synthesized beamsizes. Contour levels in all panels are $50 \text{ mJy} \times 2^{n/2}$ for $n = 0, 1, 2, \dots, 13$. The RMS noise levels of the images are 17, 13, 10, and 11 mJy. The number of sources seen decreases as frequency increases due to both the shrinking primary beam of the telescope and the general downward slope of the radio spectra of celestial sources.

observation, including observations of different calibrators. To allow further consistency checks, if multiple calibrator observations could be applied to a given snapshot, it was duplicated and a different calibration was applied to each copy. The copy with the smallest time differential between the two observations was treated as the “official” version of the observation, with the other copies not being included in the final analysis. The median time differential of the “official” observations was 21 minutes, with values ranging from 3 to 38 minutes. The stability of the ATA over such timescales was qualitatively very good. Such stability is particularly noteworthy because of the large angles on the sky separating the calibrator and target sources (see, *e.g.*, Table 4.2).

Once the observations of the science targets were flux and bandpass calibrated, system temperatures were computed and inserted into the datasets using `calctsys`. If the science targets had bright, off-center emission, the `calctsys` technique would have computed incorrect T_{sys} values due to significant visibility variations across the passband, but this was not an issue for our selected sources.

Reference images of the science targets were then generated from the calibrated observations in a manner similar to that described above. A reference image was generated for each observed combination of source, sky frequency, and calibrator. Before imaging, the merged data were phase self-calibrated against a primary-beam-corrected model generated from the NRAO VLA Sky Survey (NVSS; Condon et al. (1998)) catalog. The purpose of this step was to compensate for occasional antenna-based phase glitches that could occur between the calibrator observation and the observation it was applied to. To roughly correct for the fact that the NVSS data only comprise a single frequency, each NVSS model image was scaled by $S_\nu(\nu)/S_\nu(1.4 \text{ GHz})$ before the phase self-calibration, where ν was sky frequency of the ATA observations and $S_\nu(\nu)$ was a simplistic model of the spectrum of the source being observed. The NVSS catalog sources corresponding to NGC 253 was found to poorly reflect its extended structure, so a bootstrapped model of that source’s spatial emission was used rather than the NVSS model. After this first round of phase self-calibration, the u - v data were imaged and recalibrated as described above.

4.3.3 u - v modeling

Fluxes were determined by fitting the u - v data with point-source and elliptical Gaussian models based on the reference images. Our measurements are listed in Tables 4.3–4.5. The calibrators and Arp 220 were all well-modeled as point sources. M82 was modeled as an elliptical Gaussian. The deconvolved source sizes determined for M82 varied slightly over the range of frequencies observed. The major (minor) axis of M82 measured approximately 60 (40) arcsec at 1 GHz and 35 (20) arcsec at 7 GHz, with a roughly constant position angle (PA) of $\sim 80^\circ$ (east from north). The higher-frequency values are consistent with observations made at 1.48, 4.87, 32, and 92 GHz (Seaquist et al. 1985; Klein et al. 1988; Carlstrom & Kronberg 1991). The lower-frequency values are likely biased by the large ($\sim 250''$) synthesized beam of the ATA at low frequencies.

Table 4.3: Flux density measurements of M82

ν (GHz)	Flux (Jy)	Calibrator
1.0	8.92 ± 0.27	3C 48
	8.96 ± 0.16	3C 147
1.1	8.72 ± 0.26	3C 48
	8.46 ± 0.15	3C 147
1.4	7.25 ± 0.11	3C 48
	7.37 ± 0.10	3C 286
1.5	6.85 ± 0.15	3C 48
1.7	6.60 ± 0.09	3C 286
1.8	6.34 ± 0.09	3C 286
2.0	6.06 ± 0.13	3C 48
	6.13 ± 0.08	3C 286
2.1	5.87 ± 0.13	3C 48
2.5	5.39 ± 0.12	3C 48
2.6	5.55 ± 0.17	3C 48
	5.21 ± 0.09	3C 147
2.9	4.89 ± 0.11	3C 48
	4.88 ± 0.11	3C 147
3.0	4.33 ± 0.09	3C 48
3.1	4.70 ± 0.10	3C 48
3.2	4.54 ± 0.07	3C 147
3.3	4.50 ± 0.07	3C 147
3.4	4.44 ± 0.06	3C 286
3.5	4.32 ± 0.06	3C 286
3.6	4.15 ± 0.06	3C 286
4.3	3.85 ± 0.09	3C 48
4.4	3.73 ± 0.06	3C 147
4.5	3.67 ± 0.06	3C 147
4.7	3.55 ± 0.08	3C 48
4.8	3.53 ± 0.08	3C 48
	3.60 ± 0.09	3C 147
5.0	3.38 ± 0.05	3C 48
	3.50 ± 0.07	3C 147
	3.46 ± 0.05	3C 286
5.1	3.16 ± 0.10	3C 48

Table 4.3 (cont'd): Flux density measurements of M82

ν (GHz)	Flux (Jy)	Calibrator
	3.52 ± 0.11	3C 147
5.2	3.34 ± 0.05	3C 286
5.5	3.34 ± 0.05	3C 147
5.6	3.20 ± 0.06	3C 147
6.0	2.96 ± 0.05	3C 286
6.1	2.99 ± 0.10	3C 147
6.2	2.91 ± 0.10	3C 147
6.3	2.98 ± 0.10	3C 147
6.4	2.90 ± 0.10	3C 147
6.6	2.76 ± 0.06	3C 147
6.7	2.86 ± 0.06	3C 147

Table 4.4: Flux density measurements of NGC 253

ν (GHz)	Core	Flux (Jy) Extended	Total	Calibrator
1.1	2.98 ± 0.05	3.92 ± 0.07	6.90 ± 0.08	3C 147
1.2	2.93 ± 0.05	3.83 ± 0.09	6.75 ± 0.10	3C 147
1.3	2.78 ± 0.04	3.51 ± 0.06	6.29 ± 0.07	3C 147
1.4	2.76 ± 0.09	3.17 ± 0.10	5.93 ± 0.13	3C 48
	2.75 ± 0.06	3.22 ± 0.07	5.97 ± 0.09	3C 147
1.5	2.54 ± 0.08	2.59 ± 0.09	5.13 ± 0.13	3C 48
1.7	2.48 ± 0.05	2.68 ± 0.07	5.16 ± 0.09	3C 147
1.8	2.39 ± 0.05	2.50 ± 0.07	4.89 ± 0.08	3C 147
1.9	2.36 ± 0.05	2.42 ± 0.07	4.78 ± 0.08	3C 147
2.0	2.35 ± 0.08	2.29 ± 0.10	4.64 ± 0.12	3C 48
2.1	2.25 ± 0.08	1.96 ± 0.09	4.21 ± 0.12	3C 48
2.6	1.95 ± 0.03	1.80 ± 0.05	3.75 ± 0.06	3C 147
2.7	2.01 ± 0.02	1.74 ± 0.03	3.75 ± 0.04	3C 147
3.0	1.98 ± 0.07	1.76 ± 0.09	3.74 ± 0.12	3C 48
3.1	1.89 ± 0.06	1.70 ± 0.08	3.59 ± 0.11	3C 48
	1.83 ± 0.02	1.51 ± 0.03	3.34 ± 0.04	3C 147
3.3	1.76 ± 0.02	1.40 ± 0.04	3.16 ± 0.04	3C 147
3.6	1.75 ± 0.04	1.38 ± 0.07	3.13 ± 0.08	3C 147
3.7	1.75 ± 0.08	1.15 ± 0.17	2.90 ± 0.19	3C 147
4.5	1.50 ± 0.02	0.97 ± 0.04	2.47 ± 0.04	3C 147
4.6	1.46 ± 0.04	1.06 ± 0.07	2.51 ± 0.08	3C 147
4.7	1.47 ± 0.03	1.00 ± 0.09	2.47 ± 0.10	3C 147
4.9	1.35 ± 0.04	1.01 ± 0.07	2.35 ± 0.08	3C 147
5.0	1.37 ± 0.10	0.65 ± 0.15	2.01 ± 0.18	3C 48
	1.45 ± 0.04	0.81 ± 0.09	2.26 ± 0.10	3C 147
5.1	1.34 ± 0.07	0.90 ± 0.15	2.23 ± 0.17	3C 48
5.6	1.30 ± 0.03	0.90 ± 0.09	2.20 ± 0.10	3C 147
5.7	1.26 ± 0.03	0.75 ± 0.09	2.01 ± 0.09	3C 147
5.8	1.26 ± 0.04	0.61 ± 0.07	1.87 ± 0.08	3C 147
5.9	1.21 ± 0.04	0.76 ± 0.08	1.97 ± 0.09	3C 147
6.0	1.20 ± 0.03	0.75 ± 0.07	1.95 ± 0.08	3C 147
6.1	1.16 ± 0.03	0.60 ± 0.07	1.76 ± 0.07	3C 147
6.3	1.16 ± 0.03	0.52 ± 0.07	1.68 ± 0.08	3C 147

Table 4.4 (cont'd): Flux density measurements of NGC 253

ν (GHz)	Core	Flux (Jy) Extended	Total	Calibrator
6.4	1.18 ± 0.04	0.54 ± 0.09	1.72 ± 0.09	3C 147
6.7	1.13 ± 0.04	0.73 ± 0.15	1.87 ± 0.15	3C 147
7.0	1.04 ± 0.04	0.51 ± 0.07	1.55 ± 0.09	3C 147

Table 4.5: Flux density measurements of Arp 220

ν (GHz)	Flux (Jy)	Calibrator
1.1	0.36 ± 0.02	3C 48
	0.34 ± 0.01	3C 286
1.2	0.37 ± 0.02	3C 48
	0.33 ± 0.01	3C 286
1.4	0.32 ± 0.01	3C 286
1.5	0.26 ± 0.02	3C 286
1.6	0.31 ± 0.03	3C 286
1.7	0.30 ± 0.01	3C 286
1.8	0.31 ± 0.01	3C 286
1.9	0.30 ± 0.01	3C 286
2.0	0.30 ± 0.01	3C 286
2.5	0.27 ± 0.02	3C 48
	0.29 ± 0.01	3C 286
2.6	0.28 ± 0.02	3C 48
	0.28 ± 0.01	3C 286
2.9	0.27 ± 0.01	3C 286
3.2	0.27 ± 0.01	3C 286
3.3	0.27 ± 0.01	3C 286
3.4	0.28 ± 0.01	3C 286
3.5	0.27 ± 0.01	3C 286
3.6	0.24 ± 0.01	3C 286
4.3	0.25 ± 0.02	3C 286
4.4	0.23 ± 0.02	3C 286
4.5	0.24 ± 0.02	3C 286
4.6	0.24 ± 0.02	3C 286
4.7	0.26 ± 0.02	3C 286
5.0	0.24 ± 0.02	3C 48
	0.23 ± 0.01	3C 286
5.2	0.23 ± 0.01	3C 286
5.4	0.24 ± 0.01	3C 286
5.7	0.20 ± 0.02	3C 48
	0.22 ± 0.01	3C 286
6.0	0.22 ± 0.01	3C 286
6.2	0.29 ± 0.03	3C 286

Table 4.5 (*cont'd*): Flux density measurements of Arp 220

ν (GHz)	Flux (Jy)	Calibrator
6.3	0.20 ± 0.03	3C 48
	0.20 ± 0.01	3C 286
6.7	0.23 ± 0.03	3C 286
7.0	0.19 ± 0.02	3C 286

NGC 253 was modeled as the superposition of an elliptical Gaussian and a more compact source. At lower frequencies, the field of view contained bright sources away from the pointing center which were also modeled. The size of the extended emission is strongly frequency-dependent, having a major (minor) axis of 480 (200) arcsec at 1 GHz and 200 (50) arcsec at 7 GHz, with roughly constant PA of $\sim 50^\circ$. The morphology of this region at 1.5 GHz is consistent with the “inner disk” identified in Hummel et al. (1984). One should bear in mind that the spectra presented in Table 4.4 and Figure 4.10 manifest not only the changing brightness of this region but its changing size as well. The more compact source was unresolved at lower frequencies and modeled as a point source. Above 3.5 GHz, however, it was possible for some structure in the compact source to be discerned, and it was found that an elliptical Gaussian with major (minor) axes of ~ 16 (~ 10) arcsec at a PA of $\sim 60^\circ$ was a practical model. This morphology is consistent with the “nuclear peak” seen in the near IR by Scoville et al. (1985).

The residuals to the fitted u - v models were imaged to check for missed sources, poorly calibrated baselines, and so on. Further flagging was usually sufficient to solve any issues. Some datasets with severe RFI were not recoverable and had to be discarded.

Figure 4.6 shows a histogram of the fractional flux residuals of the calibrator snapshots as compared to the analytic calibrator models. The calibrator fluxes are those derived from u - v modeling of calibrator snapshots that have been self-calibrated to the calibrator reference images. The vast majority of the measurements agree to within 1%, demonstrating that the reference-image calibration pipeline propagates fluxes with high fidelity.

4.3.4 Uncertainties

The MIRIAD u - v fitting task, `uvfit`, reports statistical uncertainties, but these fail to account for systematics and are only lower limits. We augmented these uncertainties with both relative and absolute terms:

$$\sigma_i^2 = s_i^2 + (\eta S_{\nu,i})^2 + (\sigma_a \cdot 1 \text{ Jy})^2, \quad (4.3)$$

where σ_i is the uncertainty used for the i^{th} measurement, s_i is the uncertainty for that measurement reported by the fitting routine, $S_{\nu,i}$ is the value of the measurement, and η and

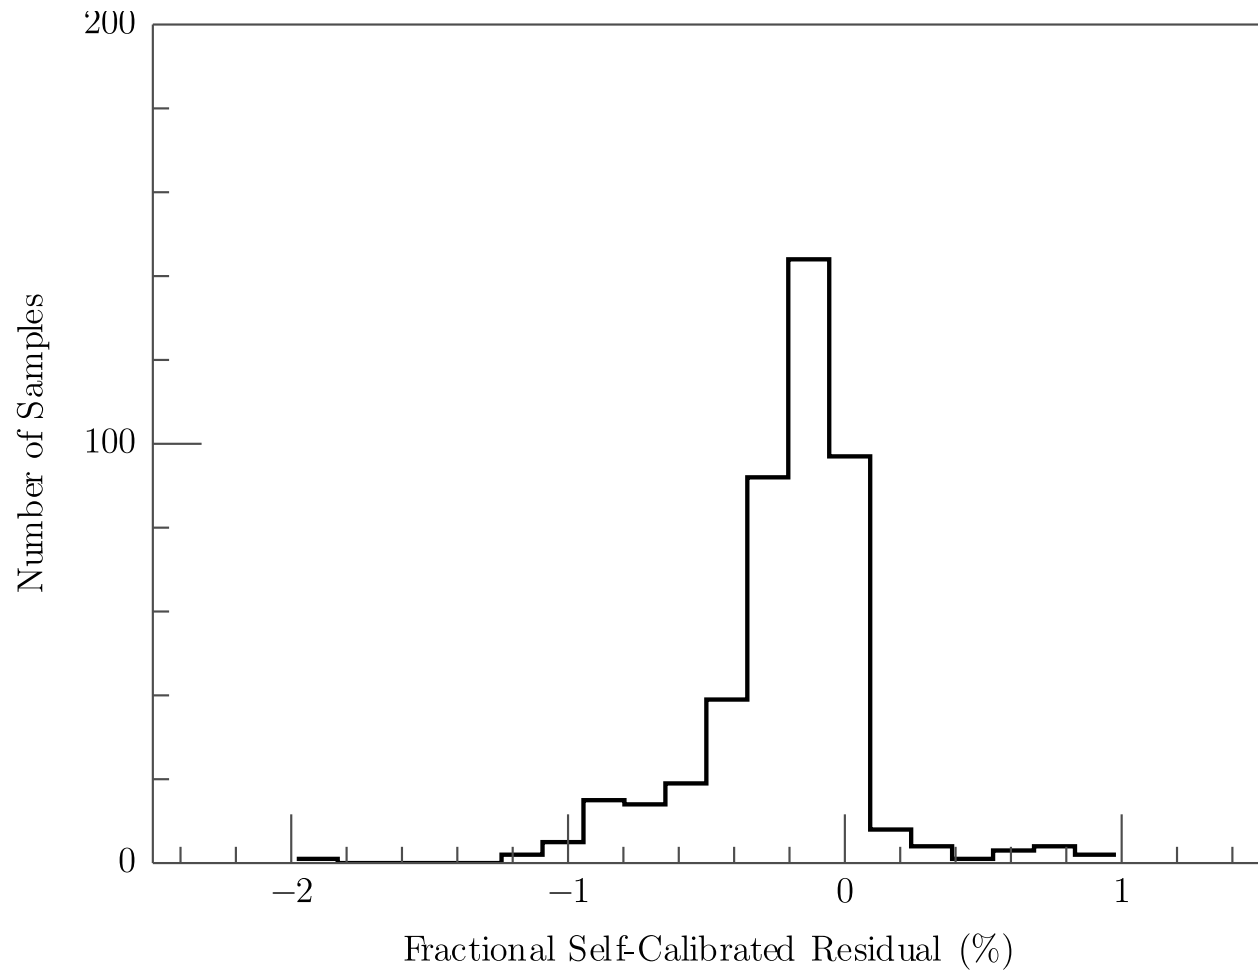


Figure 4.6: Histogram of fractional flux residuals between calibrator observations and analytic calibrator models. The fluxes for the calibrator observations are derived from u - v modeling of calibrator data that have been self-calibrated to the calibrator reference images.

σ_a are tunable parameters.

The choice of these parameters was derived from the statistics of the ensemble of observations. For each observed combination of source, frequency, and calibrator, the weighted mean of the contributing observations was computed:

$$\bar{S}_\nu = \sum_{i=1}^n S_{\nu,i} \sigma_i^{-2} / \sum_{i=1}^n \sigma_i^{-2}. \quad (4.4)$$

Let there be N observed combinations of source, frequency, and calibrator, with n_j observations of the j^{th} combination. Let $\bar{S}_{\nu,j}$ be the weighted mean of the n_j observations, and $S_{\nu,j,i}$ and $\sigma_{j,i}$ the values and uncertainties of the individual observations. If the uncertainties have been correctly assessed and the measurements are drawn from normal distributions, the sum of the squares of the normalized residuals,

$$K^2 = \sum_{j=1}^N \sum_{i=1}^{n_j} \left(\frac{S_{\nu,j,i} - \bar{S}_{\nu,j}}{\sigma_{j,i}} \right)^2, \quad (4.5)$$

should be taken from a χ_k^2 distribution, where $k = \sum_{j=1}^N n_j - N$ is the number of degrees of freedom. The expectation value of such a distribution is k .

The value of K^2 was computed on a grid of η and σ_a for three groups of measurements: the full ensemble, those with $\bar{S}_\nu < 1$ Jy (“faint”), and those with $\bar{S}_\nu > 3$ Jy (“bright”). Figure 4.7 shows contours of $K^2 = k$ (*i.e.*, uncertainties agreeing with observed scatter) in the η - σ_a plane. As expected, the contours of constant K^2 are primarily along the η axis for the faint group (for which σ_a dominates the additional uncertainty) and primarily along the σ_a axis for the bright group (for which η dominates the additional uncertainty). The $K^2 = k$ contours for the three groups intersect at an approximate location of $\eta = 0.030$ and $\sigma_a = 0.014$, *i.e.* an additional 3.0% fractional uncertainty and 14 mJy absolute uncertainty. These are the values of η and σ_a that were adopted. The value of σ_a is approximately equal to the thermal noise in good conditions and is about half of the thermal noise at high frequencies, where T_{sys} generally increases.

If the uncertainties are correctly assessed and measurements are drawn from normal distributions, the sum of the normalized residuals of each measurement is drawn from a normal distribution with mean 0 and variance $\sigma^2 = \sum_{j=1}^N n_j$. For the chosen parameter values, the sum of the residuals was 0.52σ . This moderate value suggests that the adopted uncertainties characterize well the observed scatter in the measurements and that the weighted-mean approach does not mask large variations in the individual measurements.

4.3.5 Spectra

We constructed broadband, high-resolution spectra out of the flux measurements described in §4.3.3. Repeated measurements were combined by taking a weighted mean as described

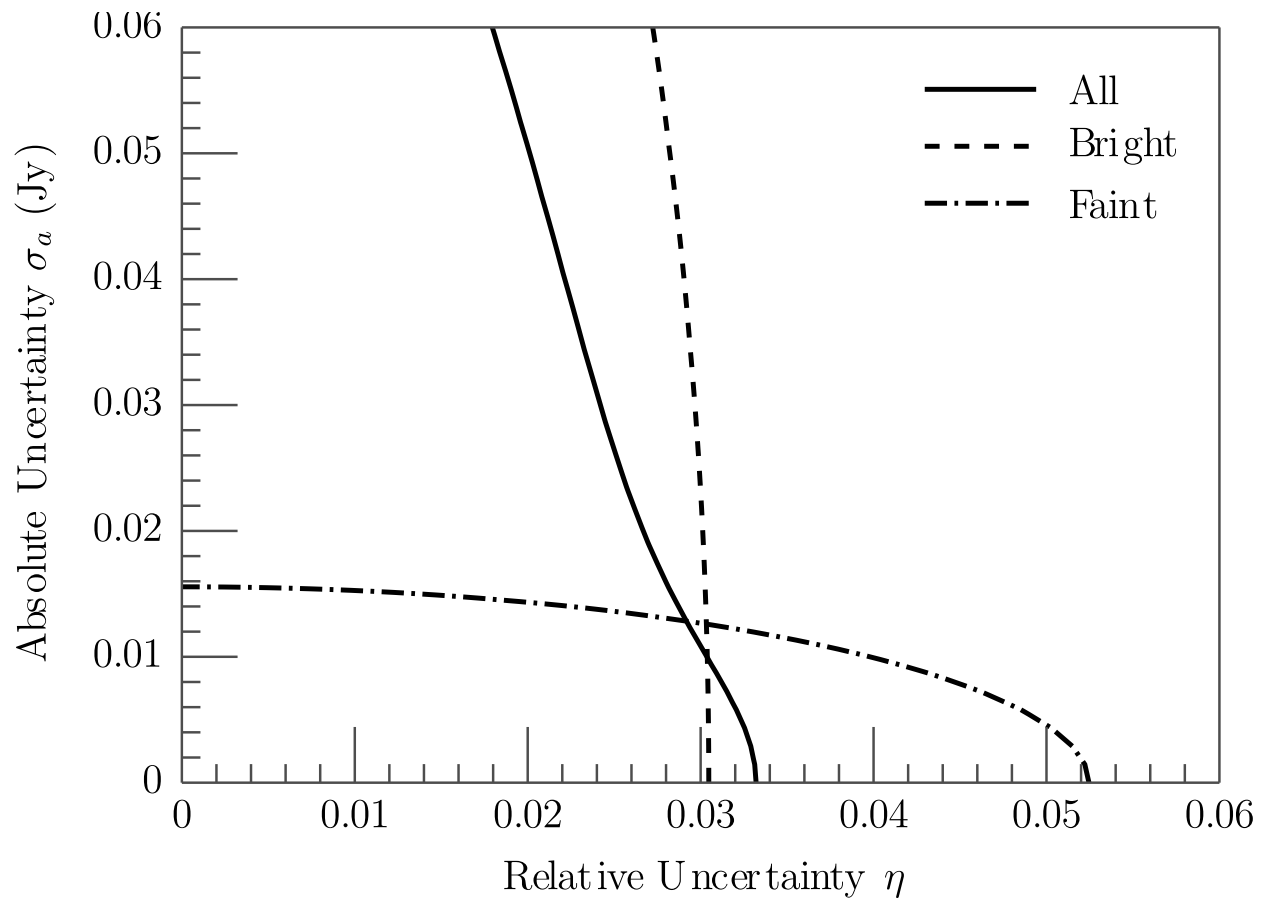


Figure 4.7: Contours of $K^2 = k$ in the η - σ_a space of uncertainty augmentation parameters. See §4.3.4 for discussion. *Solid line*: the contour for K^2 values derived from all repeated flux measurements. *Dashed line*: the contour for values derived from measurements of sources with $S_\nu > 3$ Jy. *Dot-dashed line*: the contour for values derived from measurements of sources with $S_\nu < 1$ Jy.

in the previous section. Twenty-five percent of the spectral data are based on a single observation, while 63% are based on at least three observations. The maximal number of contributing observations is 13, of M82 at 5.0 GHz. The uncertainty of each spectral point is the propagated uncertainty of the weighted mean,

$$\sigma_{S,j}^2 = \left(\sum_{i=1}^{n_j} \sigma_{j,i}^{-2} \right)^{-1}. \quad (4.6)$$

The spectra are shown in Figures 4.8-4.11 along with measurements from the literature and model results as discussed in §4.4.

4.3.6 Self-consistency Checks

We performed no absolute flux calibration of the ATA. The measurements we have made are thus measurements of flux density ratios, with absolute flux densities being derived from *a priori* spectral models of the calibrator sources. We cannot verify the models themselves, but we can assess whether our measured flux ratios are consistent with them.

Let the “true” flux density of source X (*i.e.*, the value that would be measured by a perfect, absolutely calibrated antenna at the time of observation) be T_X , the *a priori* modeled flux density of X be M_X , and the flux density of X observed by the ATA after calibration to source Y be $O_{X,Y}$. All of these values vary with frequency. In the simplest possible model, at a fixed frequency $O_{X,Y}$ and T_X are directly related by a calibrator-dependent gain factor g_Y :

$$O_{X,Y} = g_Y T_X. \quad (4.7)$$

The gain is self-calibrated such that $O_{Y,Y} \equiv M_Y$, so that ideally

$$O_{X,Y} = M_Y \frac{T_X}{T_Y} \equiv M_Y R_{X,Y}^T, \quad (4.8)$$

where $R_{X,Y}^T$ is the true flux density ratio of X to Y . Not knowing whether this simple model is correct, we will define the observed flux ratio as

$$R_{X,Y}^O \equiv \frac{O_{X,Y}}{M_Y}. \quad (4.9)$$

If this simple calibration model holds, the observed ratios should demonstrate consistency and closure:

$$R_{X,Y}^O \cdot R_{Y,X}^O = 1, \quad (4.10)$$

$$R_{X,Y}^O \cdot R_{Y,Z}^O \cdot R_{Z,X}^O = 1. \quad (4.11)$$

Furthermore, if the spectral models are accurate, it should be found that

$$R_{X,Y}^O = R_{X,Y}^M \equiv M_X/M_Y. \quad (4.12)$$

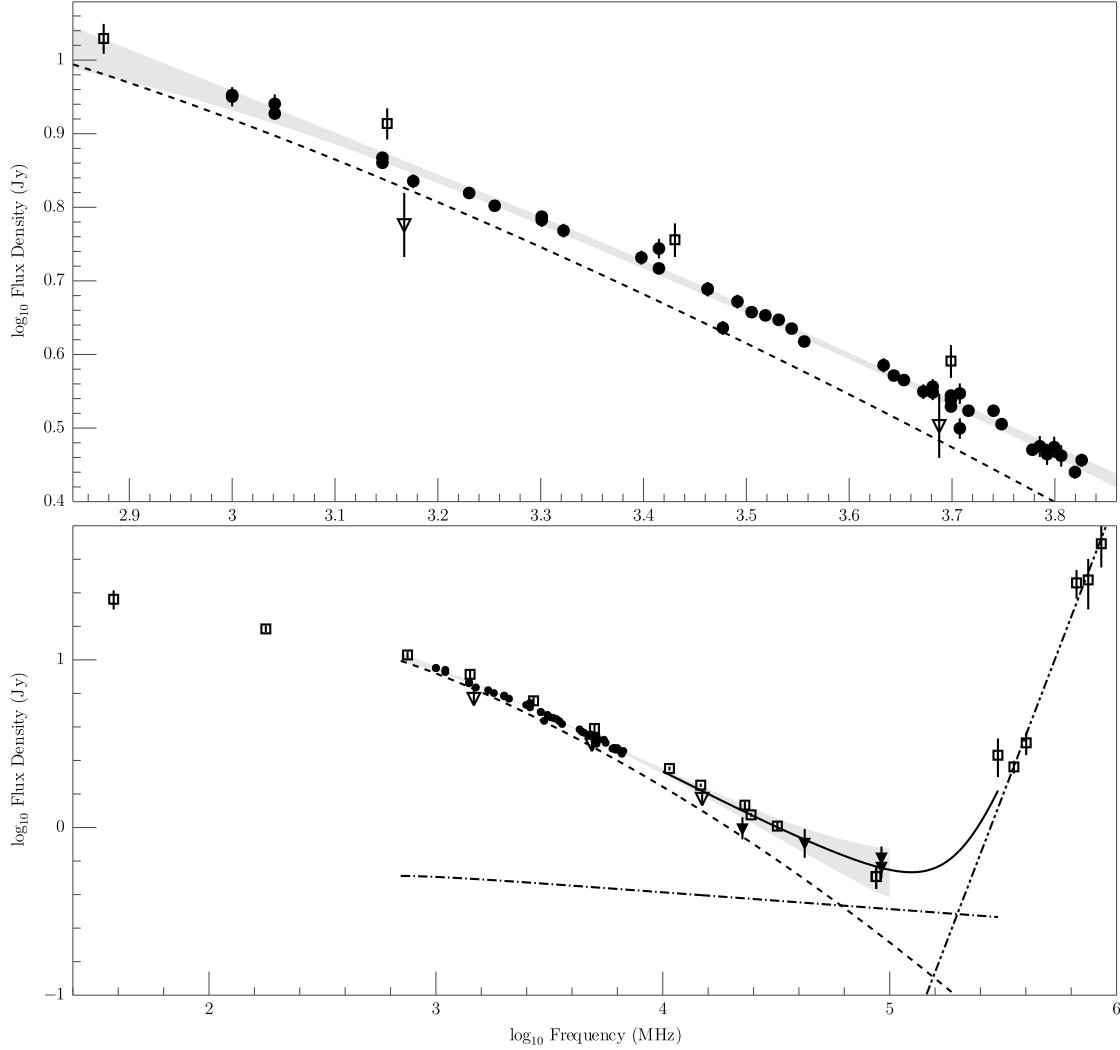


Figure 4.8: Spectrum of M82 with model. The two panels show different views of the same data. *Circles:* new measurements presented in this chapter. Other symbols indicate data from the literature: [Carlstrom & Kronberg \(1991\)](#) and citations therein; [Klein et al. \(1988\)](#) and citations therein; [Jura et al. \(1978\)](#); [Kronberg et al. \(1979\)](#); [Leeuw & Robson \(2009\)](#); [Rodríguez-Rico et al. \(2004\)](#); [Seaquist et al. \(1996\)](#). These measurements are made with various telescopes in various configurations and, due to the properties of interferometers, are not always directly comparable. *Squares:* data obtained with single dishes. *Triangles:* data obtained with interferometers. *Filled symbols:* data used in the fits described in §4.4.1. *Open symbols:* data not included in the fits. *Shaded region:* the 3σ uncertainty region of the model spectrum as described in §4.4.2. *Dashed line:* the model nonthermal contribution. *Dash-dotted line:* the model thermal contribution. *Dash-double-dotted line:* a power-law fit to the FIR data. *Solid line:* the sum of the thermal, nonthermal, and FIR power law contributions. See §4.4.2 for further discussion of the model results.

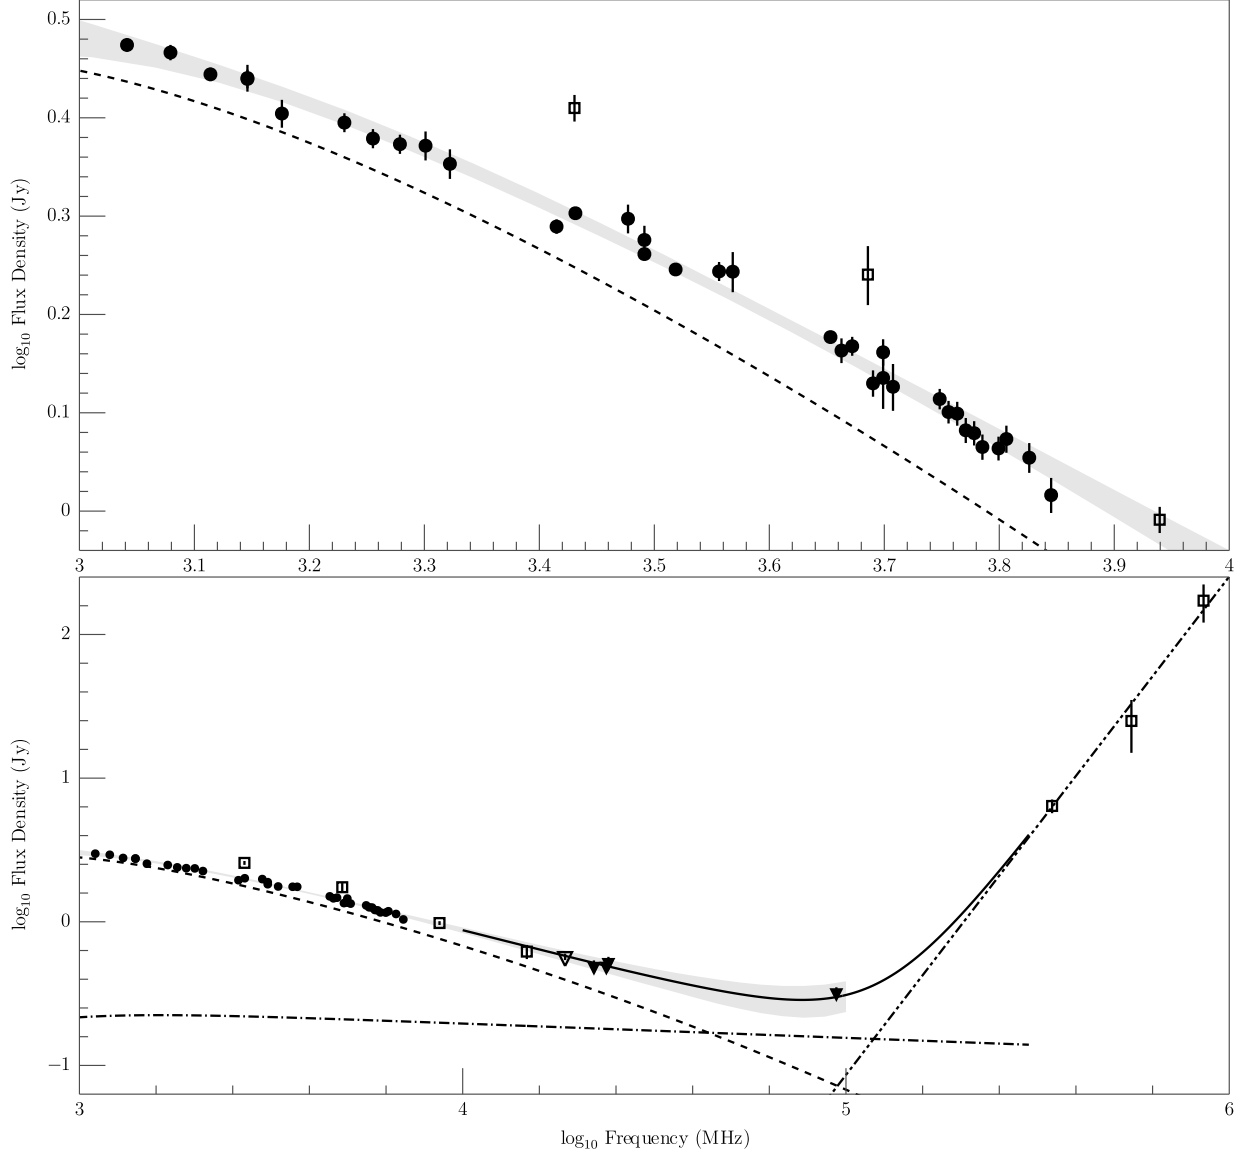


Figure 4.9: Spectrum of the core component of NGC 253 with model. The symbols are as in Figure 4.8. External data are found in [Beck et al. \(1979\)](#); [Hildebrand et al. \(1977\)](#); [Ott et al. \(2005\)](#); [Peng et al. \(1996\)](#); [Ricci et al. \(2006\)](#); [Rieke et al. \(1973\)](#); [Takano et al. \(2005\)](#); [Weiss et al. \(2008\)](#). See §4.4.2 for discussion.

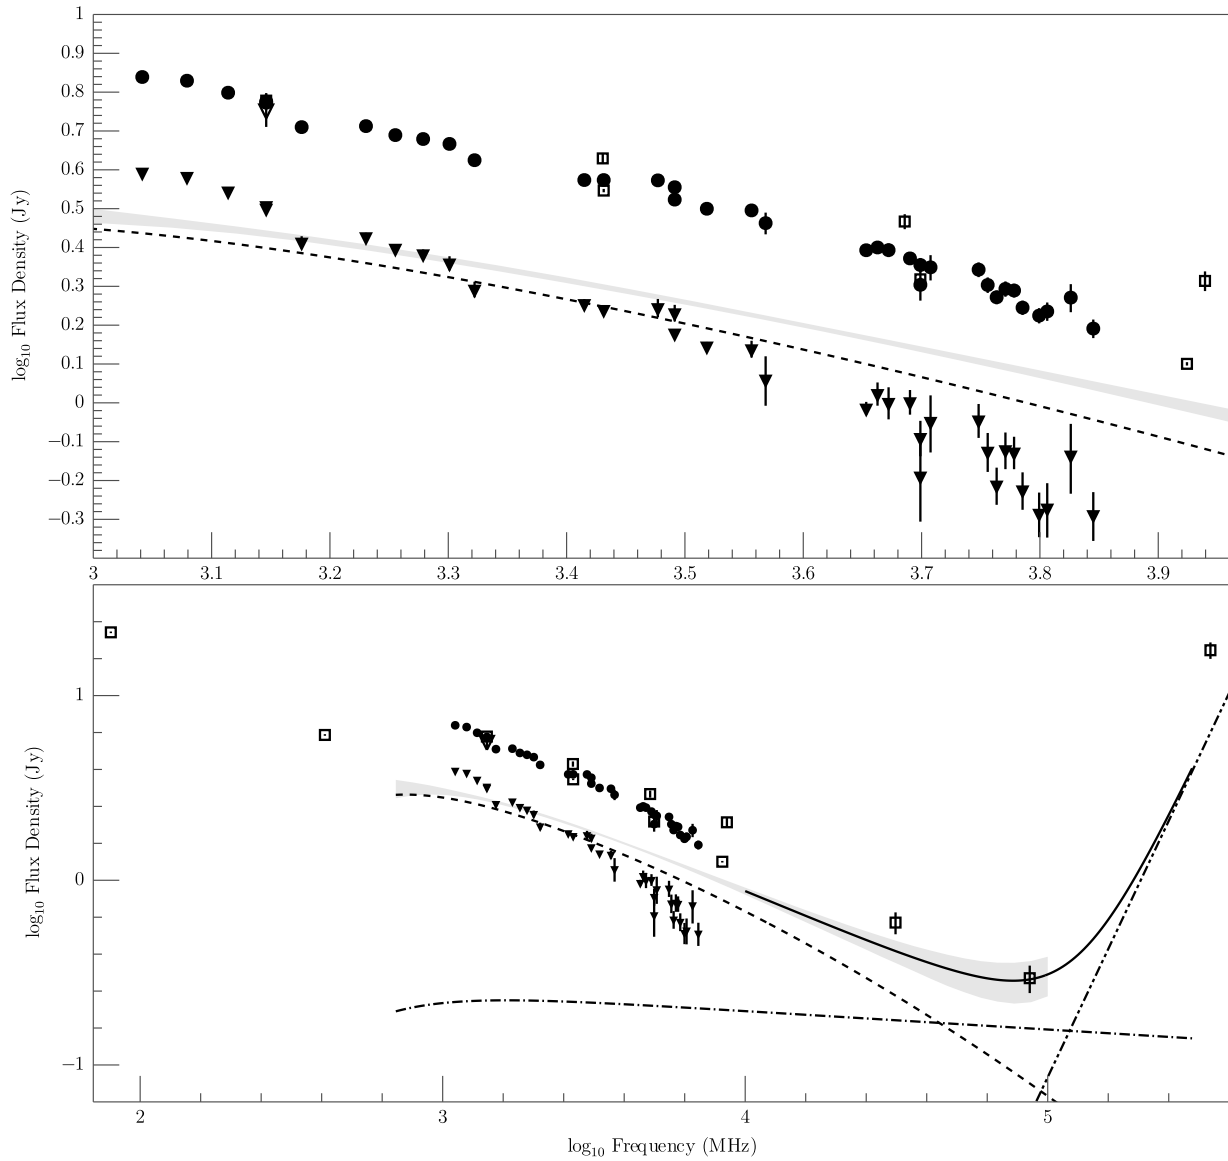


Figure 4.10: Spectra of the extended component and total flux of NGC 253. The model used in Figure 4.9 is overlaid to ease comparison. External data are found in [Beck et al. \(1979\)](#); [Geldzahler & Witzel \(1981\)](#); [Jura et al. \(1978\)](#); [Wright & Otrupcek \(1990\)](#); [Strickland et al. \(2004\)](#); [Weiss et al. \(2008\)](#). *Triangles:* fluxes of the extended component of NGC 253. *Circles:* sum of extended and core component fluxes. Other symbols are as in the previous plot.

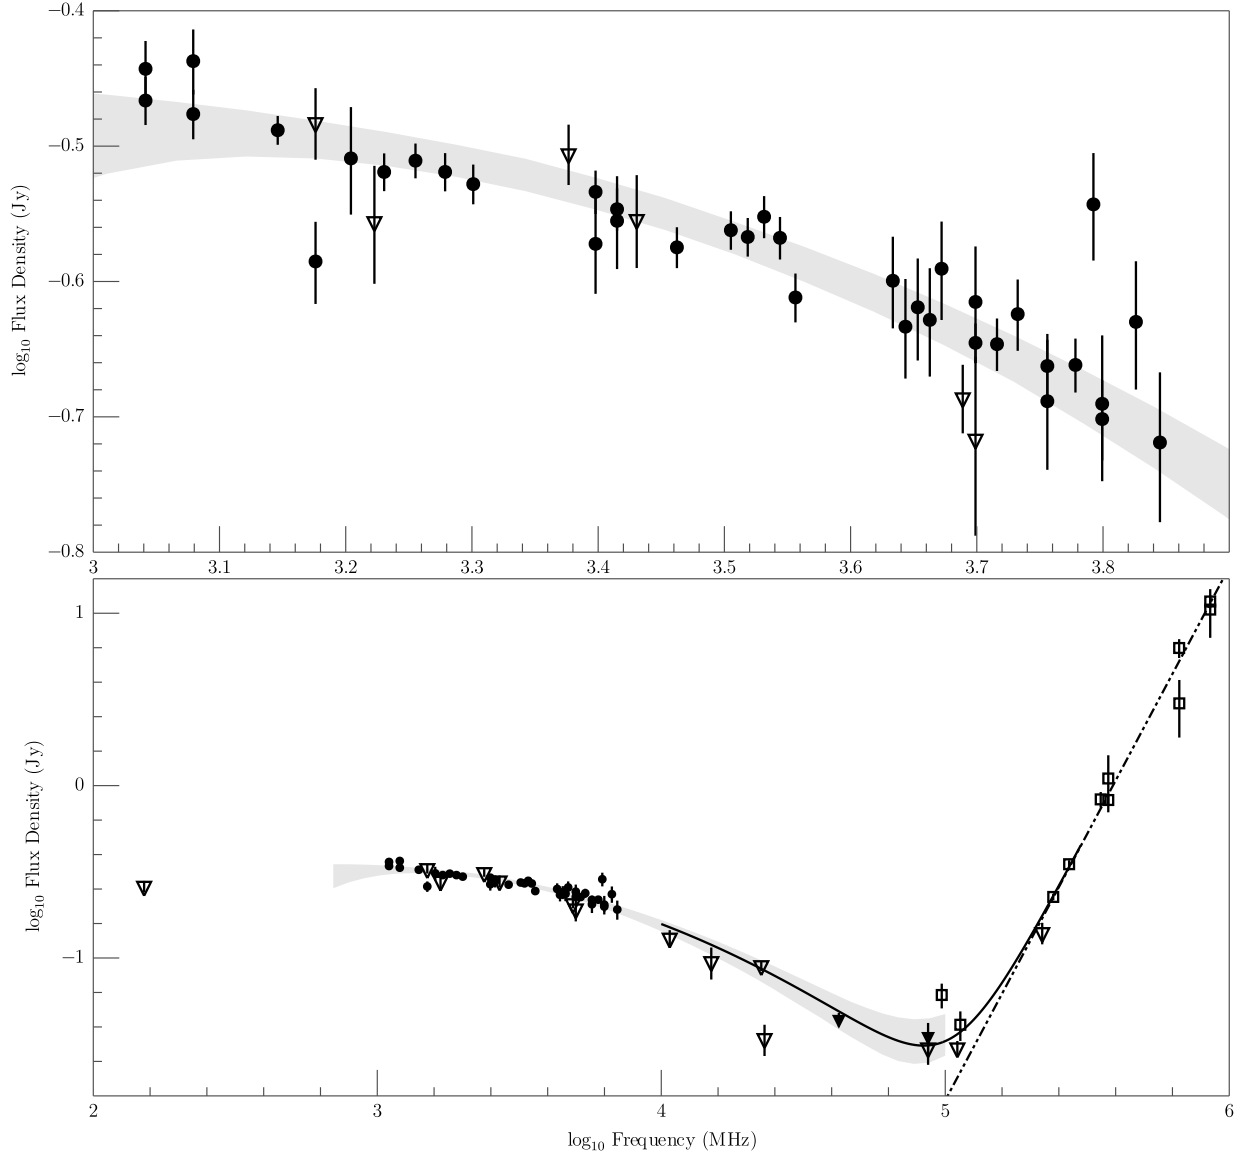


Figure 4.11: Spectrum of Arp 220 with model. The symbols are as in Figure 4.8. See §4.4.2 for further discussion. External data are found in [Anantharamaiah et al. \(2000\)](#); [Carico et al. \(1992\)](#); [Dunne & Eales \(2001\)](#); [Eales et al. \(1989\)](#); [Imanishi et al. \(2007\)](#); [Norris \(1988\)](#); [Radford et al. \(1991\)](#); [Rigopoulou et al. \(1996\)](#); [Rodríguez-Rico et al. \(2005\)](#); [Scoville et al. \(1991\)](#); [Sopp & Alexander \(1991\)](#); [Woody et al. \(1989\)](#); [Zhao et al. \(1996\)](#).

(If M_X is not known precisely, two measurements $R_{X,Y}^O$ and $R_{X,Z}^O$ can still be used to check the model ratio $R_{Y,Z}^M$ by computing $R_{Y,Z}^O = R_{X,Z}^O/R_{X,Y}^O$.) If all three of these properties are observed, there is good reason to believe that $R_{X,Y}^O = R_{X,Y}^T$ and $O_{X,Y} = T_X$, though it is impossible to demonstrate the latter without absolute flux measurements.

We compute these ratios and plot them in Figure 4.12. Our measurements do not have enough precision for us to believably solve for spectral models of the calibrators from scratch. We do, however, solve for a change in the overall scaling of the VLA 1999.2 models, *i.e.*, for three parameters f_{48} , f_{147} , and f_{286} such that $M_X = f_X V_X$, where V_X is the VLA model of source X and we denote the calibrators by their 3C numbers. Given the constraints implied by Equations 4.10 and 4.11, there are only two independent $R_{X,Y}^O$, so a normalizing assumption must be applied to derive the three f_X from observations. We choose to fix $f_{48}f_{147}f_{286} = 1$, which is equivalent to requiring that the total relative change to the models, $\ln^2 f_{48} + \ln^2 f_{147} + \ln^2 f_{286}$, be minimized.

To solve for the f_X from the data, we define $a = f_{48}/f_{147}$ and $b = f_{48}/f_{286}$. Minimizing the logarithms of $R_{X,Y}^M/R_{X,Y}^O$ in a least-squares sense leads to $a = 1.018$ and $b = 1.007$ with a root reduced χ^2 parameter ($\sqrt{\chi_r^2}$) of 1.07. This minimization included two precise measurements made at the VLA in 2008 September: $R_{48,286}^O = 1.066 \pm 0.003$ at 1.465 GHz and $R_{48,286}^O = 0.7349 \pm 0.002$ at 4.885 GHz (R. Perley, private communication). Using our normalizing assumption, we find $f_{48} = 1.011$, $f_{147} = 0.989$, and $f_{286} = 1.001$. This is consistent with evidence from the VLA that the flux density of 3C 48 has recently risen slightly (R. Perley, private communication).

Our measured flux ratios are well-described within their uncertainties by the rescaled VLA models with the exception of the $R_{147,286}^O$ data. These observations were made in June 2009, well after the main observing for this work was completed, after it was realized that there were only a few data comparing these calibrators. Unlike the 2008 observations, these occurred in the daytime, and solar interference was a significant problem with the low-frequency data. Nevertheless the high-frequency ratios, for which solar interference was much less relevant, disagree significantly with the adopted models. We have not been able to identify the cause of this disagreement, but we do not believe that it is due either to genuine inaccuracy in the models or to problems in the calibration and u - v modeling pipeline. If the former were the case, it would represent an unprecedented variation in the flux of some combination of 3C 286 and 3C 147. If the latter were the case, the consistency property of Equation 4.10 would not hold as well as it does in the observations. We believe the robustness of the rest of the results provide ample reason to believe that the models and reduction pipeline are generally robust and accurate.

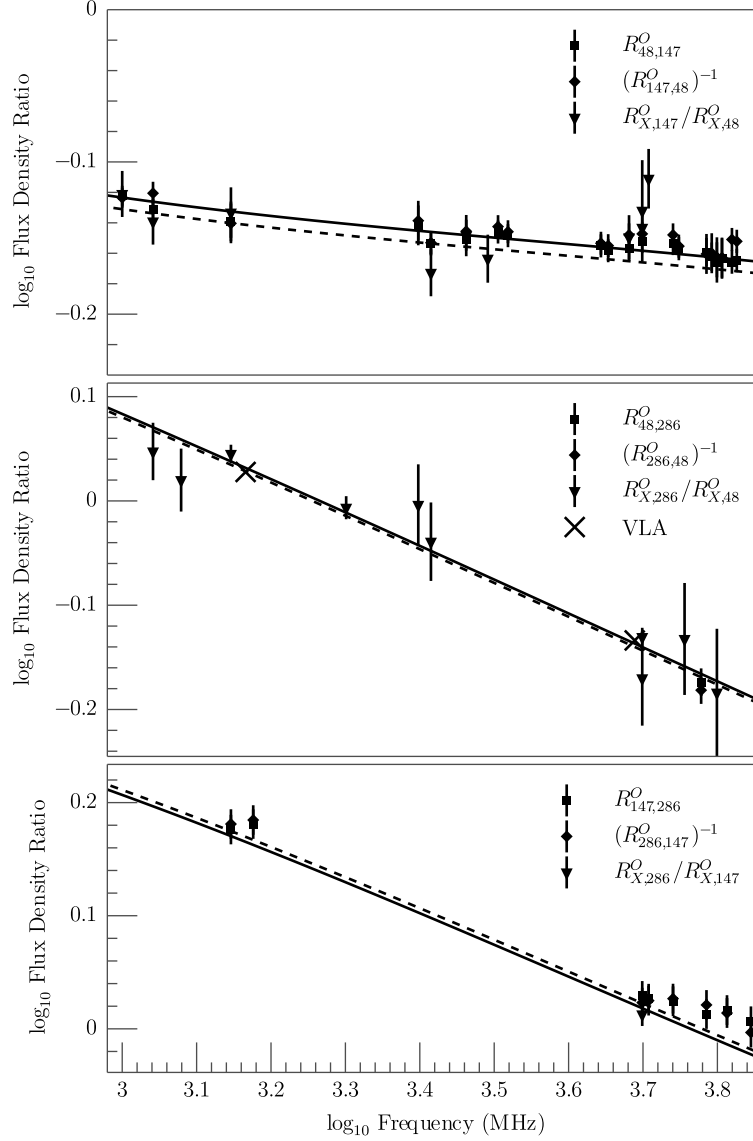


Figure 4.12: Flux density ratios of calibrator observations and spectral models. See §4.3.6 for definitions. *Dashed lines*: ratios of the VLA 1999.2 models. *Solid lines*: ratios of the spectral models adopted in this work. *Squares and diamonds*: directly observed calibrator ratios. The good agreement between the original and inverted ratios indicates that there is negligible flux loss in the secondary calibration process (§4.3.2). *Triangles*: calibrator ratios computed indirectly from the ratio of two observations of a science target. *Xs*: measurements of $R_{48,286}^O$ made at the VLA in 2008 September. The observations of $R_{147,286}^O$ were made in 2009 June and exhibit discrepant behavior; see §4.3.6 for discussion.

4.4 Spectral Modeling

In order to assess the steepening of the nonthermal components of the observed spectra, we modeled them with a two-component model similar to that used in [Niklas et al. \(1997\)](#):

$$S_\nu = S_{\text{th}}(\nu) + S_{\text{nt}}(\nu), \quad (4.13)$$

where $S_{\text{th}}(\nu)$ is the thermal contribution and $S_{\text{nt}}(\nu)$ is the nonthermal contribution. Free-free absorption (FFA) is likely to play a significant role in determining the spectra of our target galaxies. We attempt to account for its effects by including as a parameter in our model τ_1 , the free-free optical depth at 1 GHz. The thermal component is then

$$S_{\text{th}}(\nu) = B_\nu(T_e)\Omega(1 - e^{-\tau_{\text{ff}}}), \quad (4.14)$$

where $T_e \simeq 10^4$ K is the temperature of the thermal electrons, Ω is the angular size of the emitting region, and τ_{ff} is the free-free optical depth as a function of frequency. We use the canonical frequency dependence of the free-free optical depth,

$$\tau_{\text{ff}} = \tau_1 \nu_g^{-2.1}, \quad (4.15)$$

where $\nu/\text{GHz} = \nu_g$. In typical cases, the large-scale thermal emission in galaxies at cm wavelengths is optically thin (*e.g.* [Duric et al. 1988](#); [Price & Duric 1992](#)), $\tau_{\text{ff}} \ll 1$, in which case the thermal emission has the well-known frequency dependence of $S_{\text{th}} \propto \nu^{-0.1}$. The galaxies we consider are, however, decidedly atypical, with dense, compact emission regions, so we avoid this approximation in our modeling. Our results indicate, however, that it broadly holds even for the galaxies in our sample (cf. Figures 4.8-4.11).

For simplicity we assume that the thermal and nonthermal components are well-mixed and hence share the same τ_{ff} . If this is not the case (as is true of the Milky Way, *e.g.* [Beck & Reich 1985](#)), this assumption would require loosening in a more detailed analysis, although at small values of τ_{ff} its impact is minimal. In the well-mixed case, the nonthermal emission is attenuated so that

$$S_{\text{nt}}(\nu_g) = \tilde{S}_{\text{nt}}(\nu_g) \cdot \frac{1 - e^{-\tau_{\text{ff}}}}{\tau_{\text{ff}}}, \quad (4.16)$$

where \tilde{S}_{nt} is the unabsorbed nonthermal spectrum. Because our fits model the thermal and nonthermal contributions simultaneously and we do not approximate any FFA-related effects, we should obtain good results even when τ_1 is large. We model the unabsorbed nonthermal component as

$$\tilde{S}_{\text{nt}}(\nu_g) = 10^{A+B \log_{10} \nu_g + C[\log_{10} \nu_g]^2} \text{Jy}, \quad (4.17)$$

a parabola in log-log space. Our parabolic spectral model is not intended to be well-motivated physically, but rather to be simple and flexible. The theoretical spectra modeled in [T+06](#), to which we compare our results in §4.5, are derived numerically and hence do not suggest a more appropriate analytic expression for us to employ. While [Condon \(1992\)](#) and others

typically use broken power laws to model aged synchrotron spectra, one of the key contentions of T+06 is that other loss processes will significantly modify raw synchrotron spectra in starbursts. We note that functional forms analogous to ours have proven to be good empirical models in works such as Baars et al. (1977) and the VLA broadband calibrator models. The parabolic model has the additional benefits of being linear in its parameters (in log-space) and being reducible to a pure power-law spectrum (ignoring FFA) simply by fixing $C = 0$.

An alternative model for the effect of FFA upon the nonthermal emission is one in which the absorption screens but is not mixed with the nonthermal emitting medium,

$$S_{\text{nt}}(\nu_g) = \tilde{S}_{\text{nt}}(\nu_g) \cdot e^{-\tau_{\text{ff}}}. \quad (4.18)$$

We disfavor this model for our targets because they are starbursts, and their diffuse synchrotron emission propagates through a compact ISM riddled with H II regions and supernova remnants (*e.g.* McDonald et al. 2002; Rodríguez-Rico et al. 2004, regarding M82). In the analysis that follows we will, however, occasionally compare results for the two absorption models.

4.4.1 Isolating the thermal emission

Excepting those with active nuclei, most galaxies have a nonthermal spectrum much steeper than the aforementioned $\alpha = -0.1$ of optically-thin thermal free-free emission. Because the nonthermal emission dominates at cm wavelengths, good higher-frequency data are needed to usefully constrain the thermal component. However, above $\nu \simeq 100$ GHz, thermal emission from cold dust begins contributing non-negligibly to the total flux. We incorporated flux densities from the literature with $20 < \nu/\text{GHz} < 100$ into our fits for this reason. Table 4.6 summarizes the parameters of some of the relevant observations found in the literature. Those observations, and others not listed in table 4.6, are shown in Figures 4.8-4.11 for comparison to our observations.

Comparing fluxes in this way requires care because interferometers resolve out spatially extended emission. The largest angular scale (LAS) to which an interferometer observation is sensitive is inversely proportional to its shortest baseline as measured in wavelengths. In a very real way, two different interferometer observations, or an interferometer observation and a single-dish observation, are not “looking at” the same source. We deal with these considerations on a source-by-source basis below.

M82

M82 is well-observed in the 20–100 GHz range with both interferometers and single dishes. The interferometric fluxes are systematically smaller than the single-dish fluxes by about 25%, with the exception of the unusually low 87.2 GHz single-dish measurement of Jura et al. (1978). Unlike the other single-dish data, which come from maps, this observation was made with a single pointing of an 11 m dish (HPBW 75"). It furthermore needs conversion to the flux density scale of Baars et al. (1977). We regard it as an outlier and discard it.

Table 4.6: High-frequency observations. See §4.4.1 for discussion.

ν (GHz)	Flux (Jy)	Telescope	Tel. Kind ^a	θ_{LAS}^b (arcsec)	Reference ^c
M82					
22.4	1.00 ± 0.15	VLA D-cfg.	I	45	1
23	1.36 ± 0.07	Effelsberg 100m	SD		2
24.5	1.19 ± 0.03	Effelsberg 100m	SD		3
32	1.02 ± 0.06	Effelsberg 100m	SD		3
42.2	0.82 ± 0.16	VLA C-cfg.	I	43	4
87.2	0.51 ± 0.08^d	Kitt Peak 11m	SD		5
92	0.59 ± 0.09	BIMA	I	50	1
92	0.67 ± 0.10	OVRO mm Arr.	I	55	6
NGC 253 Core Component					
22	0.49 ± 0.05	ATCA EW367-cfg.	I	61	7
23.7	0.56 ± 0.06	VLA DnC-cfg.	I	60	8
24	0.52 ± 0.05	ATCA EW367-cfg.	I	56	9
94	0.32 ± 0.03	BIMA A/B/C-cfg.	I	103	10
Arp 220					
22.5	0.090 ± 0.006	VLA A-cfg.	I	2	11
23.1	0.034 ± 0.007^e	VLA A-cfg.	I	2	12
42.2	0.044 ± 0.004	VLA C-cfg.	I	43	13
87	0.030 ± 0.006^e	RAINBOW AB-cfg.	I	8	14
87	0.035 ± 0.007^e	IRAM PdBI	I	40	15
97.2	0.061 ± 0.010	IRAM 30m	SD		16

^aThe letter “T” denotes an interferometer, “SD” denotes a single dish.^bLargest angular scale to which the observation is sensitive.^c1: [Carlstrom & Kronberg \(1991\)](#), 2: [Kronberg et al. \(1979\)](#), 3: [Klein et al. \(1988\)](#), 4: [Rodríguez-Rico et al. \(2004\)](#), 5: [Jura et al. \(1978\)](#), 6: [Seaquist et al. \(1996\)](#), 7: [Ricci et al. \(2006\)](#), 8: [Takano et al. \(2005\)](#), 9: [Ott et al. \(2005\)](#), 10: [Peng et al. \(1996\)](#), 11: [Zhao et al. \(1996\)](#), 12: [Norris \(1988\)](#), 13: [Rodríguez-Rico et al. \(2005\)](#), 14: [Imanishi et al. \(2007\)](#), 15: [Radford et al. \(1991\)](#), 16: [Anantharamaiah et al. \(2000\)](#)^dConverted to flux density scale of [Baars et al. \(1977\)](#).^eUncertainty of 20% adopted.

All of the interferometer observations we consider (Carlstrom & Kronberg 1991; Rodríguez-Rico et al. 2004; Carlstrom & Kronberg 1991; Seaquist et al. 1996) have LASs of $\sim 50''$. The deconvolved size of the radio emission of M82 is $\sim 35'' \times \sim 10''$ from cm wavelengths (this work) up to 92 GHz (Klein et al. 1988; Carlstrom & Kronberg 1991), so significant flux from M82 should not be resolved out; the source of the excess emission detected by the single-dish observations is unclear. For safety, we compare the ATA fluxes only to other interferometric measurements. This criterion is coincidentally equivalent to the selection of only post-1990 observations.

NGC 253

The extended spatial structure of the emission of NGC 253, as detected by the ATA, makes the modeling of its spectrum more difficult. We do not expect the high-frequency interferometer observations of NGC 253 listed in Table 4.6 to be sensitive to the extended component because of its steep spectral index and the constrained LASs of the observations. Simple simulations of high-frequency VLA observations of the galaxy support this conclusion. Two of the FIR observations of NGC 253 (Rieke et al. 1973; Hildebrand et al. 1977) are made with single dishes with $\sim 65''$ beams that could be sensitive to the extended emission region; however, Weiss et al. (2008) resolve the central component from the extended source and find that emission from the former dominates the latter. Because of these facts, our spectra link the FIR and high-frequency measurements with fluxes for only the central component of NGC 253.

Arp 220

The observations of Arp 220 made by Imanishi et al. (2007), Norris (1988) and Zhao et al. (1996) have LASs that are too small for our purposes and are not used in our fits. The 97 GHz single-dish observation of Anantharamaiah et al. (2000) is significantly discrepant from 87 GHz observations made with the Plateau de Bure Interferometer (Radford et al. 1991), the Nobeyama Millimeter Array (Imanishi et al. 2007), and a 110 GHz measurement made with the Owens Valley Radio Observatory (OVRO) mm Array (Scoville et al. 1991). The 113 GHz measurement also presented in Anantharamaiah et al. (2000) clusters much better with these other observations and a power-law extrapolation of FIR observations of Arp 220 (Carico et al. 1992; Dunne & Eales 2001; Eales et al. 1989; Rigopoulou et al. 1996; Woody et al. 1989). Once again, we use the interferometric observation but not the single-dish one.

4.4.2 Modeling results

We performed least-squares fits between the averaged ATA data and the chosen high-frequency measurements for each source, using a Levenberg-Marquardt algorithm with τ_{ff} constrained to be nonnegative. In the case of M82 we used $\Omega = 2400 \text{ arcsec}^2$, its modeled angular size at 1 GHz. For the core component of NGC 253, we used its modeled size

Table 4.7: Results of spectral modeling.

Parameter ^a	M82	NGC 253	Arp 220
$\sqrt{\chi_r^2}$	1.21	1.02	1.17
A	0.93 ± 0.01	0.50 ± 0.01	$-0.48^{+0.007}_{-0.011}$
B	-0.56 ± 0.02	-0.51 ± 0.03	$-0.01^{+0.03}_{-0.04}$
C	-0.12 ± 0.03	-0.16 ± 0.03	$-0.31^{+0.02}_{-0.04}$
τ_1	$0.04^{+0.01}_{-0.01}$	0.27 ± 0.06	< 0.06
α_5	-0.73 ± 0.04	-0.74 ± 0.04	$-0.44^{+0.01}_{-0.05}$
Δ	-0.24 ± 0.07	-0.33 ± 0.06	$-0.62^{+0.04}_{-0.08}$
f_{th}	0.06 ± 0.02	$0.07^{+0.03}_{-0.01}$	< 0.02

^aUncertainties in the parameters are marginalized 1σ values determined via Monte Carlo simulations with 10000 realizations for each source.

as resolved at higher frequencies, 160 arcsec^2 . The ATA did not resolve Arp 220 but in [Scoville et al. \(1991\)](#) it is reported that the emission comes from an extended component of $\sim 7'' \times \sim 15''$ and a stronger central component of $\sim 2 \text{ arcsec}^2$. We adopt an intermediate value of 17 arcsec^2 . In all cases we used $T_e = 10^4 \text{ K}$. The numerical results of our fits are reported in Table 4.7. The uncertainties shown are marginalized values determined via Monte Carlo simulations with 10000 realizations for each source. To provide a more physical interpretation of B and C , we also transform them, respectively, into α_5 , the spectral index of the nonthermal model component at 5 GHz, and $\Delta \equiv d\alpha/d(\log_{10} \nu)$ at the same frequency. We also compute f_{th} , the thermal fraction (not correcting for FFA) at 1 GHz. We computed 3σ confidence limits in the total flux of the best-fit models at 40 frequencies logarithmically spaced between 0.7 and 100 GHz. Finally, we computed 1σ confidence regions in the B - C plane of parameter space. This was done by tracing the contour of $\chi^2 = \chi_{\text{min}}^2 + 2.3$, with χ_{min}^2 being the minimum χ^2 achieved in the overall fit and the offset parameter coming from [Avni \(1976\)](#). The value of χ^2 assigned to each location in the B - C plane is the minimum value of χ^2 attained in a fit with B and C fixed and A and τ_1 free.

The parameter values that are obtained depend strongly on the highest-frequency data points. Fits that were run without any very high-frequency ($\nu > 50 \text{ GHz}$) data points often converged to the physically-suspect lower bound of $\tau_1 = 0$. To investigate the possible effects of inaccurate high-frequency observations on our results, we ran a set of Monte Carlo simulations in which the uncertainty associated with each high-frequency measurement was doubled. In this case, the resulting uncertainties in the modeled parameter values increased by a factor of ~ 1.5 , demonstrating the extent to which the results are affected by those few data points. While we have done our best to choose our high-frequency data points carefully,

it remains true that our results are dependent on the choice of which measurements to include and, of course, the accuracy of those measurements.

M82

The M82 fit has $\sqrt{\chi_r^2} = 1.21$. The ATA data and fit results are presented in Figure 4.8. The value of α_5 that we obtain agrees well with that found in Klein et al. (1988), $\alpha = -0.68 \pm 0.02$. We, however, find a higher thermal content at 32 GHz, because the possibility of a steepening nonthermal spectrum allows for a larger thermal contribution.

A fit with a fixed value of $C = 0$ (*i.e.*, a pure power-law nonthermal spectrum) had a comparable $\sqrt{\chi_r^2} = 1.34$, but also $\tau_1 = 0$, a physically implausible result. In the $C = 0$ case, the many ATA data points tightly constrain the power-law slope of the nonthermal component, but the high-frequency points lie significantly below the nonthermal line. The thermal component, being spectrally shallower than the nonthermal component, can only make this situation worse, so the fit converges to a solution with no thermal component at all. If any thermal component is present, the nonthermal spectrum must steepen so that the high-frequency points can be modeled successfully.

NGC 253

The NGC 253 fit has $\sqrt{\chi_r^2} = 1.02$. The ATA data and model results for the core component of NGC 253 are presented in Figure 4.9. The symbols used are analogous to those in Figure 4.8. Figure 4.10 shows the extended and total fluxes in context with total-flux data from the literature. The same models as in Figure 4.9 are overlaid to ease comparison.

A fit with a fixed value of $C = 0$ had $\sqrt{\chi_r^2} = 1.51$ and $\tau_1 = 0.25$. Here, a change in the spectral index is detectable in the ATA data alone, so when the nonthermal component has a constant slope, τ_1 must increase to induce more curvature at low frequencies via FFA. In this case, once again the high-frequency data points lie well below the points attained by the model.

Arp 220

The Arp 220 fit has $\sqrt{\chi_r^2} = 1.23$ and $\tau_1 = 0$ and is shown in Figure 4.11. This result is due to the dramatic steepening of the spectrum of Arp 220 at $\nu \sim 10$ GHz. A purely nonthermal fit that accommodates the high-frequency observations must underestimate the low-frequency ($\nu \lesssim 2$ GHz) data points significantly. The addition of any thermal contribution would require an even steeper nonthermal component to match the high-frequency measurements, which would make the discrepancies on the low end even worse. Adjustments to the assumed angular size of Arp 220, which affect the overall scaling of the thermal component, do not improve the outcome. We report the results of our fit in Table 4.7 but caution that they are unsatisfactory. We exclude the Arp 220 results from further analysis in the following sections.

If the nonthermal curvature parameter C is fixed to zero, the model cannot accommodate the data, with $\sqrt{\chi_r^2} = 2.77$. This result quantitatively confirms the supposition of [Condon \(1992\)](#) that the nonthermal spectrum of Arp 220 steepens in the cm regime. Our data cause serious problems for the empirical spectral model of Arp 220 presented in [Sopp & Alexander \(1991\)](#), which predicts a much steeper spectrum at frequencies of a few GHz than we find. Their model was inconsistent with several measurements at $\nu \sim 2$ GHz shown in the paper, and our ATA data make it clear that those measurements are generally accurate and the simple power-law used in that work is insufficient to match the data.

4.5 Discussion

We test whether our observational results for M82 and NGC 253 are consistent with the predictions of [T+06](#) by comparing the curvature of the FFA-corrected nonthermal spectra with the output of the numerical model presented in that work. We obtained an updated version of the relevant code (E. Quataert, private communication) and modified it to allow more convenient tuning of the input parameters. These parameters are (1) Σ_g , the surface density of the gas disk; (2) h_{eff} , the effective scale height of the gas disk; (3) p , the power-law index of the injected CR energy spectrum, $dn \propto E^{-p}dE$; and (4) η , the ratio of the energy density of the magnetic field to the pressure of the ISM, $B^2 = \eta \cdot 8\pi^2 G \Sigma_g^2$. The value of h_{eff} affects the characteristic timescales of CR cooling via ionization and free-free losses by setting the density of the neutral ISM as encountered by the CRs, $n_{\text{CR}} = \Sigma_g / 2m_p h_{\text{eff}}$. (In [T+06](#), h_{eff} is specified by two parameters, the physical scale height of the gas disk h and an ISM/CR interaction efficiency parameter f : $h_{\text{eff}} = h/f$. They are degenerate for our purposes.) Because the mean value of the ISM density may not correspond well to the value of n_{CR} due to clumping and other effects, h_{eff} may likewise not correspond well to the actual scale height of the gas disk. We take the surface density of the disk of M82 to be 0.081 g cm^{-2} and that of NGC 253 to be 0.067 g cm^{-2} ([Thompson et al. 2009](#), and references therein). Note that these recently-derived values are significantly different than the ones given in [Kennicutt \(1998\)](#).

Given a choice of Σ_g , h_{eff} , p , and η , we used a simple least-squares fit to determine values for the spectral parameters B and C . We used the [T+06](#) model to compute flux densities at 25 frequencies logarithmically spaced between 1 and 7 GHz and then fit a bent power law (as in Equation 4.17) to the data. All of the spectra generated in this way were very well-described by such a model. The derived parameters were not sensitive to the details of the choice of frequencies at which to sample the model.

We computed such B and C values for the [T+06](#) model on two grids of parameters. The first had Σ_g set to the observed M82 value, h_{eff} ranging between 10 and 250 pc, p ranging between 2 and 2.5, and η ranging between 0.02 and 1, all stepping through 10 values logarithmically. The lower bound of η is the approximate value for which the synchrotron cooling time, τ_{syn} , and the timescale for escape from the galaxy via advection by a galactic

wind, τ_{esc} , are equal. For η (and hence B) smaller than this, the assumptions of the T+06 code are no longer upheld. The largest plausible η is ~ 1 , at which point the magnetic field and ISM pressure are in equipartition (Parker 1966). The second grid was identical except that Σ_g was set to the NGC 253 value and the lower bound of η was 0.04, the analogous limiting value. Figure 4.13 shows the values of B and C attained by the model and compares them to the 1σ and 2σ confidence intervals associated with our observations. The T+06 models can accommodate a wide range of spectral indices and exhibit spectral steepening for all parameter values. Our observed spectra, however, exhibit more steepening (more negative C) than the models. The models are consistent with observations of both galaxies at the 2σ level but only with M82 at the 1σ level. If the source of the FFA is configured as a screen in front of the nonthermal emission rather than a medium mixed within it (Equation 4.18), the C values derived for our observations are less negative, ~ -0.10 , and the T+06 models have large areas of overlap with the observations of both galaxies even at the 1σ level. The $\sqrt{\chi_r^2}$ values of the fits in this case are virtually identical to those of the fits in the well-mixed case.

We quantify this agreement in Figures 4.14 and 4.15, which show the regions in the η - h_{eff} plane where the results of the T+06 model are consistent with the observed properties of M82 and NGC 253, respectively, at the 2σ level, for $p = 2$ and 2.3. Because our data, unlike the outputs of the T+06 models, are affected by FFA and thermal emission, they cannot be compared to the models directly without taking into account these effects. To do so requires some sort of spectral modeling, so we compare the spectral curvature parameters B and C found for our data with ones derived for the outputs of the T+06 models. (If we corrected our observations for the modeled FFA and thermal emission and then compared them to the T+06 models, we would be effectively hiding the model-dependence of the correction and discarding the information that the data have been corrected assuming an underlying parabolic nonthermal spectral model.) We assumed a bivariate normal probability distribution for the observed values of B and C , with the parameters of the distribution estimated by applying the maximum likelihood method to the results of the Monte Carlo simulations. The agreement between the shape of the analytic distribution function and contours of constant $\Delta\chi^2$ was very good. Using the bounds on h_{eff} and η given above, $p = 2.5$ is rejected at the 99.96% confidence limit for M82 and the 99.99% confidence limit for NGC 253; larger values of p become acceptable if the lower limit on h_{eff} is decreased. For NGC 253, $p = 2.3$ is also rejected at the 99.99% confidence limit.

For M82, virtually all physically plausible values of η are allowed. For $p = 2$, the allowed values of h_{eff} are of order 100 pc, broadly consistent with observed radio sizes. For $p = 2.3$, agreement requires $h_{\text{eff}} \approx 20$ pc. Such sizes have been proposed for the neutral gas disks of some ultraluminous infrared galaxies (ULIRGs; Scoville et al. 1997; Carilli et al. 1998), but are unlikely to be found in less extreme galaxies. Furthermore, if the CRs interact with the ISM at somewhat less than its nominal mean density (the parameter f of T+06 is <1), as is suggested by measurements of CRs in the Milky Way (Connell 1998), the actual scale height of the disk becomes even less physically plausible. As η increases, the allowed h_{eff} decreases for both values of p . This is because the stronger magnetic field steepens the synchrotron

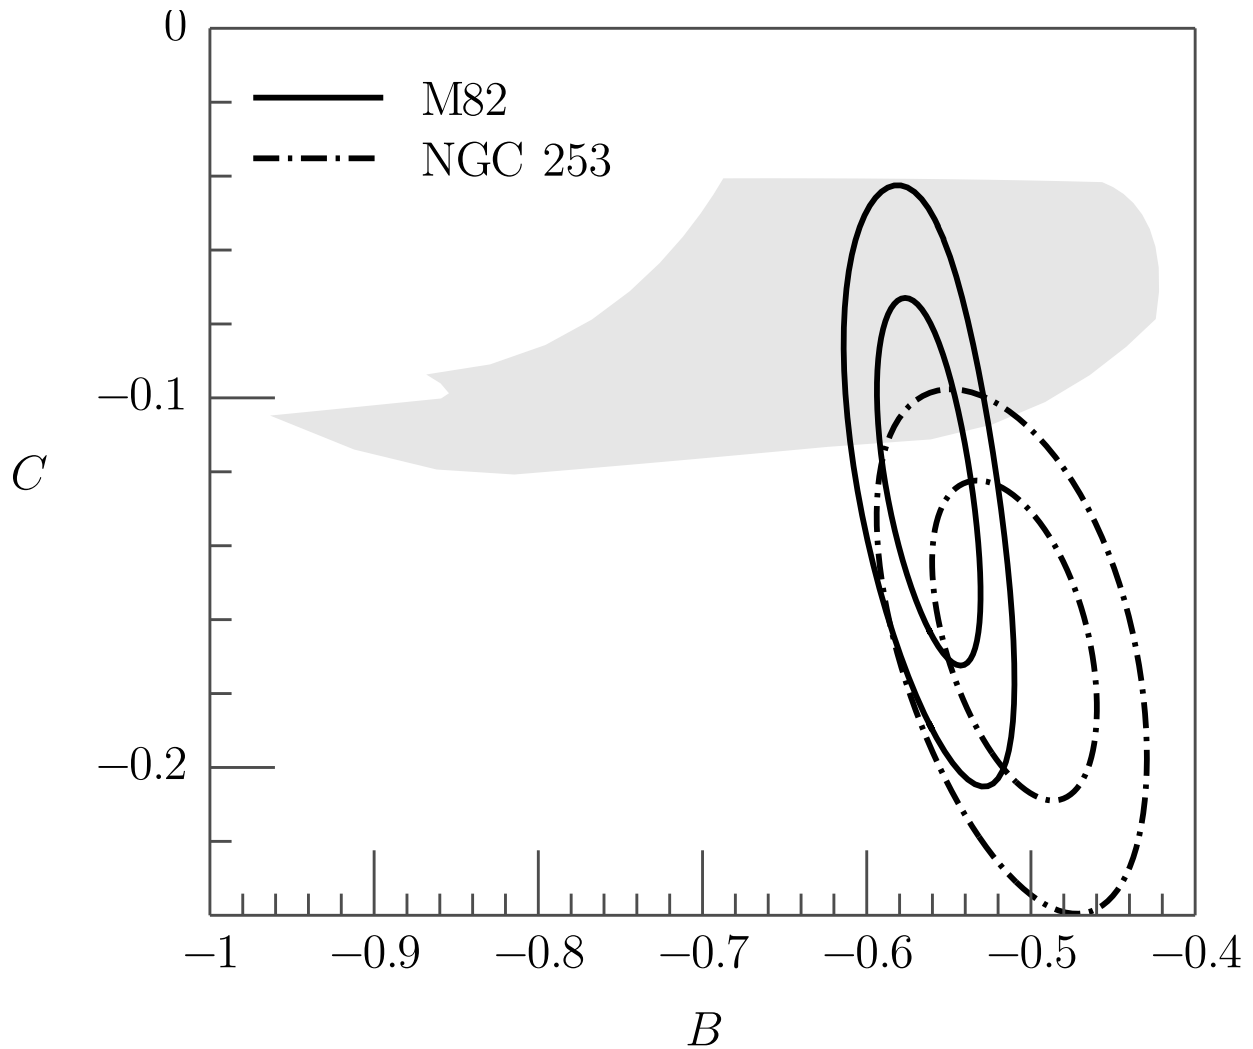


Figure 4.13: Observationally and theoretically derived values of the spectral curvature parameters B and C . *Contours*: 1σ (*inner*) and 2σ (*outer*) confidence regions for the parameters as derived from the modeling described in §4.4. *Shaded region*: the range of B and C values taken by the model of T+06 over a range of parameters appropriate for M82 as described in §4.5. The equivalent region for the NGC 253 parameters is virtually identical to the one shown.

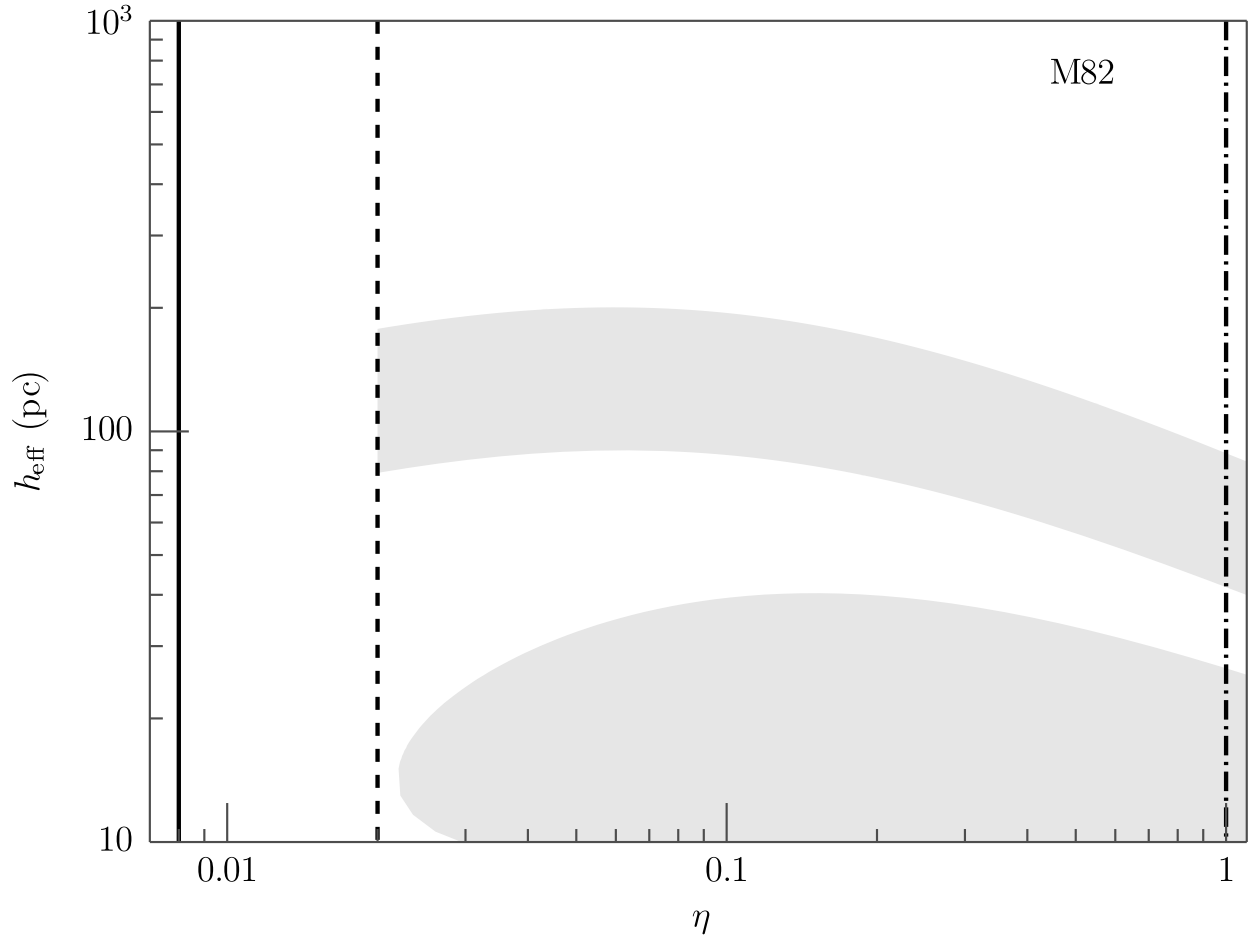


Figure 4.14: Values of η and h_{eff} for which spectral curvature parameters B and C derived from the T+06 model are consistent with observations of M82 to within 2σ . *Upper shaded region*: the allowed parameters for $p = 2$. *Lower shaded region*: the allowed parameters for $p = 2.3$. *Solid line*: the value of η such that $B = B_{\text{min}}$. *Dashed line*: the value of η such that the synchrotron cooling timescale τ_{syn} and the escape time τ_{esc} of T+06 are equal, assuming $\nu = 1.43$ GHz, $h_{\text{eff}} = 100$ pc, and the galactic wind speed $v_w = 500$ km s $^{-1}$. This defines the approximate lower limit in η of the validity of the T+06 numerical model. *Dot-dashed line*: the value of η such that $B = B_{\text{eq}}$.

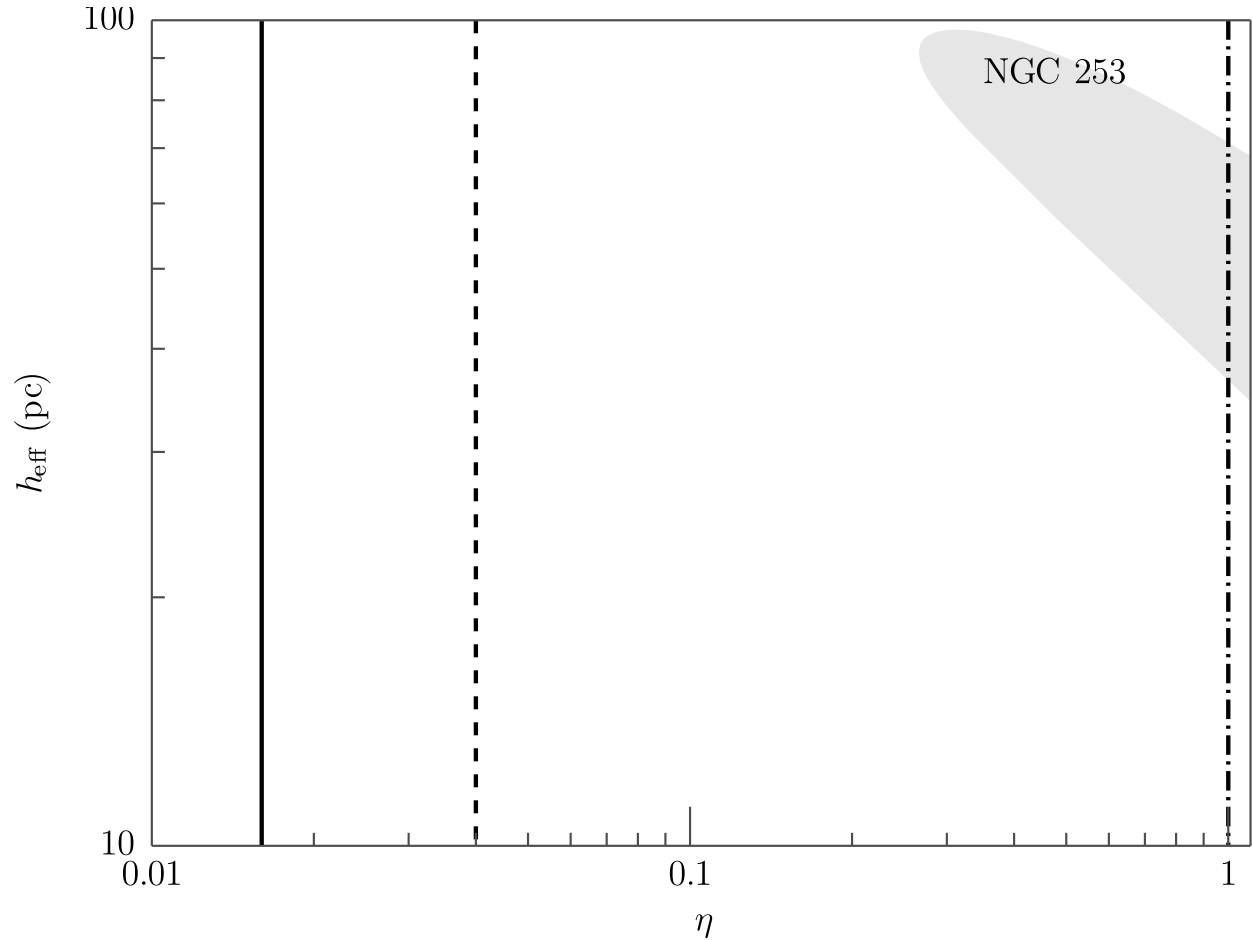


Figure 4.15: Values of η and h_{eff} for which spectral curvature parameters B and C derived from the T+06 model are consistent with observations of NGC 253 to within 2σ . The symbols are as in the previous figure. *Shaded region:* the allowed parameters for $p = 2$. The T+06 model with $p = 2.3$ is excluded to a 99.99% confidence limit.

spectrum, requiring a smaller h_{eff} to increase n_{CR} and consequently the amount of spectral flattening due to ionization and free-free emission off of neutral hydrogen. As p increases, the allowed values of h_{eff} rapidly decrease for the same reason.

The NGC 253 data are only marginally consistent with the T+06 model within the range of parameters we consider, and the statistically acceptable results have $h_{\text{eff}} \approx 50$ pc. If the source of FFA is configured as a screen rather than mixed with the nonthermal emission, much more of the T+06 model parameter space is allowed, with a general configuration of allowed parameters similar to that obtained for M82.

The relationship between h_{eff} and the observed radio sizes of starburst galaxies h_{obs} is significant within the T+06 model. Estimates in T+06 indicate that if $h_{\text{eff}} \sim h_{\text{obs}}$ and $B \sim B_{\text{min}}$, free-free emission off of neutral ISM nuclei becomes a significant sink of CR energy for $\Sigma_g \gtrsim 1 \text{ g cm}^{-2}$, which would disrupt the observed linearity of the FIR-radio correlation at high surface densities. Since the free-free timescale is independent of the magnetic field while the synchrotron timescale goes as $\tau_{\text{syn}} \propto B^{-3/2}$, a stronger magnetic field, $B \sim B_{\text{eq}}$, avoids this inconsistency. Our observations have $h_{\text{eff}} \sim h_{\text{obs}}$ in the $p = 2$ case and $h_{\text{eff}} \ll h_{\text{obs}}$ in the $p = 2.3$ case. A precise independent measurement of p could thus provide a valuable test of this aspect of the T+06 model.

4.6 Conclusions

Our observations of M82 are consistent with the scenario proposed in T+06. Our observations of NGC 253, on the other hand, are only marginally consistent with it. We decisively detect the spectral steepening at $\nu \simeq 1$ GHz predicted in that work. The uncertainties in our spectral modeling, due to uncertainties in the individual observations, the difficulties of untangling the thermal and nonthermal emission, and the inaccuracies inherent in our spatially-unresolved approach, make it difficult to significantly constrain all of the parameters of the T+06 spectral model or to reject its entire parameter space. We do, however, reject with high confidence the T+06 model with a steep electron energy injection spectrum power-law index of $p = 2.5$ for M82 and NGC 253. We also reject with high confidence the T+06 model with $p = 2.3$ for NGC 253. We are unable to place limits on the strength of the magnetic field required by the T+06 model with our current data. Further work to constrain the efficiency with which CRs interact with the ISM or the allowed values of p would significantly shrink the model parameter space. Observations of a larger sample of galaxies would allow a statistical treatment of the topics that we have dealt with here on a case-by-case basis. It would likely also be fruitful to compare our detailed radio results with FIR observations. Such an approach would be a frequency-space analog to the recent works investigating the connections between FIR and radio emission at high spatial resolutions (Marsh & Helou 1995; Paladino et al. 2006).

The broadband, continuous spectra that we have obtained with the ATA are precise and repeatable. The quality of our measurements is limited by several factors. The u - v

coverage of our snapshot observations is good but suboptimal for complex sources, which made the imaging of targets such as NGC 253 more difficult. The sensitivity the system is a limiting factor for snapshot observations of fainter sources such as Arp 220. Both of these quantities were limited by the limited number of antennas available for observing during the commissioning of the ATA-42. The planned expansion of the ATA to ~ 350 antennas will do much to improve both of them. The fact that the ATA was undergoing commissioning at the time of these observations meant that several important aspects of observing and data reduction, such as polarization calibration and RFI mitigation, were incomplete. Ongoing development since these observations were made has led to improvements in virtually all aspects of the system performance. Of particular interest is the deployment of a second independently-tunable correlator, which doubles the already-fast throughput of our spectral observations and further improves the prospects for fast, precise, broadband spectral monitoring of sources.

Acknowledgments

I thank Eliot Quataert for helpful discussions. This work was done with the support of a NSF Graduate Research Fellowship.

Chapter 5

ASGARD, A Large Survey for Slow Galactic Radio Transients: Overview and First Results

New and upgraded radio observatories provide an order-of-magnitude increase in survey speed compared to their predecessors, enabling a series of new investigations of the poorly-understood dynamic radio sky. Searches for slow radio transients and variables, however, have generally focused on extragalactic fields, which are generally much simpler to image and analyze than Galactic fields. The basic parameters of the populations of slow Galactic radio transients and variables thus remain uncharacterized. We present ASGARD, a large 3 GHz survey performed with the Allen Telescope Array that aims to improve this situation. ASGARD observations spanned 2 years with weekly visits to 23 deg^2 in two fields in the Galactic Plane, totaling 900 hr of integration time on science fields and making it significantly larger than previous efforts. We describe the observations and data analysis techniques in detail, demonstrating our ability to create accurate wide-field images while effectively modeling and subtracting large-scale radio emission, allowing standard transient-and-variability analysis techniques to be used. We present early results from the analysis of two pointings: one centered on the microquasar Cygnus X-3, and one overlapping the *Kepler* field of view, including images, catalog statistics, completeness functions, variability measurements, and a transient search. Out of 134 sources detected in these pointings, the only compellingly variable one is Cygnus X-3, and no transients are detected. We estimate number counts for potential Galactic radio transients and compare our current limits to previous work and our projection for the fully-analyzed ASGARD dataset.¹

¹This chapter and Appendix B were subsequently published, with some modifications, as Williams *et al.*, 2013 ApJ 762 85.

5.1 Introduction

The technological developments of the past few decades have led to an explosion of interest in the astronomical time domain. There has been a recent boom at radio wavelengths, where the survey capabilities of new and upgraded facilities represent order-of-magnitude improvements over their predecessors (*e.g.* Perley et al. 2011; Verheijen et al. 2008; Giovanelli et al. 2005; DeBoer et al. 2009; Jonas 2009; Carilli & Rawlings 2004). Among the motivators for the construction of these facilities are known or strongly-supported classes of highly-variable slow extragalactic radio emitters such as active galactic nuclei (Lister et al. 2011), orphan γ -ray burst afterglows (Frail et al. 2001; Levinson et al. 2002), radio supernovae (Weiler et al. 2002; Brunthaler et al. 2009; Muxlow et al. 2010), and tidal disruption events (Rees 1988; Bloom et al. 2011; Berger et al. 2012). We follow other authors in defining slow variables as those whose emission evolves on timescales $\gtrsim 1$ s; this approximately corresponds to those that emit via incoherent, rather than coherent, processes, and are typically identified using image-domain techniques. The relatively sparse prior exploration of the dynamic radio sky additionally highlights it as an exciting discovery space (Cordes et al. 2004; Fender & Bell 2011; Frail et al. 2012).

Much work has recently gone into the characterization of the population of highly-variable slow extragalactic sources. Archival studies have used existing large-area surveys to search for rare events (Levinson et al. 2002; Bower et al. 2007; Bell et al. 2011; Bower & Saul 2011; Ofek & Frail 2011; Bannister et al. 2011a,b; Thyagarajan et al. 2011, 2012; Frail et al. 2012). Followup observations of candidate or confirmed highly-variable sources have been used to characterize the detailed properties of individual objects (de Vries et al. 2004; Gal-Yam et al. 2006; Muxlow et al. 2010; Berger et al. 2012). Finally, dedicated surveys have been devoted to the systematic discovery of highly variable radio sources. Several have been undertaken with the Allen Telescope Array (ATA; Welch et al. 2009), one of the first radio observatories explicitly designed to be an efficient survey instrument, including the ATA Twenty-centimeter Survey (Croft et al. 2010a,b, 2011) and the Pi GHz Sky Survey (Bower et al. 2010, 2011, PiGSS-I and PiGSS-II hereafter). Ofek et al. (2011) describe most of these previous studies in greater depth. Frail et al. (2012) summarize the state of the field and find that reliable detection of these sources remains both a challenge and an opportunity.

In contrast, there has been relatively little work to investigate the population of highly-variable slow *Galactic* radio emitters. Source classes contributing to this population include X-ray binaries (Waltman et al. 1995; Marscher & Brown 1975), active stellar binaries (Hall 1976; Bopp & Fekel 1977; Eker et al. 2008), cool dwarfs (Berger et al. 2001; Berger 2002; Hallinan et al. 2007), and flare stars (Güdel 2002; Jackson et al. 1989; Richards et al. 2003). Towards the Galactic Center (GC), several intriguing sources of ambiguous nature have been discovered, including the Galactic Center Transient (Zhao et al. 1992), A1742–28 (Davies et al. 1976), CXOGC J174540.0–290031 (Bower et al. 2005), GCRT J1745–3009 (Hyman et al. 2005), and GCRT J1742–3001 (Hyman et al. 2009). An overall increase in the prevalence of apparent radio variability is expected towards the Galactic plane (GP),

and the GC in particular, due to interstellar scintillation (e.g., Spangler et al. 1989; Rickett 1990; Ghosh & Rau 1992; Gaensler & Hunstead 2000; Lovell et al. 2008; Ofek & Frail 2011). Pulsars and several other well-known classes of variable Galactic sources are not discussed here because they do not fall into the “slow” category; see Cordes et al. (2004) and references therein.

For many years, the best available data on Galactic radio variability came from the work of Taylor & Gregory (1983) and Gregory & Taylor (1986), who used the NRAO 91-m transit telescope to repeatedly survey the GP. Recent work has begun to expand and update these results. Hyman et al. (2002, 2003) have periodically monitored the GC with the VLA and GMRT, and as indicated above they have discovered several unusual transient sources. Becker et al. (2010) conducted an archival search for GP transients by comparing the VLA survey described in Becker et al. (1994) with CORNISH, the “Co-Ordinated Radio ‘N’ Infrared Survey for High-mass star formation” (Purcell et al. 2008), discovering a population of highly variable sources and analyzing their statistical properties. Most recently, Ofek et al. (2011) conducted a VLA variability survey at low Galactic latitudes with rapid multiwavelength followup, discovering a candidate transient source and measuring more statistical properties of the population.

In this chapter we present ASGAR: the ATA Survey for Galactic Radio Dynamism. Its primary goals are to perform a deep search for Galactic radio transients and to measure the variability properties of a wide variety of Galactic radio sources on day–year timescales. To this end, we repeatedly observed 24 pointings near the GP at 3 GHz with the ATA, visiting most pointings on a weekly cadence over the course of ~ 2 yr and obtaining ~ 900 hr of integration on our science pointings. The large field of view (FOV) of the ATA allows us to cover a relatively large footprint on the sky, and our frequent visits provide thorough sampling of the variability structure function. The compact configuration of the ATA also makes our observations sensitive to static extended structures in the GP, such as nonthermal radio filaments (Yusef-Zadeh et al. 1984; Law et al. 2008), H II regions (Brogan et al. 2003; Nord et al. 2006), and supernova remnants (Gray 1994).

We proceed by describing the ASGAR observations (§5.2) and data processing (§5.3) in detail. This work concerns itself with a subset of the whole dataset (§5.3.1), representing $\sim 10\%$ of the expected usable observations, while another portion of the dataset will be described elsewhere (Chapter 6). We present first results (§5.4) derived from analysis of this subset, including deep images, source catalogs, a transient search, and variability statistics. Finally we discuss our current conclusions and the prospects of the fully-analyzed survey (§5.5), populating a “log- N /log- S ” plot for Galactic radio transients.

5.2 Observations

We used the ATA to monitor two fields: an area around the Galactic center (GC) spanning $-4.5^\circ < \ell < 8^\circ$, $|b| < 2^\circ$; and ~ 5 deg² towards Cygnus including the highly radio-variable

microquasar Cyg X-3 (Gregory & Kronberg 1972) and a portion of the FOV of the *Kepler* (Koch et al. 2010) mission. Coordinates of the ASGARD pointing centers are listed in Table 5.1, along with the observing time devoted to each pointing over the course of the survey. Throughout this work we use the word “field” to refer to either of the two survey regions (GC and Cygnus) and the word “pointing” to refer to a specific pointing center that was observed.

While some ASGARD observations were conducted with complete control of the ATA, the vast majority were conducted commensally with a SETI (Search for Extraterrestrial Intelligence) survey (Blair & ATA Team 2009). Both of these projects were designed to increase the likelihood of discovering rare Galactic events by targeting regions with high source densities, as well as to take advantage of the data stream provided by the *Kepler* mission, which motivated the basic choice of survey fields. The division of the survey area into two fields was also a practical choice given year-round observations, since the GC region is seasonally difficult or impossible to observe from the ATA site (latitude $+40.8^\circ$). The Cygnus region, on the other hand, can be observed year-round, also has a very high source density, and contains the benchmark source Cyg X-3. The observed GC field is as centered on Sgr A* as possible given visibility constraints. The overall footprint of each field was determined by the joint needs of the commensal SETI search and ASGARD. The former required a certain dwell time per pointing to sequentially survey targets using the ATA digital beamformer backends. The latter aimed for a weekly revisitation cadence. Combining these requirements with the typical weekly survey time allocation determined the total footprint that could be observed. The system used to organize and execute the commensal observations is described in Chapter 3.

The pointings in the GC field fall on an 11×2 square grid in Galactic coordinates with a spacing of 1° and a northeast corner located at $\ell = 6.5^\circ$, $b = 0.5^\circ$. The half-power beam width (HPBW) of the ATA is approximately $3.5 \text{ deg GHz}/\nu_{\text{obs}}$ (Hull et al. 2010; Harp et al. 2011, but see §5.3.5), so that the pointing centers range from slightly oversampled to critically sampled for ASGARD observing frequencies of $\sim 1\text{--}3 \text{ GHz}$. The Cygnus field consists of an analogous 2×2 grid, with Cyg X-3 located at the northwest pointing center ($\ell = 79.8^\circ$, $b = 0.7^\circ$), as well as a disjoint pointing toward $\ell = 75.8^\circ$, $b = 13.5^\circ$ ($\alpha = 19^{\text{h}}21^{\text{m}}24^{\text{s}}$, $\delta = 44^\circ00'00''$, ICRS J2000), which lies within the *Kepler* mission FOV. The total footprint of the survey on the sky is $\sim 23 \text{ deg}^2$ if each pointing is taken to image a circle with a diameter of the nominal HPBW.

At the time of the observations the ATA frontend consisted of forty-two 6.1-m offset Gregorian dishes. The backends used for this work were two FX correlators with bandwidths of 104.9 MHz divided into 1024 channels and a dump time of 10 s. The correlators accepted 64 “antpol” inputs, meaning that they could perform full-Stokes correlation of 32 antennas or, hypothetically, single-Stokes correlation of 64 antennas. Because of the desirability of full-Stokes coverage and an ongoing program of feed retrofits, the correlators generally accepted data from ~ 32 distinct antennas. The precise set of antennas used varied with time due to availability changes stemming from maintenance or hardware failures, and the two correlators

Table 5.1: ASGARD pointing centers and 3 GHz summary statistics

Field	Identifier ^a	Galactic Coordinates (ℓ , b)		# Epochs ^b	Integ. Time (hr) ^b
GC	gc-07+1	356.500	+0.500	66	45.29
	gc-07-1	356.500	-0.500	0	0.00 ^c
	gc-05+1	357.500	+0.500	66	46.28
	gc-05-1	357.500	-0.500	0	0.00 ^c
	gc-03+1	358.500	+0.500	64	14.20
	gc-03-1	358.500	-0.500	59	10.84
	gc-01+1	359.500	+0.500	67	14.72
	gc-01-1	359.500	-0.500	76	22.44
	Sgr A*	359.944	-0.046	3	5.67
	gc+01+1	0.500	+0.500	70	21.70
	gc+01-1	0.500	-0.500	5	0.80
	gc+03+1	1.500	+0.500	65	16.87
	gc+03-1	1.500	-0.500	0	0.00 ^c
	gc+05+1	2.500	+0.500	68	15.92
	gc+05-1	2.500	-0.500	0	0.00 ^c
	gc+07+1	3.500	+0.500	33	4.94
	gc+07-1	3.500	-0.500	67	8.45
	gc+09+1	4.500	+0.500	70	19.73
	gc+09-1	4.500	-0.500	26	2.48
	gc+11+1	5.500	+0.500	69	11.33
	gc+11-1	5.500	-0.500	3	6.12
	gc+13+1	6.500	+0.500	72	16.55
	gc+13-1	6.500	-0.500	18	14.97
Cygnus	Kepler	75.756	+13.491	47	96.72
	Cyg X-3	79.845	+0.700	86	153.05
	x3+0-2	79.845	-0.300	13	10.87
	x3+2+0	80.845	+0.700	20	26.58
	x3+2-2	80.845	-0.300	10	3.96

^aIdentifiers with numerical codes refer to offsets from a reference point measured in half-degree increments.

^bSummary statistics refer to the 3 GHz subset of data as described in §5.3.1.

^cThese pointing centers were completely observed by the end of the first GC season, before the switch to 3 GHz observing frequencies, and so have no coverage in the 3 GHz subset.

did not necessarily have identical antpol inputs.

Each observing session (“epoch”) generally began with a long (~ 30 min) observation of a bright, unresolved source, usually one of 3C 48, 3C 147, 3C 286, or 3C 295. These observations were necessary for delay calibration of the beamformers used in the SETI survey but also provided excellent bandpass and flux density scale calibration data for ASGARD. Science pointings were observed with periodic (every ~ 45 min) visits to a nearby phase calibrator. Because of the low declinations of the GC pointings, these were generally observed at relatively low elevations and at constrained hour angle ranges.

The ATA feeds have a high-bandwidth design using a log-periodic architecture, and as such the location of the active region on the feed may vary significantly depending on the observing frequency. Because the optics of the ATA reflectors are frequency-independent, each feed is mounted on a piston drive so that the appropriate part of the feed may be moved to the optical focus for each observation. Focus positions are identified by the optimal corresponding observing frequency, i.e., a focus position of $\nu_{\text{foc}} = 1.4$ GHz provides the best configuration for observations at that frequency. The relationship between the piston position and observing frequency is based on both theoretical and empirical analysis (Harp et al. 2011). Defocused observations are possible with some loss in system performance. For $\nu_{\text{foc}} \gtrsim 0.9\nu_{\text{obs}}$, the penalty is slight, while for lower focus frequencies (active region too close to the secondary) sensitivity degrades and the PB broadens. For ASGARD observations made in complete control of the array, the focus position was set optimally, usually at 3.14 GHz. For commensal observations, the focus position was set at 1.90 GHz.

The overall coverage statistics of ASGARD observations are recorded in Table 5.2. We divide the observations into four campaigns: three seasons of GC observations during the summers of 2009–2011, and regular Cygnus observations from 2009 November to 2011 April. For reasons discussed in the following section, we isolate the set of observations made with the ATA correlators tuned to sky frequencies of 3.04 and 3.14 GHz, which we refer to as the “3 GHz” observations. These are the same frequencies used by the PiGSS survey. The overall amount of observatory time dedicated to the project was 1650 hr, with 902 hr spent observing science targets.

5.3 Analysis

We mostly use standard techniques to calibrate and image the data, using tools from the MIRIAD data reduction package (Sault et al. 1995) with several customized steps implemented in the Python programming language via the package `miriad-python` (Chapter 2). We use our images to construct a catalog of every compact source detected in the survey and to measure fluxes (or upper limits) for every observation of every source. We perform a variability analysis on this photometric dataset. We define a criterion by which sources may be classified as “transient” or not, but the classification does not affect the variability analysis, and no transient sources are detected.

Table 5.2: Parameters of ASGARDCampaigns

Subset ^a	Field	Start Date	End Date	# Epochs	Time On-Source (h)
Complete	GC	2009 May 22	2009 Oct 14	64	186
	GC	2010 Mar 01	2010 Oct 21	84	285
	GC	2011 Mar 01	2011 Apr 11	22	54
	Cygnus	2009 Nov 24	2011 Apr 11	159	377
3 GHz	GC	2010 Apr 29	2010 Oct 21	73	245
	GC	2011 Mar 01	2011 Apr 11	22	54
	Cygnus	2010 Feb 03	2011 Apr 11	108	291
This work	Cygnus	2010 Feb 03	2011 Apr 11	30	83

^aThe group of rows labeled “Complete” gives summary statistics for the complete ASGARDCampaign dataset. The group labeled “3 GHz” gives statistics for the 3 GHz subset of data as described in §5.3.1. The group labeled “This work” gives statistics for the datasets analyzed in this work.

The chief difficulty in the particular case of ASGARDCampaign is successful imaging of the significant large-scale structure (LSS) present in virtually all ASGARDCampaign fields of view, with the notable exception of the *Kepler* pointing. Not only is LSS imaging generally challenging, but by the nature of the survey most ASGARDCampaign epochs have sparse hour angle coverage. Fortunately, the LSS emission is expected to be time-invariant, while any astrophysical transients will be unresolved by the ATA (resolution $\sim 1'$ at 3 GHz). All observations of a given pointing can therefore be combined to produce an LSS model that can then be subtracted in the visibility domain from each epoch’s observations. To ease the measurement of the variability of compact sources, we subtract these sources from the LSS model, causing them to remain in the per-epoch images. Putting aside for the moment the important issue of subtraction errors, this approach yields per-epoch images consisting only of compact sources upon which standard transient-search techniques may be used. We describe our implementation of this general approach below.

5.3.1 Subset of Data Presented in This Work

The first season of GC observations was performed at sky frequencies of 1.43 and 2.01 GHz. It was found, however, that there was substantial broadband interference at these frequencies that presented numerous challenges for data analysis (Williams 2010). Subsequent observations primarily used the PiGSS 3 GHz setup, although a few observations were made at other frequencies. The analysis presented in this work is restricted to 3 GHz observations. In Tables 5.1 and 5.2 we provide coverage statistics as computed for the 3 GHz subset of the

complete ASGARD dataset.

In this work, we analyze two particular ASGARD pointings. The first is towards the highly-variable source Cyg X-3, which allows us to demonstrate the detection of variable radio sources embedded in complex, large-scale emission. The second is the *Kepler* pointing, which allows us to investigate the performance of our techniques in a field without unusual imaging challenges. We demonstrate the analysis of 29 epochs of the Cyg X-3 pointing, representing 34% of the 3 GHz Cyg X-3 dataset by number of epochs and 17% by raw data volume, and six epochs of the *Kepler* pointing, five of which also involved visits to the Cyg X-3 pointing. These epochs nearly completely sample the timespan of the 3 GHz Cygnus campaign, ranging from 2010 February 03 to 2011 April 04, and are roughly uniformly spaced in time. We have processed and imaged epochs surrounding the Cyg X-3 radio flares of 2010 May (Bulgarelli et al. 2010) and 2011 March (Kotani et al. 2011), but do not run our full pipeline on the latter epochs, of which there are six, because the 20 Jy flare leads to severe imaging problems related to dynamic range limitations. Tests of our transient detection process confirm, however, that it succeeds in this trivial case. The ATA-42 observations of these events are described in more detail elsewhere (Chapter 6). A somewhat larger portion of our observations has been processed and analyzed but without significant multi-epoch coverage of other pointings, and so is not presented in this work to maintain a clear focus on the better-covered Cyg X-3 and *Kepler* pointings.

5.3.2 Calibration & Flagging

Radiofrequency interference (RFI) was an intermittent problem during the 3 GHz observations. The forms of RFI most commonly encountered were narrow-bandwidth (1–5 channel) tones that affected the majority of baselines with a $\sim 100\%$ duty cycle, wider-bandwidth tones (~ 5 MHz) that affected moderate numbers of baselines with a moderate duty cycle, and brief (< 1 dump) broadband bursts affecting all baselines. Problems in the digital hardware (e.g., overheating) could also cause RFI-like effects, typically manifesting themselves as complete corruption of one half or one quarter of the spectrum for certain baselines. In the results we report, RFI was primarily excised from the data manually, using standard MIRIAD tools and an interactive, graphical visibility visualizer for the RFI with more complex time/frequency structure (Chapter 2). Several approaches to automatic flagging have been pursued with ATA data (e.g. Keating et al. 2010; Bower et al. 2010) and their integration into the ASGARD pipeline is being investigated.

Standard bandpass and gain calibration techniques are used. The long calibration observations at the beginning of each epoch are used to set the flux density scale, referencing to Baars et al. (1977). For the GC observations, the gain calibrator was NRAO 530 ($\ell = 12.03, b = +10.81$), while for the Cyg X-3 observations it was usually BL Lac. Because this latter source is variable and usually $\sim 10\%$ linearly polarized, gain parameters for it are derived from the bandpass observations whenever possible, treating X and Y feeds separately. Our observing program did not allow for planned observations of polarimetric calibrators

over wide hour angle ranges, so we are unable to solve for the frequency-dependent leakage terms that would be required for polarimetric calibration of the ATA (Law et al. 2011) on an epoch-by-epoch basis. Future work will investigate the stability of the system leakages to see if polarimetric calibration can be achieved by combining observations from multiple epochs.

5.3.3 Imaging and Source Extraction

After calibration, the data are averaged down to 16 spectral channels of 6.5536 MHz bandwidth per correlator and are converted to CASA (McMullin et al. 2007) format for imaging. This is necessary because MIRIAD does not implement any wide-field imaging algorithms, which we have found to be necessary for our analysis. In particular, without the use of techniques such as polyhedral imaging (Cornwell & Perley 1992) or w -projection (Cornwell et al. 2008), sources far away ($\gtrsim 0.8^\circ$) from phase center do not deconvolve well and acquire an hour angle dependence in their position, which is severely problematic for both LSS modeling and photometric extraction. In the steps we describe below, the averaged channels are gridded using multifrequency synthesis (Sault & Wieringa 1994) and imaged using w -projection with 128 planes. Images are 2048×2048 with a pixel size of $10''$ and thus span approximately five times the HPBW at 3 GHz. Some of the technical details of the imaging process are discussed in Appendix A.

We first construct a deep image for each pointing. Because these images are used to generate several important data products, the imaging techniques vary from pointing to pointing depending on what produces the best results empirically. For the *Kepler* pointing, we use the Cotton-Schwab deconvolution method (Schwab 1984; Cornwell et al. 1999), while for Cyg X-3 we use the maximum-entropy algorithm (Gull & Daniell 1978) with a Gaussian prior. For consistent PB-uncorrected LSS sampling, the data contributing to each deep image have a consistent ATA feed focus position. Imaging artifacts, rather than thermal noise, currently limit the image quality, so the use of only part of our data to form the deep images does not significantly affect their sensitivities.

The deep images are used to construct a catalog of compact sources, the properties of which are described in §5.4.2. Sources are detected using a combination of the MIRIAD task `sfind` and manual inspection to check for missed detections and reject dubious ones. (Given the resolution, sensitivity, and footprint of ASGARD, the latter technique is scalable to the whole survey.) Source positions, shapes, and mean fluxes are cataloged. Because the primary aim of ASGARD is to study variability, our catalog does not initially include sources that are marginally detected in the deep image, since such sources will typically not be detectable in the epoch images. The epoch images, however, are searched for uncataloged sources as described below, so that any variable source that becomes detectable during an epoch will eventually be included in the catalog.

Each ASGARD pointing is also associated with an LSS model. In the *Kepler* pointing, this model is blank; for the rest, the model is derived from the deep image and the compact source catalog, using either deconvolution of a variant deep image in which the compact

sources have been subtracted, or source-fitting techniques on the model image. The latter approach can be helpful because maximum-entropy deconvolution tends to model unresolved sources as Gaussians about the size of the synthesized beam. Because there are substantial numbers of Cyg X-3 observations made at focus positions of both 1.90 and 3.14 GHz, we generate one LSS model of this pointing for each focus position, so that the approximations of our PB modeling scheme (§5.3.5) do not lead to avoidable LSS subtraction errors.

Images from individual epochs are made the same way as the deep images, except that before imaging the appropriate LSS model is subtracted from the u - v data, and baselines shorter than 50 m are not imaged. Deconvolution of the individual epoch images is performed with 800 iterations of CASA’s “wide-field” implementation of the Högbom CLEAN algorithm (Högbom 1974). In some epochs, there is a small but discernable normalization difference between the LSS model and the epoch data, and we rescale the latter to match the LSS model. The largest such correction factor is 5%. The restoring beam was chosen automatically based on the u - v coverage of each epoch, with a typical size of $80'' \times 40''$.

Lightcurves of the known, compact sources are derived using image-domain fitting on the LSS-subtracted individual epoch images. Each fit holds the source position and shape fixed but allows the total flux to vary. Very nearby sources are fit simultaneously. Because the reality, positions and shapes of these sources are known from the deep image, detection limits are set to the relatively weak constraint of three times the local background rms. Undetected sources are cataloged with an upper limit of this value.

5.3.4 Detection of Uncataloged Sources

We use `sfind` to search for any residual sources in the individual epoch images after subtraction of both the LSS and the known, compact sources. (The LSS is subtracted in the u - v domain but the compact sources are subtracted in the image domain.) In our analysis, a “transient” is any source detectable in an individual epoch image that is not detectable in the corresponding deep image. Sources discovered via `sfind` are added to the ASGARD catalog and so are subsequently processed in the same way as all others. For those pointings in which the LSS model is derived from a deep image, such a source’s contribution to the deep image propagates into the LSS model. This mean flux is equal to the bright-epoch flux diluted by $\sim 1/n$, where n is the number of epochs. This contribution is subtracted from the per-epoch images used for residual source detection, but is not significant for n more than a few.

Previous searches for radio transients have typically been dominated by false positives (Frail et al. 2012). Our particular procedure involves multiple rounds of sky modeling and subtraction which will inevitably leave artifacts as well. Our transient-detection step therefore has stringent detection limits and cross-checks for systematic effects. We use the false-discovery rate (FDR) algorithm in `sfind` (Hopkins et al. 2002). The background rms is computed in 64×64 -pixel boxes (`sfind` keyword `rmsbox`) and the target FDR is set to 0.5% (`sfind` keyword `alpha`). To assist in computing detection limits, `sfind` was modified to report an estimated minimum detectable source flux density in the event that no sources were found in

an image, basing this value on an estimated FDR “ p -value” that would be needed to have yielded a source detection.

Sources reported by `sfind` are filtered according to several criteria. Sources for which `sfind` does not report a positional uncertainty, usually indicative of a very poor fit, are rejected. Sources in which the product of the deconvolved major and minor axes exceeds 7000 arcsec^2 are rejected, as are those in which the deconvolved major axis exceeds 130 arcsec . Sources for which the modeled primary beam attenuation exceeds 98.9% ($3\sigma_{\text{PB}}$, where $\sigma_{\text{PB}} = \text{HPBW}/\sqrt{8 \log 2}$) are rejected, although our later analysis uses much more conservative PB correction cutoffs. Finally, newly-detected sources near previously cataloged steady sources are also rejected. The match radius for this test is $50f \text{ arcsec}$ where f depends on the cataloged total flux S_k of the known source. For $S_k < 20 \text{ mJy}$, $f = 1$; for $S_k > 2.97 \text{ Jy}$, $f = 10$; and for intermediate values, $f = 2 \log(S_k/20 \text{ mJy})$. All of the above cutoff values were determined by examining the properties of the transient candidates that were both significantly detected and obviously spurious. Remaining candidates are examined manually as described in §5.4.4.

5.3.5 Primary Beam Modeling

In order to compute accurate fluxes, sensitivities, and sky models, we must account for the primary beam of the ATA. [Hull et al. \(2010\)](#) analyzed the primary beam of the ATA using data from the PiGSS survey, finding $\text{HPBW} = 1.10 \pm 0.01^\circ$ for a circular Gaussian PB model. Because many of our observations are performed with the same frequency configuration as PiGSS, we adopt this value assuming a ν^{-1} frequency dependence, i.e. $\text{HPBW}_{\text{optimal}} = 3.40 \text{ deg GHz}/\nu_{\text{obs}}$ for a mean PiGSS observing frequency of 3.09 GHz . The numerator is slightly smaller than the more generic value of 3.5 deg GHz reported by [Harp et al. \(2011\)](#).

Our PB modeling is complicated, however, by the fact that a substantial number of measurements are made with the focus set to 1.90 GHz . Although we do not have observations specifically aimed at characterizing the PB shape of 3 GHz observations at this focus setting, we measure this value from the data in two ways. Firstly, our observations of the Cyg X-3 pointing have extensive coverage in both this focus position and in the optimally focused position. By comparing compact source fluxes in deep images made from the two sets of data, we obtain $\text{HPBW}_{1900} = 3.92 \text{ deg GHz}/\nu_{\text{obs}}$. We also compare the ATA-apparent fluxes of point sources observed in the *Kepler* pointing to those from the NVSS catalog (see §5.4.2), assuming a typical source spectral index of $S_\nu \propto \nu^{-0.7}$, and find a factor of 3.96 deg GHz . We use the former value. Below, we denote the PB correction factor as applied to a particular location f_{PB} , defined such that $f_{\text{PB}} \geq 1$.

Holographic measurements of the ATA dishes suggest that the PB pattern becomes increasingly noncircular as the focus moves away from its optimal setting ([Harp et al. 2011](#)). Due to alt-az mount of the ATA dishes, the PB rotates on the sky over the course of an observation. Proper accounting for such an effect would need to occur during the Fourier inversion process, which neither MIRIAD nor the standard CASA imager are capable of doing.

We are thus unable to measure, or compensate for, this effect. (PB rotation can be dealt with approximately by imaging the data in blocks of similar hour angle, but this reduces u - v coverage and thus hampers the deconvolution of our complex fields.) The measurements of [Harp et al. \(2011\)](#) suggest that in our configuration the axial ratio should be about 10%.

We also assume that each dish has an identical PB pattern and that all dishes are pointed identically. These assumptions allow us to handle PB correction in the image domain and are relied on in some of the analysis that follows. The ability to relax these assumptions, like the ability to model PB rotation, relies on support in imaging software that is not yet widespread; two notable implementations are the MeqTrees system ([Noordam & Smirnov 2010](#)) and a derivative of the CASA imager equipped with the A -projection algorithm ([Bhatnagar et al. 2008](#)). Adapting our analysis to a pipeline in which PB correction occurred during the imaging process would require actual measurement, rather than calculation, of the spatial variation in image noise, but would not require major conceptual changes.

5.3.6 Multi-Epoch Photometry

Our imaging and source extraction pipeline yields multi-epoch photometry for our catalog sources, though for some faint sources our measurements yield mostly upper limits. Each source additionally has a flux measurement from the deep image in which it was detected.

Our images are not corrected for the attenuation of the ATA PB, so we correct our flux measurements for this effect. In the subset of data we consider, all *Kepler* observations are made with a focus setting of 1.90 GHz, and so the PB correction for a given source is the same in every epoch. The Cyg X-3 observations are made with focus settings of both 1.90 and 3.14 GHz, so the PB correction can vary from epoch to epoch. The deep Cyg X-3 image from which deep fluxes are obtained derives from the optimally-focused data. An additional factor that complicates comparison of the Cyg X-3 fluxes is the fact that we use different LSS models for each focus setting. While this is necessary to accurately subtract LSS from each epoch as well as possible, differences in the LSS model around each source can lead to an additive flux offset between measurements made at different focus settings, above and beyond any multiplicative errors caused by the limitations of our analytic PB models. In this work, we choose to consider only measurements from a consistent focus setting, choosing the one that resulted in the higher mean detection significance. For sources close to the pointing center, this is generally the optimally-focused data, whereas for sources far from the pointing center, the broad defocused PB results in more significant detections.

Some images yield particularly poor photometric results and we discard their measurements. These are generally either epochs with very few contributing data or those from 2011 March in which Cyg X-3 was undergoing a major (~ 20 Jy) flare, leading to significant dynamic range issues. Even without the inclusion of the latter flare in our analysis, Cyg X-3 is the most significantly variable source in our dataset.

We set the uncertainty on each flux density measurement to be

$$\sigma^2 = \sigma_{\text{rms}}^2 + (0.05S)^2 + (0.5 \text{ mJy})^2, \quad (5.1)$$

where σ_{rms} is the background rms of the relevant image-domain fit and S is the measured value. We investigated a correction for CLEAN bias (Becker et al. 1995; Condon et al. 1998) but did not find compelling evidence that this improved our results, so we do not include such a correction in our analysis. Likewise, we investigated but did not apply a “post-imaging calibration” correction of the kind used by Bannister et al. (2011a,b) and Ofek et al. (2011), in which a scalar flux density correction factor is computed for each image and applied to its photometric results. As was also found in PiGSS-II, we found that such a calibration sometimes had encouraging results, especially for bright sources, but that its effect on the photometry of faint sources appeared to range between neutral and negative. We also investigated photometry via parabolic fits, as used in Bannister et al. (2011a,b). The full ASGAR dataset should have enough flux measurements to be able to investigate the impact of these and other techniques in a more statistically rigorous manner. We also hope to improve the means by which we extract the raw photometry, as examination of the residual images from both the deconvolution and source-fitting processes indicates that there is room for improvement.

For those compact sources that are found within a region of more extended emission, our LSS subtraction technique naturally leads to the possibility of an additive flux bias in our measurements. We constructed a set of images in which the LSS model was added back to each epoch image after convolution with the synthesized beam. We then extracted photometry again using a specialized routine that simultaneously fit for the flux of a source in each image on top of a local, constant background term. This more elaborate procedure did not produce more consistent results, and we found that our LSS subtraction procedure did not lead to consistent flux biases.

In Figure 5.1 we compare our deep and epoch flux measurements after applying the above corrections. There is very good agreement in the scales between the two. Cyg X-3 is clearly the most significantly variable source we observe. The nominally very bright sources are all detected at very large distances from the pointing center and are subject to highly uncertain PB corrections. In order of decreasing deep image flux density, these are DR 22 (separation $110'$; PB correction ~ 340), DR 21 ($80'$; ~ 25), and 4C 44.32 ($99'$; ~ 108). The last PB factor listed obeys a different Gaussian relation than the first two because of the different focus settings used in the *Kepler* and Cyg X-3 deep images.

5.3.7 Timescales

We sample the variability of our catalog sources on timescales of days to years. If a given source is measured n times at dates t_i , there are $n(n-1)/2$ possibly-redundant intervals sampled, $|t_i - t_j|$ for $i < j$. We plot these intervals in Figure 5.2 for the three samplings in this work. Coverage is fairly uniform across the range of probed timescales. The variability metrics that we describe in §5.4.5 do not take the temporal sampling into account, however, and so they are not optimally sensitive to variability on any one given timescale. A structure-function analysis (Simonetti et al. 1985; Emmanoulopoulos et al. 2010) would take into account the

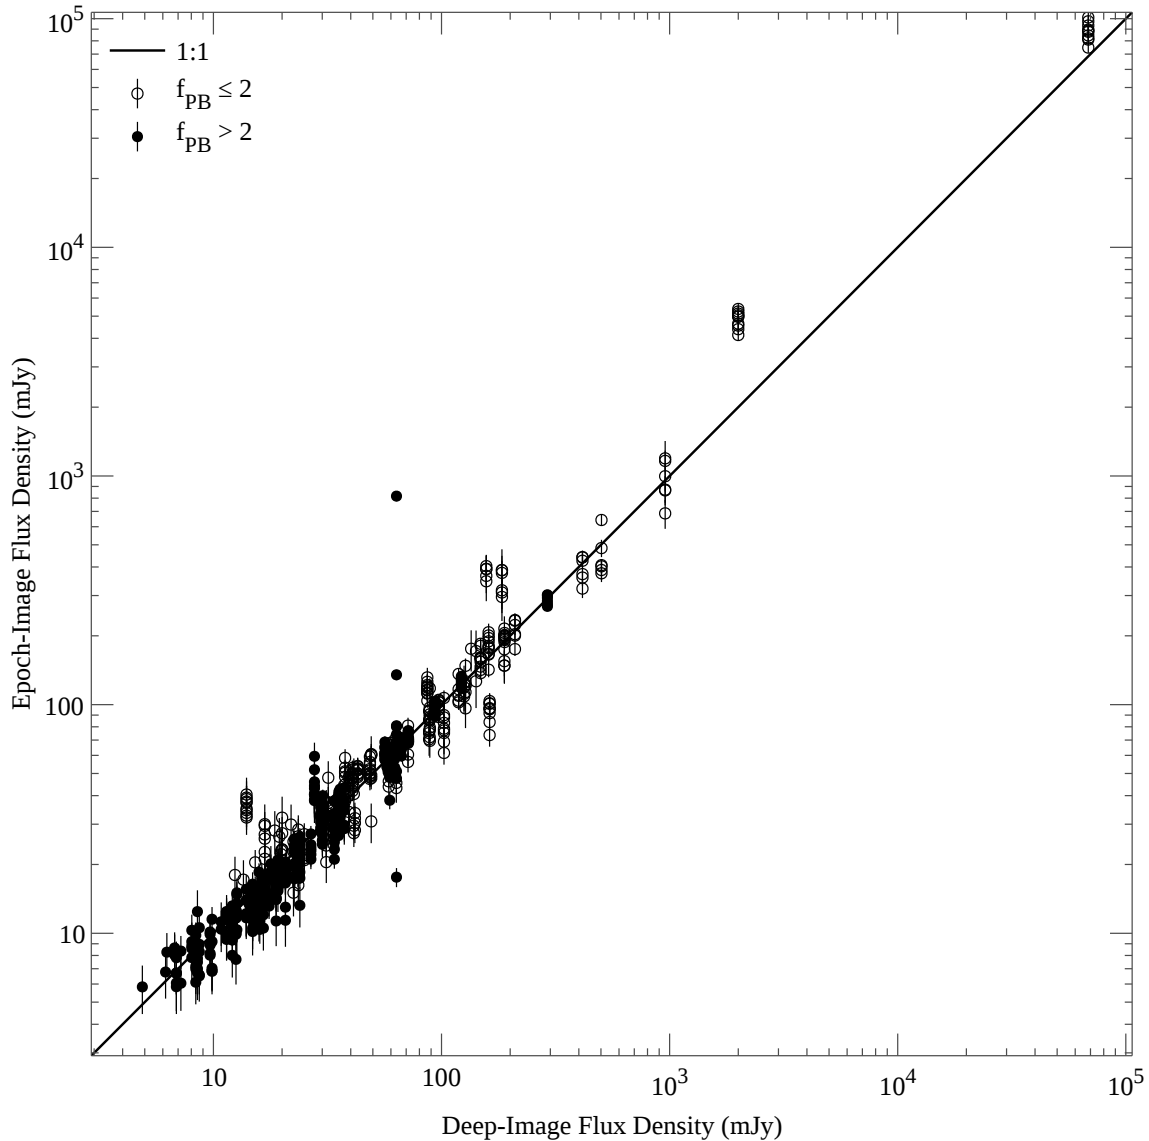


Figure 5.1: Comparison between fluxes of cataloged sources as measured in the deep images and those measured in individual epochs. Agreement is generally very good, especially for those sources for which the PB correction is not large. The vertical set of points with a large spread corresponds to the highly variable source Cyg X-3. Several sources have a discernable offset between the deep image and epoch fluxes; these are seen in Cyg X-3 pointing, and the offset arises from disagreements between the large-radius behavior of the optimally focused PB model (used to obtain the deep flux) and that of the defocused PB model (used to open the epoch fluxes). The nominally very bright sources are subject to large, uncertain PB corrections; see §5.3.6.

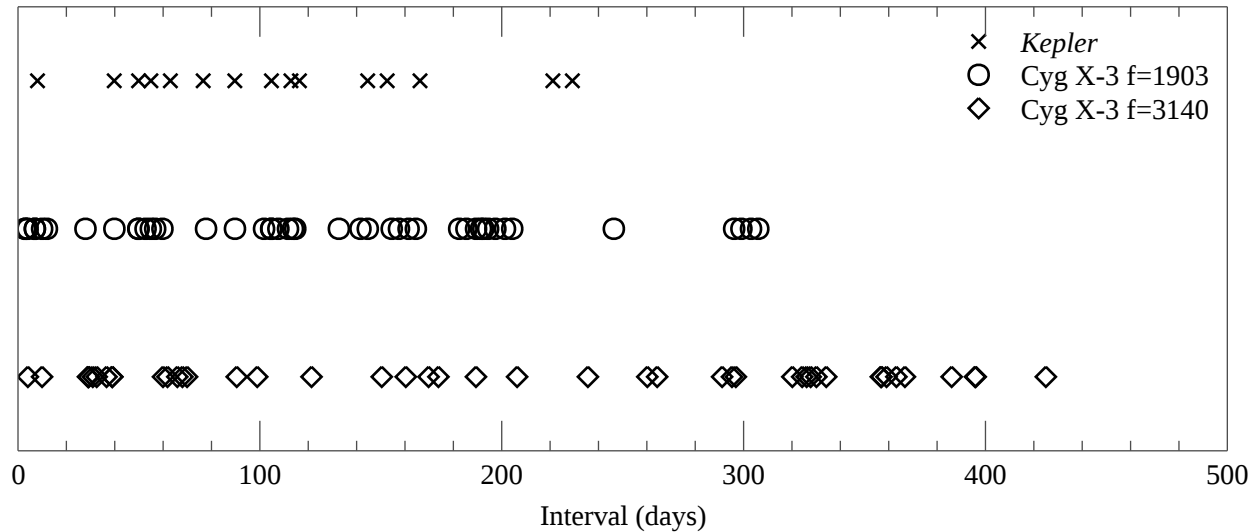


Figure 5.2: Intervals between observations for the three sampling patterns present in this work. Different samples are not independent because for n samples there are $n(n-1)/2$ intervals probed. A higher density of measurements around a certain timescale, however, generally indicates increased sensitivity to variability on that timescale.

timing of each observation but is not undertaken here (§5.4.5). Ofek et al. (2011) find that the structure function of Galactic radio variables saturates at $\tau \sim 10$ d, remaining flat to at least $\tau \sim 60$ d, suggesting that the more detailed analysis would not significantly affect our results. Ofek & Frail (2011) find a flat structure function for variable extragalactic radio sources on timescales from ~ 250 – 3250 days.

With additional analysis, our observations can also be sensitive to variability on timescales around the observation time of a typical epoch image, ~ 1 hour. Evolution at and below this threshold can be searched for by imaging subsets of the data within each epoch, with the expected tradeoff between time resolution and sensitivity. Our maximum time resolution is the ATA correlator integration time of 10 s. Any single-epoch transients discovered in the complete dataset will be examined for intra-epoch evolution in this manner, but in this work we find no sources that merit this detailed investigation. In the gap between our \sim hour integration time and our \sim week observing cadence, we are potentially sensitive to events (*e.g.*, a lucky observation may occur precisely during an hour-long stellar flare) but can only poorly constrain their duration.

Our sensitivity to slowly-evolving sources might correspondingly be improved by imaging, *e.g.*, months of data at a time, and searching for sources too faint to be detected in a single epoch but too variable to be detected in the deep image, as in Bower et al. (2007) and other works. (This is simply a rough matched-filter approach.) Because our images are not limited by thermal noise it is unclear how much of a benefit this technique would provide in practice.

We defer an investigation of this approach to future work.

5.4 First Results

In the following subsections we present the first results from analyzing the subset of ASGARD data described above.

5.4.1 Images

The deep image of the *Kepler* pointing at 3 GHz is shown in Figure 5.3. All of the observations of this pointing were made with a focus setting of 1.9 GHz, so the PB is broadened as compared to the nominal, optimally-focused configuration.

The optimally-focused deep image of the Cyg X-3 pointing at 3 GHz is shown in Figure 5.4. The LSS and FOV may be compared with the 1.4 GHz image made with the Westerbork Synthesis Radio Telescope (WSRT) presented in Setia Gunawan et al. (2003, fig. 2; S+03 hereafter). The ATA's compact configuration and large FOV make it much more sensitive to LSS than the WSRT, and it is notable that the image shown in Figure 5.4 is produced from (multiple epochs of) a single pointing. We detect significant LSS out to a radius of $\sim 45'$ from the pointing center.

Figure 5.5 shows an image made from a single Cyg X-3 epoch (2011 Feb 01) after subtraction of the LSS. The Cyg X-3 field contains three bright sources near the half-power point that are somewhat problematic in the subtraction: DR 7 (west), DR 15 (south), and 18P 61 (northeast; Wendker et al. 1991). All of these are at least marginally resolved and embedded in cuspy extended emission. The residuals due to these sources nonetheless do not significantly impair the imaging of each epoch. Not shown in the figures, but easily detectable in the Cyg X-3 epoch images, are DR 22 and DR 21, as described above.

We generated a total of 29 epoch images, 23 of the Cyg X-3 pointing (dropping six analyzed epochs with dynamic range issues caused by the 2011 March major flare), and 6 of the *Kepler* pointing. In Figure 5.6 we plot the representative rms of each image (as reported by `sfind`) as a function of integration time. In Figure 5.7, we show the number of sources detected in each image.

To assess the effectiveness of the LSS subtraction process, we compare image rms values in regions with varying levels of LSS flux. We selected several source free regions in the Cyg X-3 and *Kepler* fields of 50×100 image pixels. For each epoch image and each region, we computed the ratio of the rms in that region to the rms in an equally-sized region on the outskirts of the image, away from all source emission. Taking this ratio compensates for the varying noise baseline of each image. Figure 5.8 summarizes these ratios as a function of the mean LSS model flux associated with each source-free region. (Measurements from the *Kepler* field, in which there is no LSS subtraction, are assigned a mean LSS model flux of zero.) The rms ratios increase above the baseline as LSS becomes more significant, but

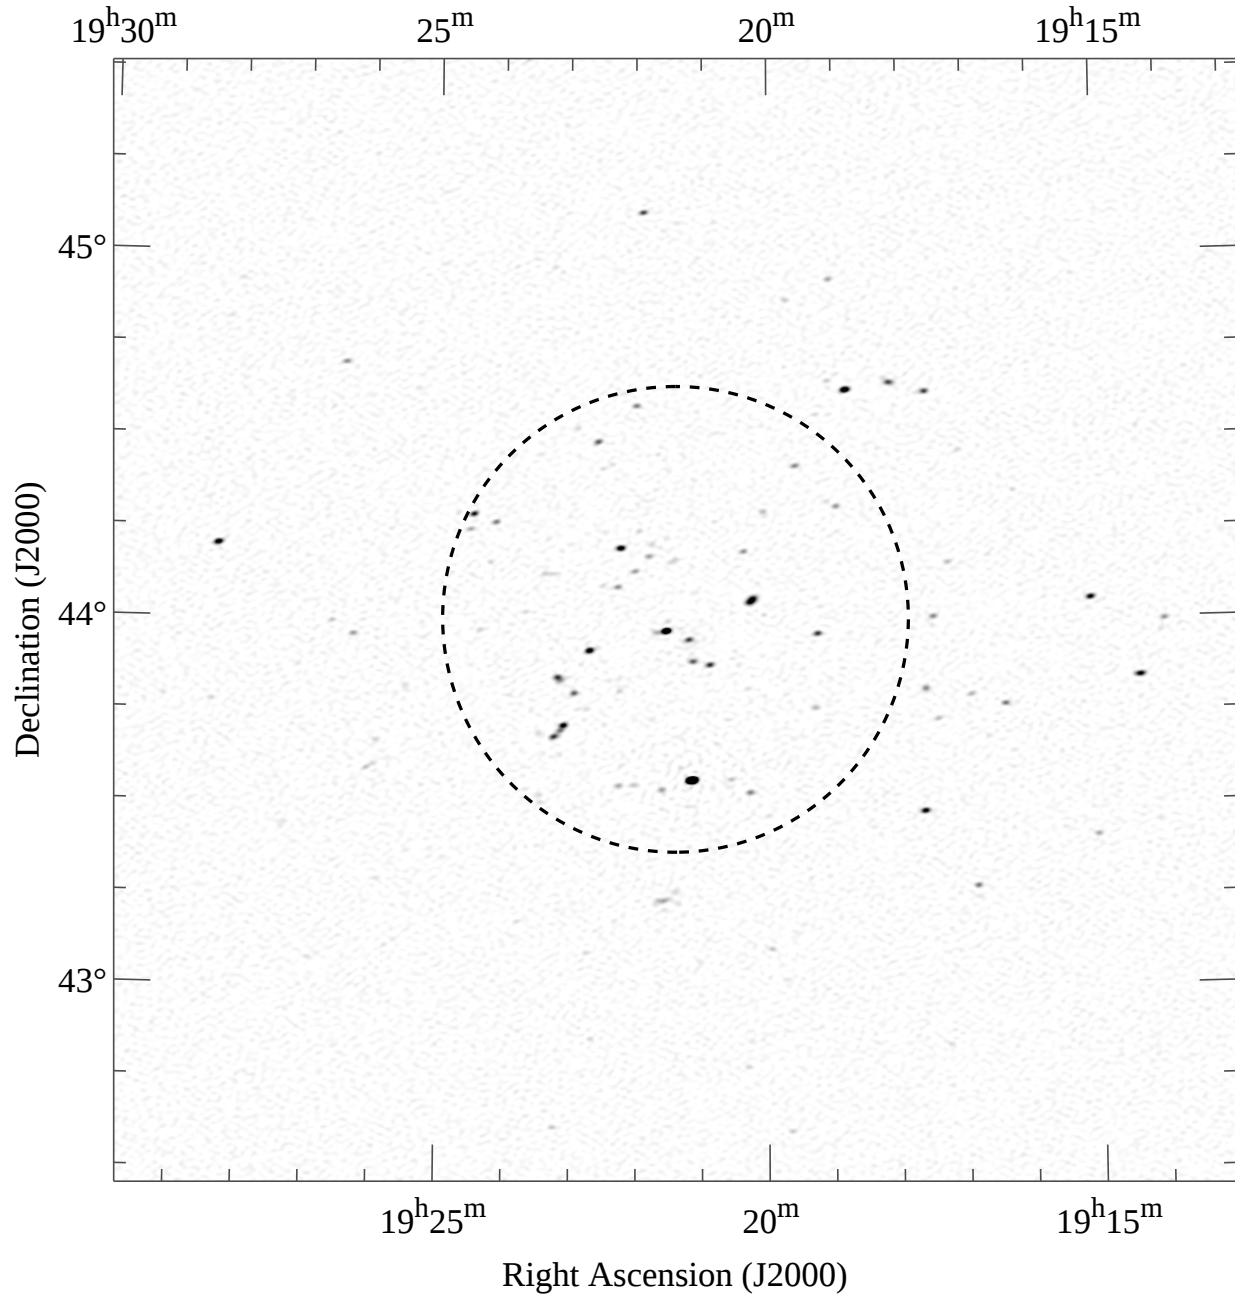


Figure 5.3: Deep image of the *Kepler* pointing, uncorrected for PB attenuation. The grayscale is linear from zero (white) to 20 mJy/beam (black). The maximum brightness in the image is ~ 200 mJy/beam. The rms deconvolution residual is 0.5 mJy/beam. The HPBW for this image is 1.27° and is denoted by the dashed circle. The synthesized beam is $65'' \times 41''$ at a position angle of 98° .

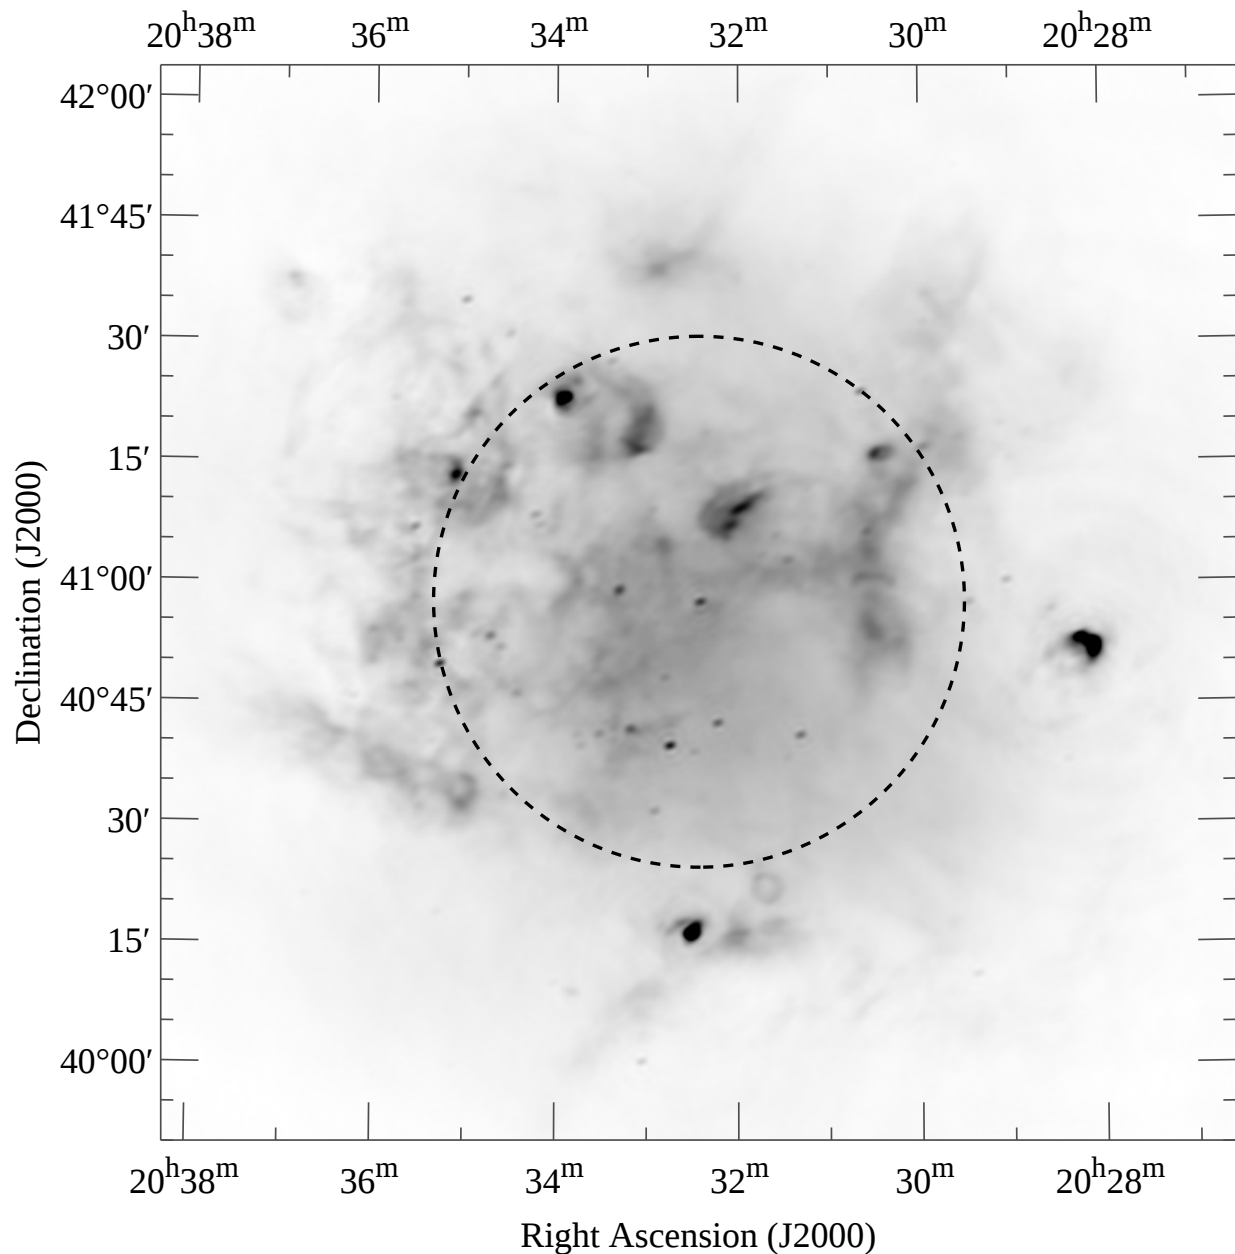


Figure 5.4: Deep image of the Cyg X-3 pointing, uncorrected for PB attenuation. The grayscale is linear from zero (white) to 150 mJy/beam (black). The maximum brightness in the image is ~ 500 mJy/beam. The rms deconvolution residual is 1 mJy/beam. The HPBW for this image is 1.1° and is denoted by the dashed circle. The synthesized beam is $60'' \times 42''$ at a position angle of 106° .

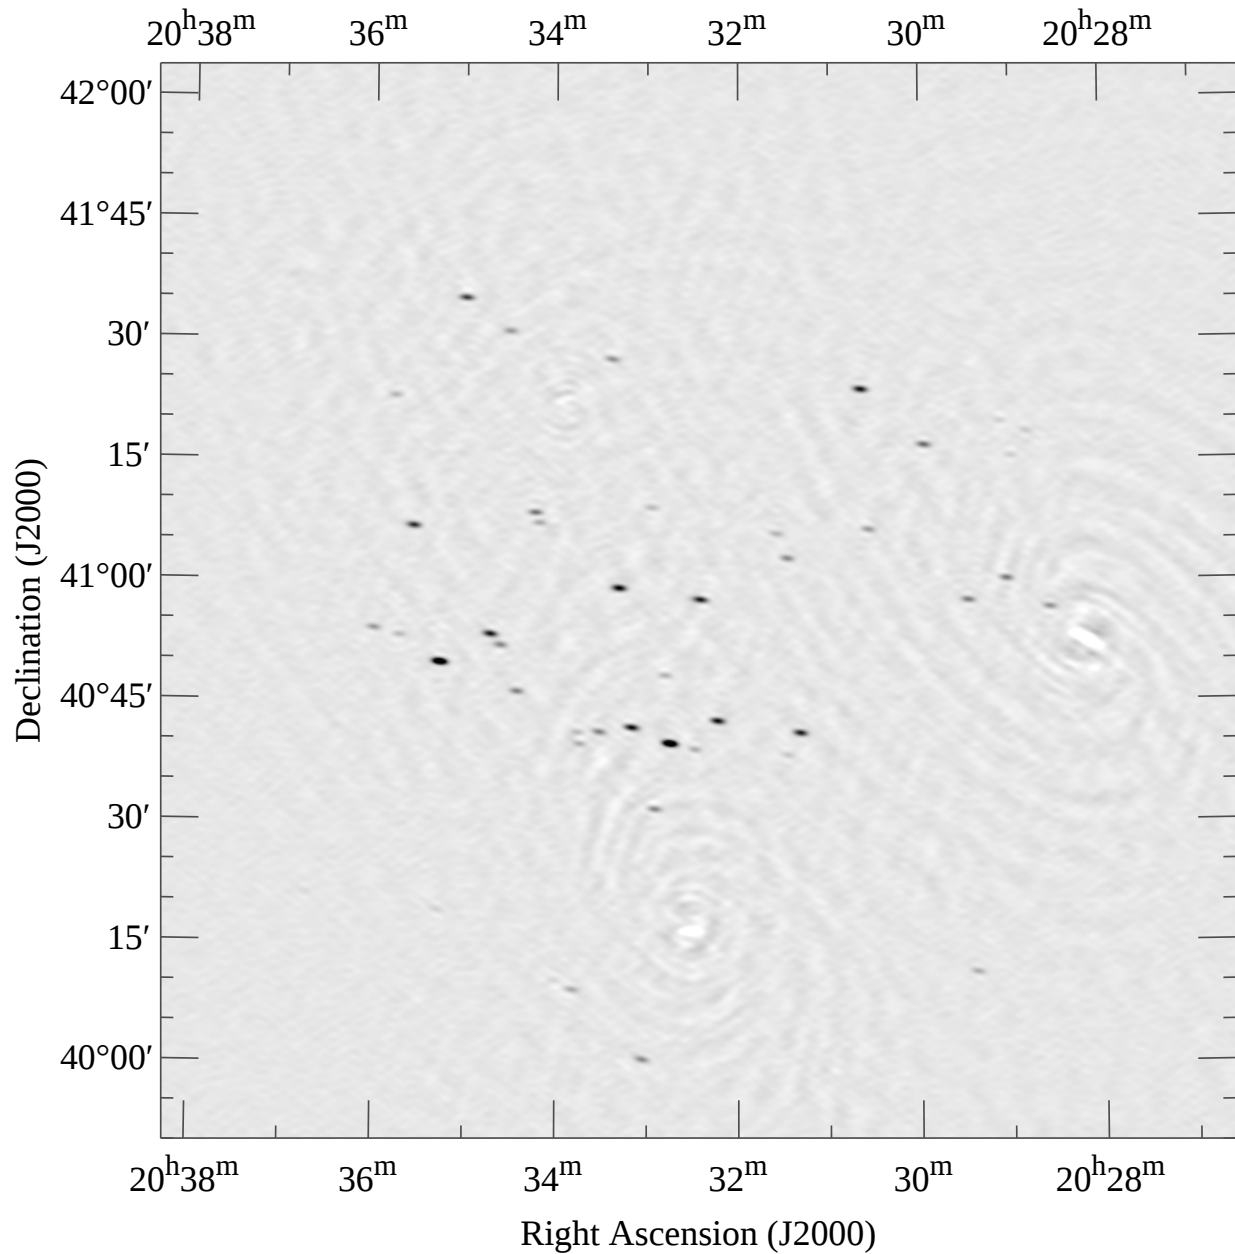


Figure 5.5: Image of the Cyg X-3 pointing from 2011 Feb 01 with LSS subtracted. The grayscale is linear from -5 mJy/beam (white) to 50 mJy/beam (black). (Note that this is different than that used in Figure 5.4.) The maximum brightness in the image is ~ 380 mJy/beam. The rms deconvolution residual is 0.9 mJy/beam. The synthesized beam is $85'' \times 36''$ at a position angle of 83° .

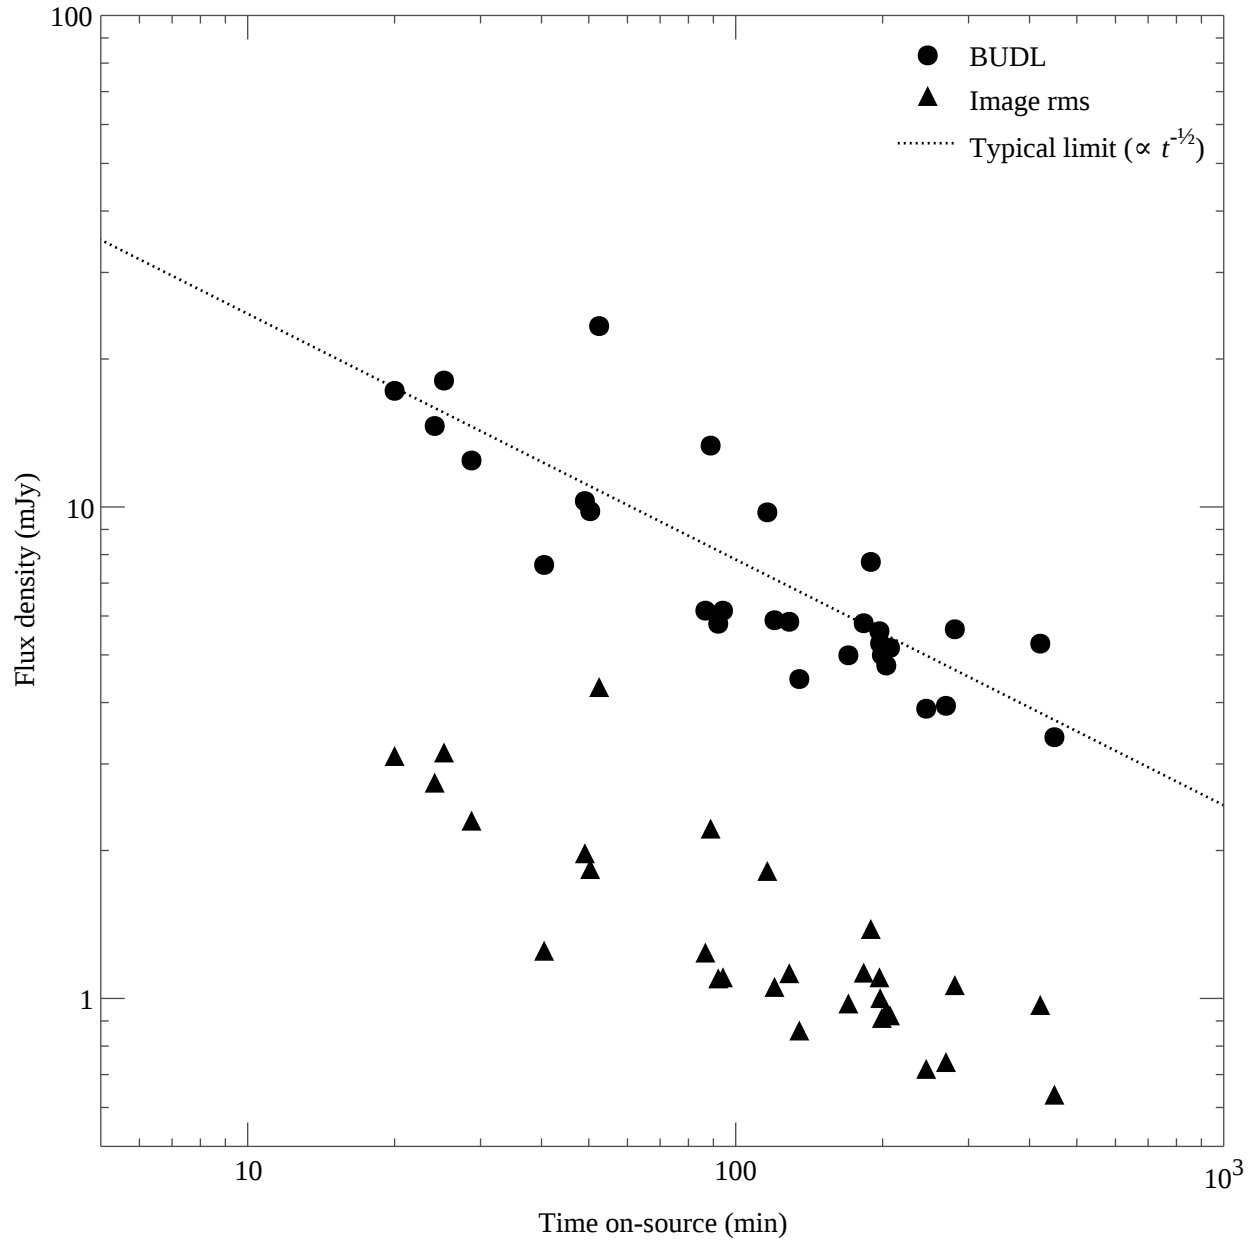


Figure 5.6: Noise statistics for epoch images as a function of on-source integration time t . *Triangles*: representative rms values as reported by `sfind`. *Circles*: blind unresolved-source detection limits (BUDLs; §5.4.3). At fixed t , achieved values will vary with system temperature, data flagging, calibration quality, and the source detection cutoff dynamically determined by the `sfind` FDR algorithm. The *dashed line* shows the typical effective BUDL assuming a scaling of $t^{-1/2}$; for $t = 10$ min, this corresponds to ~ 24 mJy. Both sets of numbers refer to PB-uncorrected images and hence *apparent* flux densities.

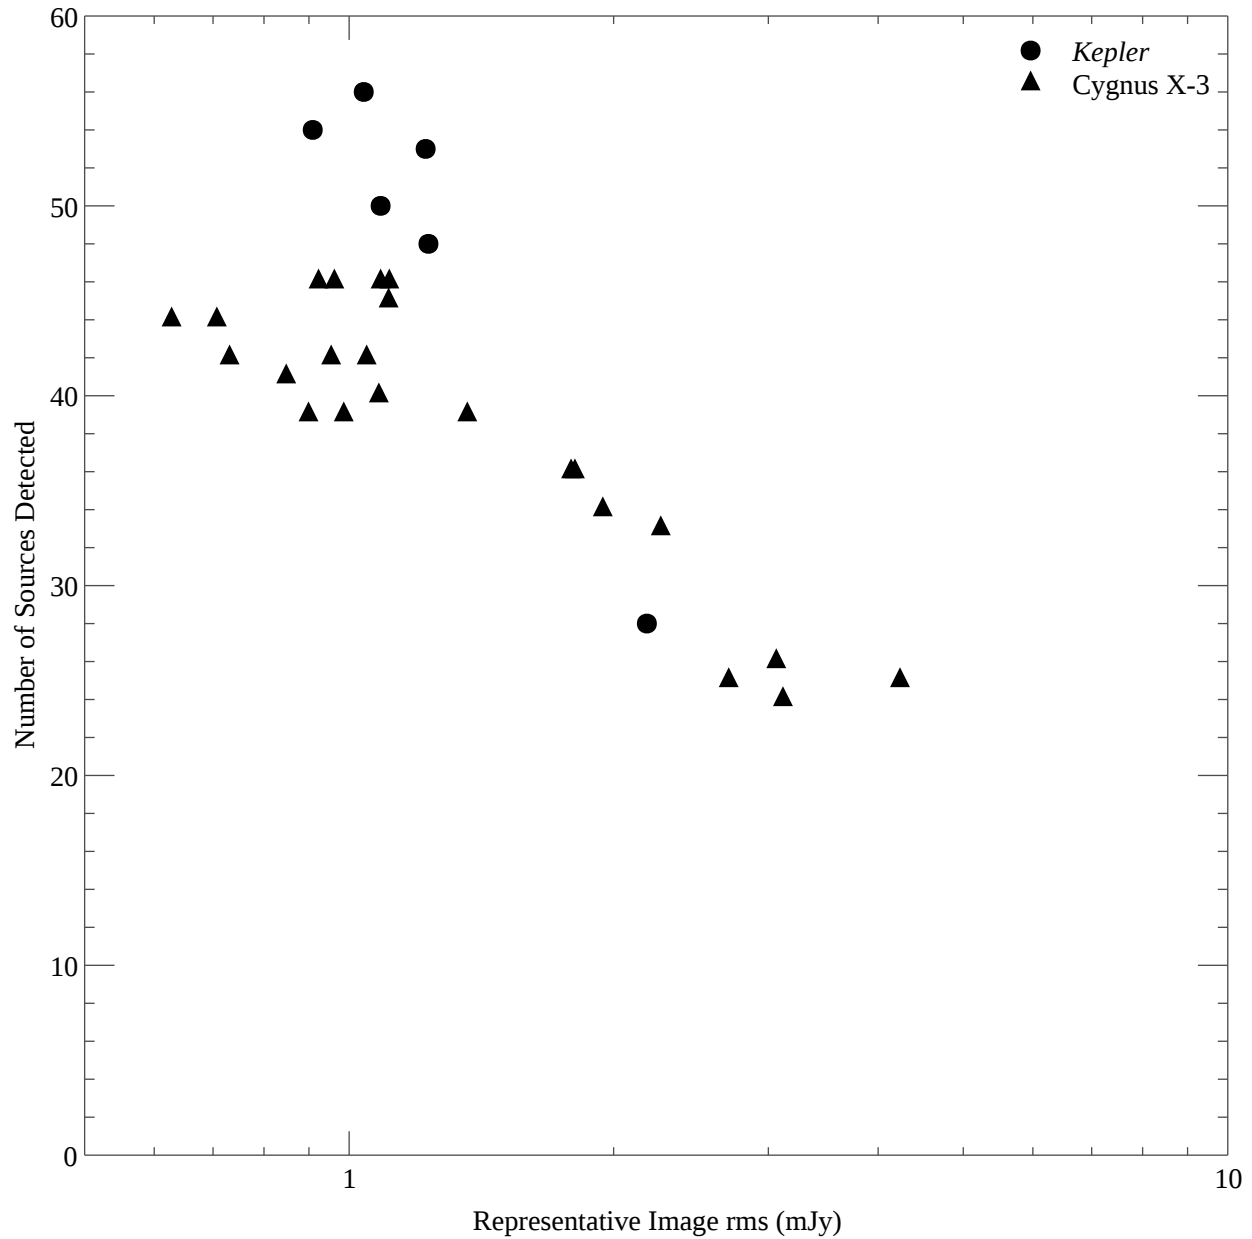


Figure 5.7: Number of sources detected in each epoch image, as a function of the representative image rms reported by `sfind`. Symbols are grouped by pointing. There are 86 cataloged sources associated with the *Kepler* pointing and 48 with the Cyg X-3 pointing. The detection rate for the *Kepler* pointing is lower because the *Kepler* deep image is much less limited by systematics than the Cyg X-3 deep image, and so contains many more faint sources that are cataloged but cannot easily be detected in the epoch images.

saturate at a factor of ~ 2 for mean LSS fluxes of ~ 0.4 mJy. We placed a group of regions especially near the three bright sources of the Cyg X-3 field. These display somewhat higher rms ratios than other regions in areas distant from these bright sources, but the difference is slight.

To give a sense of the areal coverage of the ASGARD GC field, we show a preliminary mosaicked image of 3 GHz GC data in Figure 5.9. LSS is not subtracted in this image. The missing pointings have coverage at lower frequencies but not at 3 GHz. The Sgr A complex, with structured emission reaching brightnesses of 50 Jy/beam, presents a clear challenge. The data contributing to Figure 5.9 come from only a few epochs, so the more thorough hour angle coverage of the complete dataset will make a significant difference to image quality. Joint deconvolution of some or all of the pointings might greatly improve the deep map, although this would depend strongly on how well the ATA PB can be modeled.

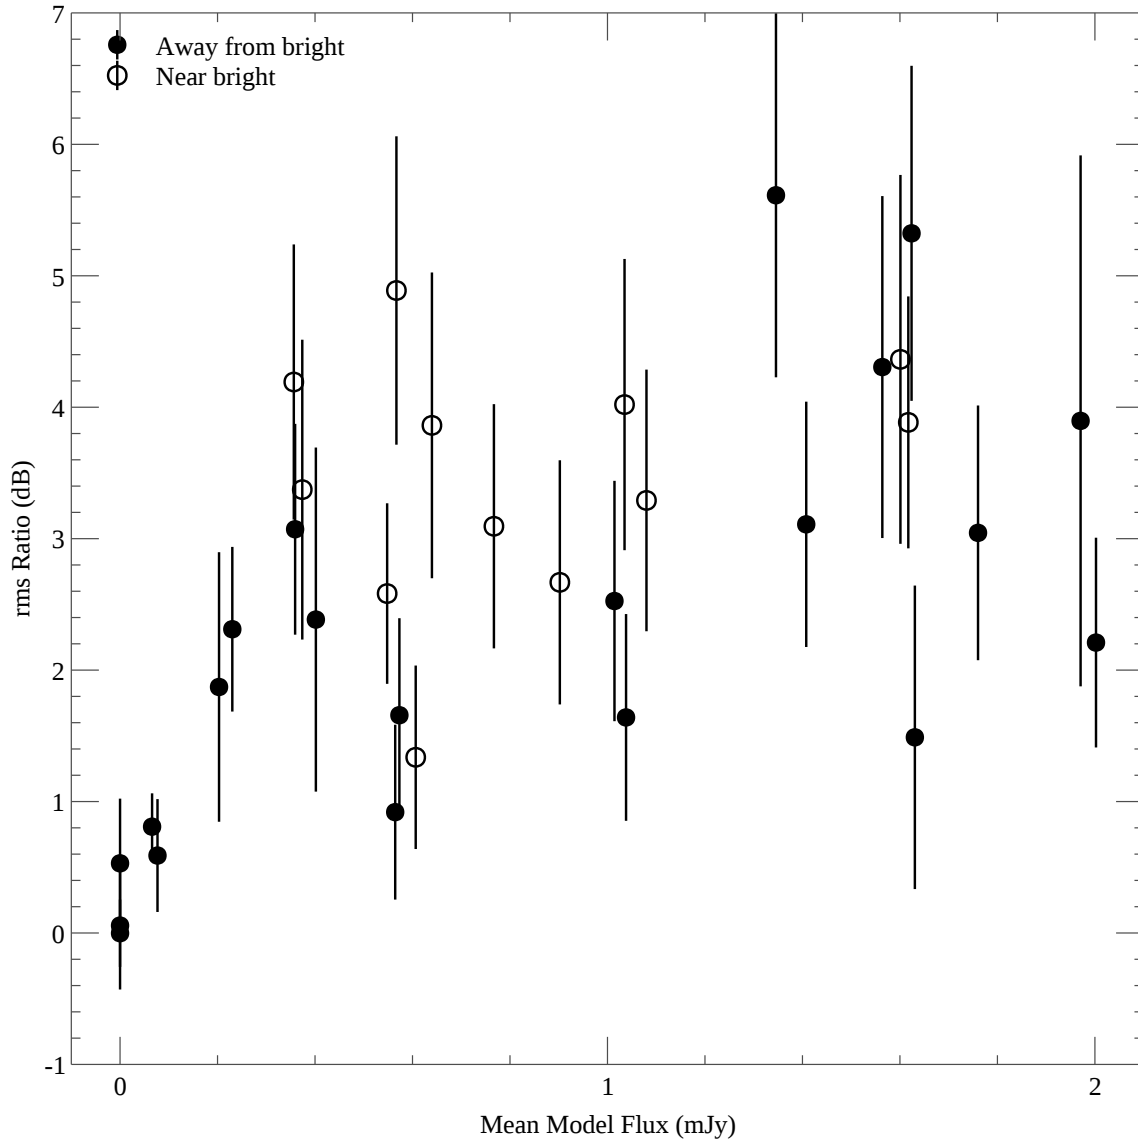


Figure 5.8: Measurements of the increase in image noise as a function of LSS flux. Source-free regions of 50×100 pixels were selected in the *Kepler* and Cyg X-3 fields. For each epoch image and each region, the rms was computed and divided by the rms in an equally-sized box in the lower left corner of each image, far away from all source emission. Each plotted point gives the mean and standard deviation of these rms ratios over all imaged epochs for a particular pointing, source-free region, and focus setting. The *abscissa* is the mean flux of the LSS model of that region. *Open circles* denote measurements from regions very near (~ 20 pixels) any of the three bright sources of the Cyg X-3 field. Areas in which there is significant LSS show rms increases of a factor of ~ 2 versus the source-free baseline. See §5.4.1 for discussion.

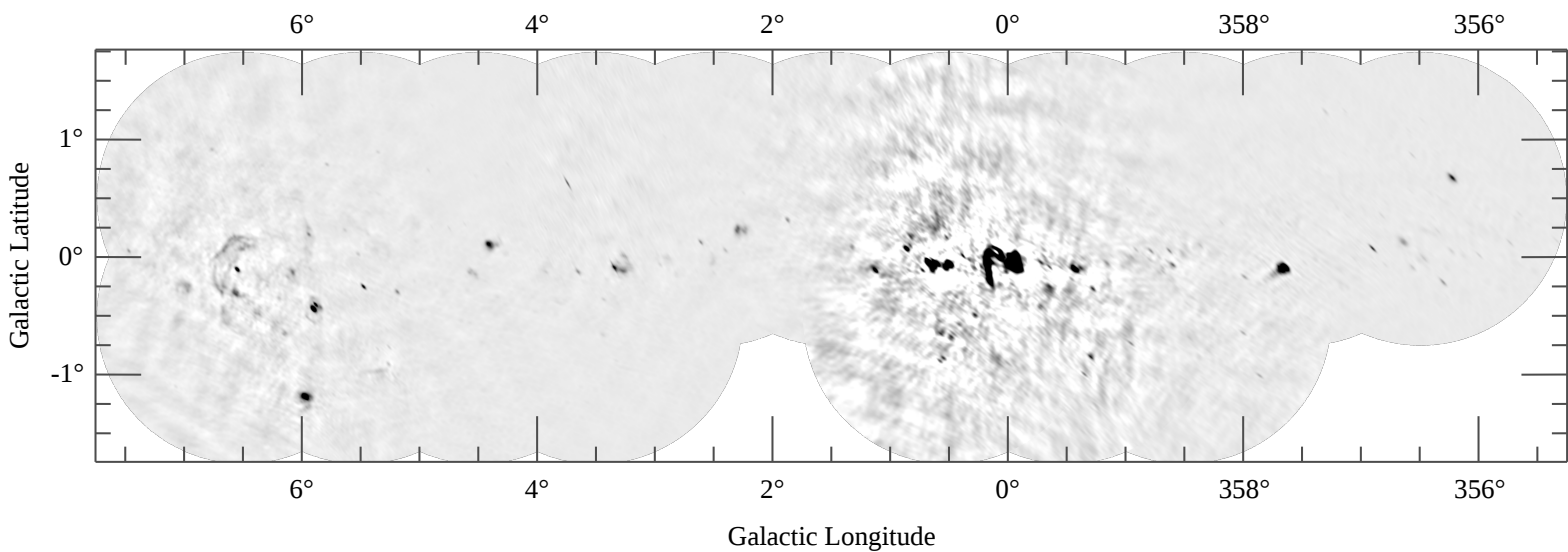


Figure 5.9: Preliminary mosaicked image of ASGARD GC field. The grayscale is linear from -70 mJy/beam (white) to 800 mJy/beam (black). The maximum brightness in the image is ~ 3920 mJy/beam. Each pointing has been imaged to a diameter of about 2.2 times the HPBW. Most pointings have only a single epoch of observations contributing to the image.

5.4.2 Source Catalog

We robustly detect 134 compact sources in the deep images of the two pointings. There are 86 sources detected in the *Kepler* region and 48 in the Cyg X-3 region, the difference between the two being primarily due to the effects of LSS in the latter pointing. The faintest source discovered in the current *Kepler* deep image has a flux of 2.6 mJy. The faintest source discovered in the Cyg X-3 image has a flux of 9.2 mJy.

Because it is 13.5° out of the Galactic plane and contains no LSS, the *Kepler* pointing provides a clean testbed for our catalog. All but one of the sources associated with the *Kepler* pointing in our catalog has a match to a source in the NVSS catalog within $20''$, and none of these match multiple NVSS catalog entries. The probability of an individual chance match between the catalogs at this positional tolerance is $\sim 0.5\%$ (Condon et al. 1998). The unmatched source is visible, but marginal, in the NVSS imagery. We define a set of NVSS sources that might be detected in the *Kepler* deep image as those that are projected to have a PB-corrected (ATA-apparent) flux greater than 2.6 mJy, so long as the PB correction factor f_{PB} is smaller than two. Two-thirds (43/66) of these sources are detected, with most of the nondetections being marginally discernable in the ATA imagery. Meanwhile an additional 41 NVSS sources are detected that do not meet the above criteria. The disjunction between the detections is a result of some combination of spectral dependence, uncertain PB modeling, incompleteness to marginal detections, and possible genuine source variability. Of the undetected NVSS sources, the maximum predicted ATA flux is 5.6 mJy; that is, no NVSS sources reduced their flux by a measurably significant amount by the time of our observations. All NVSS sources with predicted ATA fluxes ≥ 5.8 mJy and PB correction factors less than 2 are detected. If no PB correction limit is applied, the limit is 17.7 mJy, with a maximum applied PB correction factor of 21.0, corresponding to an angular separation of $79'$.

NVSS imaging of the Cyg X-3 region is quite poor and so we do not compare the catalogs for this region. Instead we use the S+03 L-band catalog of compact sources in the Cygnus OB2 region. Of our catalog entries associated with the Cyg X-3 pointing, all but 8 are matched to S+03 sources. Of those, four are outside of the S+03 FOV, one is clearly detected but is too extended to meet their selection criteria, and three genuinely do not appear to be detected in the S+03 maps, in either their L-band or 350 MHz observations. Using the same criteria as above, five S+03 sources might have been expected to be detected with predicted apparent ATA fluxes of ~ 10 mJy but were not. One of these is blended with the very bright DR 15. The remaining four all have negative spectral indices (α , where $S_\nu \propto \nu^\alpha$), which we do not account for in our initial predictions due to inconsistent availability of information in the S+03 source catalog. Taking the spectral index into account, we do not expect to detect any of these sources. All S+03 sources with predicted ATA fluxes ≥ 11.2 mJy and PB correction factors less than 2 are detected.

In Figure 5.10 we compare various source counts. In order to avoid uncertain PB corrections, we restrict our assessment to sources within the half-power point, reducing our

catalog to 43 and 27 sources within the *Kepler* and Cyg X-3 pointings, respectively. With such limited numbers, differential functions are extremely uncertain, and so we report cumulative source count functions. In the *Kepler* pointing, our results are comparable to those of the NVSS, which is expected given the results reported above. Our *Kepler* counts are somewhat higher than those found in [PiGSS-I](#), probably due to a combination of higher-quality deep imaging and manual source identification, both made possible by the much smaller ASGARD survey footprint. Our source counts in the Cyg X-3 pointing are somewhat lower than those of [S+03](#), which may reflect the different observing frequencies of the two surveys, although such an effect would also be relevant to the NVSS comparison. Both the ASGARD and [S+03](#) source counts are uncertain on the 20–50% level due to small-number statistics.

5.4.3 Detection Limits and Completeness

In order to analyze searches for transient and highly-variable sources, we must understand the completeness of these searches. The relevant detection limit for transient searches such as our **sfind** step is, by definition, the *blind* source detection limit. Furthermore, because transient sources will be unresolved, the detection limit may be expressed as a flux density limit rather than a brightness limit. As a shorthand we thus refer to this particular quantity as the “BUDL”: blind unresolved-source detection limit. An image that has not been corrected for PB attenuation has a single value of the BUDL in terms of *apparent* flux density if imaging artifacts are not significant. Figure 5.6 depicts the variation of this value as a function of integration time for the epoch images analyzed in this work. The BUDL in terms of *intrinsic* flux density varies spatially as defined by the PB model. For each image, the apparent-flux BUDL is a function of the thermal noise; the additional effective noise due to limitations in calibration, deconvolution, etc.; and the source detection cutoff determined by the **sfind** FDR algorithm, which varies from image to image but is typically $\sim 5.1\sigma$. Despite the latter effects, we generally find that $\text{BUDL} \propto \tau^{-1/2}$ as would be expected, where τ is the image integration time.

We find an upper limit to the flux-dependent completeness of our source finding method by inserting false sources into the images and attempting to recover them. This approach only finds an upper limit to the completeness because we insert the sources mid-way through our processing pipeline, after the *u-v* data have been initially calibrated. We experimented with inserting false sources in both the *u-v* and image domains and obtained equivalent results with both approaches, so for more detailed studies we used image-domain insertion, which avoids the need to rerun the computationally-expensive imaging steps.

We began by sampling the differential completeness at fixed *apparent* flux $c_a(S_a)$, that is, without accounting for PB attenuation. To sample this function, we spread ~ 500 unresolved sources of flux density S_a throughout all of the epoch images, spacing them randomly but evenly in position angle and radius relative to the pointing center to avoid overlaps. (This is fair because we expect any image to contain at most one transient.) The blind source detection routines are then run and the fraction of detected sources is reported as the completeness,

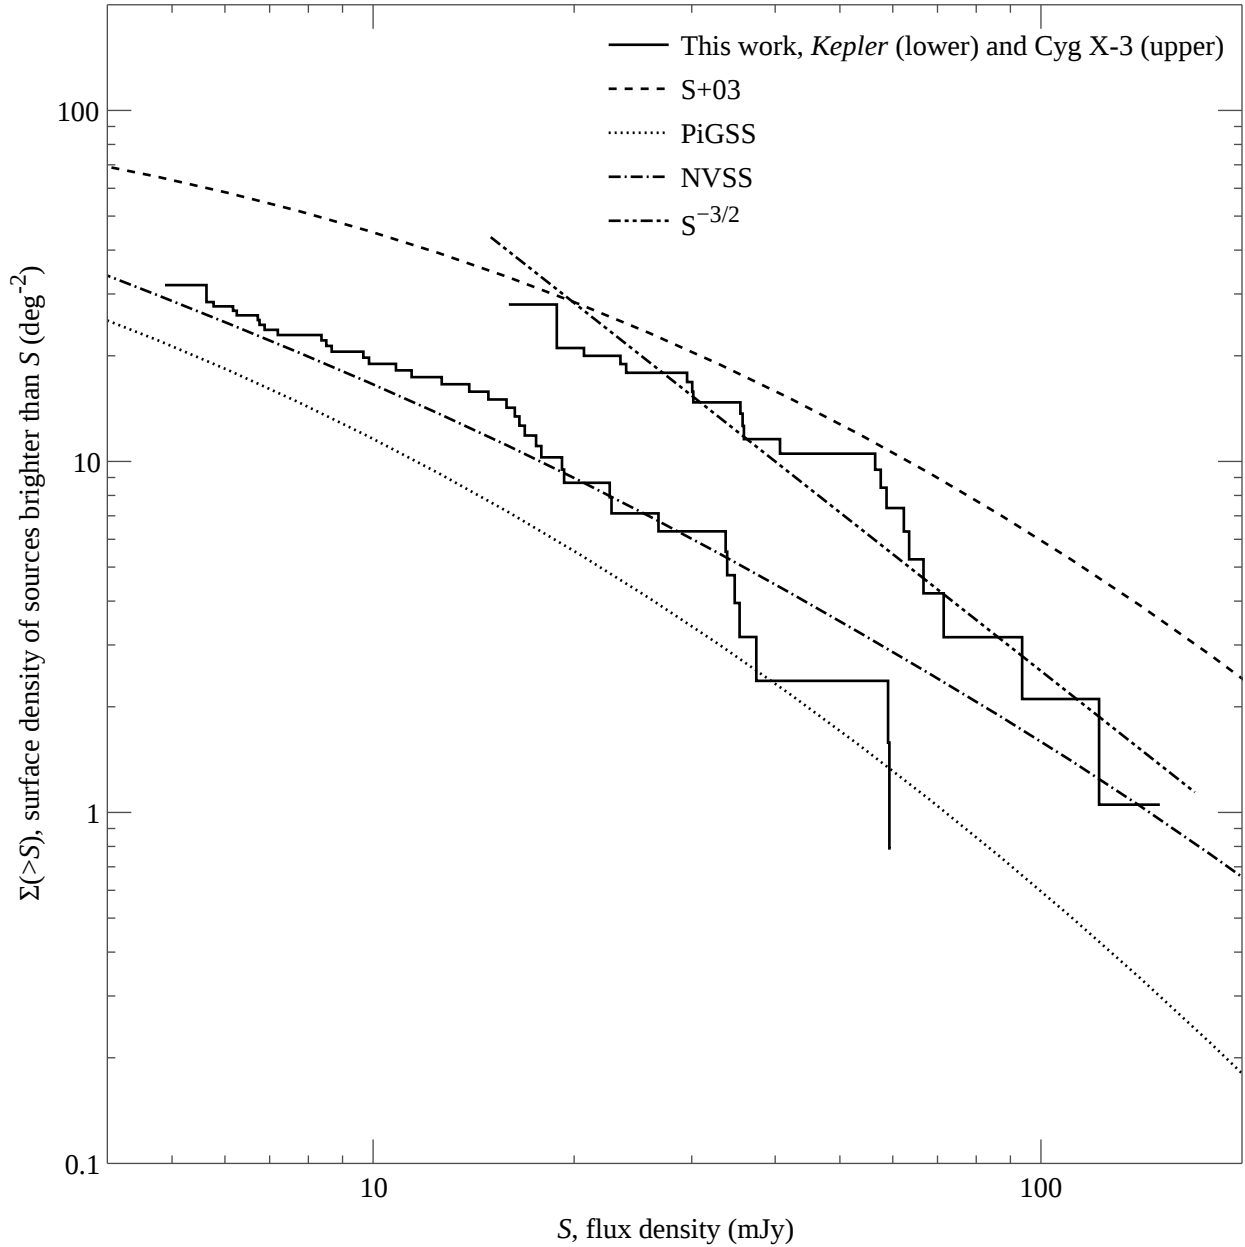


Figure 5.10: Cumulative source counts in this and other surveys. ASGARD measurements are derived from the cumulative distribution function of source flux densities as detected in the *Kepler* and Cyg X-3 deep images, combined with effective search areas based on our analytic PB models. Reference values are derived from differential source count measurements reported in S+03, the NVSS (Condon et al. 1998), and PiGSS-I. The arbitrarily-normalized $S^{-3/2}$ line shows the expected scaling for a Euclidean, volume-limited distribution. See §5.4.2 for discussion.

assigning an uncertainty to the measurement assuming a Poisson process. The top curve of Figure 5.11 shows the results of sampling $4 < S_a < 200$ mJy and a four-parameter analytic fit to these results. Our model function is

$$\tilde{c}_a(S_a) = \begin{cases} 0 & \text{if } S_a < (-D/B)^{1/C}, \\ A \tanh(BS_a^C + D) & \text{otherwise.} \end{cases} \quad (5.2)$$

This can be made to match the data well empirically, plateaus as $S_a \rightarrow \infty$, and has realistic discontinuous behavior at the completeness zero point. The fitted parameter A , the plateau value, is 99.3%, and reflects the fraction of our images that are disqualified from blind source detection by virtue of proximity to a cataloged source.

We also sample the differential *intrinsic* completeness function $c_i(S_i)$, assuming our analytic PB model, by inserting false sources with appropriately PB-attenuated fluxes. Our results are also shown in Figure 5.11, along with a corresponding fitted model of the completeness function in which the plateau value was fixed to match that of the apparent completeness function. There is substantial disagreement between our samples of this function and an estimate derived from c_a and our image BUDLs, suggesting that the assumption of a spatially-uniform BUDL in each image does not hold strongly.

Finally, we compute the *cumulative* intrinsic completeness C_i ; that is, the completeness to all sources intrinsically brighter than some limiting value. In order to do this we must assume a distribution function for source number counts as a function of flux. We use the standard Euclidean, volume-limited distribution $N(>S) \propto S^{-3/2}$, which agrees well with the observed number counts of our catalog (Figure 5.10). The cumulative intrinsic completeness is then:

$$C_i(>S_i) = \int_{S_i}^{\infty} dS'_i c_i(S'_i) (S'_i)^{-5/2} \bigg/ \int_{S_i}^{\infty} dS'_i (S'_i)^{-5/2} \quad (5.3)$$

$$= \frac{3}{2} S_i^{3/2} \int_{S_i}^{\infty} dS'_i c_i(S'_i) (S'_i)^{-5/2}. \quad (5.4)$$

We evaluate this integral numerically using the analytic fit to our samples of $c_i(S_i)$. Below the cutoff of c_i at 4.4 mJy, $C_i(>S_i) \propto S_i^{3/2}$.

5.4.4 Search for Transients

We performed a search for transient radio sources with our subsample of processed ASGARD images. Recall that by our definition, “transient” sources are merely ones that are not detected in our deep images, and so must be searched for in the individual epoch images. Sources detected in this way will, by construction, evolve on relatively short timescales, but it is important to note that we are also sensitive to sources that vary on relatively long timescales: they are just detected in the deep images. Once detected, any transient sources are entered into our catalog and their photometry is measured at every epoch, so both

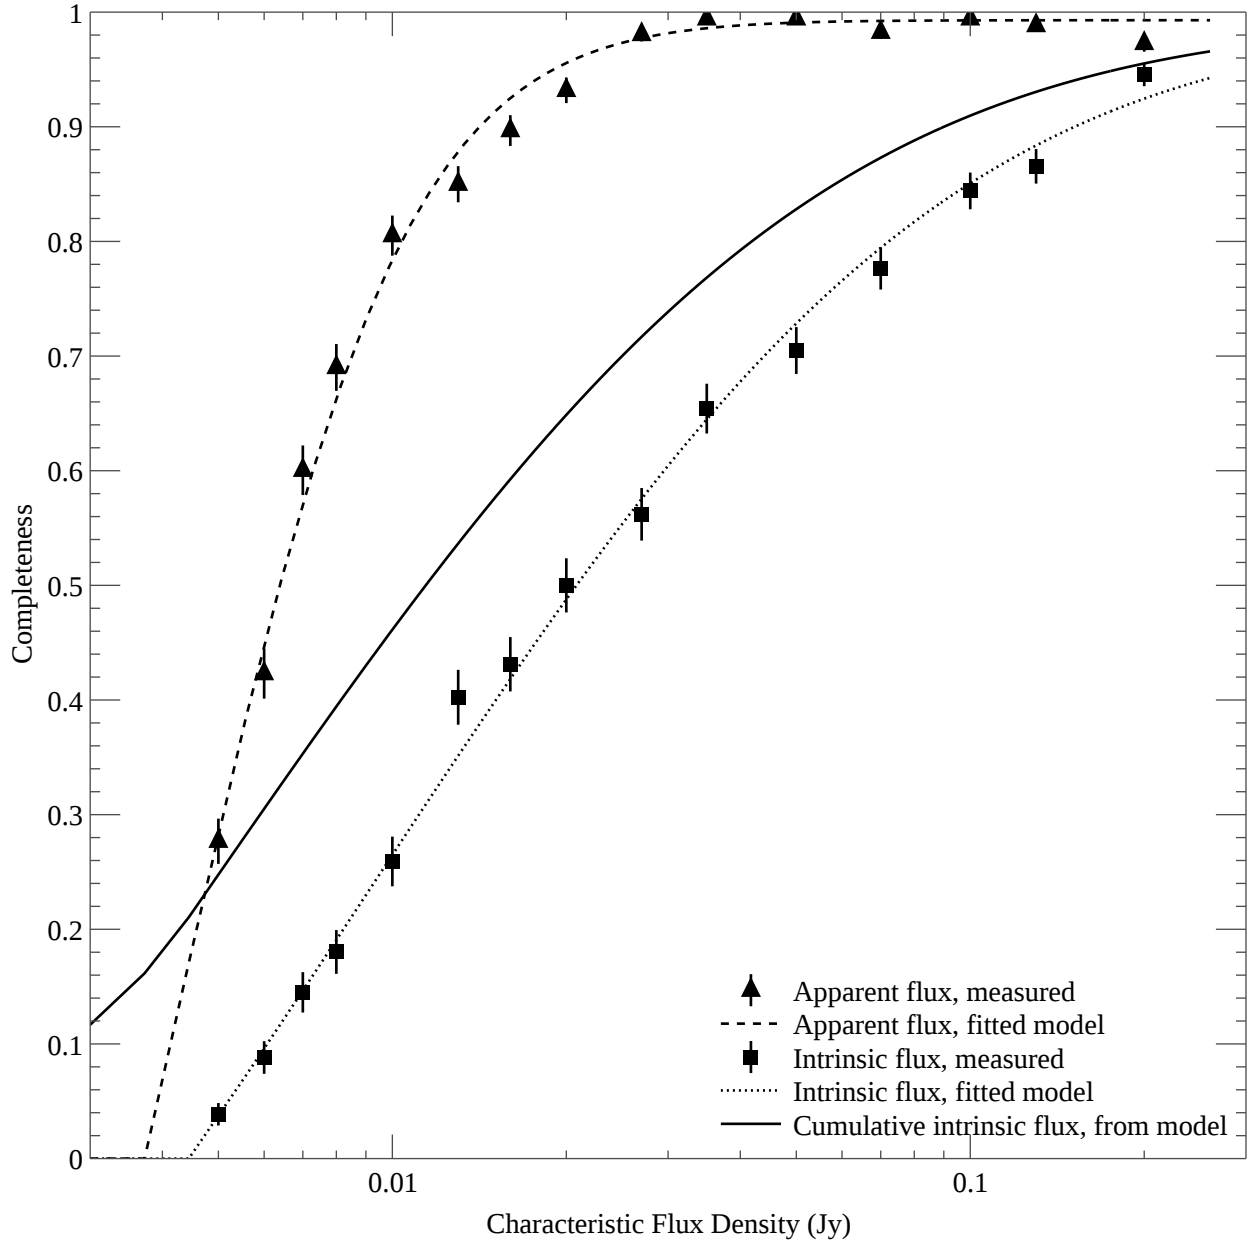


Figure 5.11: Completeness of blind source detection in various regimes, as determined from simulations with the current set of epoch images. *Apparent* completeness measures the detection fraction of sources of a given flux in images without correction for PB attenuation. *Intrinsic* completeness measures the detection fraction of sources of a single given flux after correction for PB attenuation. *Cumulative intrinsic* measures the expected detection fraction for all sources brighter than a given intrinsic flux, assuming luminosity function $N(>S) \propto S^{-3/2}$. See §5.4.3 for discussion.

classes of sources are analyzed identically, although in the stereotypical single-epoch case the “lightcurve” for a transient will consist of a numerous upper limits and a single detection.

After applying the `sfind` source detection and filtering procedure described above, we were left with a list of 17 transient candidates in our 29 epoch images (23 of the Cyg X-3 pointing excluding of the 2011 March major flare, six of the *Kepler* pointing). All but one were detected in the Cyg X-3 pointing. Visual inspection confirmed that all of the candidates were spurious, as indicated by various combinations of unphysical extended structure, a large distance from the pointing center, association with sidelobes or incompletely subtracted LSS, and/or indistinguishability from other noise fluctuations in the image. The full-scale ASGARD transient search will verify transient candidates rigorously, both by partitioning the imaged data to check for instrumental errors (cf. [Frail et al. 2012](#)) and by treating the image noise statistics quantitatively.

Although experience suggests that systematic effects are far more likely to cause false detections than thermal fluctuations, it is instructive to consider the limits imposed by noise ([Frail et al. 2012](#)). Our search examined a total area of $\sim 7 \cdot 10^5$ synthesized beams, accounting for the variable beam size and HPBW of each image. This corresponds to one expected noise event of 4.68σ . All of our candidates are above this threshold, so our search is not yet contaminated by statistical fluctuations. Extrapolating to the complete 3 GHz dataset, one noise fluctuation of $\sim 5.4\sigma$ may be expected.

We quantify the power of this early search by computing its effective search area as a function of the BUDL. Given the BUDL of an image $S_{a,\text{lim}}$, our analytic PB model, and the cumulative apparent-flux completeness function $C_i(S_i)$ determined in the previous section, the completeness-corrected effective solid angle in which sources with intrinsic flux density greater $S_{i,\text{lim}}$ may be detected is

$$\Omega_{\text{eff}}(S_{a,\text{lim}}, >S_{i,\text{lim}}) = C_i(>S_{i,\text{lim}})\pi r^2(S_{a,\text{lim}}, S_{i,\text{lim}}), \quad (5.5)$$

where r is the radius within which every source of intrinsic flux $>S_i$ will be detected, given a fixed BUDL:

$$r^2(S_a, S_i) = \begin{cases} 0 & \text{if } S_i < S_a, \\ 2\sigma_{\text{PB}}^2 \log(\frac{S_i}{S_a}) & \text{if } S_i < 90S_a, \\ (3\sigma_{\text{PB}})^2 & \text{otherwise.} \end{cases} \quad (5.6)$$

(The upper limit stems from our rejection of `sfind` sources found beyond the 98.9% attenuation point.) Following [Bower et al. \(2007\)](#), a transient survey consisting of n visits to the same field of solid angle Ω probes a total area of $\sim \Omega(n-1)$, if the time elapsed between visits is larger than the transient timescale. If the noise in each epoch varies, the effective area searched in each epoch varies as well, given a fixed $S_{i,\text{lim}}$. Denoting these areas Ω_i , where $\Omega_1 > \Omega_2 > \dots > \Omega_n$, the effective area searched in this case is $\sum_{i=2}^n \Omega_i$, if the area probed by each epoch is a strict subset of the area probed by every more sensitive epoch.

The above condition holds in the case of identical pointing centers and nonincreasing radial PB shapes, but not for mosaics with overlapping pointings. The mosaic case can be

treated numerically by explicitly mapping the areal contribution of each epoch. We take this approach to determine the effective area as a function of $S_{i,\text{lim}}$ for our processed data. Using the same approach as [Frail et al. \(2012\)](#), these measurements can be converted into upper limits on snapshot transient source areal densities. If a total area of Ω_{eff} has been searched, the 95% confidence limit (CL) density upper limit is $3/\Omega_{\text{eff}}$ to a very good approximation. Our results are plotted in Figure 5.12 and discussed in §5.5.

We also extrapolate this technique to anticipate the parameter space that will be probed by the full survey. We use the empirical relationship between integration time and BUDL shown in Figure 5.6 to estimate the limit for all epochs as yet unprocessed. When doing so we add in a scatter comparable to that seen in the figure. This approach “bakes in” typical levels of data loss due to RFI, instrumental malfunctions, and so on, assuming that the data processed thus far are not unusual in those regards. Compared to what is presented in this work, the full dataset will be more powerful at the bright/rare end because most of the longest integrations have already been processed.

5.4.5 Variability Analysis

Several different metrics are commonly used for assessing the variability of radio sources. Most techniques compute a single scalar variability metric for each source, although this approach is necessarily reductive. (For instance, a source may be highly variable on one timescale and less so on another.) Analyses such as the structure function approach (§5.3.7, below) are more sophisticated, although they still encode certain assumptions about the nature of the variability being measured. Well-designed scalar metrics can still, however, capture meaningful information about overall variability, as we describe below.

[Becker et al. \(2010\)](#) define a modulation index or “fractional variation”

$$f = \frac{S_{\text{max}}}{S_{\text{min}}}, \quad (5.7)$$

but as discussed by [Ofek et al. \(2011\)](#) this metric has irregular statistical properties depending on the number of epochs of observations. The same is true of the metric

$$V = \frac{S_{\text{max}} - S_{\text{min}}}{S_{\text{max}} + S_{\text{min}}} \quad (5.8)$$

used by [Gregory & Taylor \(1986\)](#), which is algebraically interchangeable with the above. [Ofek et al. \(2011\)](#) prefer the ratio of the standard deviation of the flux measurements to the weighted mean, written in their notation as “StD/ $\langle f \rangle$ ” (and sometimes also referred to as “the” modulation index). In our notation this would be $\sigma_S/\langle S \rangle$, where the maximum-likelihood mean flux $\langle S \rangle$ of a set of n measurements S_i with associated uncertainties σ_i is

$$\langle S \rangle = \sum_i S_i \sigma_i^{-2} / \sum_i \sigma_i^{-2}. \quad (5.9)$$

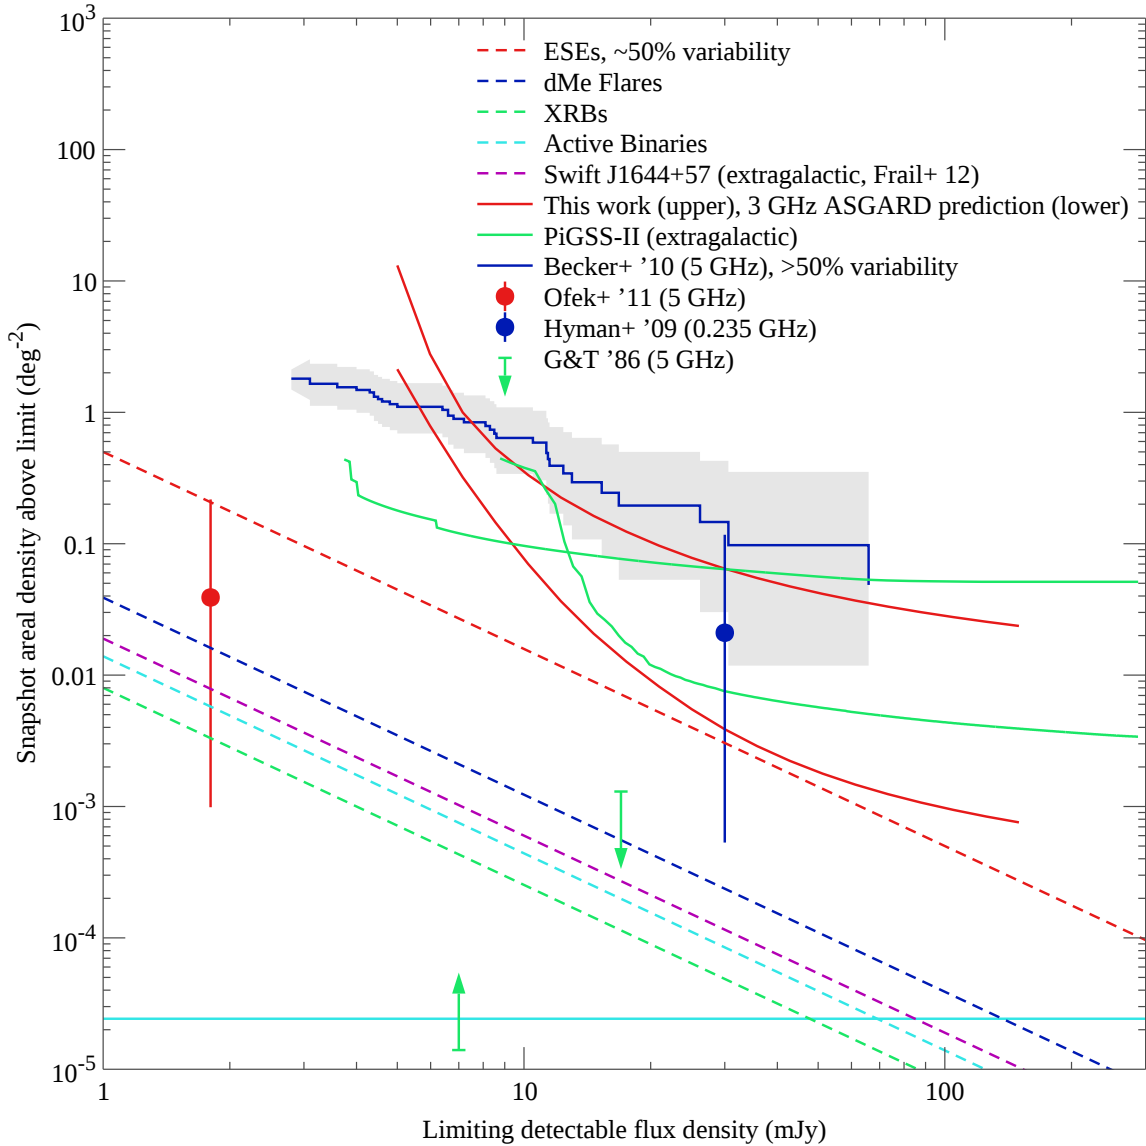


Figure 5.12: Measurements of and limits to snapshot areal densities of radio transients. See §5.5.1 for detailed discussion. All uncertainties are to 95% confidence limit. *Upper red line:* upper limit on the areal density of Galactic transients from this work alone. *Lower red line:* predicted upper limit from the complete 3 GHz ASGARD dataset, should no Galactic transients be found. *Blue line:* areal density of Galactic sources with fractional variability $> 50\%$ at 5 GHz on multi-year timescales as measured by [Becker et al. \(2010\)](#). *Points and limit arrows:* measurements of Galactic transient rates from the VLA at 5 GHz ([Ofek et al. 2011](#)), GMRT at 0.235 GHz ([Hyman et al. 2009](#)), and NRAO 91-m transit telescope at 5 GHz ([Gregory & Taylor 1986](#)), as discussed in the text. *Dashed lines:* areal density estimates for various radio transient populations as discussed in the text. *Horizontal line:* density corresponding to one event on the whole sky at any given time.

While $\sigma_S/\langle S \rangle$ is useful for comparing different studies, we disfavor it for the ranking of candidate variables because it is insensitive to the scale of the uncertainties in each measurement, and thus the confidence with which varying and nonvarying sources can be distinguished. This is easy to see if the ratio is expressed as

$$\frac{\sigma_S}{\langle S \rangle} = \frac{\sqrt{\langle S^2 \rangle - \langle S \rangle^2}}{\langle S \rangle}, \quad (5.10)$$

where the uncertainties only appear implicitly in the averages.

Building on [PiGSS-II](#) and [Bannister et al. \(2011a\)](#) we prefer χ^2 statistics for this purpose. The χ^2 statistic regarding the hypothesis that some source is unvarying is

$$\chi^2 = \sum_i \left(\frac{S_i - \langle S \rangle}{\sigma_i} \right)^2. \quad (5.11)$$

The distribution of observed χ^2 values among the observed sources is not well-defined because, as alluded to in [Bannister et al. \(2011a\)](#), they are not drawn from one parent distribution. They rather come from the family of χ_k^2 distributions, where $k = n - 1$ is the number of degrees of freedom associated with each source. Computing the reduced χ^2 does not help because that only normalizes the expectation value of the distribution, not its shape. To obtain a well-defined distribution we apply the full cumulative distribution function (CDF) of the χ_k^2 family, computing for each source P_c , the probability of accepting the hypothesis that it is constant. We give the expression for P_c in [Appendix B](#). In the case of Gaussian statistics and no variability, the observed P_c will be uniformly distributed between 0 and 1.

The metrics listed above do not use the timestamps associated with each flux density measurement. Therefore, although they can provide an overall assessment of whether a source is somehow “variable,” they cannot describe the nature of that variability. Among this set of metrics P_c is in some sense ideal because it is precisely a probabilistic assessment of this matter. [Ofek et al. \(2011\)](#) find that the structure function of Galactic radio variables saturates at $\tau \sim 10$ d, remaining flat to at least $\tau \sim 60$ d, suggesting that our basic rankings are useful. Nonetheless for particular sources structure functions can, for instance, probe the role that interstellar scintillation plays in the observed variability of our sources (*e.g.* [Rickett et al. 2000](#); [Lovell et al. 2008](#)). The many epochs of ASGARD provide a rich dataset for this analysis: the 3 GHz dataset spans a total time baseline of 1.2 yr and contains 18 pointings with more than twenty visits and 14 pointings with more than fifty.

Sources farther away from the pointing center tend to have increasing systematic effects. This can be seen by examining trends in P_c or $\sigma_S/\langle S \rangle$ against the PB correction factor, as seen in [Figure 5.13](#): although genuine variability will not increase with distance from the pointing center, apparent variability does. In [Figure 5.14](#), we compare the P_c CDFs of sources inside and outside of the half-power radius. In the latter population, there is an overabundance of both apparent variability and sources with overestimated uncertainties (cf. [Appendix B](#)). This leads us to restrict our variability analysis to sources within the

half-power point, which constitute 53 of the 97 sources that have sufficient detections to assess variability.

The excess on the low end of the P_c CDF may be indicative of the presence of genuinely variable sources, but the sample size is insufficient to allow conclusive determination. For the same reason we do not attempt to evaluate a variability confidence threshold as computed by [Bannister et al. \(2011a\)](#). More generally, the irregularity in the CDF and its divergence from the theoretically-motivated shape defined by Equations [B.4–B.6](#) indicate that our P_c values are not statistically rigorous, although they remain useful for ranking the variability of sources. For instance, in [Figure 5.15](#) we compare the dependence of P_c and $\sigma_S/\langle S \rangle$ on PB-corrected source flux. For the more reliable sources with $f_{PB} < 2$, there is a clear increase in $\sigma_S/\langle S \rangle$ as $\langle S \rangle$ decreases. As mentioned above, we attribute this to the fact that $\sigma_S/\langle S \rangle$ does not account for the overall scale of the uncertainties in a set of measurements, which are fractionally larger for fainter sources. The P_c metric does not show this flux dependence.

In [Table 5.3](#) we present positions and variability metrics of the eight most variable sources, as ranked by P_c , in the subset of sources found within the half-power point. Lightcurves for these sources are plotted in [Figure 5.16](#), and postage stamps showing the sources as imaged are shown in [Figures 5.17 and 5.18](#). We discuss the nature of these sources in [§5.5.2](#). Cyg X-3 is the most variable source and is emphatically detected as so by every metric, even without the inclusion of the 20 Jy measurements from its 2011 March major flare. None of the other sources ranked as highly-variable are obviously so. Six of the eight highly-variable sources are found in the *Kepler* pointing, and four of these have nearby companions (see [§5.5.2](#)). Although our photometry routines simultaneously fit nearby sources, it is possible that undeconvolved sidelobes are affecting our flux measurements. Epochs with low-quality images tend to be more obviously subject to systematic photometric effects. In our investigations, post-imaging calibration of these images has been unable to remove these effects: although certain sources are systematically shifted to lower fluxes, different sources are systematically shifted upward, so a simple scale factor is an insufficient correction.

5.5 Discussion & Conclusions

Our first results demonstrate the characteristics of the ASGARD dataset, the strategies we have used in our processing pipeline, and an initial search for highly variable and transient sources.

5.5.1 Galactic Radio Transient Areal Densities

[Figure 5.12](#) shows estimates of areal densities for various Galactic radio transient phenomena and the parameter spaces probed by several observational efforts. In this section, we describe the different components shown in the Figure and their derivatoinis. This standard “log N –log S ” plot inevitably collapses important distinctions between different surveys and

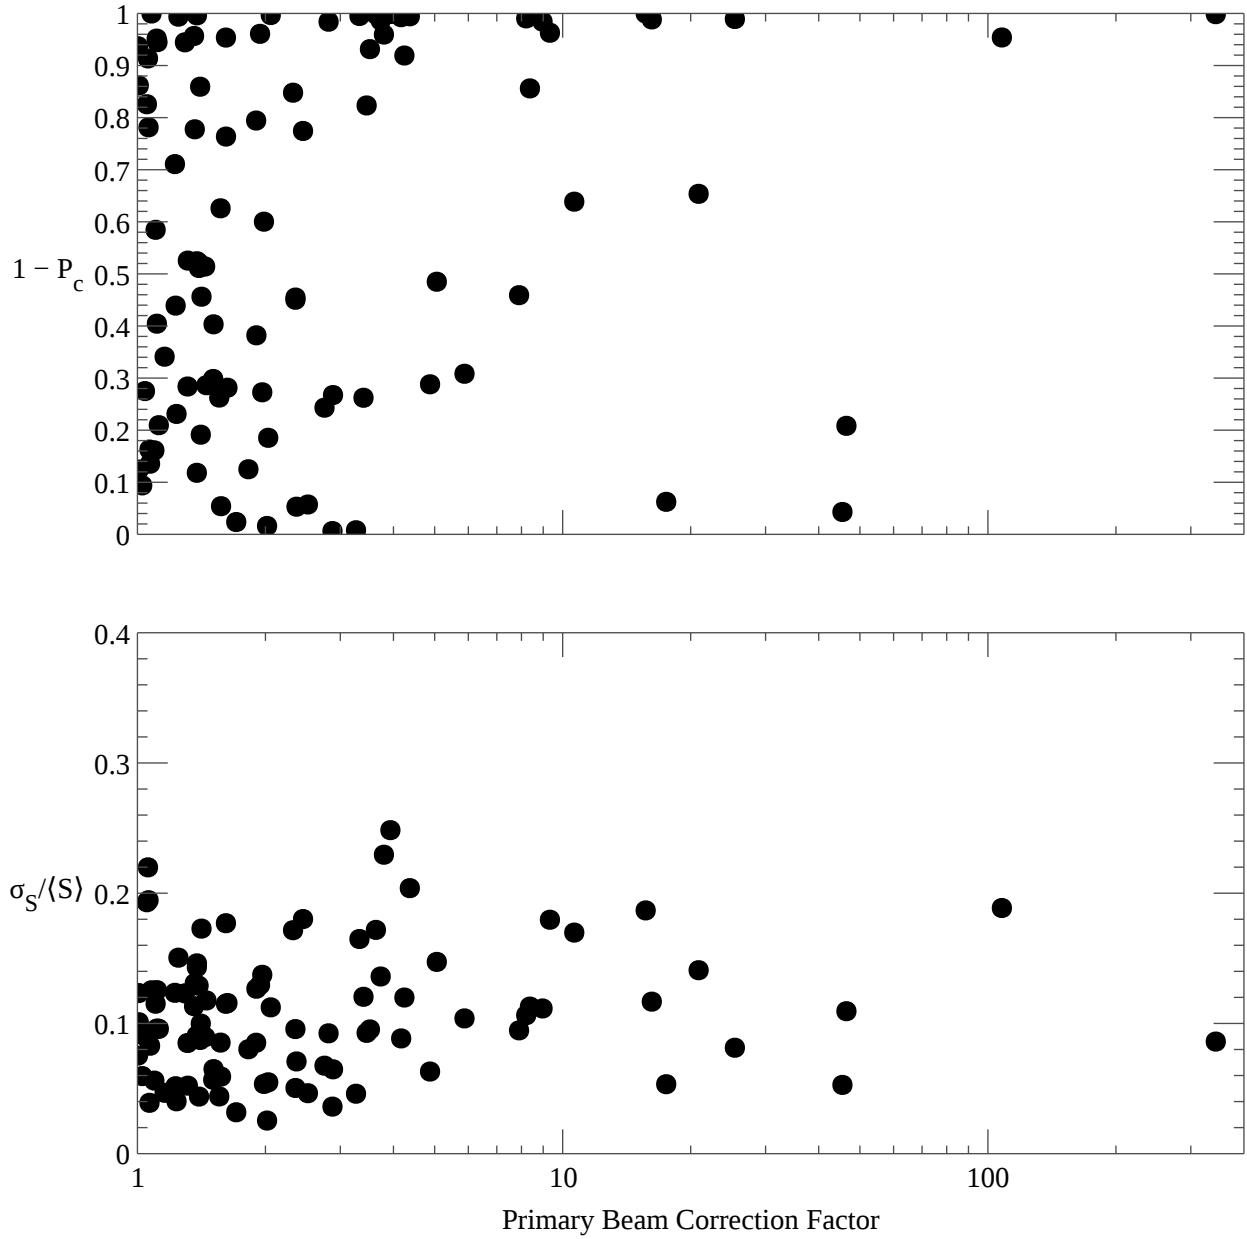


Figure 5.13: Variability metrics as a function of primary beam correction factor f_{PB} . The upper panel plots $1 - P_c$ so that a higher vertical position indicates more variability in both panels. The distributions of the metrics both skew toward higher values as f_{PB} increases, suggesting less-reliable measurements. Cyg X-3 is omitted from these plots. It has $1 - P_c \approx 1$ and $\sigma/\langle S \rangle \approx 0.69$.

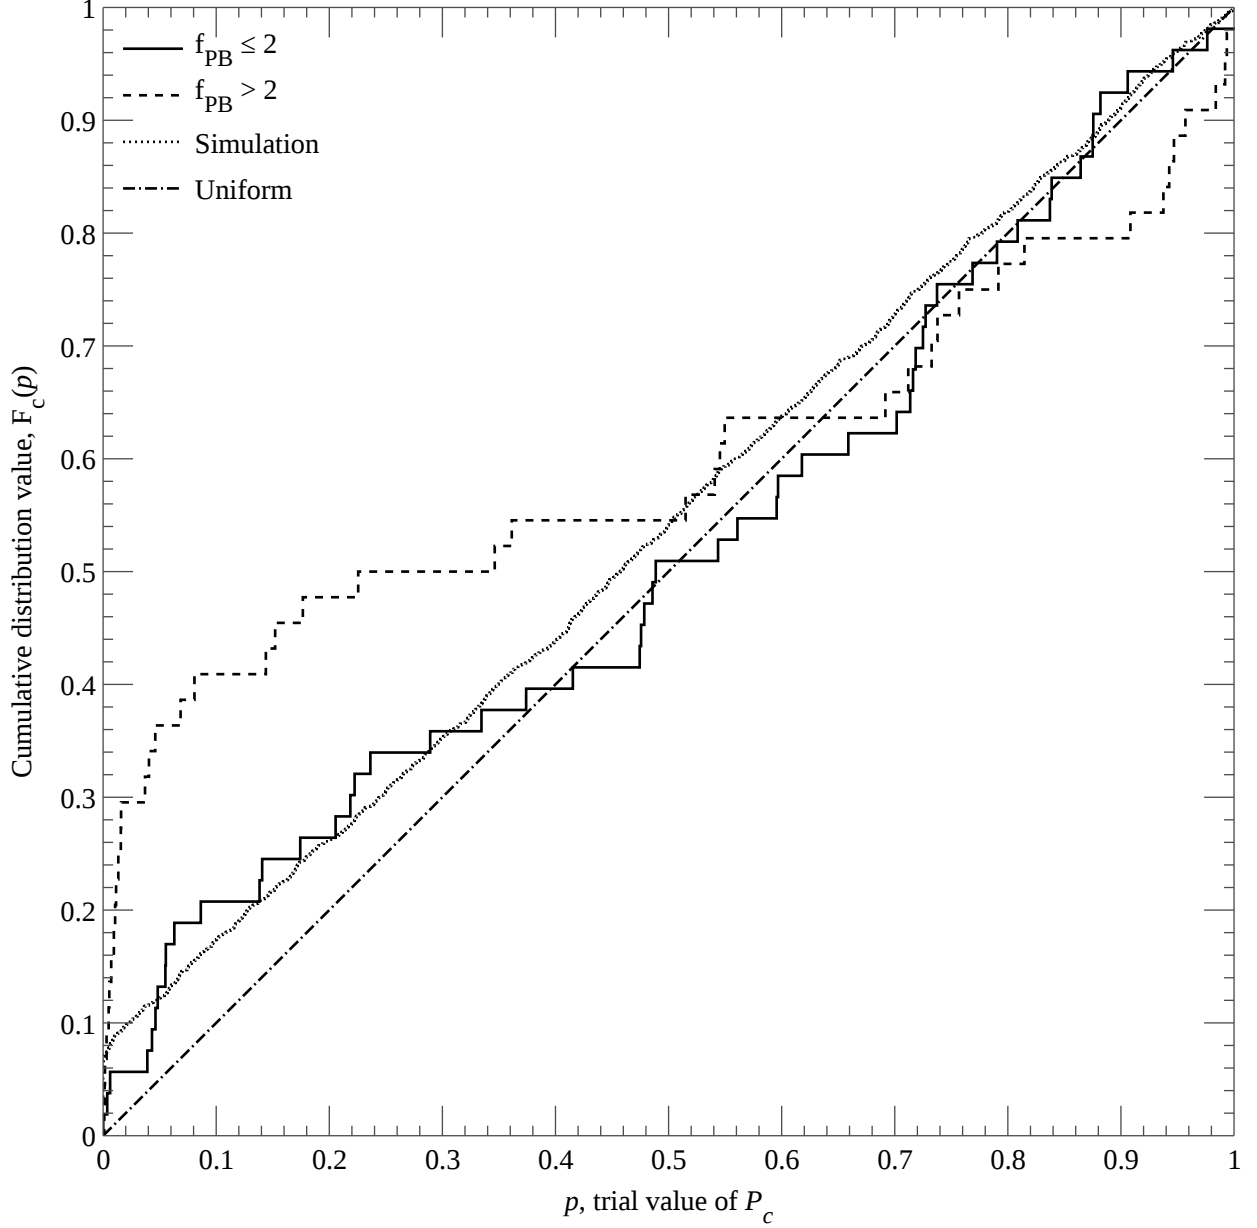


Figure 5.14: Cumulative distribution of observed P_c values for sources inside and outside the half-power point: given a randomly-chosen source in the dataset, $F_c(p)$ is the probability that its P_c value is less than p . Dash-dotted line: the uniform distribution, $F_c(p) = p$, which would be expected in the absence of variability with purely Gaussian errors. Dotted line: simulated observations based on our sampling and a population in which 10% of sources have log-normal flux variability with a scatter of 0.13 dex. Solid (dashed) line: distribution of values inside (outside) the half-power point. See §5.4.5 and Appendix B for discussion.

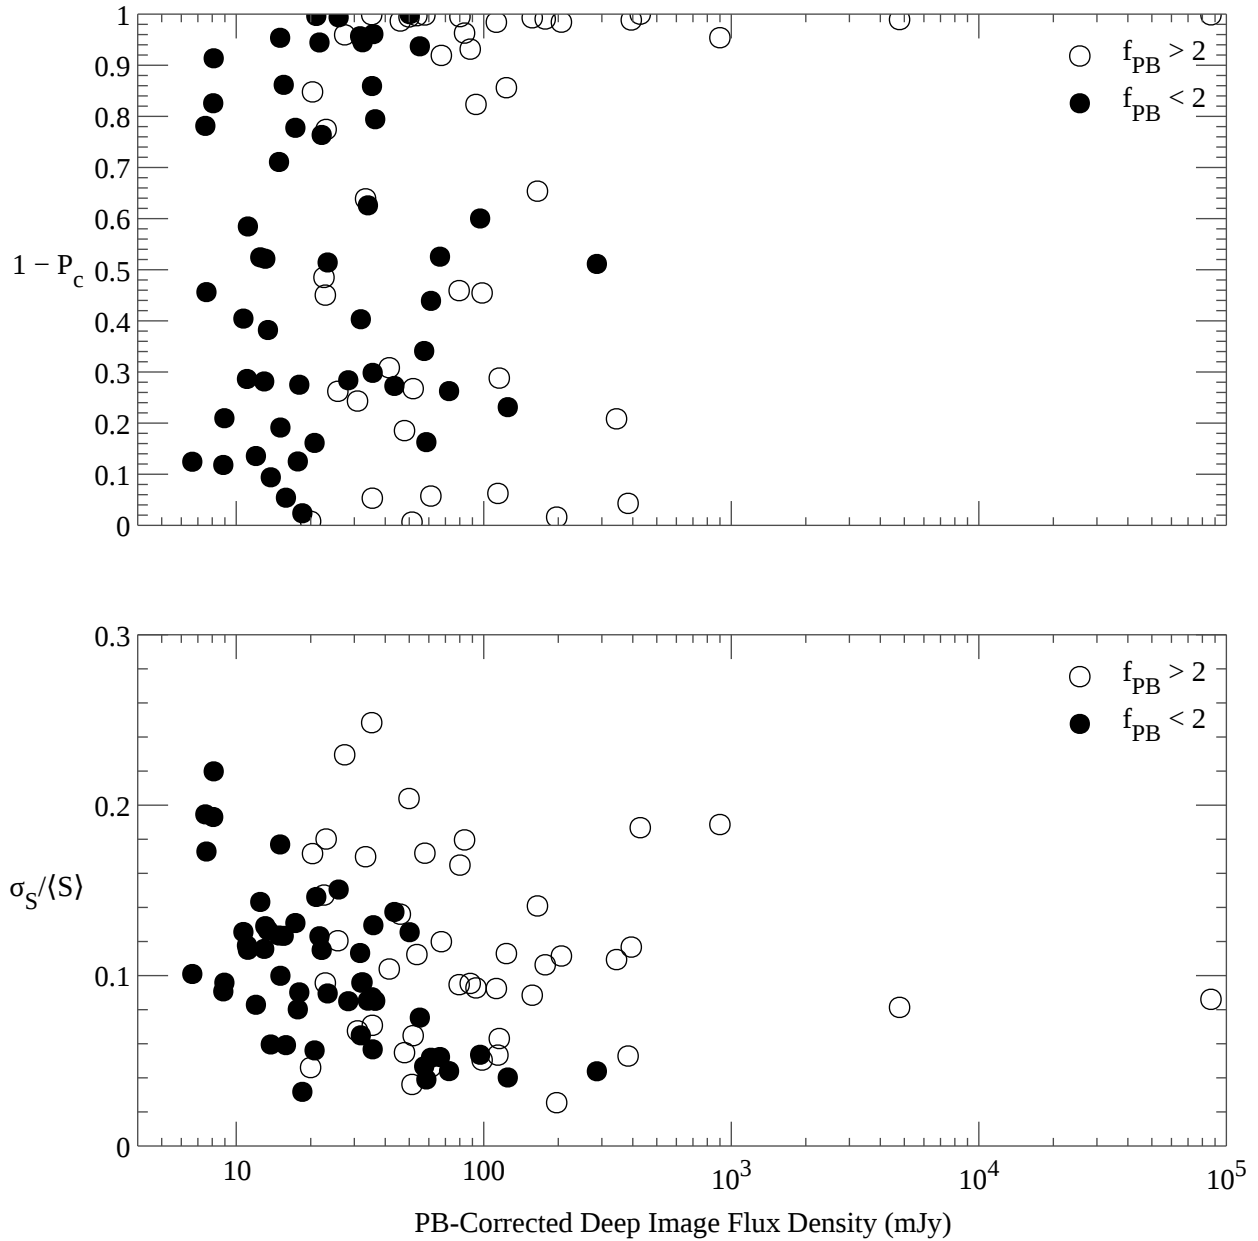


Figure 5.15: Variability metrics as a function of deep image flux density. The upper panel plots $1 - P_c$ so that a higher vertical position indicates more variability in both panels. Among the sources with $f_{\text{PB}} \leq 2$, $\sigma_S / \langle S \rangle$ increases for the less reliably-measured faint sources. The quantity $1 - P_c$ does not show an obvious trend.

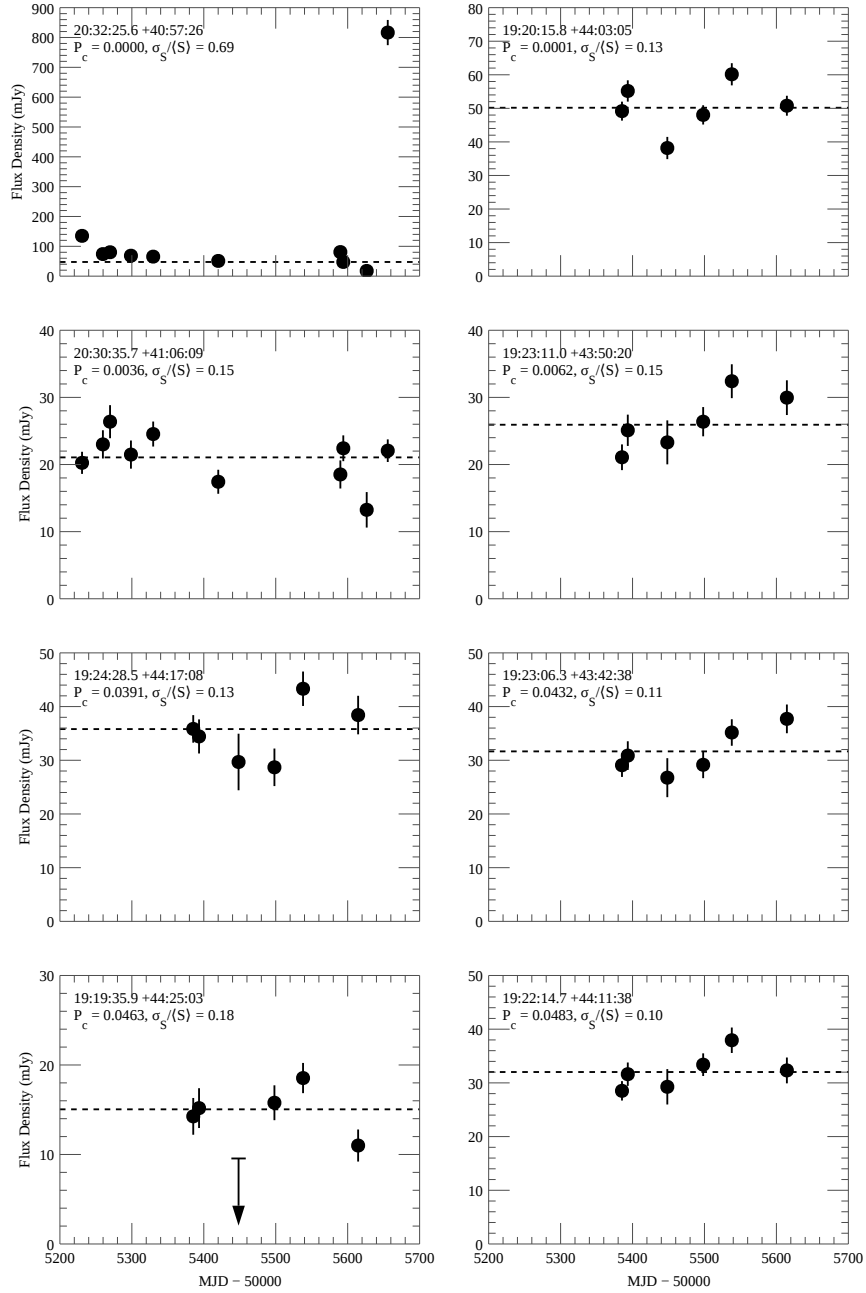


Figure 5.16: Lightcurves of the most variable sources in this study. The variability metric P_c increases from left to right, top to bottom.

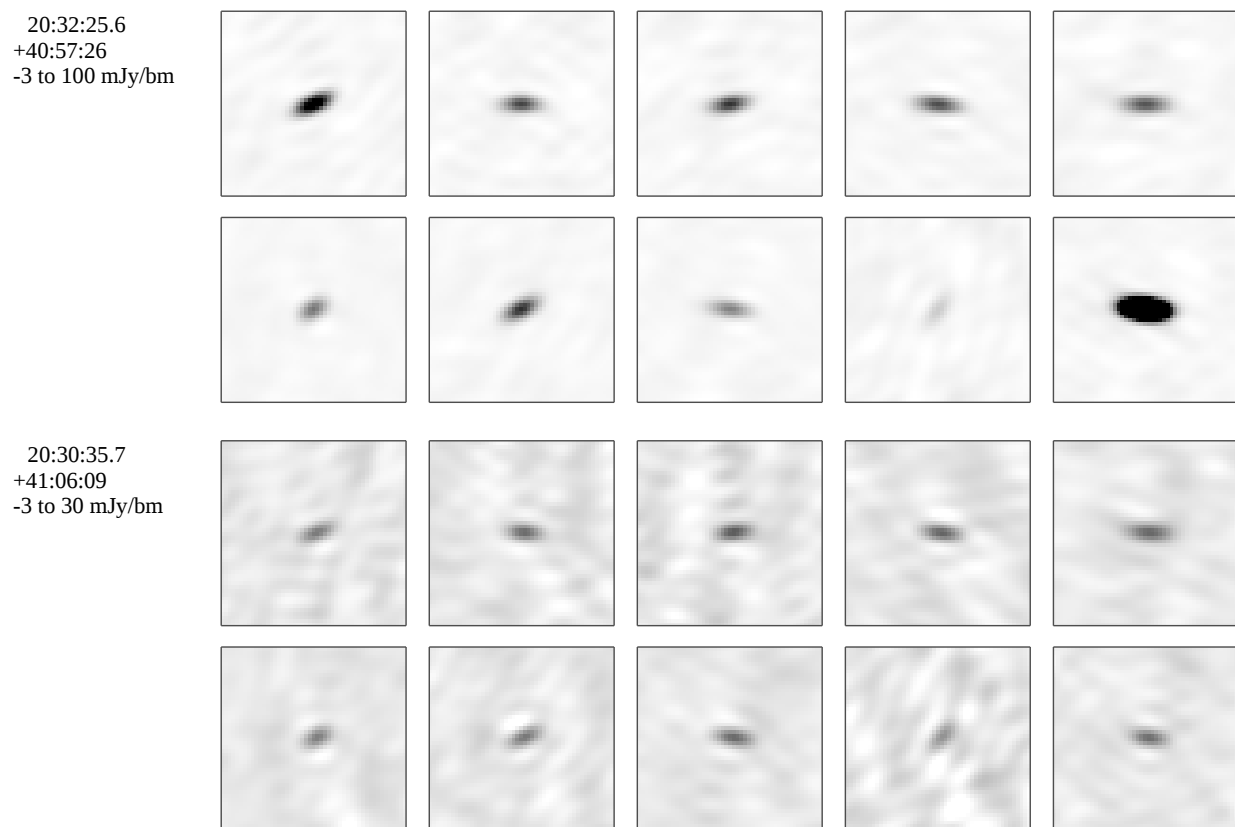


Figure 5.17: Postage stamps of the most variable sources in the Cyg X-3 pointing. Time increases left to right, top to bottom. The coordinates of each source and the linear intensity scale are annotated to the left of each set of panels. Each panel is $8.2 \times 8.2'$.

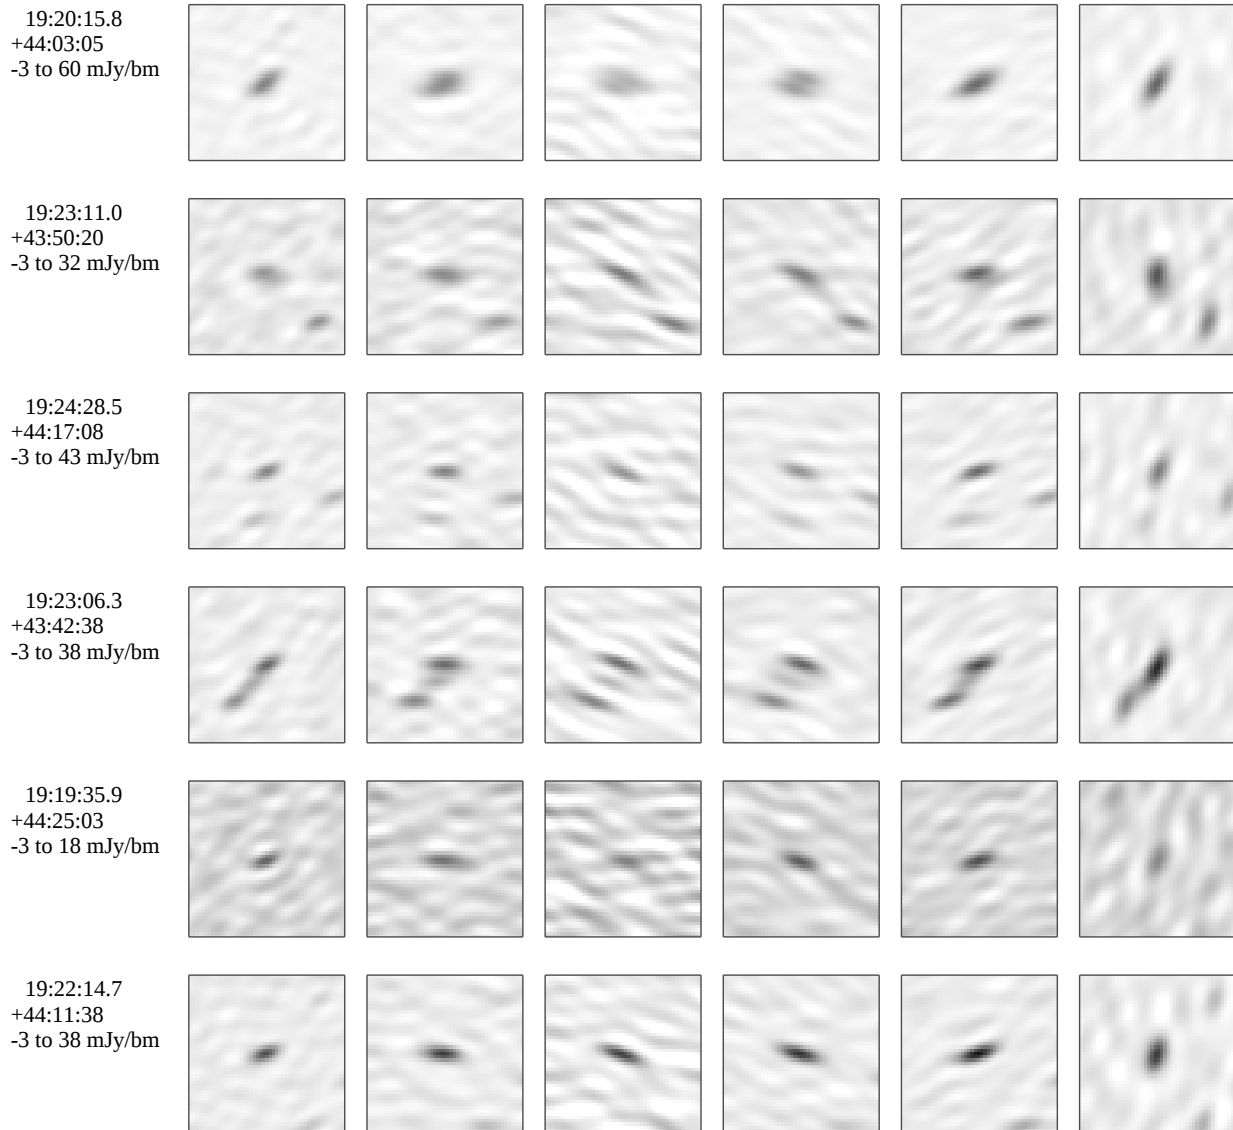


Figure 5.18: Postage stamps of the most variable sources in the Kepler pointing. Layout is the same as in Figure 5.17.

Table 5.3: Eight most variable sources in the ASGARD analysis presented in this work

R.A.	Decl.	$\langle S \rangle$ (mJy)	f_{PB}	P_c	f	$\sigma_S/\langle S \rangle$
20 32 25.6	+40 57 26	47.5	1.0	0.0000	46.39	0.69
19 20 15.8	+44 03 05	50.2	1.1	0.0001	1.58	0.13
20 30 35.7	+41 06 09	21.1	1.4	0.0036	1.99	0.15
19 23 11.0	+43 50 20	25.9	1.2	0.0062	1.54	0.15
19 24 28.5	+44 17 08	35.8	1.9	0.0391	1.51	0.13
19 23 06.3	+43 42 38	31.7	1.4	0.0432	1.41	0.11
19 19 35.9	+44 25 03	15.0	1.6	0.0463	1.69	0.18
19 22 14.7	+44 11 38	32.0	1.1	0.0483	1.33	0.10

Note. — $\langle S \rangle$ is the weighted mean flux of the source across all epochs. f_{PB} is the primary beam correction factor used to determine the intrinsic source flux. P_c is the probability that the source is constant given the measurements, assuming purely Gaussian errors. f is the modulation index of [Becker et al. \(2010\)](#). $\sigma_S/\langle S \rangle$ is the modulation index of [Ofek et al. \(2011\)](#). See discussion in §5.5.2.

populations, such as observing frequency, relevant timescale, and pointing direction ([Bell 2011](#)). When relevant we attempt to make these distinctions explicit in the discussion below. Most densities on the plot are for *transient* sources, as defined either empirically (for the observations) or by an order-of-magnitude increase in flux density (for the theoretical predictions). The exceptions are the values shown for extreme scattering events (ESEs) and the [Becker et al. \(2010\)](#) variability results, both of which correspond to $\sim 50\%$ variability.

To provide a reference for the Galactic snapshot densities we determine below, we also include on Figure 5.12 the snapshot density of extragalactic tidal disruption events similar to *Swift* J1644+57 as estimated by [Frail et al. \(2012\)](#). These events are detected at cm wavelengths and evolve over timescales of about a month.

This Work and Forecast for Complete 3 GHz Survey

We show the parameter space probed in this work and a forecast of the results of a complete analysis of the 3 GHz survey data as described in §5.4.4. The forecast comes in the form of an estimated upper limit to the snapshot areal density of transient sources should none be discovered. This limit should be comparable to the extragalactic results of [PiGSS-II](#) above ~ 10 mJy. Our calibrator pointings could also be used for transient searches, although current results suggest the odds of successful detection would be small ([Bell et al. 2011](#); [Frail et al. 2012](#)).

M-Dwarf Flares

We show the snapshot areal density of flaring M dwarfs as computed by [Osten \(2008\)](#), taking her “submillijansky” detection limit to be 0.5 mJy. We note that the $N(>S) \propto S^{-3/2}$ scaling here is well-justified because M dwarf flares are relatively faint and are only detectable out to hundreds of pc at this limit. Because coherent flares evolve quickly ($\tau \sim 60$ s) and can be fairly narrowband ($\Delta f/f \sim 0.1$), the detectability of these events in practice will depend strongly on specific survey characteristics ([Abada-Simon & Aubier 1997](#)). A survey conducted in the same manner as ASGARD would have extremely poor sensitivity to these events, even if it appeared to reach our density limit on Figure 5.12, because of the very different event timescales and steep, narrow flare spectra.

Active Stellar Binaries

We again follow the analysis of [Osten \(2008\)](#). [Drake et al. \(1989\)](#) analyzed VLA observations of 122 RS CVn-like active binary systems and detected 66, finding luminosity densities ranging from $\sim 10^{14}$ – 10^{18} erg s $^{-1}$ Hz $^{-1}$ at 6 cm wavelengths, with a median of $\sim 10^{16}$ erg s $^{-1}$ Hz $^{-1}$. We set a cutoff of a factor of >10 luminosity increase for such a system to be considered a transient rather than a variable. Approximately 10% of the systems observed had luminosity densities $>10^{17}$ erg s $^{-1}$ Hz $^{-1}$. If the observations of [Drake et al. \(1989\)](#) probe a uniform, variable population, we can take this number to be an approximate duty cycle for active binaries to appear as radio transients. We then find that the areal density of such systems brighter than 10 mJy should be $4 \cdot 10^{-4}$ deg $^{-2}$, taking the spatial density of active binary systems to be $\sim 6 \cdot 10^{-5}$ pc $^{-3}$ ([Favata et al. 1995](#)). As with flaring M dwarfs, these systems are only detectable out to hundreds of pc at mJy sensitivities and thus can be treated as an isotropic population for our purposes. Even with our somewhat generous treatment of both the duty cycle and radio luminosity dynamic range of these systems, they will be very difficult to detect blindly. Unlike M dwarf flares, however, active binary flares evolve on day timescales and have flat, broadband spectra that are amenable to detection at cm wavelengths ([Owen & Gibson 1978](#)).

X-Ray Binaries

[Lutovinov et al. \(2007\)](#) used the INTEGRAL dataset to find 74 XRBs with $|b| < 2^\circ$, with 41 of those being high-mass systems. Although the distribution of these systems is nonuniform and, in the case of high-mass systems, appears to be linked to Galactic spiral structure ([Bodaghee et al. 2012](#)), we compute a characteristic areal density in the GP by assuming these sources are distributed in a region of $|\ell| < 90^\circ$, $|b| < 2^\circ$. We compute typical flaring flux densities of these systems as follows. For black hole XRBs (BHXBs), we use radio observations of 20 systems presented by [Gallo et al. \(2012\)](#). The radio luminosities of systems with multiple observations vary by factors of ~ 10 , so that marginally-detected flaring systems may appear as transients. The maximum observed flux densities of these systems

range from ~ 0.1 – 400 mJy, with a median of 8 mJy, assuming flat spectra and an observing frequency of 3.09 GHz. For neutron star XRBs (NSXBs), [Migliari & Fender \(2006\)](#) find a typical flux density of ~ 0.4 mJy and a typical radio luminosity dynamic range of 5, so that flaring systems typically reach ~ 2 mJy. For both classes of systems, the flaring duty cycle appears to be $\sim 1\%$ ([Fragos et al. 2009](#); [Körding et al. 2005](#)). Combining these results, we arrive at an areal density of $\sim 10^{-3}$ deg $^{-2}$ for XRBs flaring above ~ 4 mJy, where the latter value is intermediate between the ones mentioned above with a bias towards a smaller value to reflect the smaller luminosity dynamic range of the NSXBs and the likelihood that the brighter flaring systems have already been discovered. For consistency we plot this estimate using $S^{-3/2}$ scaling but we warn that this is not as well-motivated as in the previous cases. Although XRBs are brighter than the aforementioned stellar systems, their rarity makes their blind detection difficult. As will be discussed in more detail in Chapter 6, flares from these systems evolve on day–week timescales and have spectra generally peaking in the cm regime ([Hjellming & Han 1995](#)).

Extreme Scattering Events

Most examples of interstellar scintillation affect source fluxes by $\sim 10\%$ ([Lovell et al. 2008](#)) and so would not be associated with radio transience. Extreme scattering events, however, can cause variations of order 50% at the ~ 3 GHz frequencies we consider ([Fiedler et al. 1987](#); [Fiedler et al. 1994](#)). These events last months and so are well-suited to detection with the ASGAR observing cadences and analysis method. [Fiedler et al. \(1987\)](#) find an ESE duty cycle of $\sim 7 \cdot 10^{-3}$. Considering previous indications of increased scintillation in the GP ([Becker et al. 2010](#); [Ofek et al. 2011](#)) and an association between ESEs and Galactic structure ([Fiedler et al. 1994](#); [Lazio et al. 2000](#)), we assume a doubled ESE duty cycle of $1.4 \cdot 10^{-2}$ in the GP, which is consistent with pulsar observations ([Pen & King 2012](#)). The areal density of blazars brighter than 100 mJy is ~ 0.03 deg $^{-2}$ ([Padovani et al. 2007](#)), and [Kraus et al. \(2003\)](#) detected intraday variability (IDV) in 86% (25/29) of a blazar sample they observed. If every source subject to IDV may experience an ESE, we find an areal density of $4 \cdot 10^{-4}$ ESE-affected sources intrinsically brighter than 100 mJy per square degree in the GP. We note again that these sources may not be detected as traditionally-defined transients because of both the typical scale of the effect and the fact that ESEs involve significant dimmings, not brightenings, of sources. The $S^{-3/2}$ scaling is of course appropriate for extragalactic sources, but the rate of ESE incidence may vary significantly by line of sight.

50% Variables from Becker et al. (2010)

[Becker et al. \(2010\)](#) defined strong Galactic variables as sources with $f > 1.5$ over the multi-year time baseline of their observations, approximately equivalent to $\sigma_S / \langle S \rangle = 1/3$ ([Ofek et al. 2011](#)), where the measurements in question are peak, not integrated, flux densities. After correcting for an estimated extragalactic contribution, [Becker et al. \(2010\)](#) find that these sources have an areal density of 1.6 deg $^{-2}$ in their survey, not accounting for the

completeness of the underlying catalogs. We combine this normalization with the CDF of the brightest measurements of each applicable source and an approximation of the completeness function of the underlying catalogs as determined by [White et al. \(2005\)](#), which we assumed applied to both catalogs used by [Becker et al. \(2010\)](#). Our observations are not directly comparable to those of [Becker et al. \(2010\)](#) for two major reasons: firstly, the *Kepler* pointing that we analyze is at a much higher Galactic latitude ($b \approx +13.5$); secondly, our observations lack the 2–15 yr time baseline over which much of the variability found by [Becker et al. \(2010\)](#) occurred, although the structure function results of [Ofek & Frail \(2011\)](#) suggest that the level of variability on 1–10 yr timescales is approximately constant. The difference between the observing frequencies of the two surveys may also be relevant due to the frequency dependence of scintillation effects in the GHz regime ([Hjellming & Narayan 1986](#)). Applying the [Becker et al. \(2010\)](#) estimates to our sensitivity limit and the footprint (not effective area) of the Cyg X-3 pointing, we would expect to detect ~ 5 highly variable Galactic sources. In our catalog Cyg X-3 is the only source with $\sigma_S / \langle S \rangle > 1/3$. We tentatively attribute this discrepancy to the comparatively short time baseline of our study, but defer a deeper analysis until later work.

GC Radio Transients from Hyman et al. (2009)

[Hyman et al. \(2009\)](#) discovered one robust GC radio transient in 22 epochs of observations with the GMRT. (Although these observations were made to follow up the detection of a separate transient, GCRT J1745–3009 [[Hyman et al. 2005](#)], we do not consider this first transient here because the parameters of the survey that led to its discovery are unclear.) [Hyman et al. \(2009\)](#) observed on a \sim monthly cadence with several-hour integrations and so are sensitive to similar but somewhat longer timescales than ASGARD. Because [Hyman et al. \(2009\)](#) do not specify their detection criteria in detail, we make several assumptions in our derivation: that no epochs were searched for transients besides the 22 reported in the paper; that in each epoch the FOV was searched out to the half power point (with $\text{HPBW}_{\text{GMRT}} = 2^\circ$); and that each epoch had a blind detection limit of 30 mJy. From these parameters we find an effective search area of 47.6 deg^2 . The full ASGARD dataset will probe a significantly larger area at this detection limit. As with the comparison to [Becker et al. \(2010\)](#), however, a direct extrapolation of the areal density would be inadvisable: [Hyman et al. \(2009\)](#) observed only toward the GC itself, where the source density is larger, and the observing wavelengths of the GMRT and the ATA differ by approximately an order of magnitude.

Galactic Radio Transients from Gregory & Taylor (1986)

[Ofek et al. \(2011\)](#) and [Frail et al. \(2012\)](#) have reported a density measurement of 10^{-3} deg^{-2} for this survey (in their Figures 1 and 6, respectively) but this number is erroneous (E. Ofek, 2012, private communication). [Gregory & Taylor \(1986\)](#) surveyed a footprint of 500 deg^2 over 16 epochs and discovered one bright ($>1 \text{ Jy}$) transient, GT 0351+543a, that appeared in one

epoch, a phenomenology similar to that of the events detected at Nasu Observatory (Kuniyoshi et al. 2007; Matsumura et al. 2007; Niinuma et al. 2007; Kida et al. 2008; Matsumura et al. 2009). Both day and year timescales were sampled but intermediate ones were not; about 2/3 of the detected variables evolved on day timescales. The single transient detection straightforwardly gives a density measurement of $(1.33_{-1.30}^{+6.10}) \cdot 10^{-4} \text{ deg}^{-2}$ (95% confidence limit), but the effective search area of the survey as a function of limiting detectable flux density is unclear due to variable survey sensitivity. Taylor & Gregory (1983) report that over the whole survey the worst-case 3.5σ detection limit, corresponding to an empirical 50% survey completeness, was 17 mJy. At this flux density limit the effective area of the survey is thus no less than 3750 deg^2 , and a 95% CL upper limit on the Galactic transient surface density is thus $1.3 \cdot 10^{-3} \text{ deg}^{-2}$. Taylor & Gregory (1983) also report a best-case 3.5σ detection limit of 7 mJy, so that at this limit the effective area of the survey is no more than 3750 deg^2 , and the Galactic transient surface density at that limit has a 95% CL lower bound of $1.4 \cdot 10^{-5} \text{ deg}^{-2}$ or $\sim 0.6 \text{ sky}^{-1}$. We advise caution in the use of these limits because it seems unlikely that the only transient detected by Gregory & Taylor (1986) would have a flux $> 1 \text{ Jy}$, when their survey should have easily been sensitive to events $\lesssim 50$ times fainter. Gregory & Taylor (1986) investigated and deemed unlikely the possibility that this event was due to interference, however. No source was detected in later VLA followup of the event (Tsutsumi et al. 1995).

5.5.2 The Nature of the Most-Variable Sources

We investigated the nature of the most-variable sources listed in Table 5.3 and plotted in Figures 5.16–5.18. Here we briefly describe our findings, omitting the most variable source, Cyg X-3. All of the other sources are consistent with being AGN. As mentioned in §5.4.5, none of the sources besides Cyg X-3 are obviously variable upon visual inspection.

19:20:15.8 +44:03:05: This source is a marginally-resolved double, as confirmed by observations from the Cambridge One-Mile telescope (Birkinshaw 1978) and the VLA (O’Dea & Owen 1985). Our modeling of this source as a single component likely leads to its apparent variability. It is a head-tail radio source found near the center of the cluster Abell 2319. The position is consistent with that of the infrared source WISE J192015.63+440307.9 (Cutri et al. 2012), which has colors consistent with an AGN (Wright et al. 2010).

20:30:35.7 +41:06:09: This source is also detected in the NVSS (J203032+410634) and the WSRT 327 MHz GP survey (Taylor et al. 1996, as 2028+4055). The WSRT survey included two epochs of observations separated by several years. Taylor et al. (1996) searched for variability between the epochs, and this source was not detected as a variable. There were no counterparts in other bands found in standard catalogs. The fluxes at 0.327, 1.4, and 3.09 GHz are 150, 12, and 21 mJy, respectively. Although variability or an unusual spectrum are possible explanations, we interpret the L-band measurement as suggesting that much of this source’s emission is resolved out by the VLA.

19:23:11.0 +43:50:20: This source is also detected in the NVSS (J192310+435016), the

Westerbork Northern Sky Survey (WENSS, as B1921.6+4344; [Rengelink et al. 1997](#)), and early Westerbork observations of Abell clusters ([Harris & Miley 1978](#)). It is 12% linearly polarized at 5 GHz ([Broten et al. 1986](#)) which may contribute to some apparent variability. There were no counterparts in other bands found in standard catalogs.

19:24:28.5 +44:17:08: This source is also detected in the NVSS (J192428+441708), WENSS (B1922.9+4411), and the Green Bank 6 cm survey (GB6, [Gregory et al. 1996](#), as J1924+4417). It has a generally decreasing spectrum over this frequency range. The position is consistent with that of WISE J192428.35+441709.0, which has colors consistent with an AGN.

19:23:06.3 +43:42:38: This source appears to be the northwest lobe of a double source with a fainter central component. It is detected in the NVSS (J192306+434233). The central component is detected in the IR (2MASS J19230940+4341460) and X-rays by *Einstein* (2E 1921.5+4335) and ROSAT ([Brinkmann et al. 1997](#)). These properties suggest that the ASGARD source in question is a lobe of an FR II active radio galaxy ([Fanaroff & Riley 1974](#)). If that is the case, its flux should be virtually constant. We do not model the faint central radio component, which will increase the source’s apparent variability.

19:19:35.9 +44:25:03: This source is also detected in the NVSS (J191936+442503) and possibly WENSS (B1918.0+4419, separation $\sim 20''$), but not GB6. It has a generally decreasing spectrum over the range of frequencies in which it is detected. There were no counterparts in other bands found in standard catalogs.

19:22:14.7 +44:11:38: This source is also detected in the NVSS (J192214+441137), WENSS (B1920.6+4405), GB6 (J1922+4411), and by [Harris & Miley \(1978\)](#). It has a strongly decreasing spectrum over this range. There were no counterparts in other bands found in standard catalogs.

5.5.3 Outlook

There is substantial evidence that the prevalence of apparent variability increases with proximity to the Galactic plane ([Spangler et al. 1989](#); [Ghosh & Rau 1992](#); [Gaensler & Hunstead 2000](#); [Becker et al. 2010](#); [Ofek & Frail 2011](#)). [Ofek et al. \(2011\)](#) suggest that there is also a significant increase in the typical variability amplitude, based on the difference between their findings from observations at $|b| \approx 7^\circ$ and those of [Becker et al. \(2010\)](#) at $|b| \lesssim 1^\circ$. Our observations, with $|b| \lesssim 0.5^\circ$, will allow a strong test of this conjecture if the source variability function is flat over 1–3 yr timescales. We also probe Galactic longitudes from $-3^\circ < \ell < 7^\circ$, investigating a possible increase in radio variability with proximity to the GC as has been tentatively measured ([Becker et al. 2010](#)) and might be expected simply from the increased source density towards Sgr A* ([Genzel & Townes 1987](#)).

As demonstrated by Figures 5.4 and 5.9, the ASGARD dataset allows for sensitive moderate-resolution imaging of static large-scale radio structure in the Galactic plane. Extended structures in our deep images will include non-thermal radio filaments ([Yusef-Zadeh et al. 1984](#); [Law et al. 2008](#)), H II regions ([Brogan et al. 2003](#); [Nord et al. 2006](#)), and supernova

remnants (Gray 1994) at a wide range of GC separations, with implications for WIMP dark matter models motivated by *Fermi* γ -ray observations (Hooper & Goodenough 2011; Linden et al. 2011), the acceleration and composition of Galactic cosmic rays (Protheroe et al. 2008; Crocker et al. 2011), and the energetics of the interstellar medium in the GC region (Morris & Serabyn 1996; Crocker et al. 2010). Although there are formidable technical challenges to full polarimetric calibration and imaging of the data, these would provide a wealth of information on the structure of the GC magnetic field, ionization content, and outflows (Nishiyama et al. 2010; Law et al. 2011a).

Acknowledgments

I thank Peter Backus, Tom Kilsdonk, and Jon Richards for designing and executing the SETI component of the commensal observing campaign, which involved responsibility for essential matters such as scheduling and frontend hardware control. This chapter makes use of data products from the Wide-field Infrared Survey Explorer, which is a joint project of the University of California, Los Angeles, and the Jet Propulsion Laboratory/California Institute of Technology, funded by the National Aeronautics and Space Administration.

Chapter 6

The 2010 May Flaring Episode of Cygnus X-3 in Radio, X-Rays, and γ -Rays

In 2009, Cygnus X-3 (Cyg X-3) became the first microquasar to be detected in the GeV γ -ray regime, via the satellites *Fermi* and *AGILE*. The addition of this new band to the observational toolbox holds promise for building a more detailed understanding of the relativistic jets of this and other systems. We present a rich dataset of radio, hard and soft X-ray, and γ -ray observations of Cyg X-3 made during a flaring episode in 2010 May. We detect a ~ 3 -d softening and recovery of the X-ray emission, followed almost immediately by a ~ 1 -Jy radio flare at 15 GHz, followed by a 4.3σ γ -ray flare ($E > 100$ MeV) ~ 1.5 d later. The radio sampling is sparse, but we use archival data to argue that it is unlikely the γ -ray flare was followed by any significant unobserved radio flares. In this case, the sequencing of the observed events is difficult to explain in a model in which the γ -ray emission is due to inverse Compton scattering of the companion star’s radiation field. Our observations suggest that other mechanisms may also be responsible for γ -ray emission from Cyg X-3.¹

6.1 Introduction

The X-ray binary (XRB) system Cygnus X-3 (hereafter Cyg X-3), discovered by [Giacconi et al. \(1967\)](#), is notable for its Wolf-Rayet (WR) companion ([van Kerkwijk et al. 1996](#)), short (4.8 h) orbital period ([Bonnet-Bidaud & Chardin 1988](#)), and dramatic radio variability ([Waltman et al. 1994](#)). The relativistic jets that it produces ([Geldzahler et al. 1983](#)) classify the source as a “microquasar” system. The distance to Cyg X-3 is ~ 7 or ~ 9 kpc, depending on the method used ([Predehl et al. 2000](#); [Ling et al. 2009](#)). It is also uncertain whether the compact object is a black hole (*e.g.* [Cherepashchuk & Moffat 1994](#)) or neutron star (*e.g.* [Stark & Saia 2003](#)).

¹This chapter has been previously published as Williams *et al.*, [2011 ApJ 733 L20](#), and is reproduced with the permission of all coauthors and the copyright holder. Copyright 2011 American Astronomical Society.

Cyg X-3 became more notable in 2009 when it became the first microquasar to be detected in the GeV γ -ray regime, via the satellites *Fermi* (Abdo et al. 2009a) and *AGILE* (Tavani et al. 2009). A 4.8-h modulation in the *Fermi* data makes the detection definitive. Emission of γ -rays has been detected by both observatories during every “high/soft” X-ray state to occur since they began science operations (Abdo et al. 2009a; Tavani et al. 2009; Corbel & Hays 2010; §6.3). The high/soft state is also associated with radio flaring and is closely tied to the launching of relativistic jets (*e.g.* Fender et al. 2004; Koljonen et al. 2010).

Various physical processes are expected to be involved in the emission and propagation of high-energy photons in Cyg X-3. Inverse Compton (IC) upscattering of the intense ($\sim 10^5$ erg cm $^{-3}$) radiation field of the WR companion by high-energy ($\gamma \gtrsim 10^3$) electrons in the relativistic jet is the most efficient source of GeV emission (Abdo et al. 2009a; Dubus et al. 2010). If the density of relativistic protons in the jet is sufficiently high, inelastic collisions between these and protons in the stellar wind will produce neutral pions that decay into γ -rays (among other products) as well (Romero et al. 2003). In this case, there will be further γ -ray emission due to secondary leptons created during the hadronic interactions (Orellana et al. 2007). The typical hadronic content of microquasar jets is unknown: large numbers of hadrons impose significant energetic constraints, but there is direct evidence of their presence in at least one system, SS 433 (Kotani et al. 1994, 1996; Migliari et al. 2002).

A γ -ray flare due to IC upscattering by jet electrons will be followed by a radio flare a few days later as those electrons cool, rarify, and become transparent to their radio synchrotron emission. This phenomenon has been observed repeatedly but not invariably (Abdo et al. 2009a; Tavani et al. 2009). If the γ -ray emission is primarily hadronic, the result will be substantially the same, as secondary processes will create a similar population of energetic leptons. Abdo et al. (2009a) determine an approximate 5 ± 7 d lag between radio and γ -ray emission. Inspection of their Fig. 2 and the data presented by Tavani et al. (2009) suggests that this uncertainty is primarily systematic rather than statistical; *i.e.*, that the radio/ γ -ray lag varies significantly from event to event. Blazar observations often show lag variability, both between sources and in the same source over time, with theoretical support for sign changes in the lag between γ -ray and other bands (Aller et al. 2010).

In 2010 May, multiple monitoring programs detected that Cyg X-3 was becoming more active in γ -rays and was entering the high/soft state (Bulgarelli et al. 2010; Corbel & Fermi Large Area Telescope Collaboration 2010; Kawai et al. 2010; Corbel & Hays 2010). We describe multiband observations of Cyg X-3 made during this episode (§6.2) and present our results (§6.3). Our data show a γ -ray flare apparently lagging a radio flare, a sequencing opposite of that expected in the typical interpretation. We discuss and interpret this result (§6.4) and, finally, present our conclusions (§6.5).

6.2 Observations & Data Analysis

We observed Cyg X-3 during its 2010 May flaring episode with a variety of instruments. Below, we describe the observations and their analysis. Many dates and times we mention are near MJD 55,340. For brevity, we express these in an offset MJD, defining $\text{OJD} = \text{MJD} - 55,300$. Thus OJD 0 is JD 2,455,300.5 or 2010 Apr 14.0 UT, and the γ -ray flare was detected on OJD 43 = 2010 May 27 UT.

6.2.1 Radio

The Arcminute Microkelvin Imager (AMI) arrays (Cambridge, UK) are two aperture synthesis telescopes mainly used to study the cosmic microwave background (Zwart et al. 2008). The observations described herein were made with the Large Array (AMI-LA), the reconfigured and reequipped Ryle Telescope, consisting of eight 13-m antennas with a maximum baseline of about 120 m, observing in the band 12–18 GHz. The angular resolution is typically $25''$. Monitoring of small-diameter sources is undertaken as described in Pooley & Fender (1997); observations are interleaved with those of a phase-reference calibrator, and after calibration the data for individual baselines are vector-averaged. The in-phase component then provides an unbiased estimate of the target source’s flux density. The amplitude scale was calibrated by (at least) daily observations of 3C 48 and 3C 286, both of which are believed to be very nearly constant on long timescales.

Should there be emission on a scale resolved by some of the baselines, the flux density estimate would be incorrect. There is an extended region of low-brightness emission around Cyg X-3 (Sánchez-Sutil et al. 2008) which is detected on the shortest baselines at extreme hour angles with the AMI-LA, and is sufficiently bright to need separate treatment when the central source in Cyg X-3 is “quenched” (flux density $\lesssim 10$ mJy at 15 GHz). No such correction was required in the observations described here.

Observations of Cyg X-3 were also made with the RATAN-600 telescope (Nizhnij Arkhyz, Russia) as part of an ongoing microquasar monitoring campaign (Trushkin 2000; Trushkin et al. 2006). The observations and data analysis were performed as described in Trushkin et al. (2006). In this work we show results at 11.2 GHz; near-simultaneous observations were made at 4.8 and 7.7 GHz as well.

Cyg X-3 was observed by the Allen Telescope Array (ATA; Northern California, USA; Welch et al. 2009) six times in the period OJD 26–43 as part of a larger transient search, ASGARD (Chapter 5). Continuum images at 3.09 GHz were made after subtraction of model of the static sky. Source fluxes were determined in ~ 10 -minute segments by fitting point source models in the image domain.

6.2.2 X-Ray

Our study includes hard X-ray observations made with the *International Gamma-Ray Astrophysics Laboratory* (*INTEGRAL*) satellite (Winkler et al. 2003). We used the *INTEGRAL* Soft Gamma-Ray Imager (ISGRI; Lebrun et al. 2003) instrument and included two observations that were part of a program to observe Cyg X-1 (PI: Wilms). Although the primary target of the observations was Cyg X-1, the ISGRI FOV is large enough to include the entire Cygnus region, and our target was in the FOV throughout the observations. The first observation occurred during revolution 929, starting at OJD 40.625 and ending at OJD 41.940. The second observation occurred during revolution 938, starting at OJD 67.469 and ending at OJD 68.805. A preliminary report of the first observation was given in Tomsick et al. (2010).

We reduced the ISGRI data using the Off-line Scientific Analysis (OSA v9.0) software package. The program `ibis_science_analysis` is the primary tool for extracting the standard data products. We produced and inspected the ISGRI image in the 20–40 keV band, and then made energy spectra and light curves for all of the bright sources in the field. After correcting for instrumental deadtime, the total exposure times for Cyg X-3 from the first and second observations are 62,420 s and 66,540 s, respectively. We produced 20–40 keV light curves with a time resolution of 1 ks.

We also obtained one-day average quick-look measurements from the All-Sky Monitor on the *Rossi X-Ray Timing Explorer* (*RXTE*/ASM; Levine et al. 1996) and transient monitor results from the Burst Alert Telescope on the *Swift* γ -ray burst mission (*Swift*/BAT; Barthelmy et al. 2005). These results are provided publicly by the *RXTE*/ASM and *Swift*/BAT teams, respectively.

6.2.3 γ -Ray

Fermi Science Tools 9.17 and HEASoft 6.9 were used to reduce and analyze all *Fermi*-LAT (Atwood et al. 2009) observations within 10° of Cyg X-3 that took place in the range OJD -59 – 150 (2010 Feb 14 – 2010 Sep 11 UT), providing a baseline of ~ 100 d before and after the flare of OJD 43–44. Cyg X-3 is within the LAT FOV for ~ 15 ks per day. Given this and the typical flux of Cyg X-3 in the LAT band, we were compelled to use 1-day time bins in the following analysis.

Thirty arcminutes away from Cyg X-3 is a comparatively bright pulsar, PSR J2032+4127. Using `gtempem` and `gtpphase`, we added the phase from the most recent ephemeris available (Abdo et al. 2009a) to the events file. Then, we extracted events corresponding to the off-pulse phases (0–0.12, 0.2–0.6, and 0.72–1; Camilo et al. 2009). This removed 20% of the live time, so the exposure time of the off-pulse events file was corrected accordingly.

With `gtselect` and `gtmktime`, good time intervals from the off-pulse events file were selected in the 0.1–10 GeV energy range. At high energies ($\gtrsim 10$ GeV), Cyg X-3 is expected to have negligible emission. The events class was set to “3” which selects only high-quality

diffuse-background photons. To minimize background albedo photons from the Earth’s limb, zenith and rocking angles were restricted to less than 105° and 52° , respectively.

Exposure maps were generated by `gtexpcube` and `gtexpmap` while `gtbin` created photon counts maps in the region of interest. Emission from Cyg X-3 was not apparent on the counts maps, which is unsurprising given that Cyg X-3 is a faint source whose peak emission is at the low end of the LAT band where the point spread function is $5\text{--}10^\circ$. Complicating matters is the high level of diffuse background emission.

Following the procedure described in [Abdo et al. \(2009a\)](#), unbinned likelihood analysis was performed for the daily bins with `gtlike` considering photons inside a 7° radius of Cyg X-3 from all 1FGL ([Abdo et al. 2010](#)) sources up to 5° away, and all bright (detection significance $> 7\sigma$ and flux $[>100 \text{ MeV}] > 5 \times 10^{-8} \text{ ph cm}^{-2} \text{ s}^{-1}$) 1FGL sources up to 20° away. The instrument response function was “Pass 6 v3” (P6_V3_DIFFUSE), and the convergence relied on the NEWMINUIT method. Spectral models and input parameters from the 1FGL catalog were used except for pulsars 1FGL J2021.0+3651 and 1FGL J2021.5+4026 which had exponential cutoffs in addition to their power laws. Besides Cyg X-3, whose photon index and normalization were free, all sources had spectral parameters fixed to the 1FGL values. We included models for the galactic (`gll_iem_v02`) and extragalactic (`isotropic_iem_v02`) diffuse emission. The normalizations were left free to vary, as was the photon index of the galactic component.

6.3 Results

Long-term lightcurves of Cyg X-3 in soft and hard X-ray bands are presented in the top panel of Figure 6.1. The system was in a soft X-ray state during OJD $\sim 20\text{--}55$, with an episode of particular softness occurring during OJD 39–41.

The *Fermi* source flux integrated over 0.1–10 GeV is presented in the lower panel of Figure 6.1. In bins where the test statistic ($\text{TS} \sim -2 \ln L$, where L is the ratio of the likelihood of models without and with the source) is less than one, we compute and plot a 1σ upper limit. A peak is seen starting on OJD 43. The TS value in this bin is 18.4, which translates approximately to 4.3σ , and the flux is $(4 \pm 1) \times 10^{-6} \text{ ph cm}^{-2} \text{ s}^{-1}$. We are unable to achieve a more precise timing of the flare. Shifting the binning by half a day results in lower-significance and lower-flux detections in the range OJD 42.5–44.5. Analysis using higher-cadence binning does not yield significant detections of the peak.

Radio measurements from just before the γ -ray flare are presented in the top panel of Figure 6.2. A ~ 1 Jy radio flare occurred on OJD 42.25. There is evidence of an earlier flare on OJD 41.11, with a maximum observed flux density of 366 mJy, but only the rising portion of the flare is observed, so its properties are poorly-constrained. As is shown in Figure 6.3, the 15 GHz flux density of Cyg X-3 is depressed in the seven days following the γ -ray flare, with a mean flux density of 48 mJy and an 80th-percentile observed flux density of 66 mJy. (That is, 80% of the measurements during this period are < 66 mJy.) Starting on OJD 51,

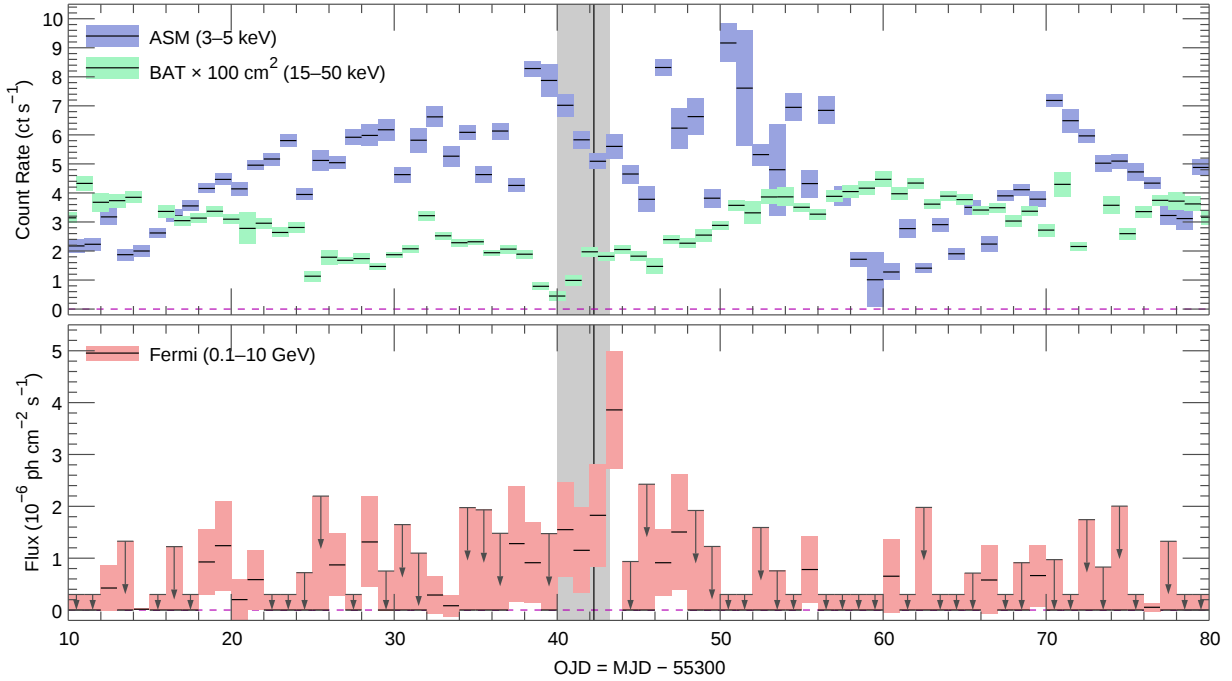


Figure 6.1: Long-term lightcurve of Cyg X-3 in X- and γ -rays. Shaded region: timespan depicted in Figure 6.2. Vertical black line: time of maximum observed 15 GHz flux density. On some days, our derived *Fermi* upper limits are smaller than the canonical background rate, $0.3 \times 10^{-6} \text{ ph cm}^{-2} \text{ s}^{-1}$; on those days, we use the background rate as the upper limit.

the typical radio flux density increases by a factor of ~ 2 to ~ 100 mJy.

The bottom panel of Figure 6.2 shows high-cadence hard X-ray measurements from before the γ -ray flare, made as Cyg X-3 exited the softest portion of its soft state. They show the recovery of the hard X-ray flux, subject to orbital modulation. The modulation makes it difficult to determine the significance of the drop in flux seen in the last three measurements; it may be evidence of a significant reduction in hard X-ray emission precursing the radio flare. Combining the *Swift*/BAT and *INTEGRAL* measurements, the hard X-ray flux has recovered by OJD ~ 42.0 , about 5 h before the observed radio flare. The γ -ray flare then occurred ~ 1 – 2 d later.

6.4 Discussion

If our observations are taken at face value — that is, one assumes no significant radio activity during coverage gaps and that the flaring stems from one ejection event — the sequencing of the radio and γ -ray emission is inconsistent with the companion IC model discussed in §6.1. Emission in γ -rays following a radio flare could be interpreted as a

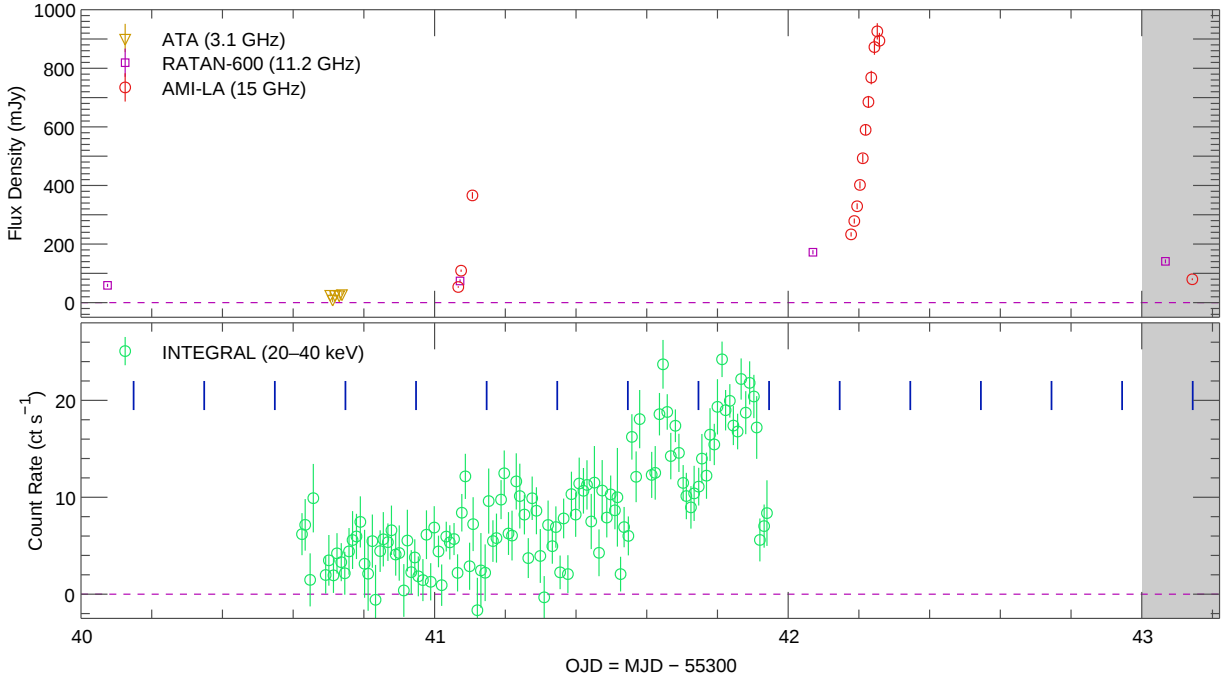


Figure 6.2: High-cadence lightcurve of Cyg X-3 in radio and hard X-rays in the time immediately before the γ -ray peak (= OJD 43–44). The X-rays are subject to orbital modulation with a 4.8-h periodicity. Vertical blue markers: times of X-ray minima according to parabolic ephemeris of Singh et al. (2002). Shaded region: partial timespan of *Fermi* peak.

reenergization of the relativistic jet electrons by a shock (Dubus et al. 2010). With a typical jet speed of $\sim 0.5 c$ (Tudose et al. 2007) and a delay of ~ 1 d, the reenergization would occur ~ 100 AU from the system. This is much larger than the orbital separation but three orders of magnitude smaller than the distance at which the jet is expected to interact significantly with the ISM (~ 1 pc). Such reenergization could be due to a collision between the jet and a dense clump of the WR wind (Araudo et al. 2009). In this scenario, the absence of γ -ray emission at the time of ejection can be explained by absorption within the system (Cerutti et al. 2011), although this model must be reconciled with the observations of Tavani et al. (2009) and Abdo et al. (2009a). The reenergization would lead to additional radio emission as the γ -ray-emitting electrons cool, although the ambient magnetic field, and hence synchrotron luminosity, will be much weaker than that found close to the system.

If the γ -ray flare is interpreted as the result of a discrete ejection event, the lack of a notable subsequent radio flare could be explained by the ejected material being largely hadronic. While hadronic γ -ray emission is not as efficient as IC upscattering, it does yield a higher ratio of γ -ray to radio luminosity (Romero & Vila 2008). The secondary leptons due to hadronic interactions would, however, radiate, and detailed studies typically find that their

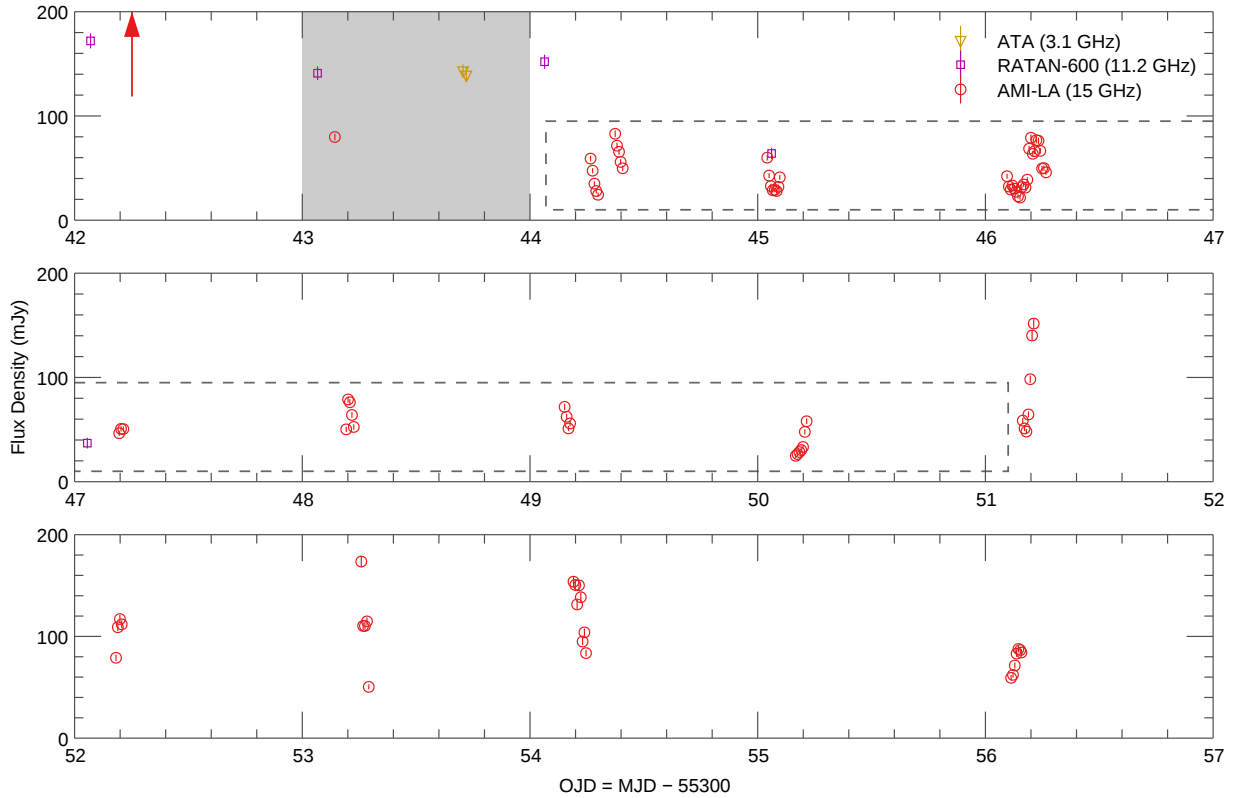


Figure 6.3: Radio lightcurve of Cyg X-3 during and after the γ -ray flare. The three panels proceed in chronological order from top to bottom. *Red arrow*: time of maximum observed 15 GHz flux density ($= 926$ mJy). *Shaded region*: timespan of the *Fermi* peak ($= (4 \pm 1) \times 10^{-6}$ ph cm $^{-2}$ s $^{-1}$). *Dashed box*: period of quiescent radio emission.

bolometric luminosity is comparable to that of the primaries (*e.g.* Vila & Romero 2010).

Should our data be taken at face value? Application of a simple synchrotron-cloud model (van der Laan 1966) supports the intuition that radio flares with sizes comparable to that of the largest observed could have occurred without detection, with the modeled lifetimes being ~ 0.2 d. While this model has had ambiguous success when applied to Cyg X-3 (Fender et al. 1997), its simplicity is advantageous for our sparse data, and it has been successfully applied to observations of other systems (*e.g.* Fender et al. 1997; Wilms et al. 2007; Prat et al. 2010). To obtain a better understanding of the radio behavior of Cyg X-3 we compared the 2010 May data to 15 years of archived AMI-LA / Ryle Telescope observations. The ~ 1 Jy radio flare is unusual: no flux densities > 400 mJy are detected in a 500 d span around the observations we describe.

In light of the predictions of the companion IC model, we consider particularly the radio observations after the observed γ -ray flare. The duty cycle of the observations is only 6.8%

in the seven-day post-flare quiescent interval. We identified 51 epochs similar to this in the archives, each lasting at least 7 days and having an 80th-percentile observed flux less than 70 mJy (cf. §6.3). These archival measurements have a duty cycle comparable to that of the 2010 May data (4.9%), but a much larger time on-source (55.5 days). Because this time is large compared to the total duration of the 2010 May post- γ -flare quiescence, a search for flares in the archival dataset can constrain the likelihood of there having been an unobserved flare in the 2010 May dataset, assuming no long-term evolution in the behavior of the quiescent state and stochastic flaring.

The epochs were identified by exhaustively searching for seven-day segments meeting the aforementioned criteria, then lengthening these segments as much as possible without violating the 80th-percentile constraint. Segments separated by less than three days were merged, in a few cases shortening the segment somewhat to preserve the statistical constraint. This method of construction does not bias against epochs containing rapid flares, no matter their size. Qualitatively, the lightcurve of the typical epoch starts high, drops to very low flux densities, and then becomes high again, possibly with rapid flares in the middle.

The largest flux density seen in the selected archival data was 524 mJy, in the context of a single rapidly-evolving (risetime ~ 0.1 d) flare that would not have necessarily been detected in the 2010 May observing. (Here we exclude slow, large, epoch-terminating flares that would have been easily seen in the 2010 May data.) Approximating the rate of such flares as one per 55.5 days of observing, we derive an 11% chance that such a flare occurred in the 7 days after the gamma-ray peak without being observed. Ignoring the epoch beginnings and endings, which contain high flux values by construction, we find that the flux density of Cyg X-3 is larger than 250 mJy only 0.3% of the time during its quiescent state.

6.5 Summary & Conclusions

We have presented observations of a 2010 May Cyg X-3 flaring episode in the radio, soft and hard X-ray, and γ -ray bands. Our data show a sequence of three notable events: a particular softening and partial recovery of the X-ray emission, a rapid ~ 1 Jy radio flare, and a γ -ray flare. Interpretation of the data is made more challenging by the sparse sampling of the radio data and the low cadence of the high-energy observations as compared to the rapidity with which the radio emission can evolve. Nonetheless, the observations we do have and comparisons to historical data challenge the interpretation that the γ -ray emission is due to IC upscattering of the companion radiation field by high-energy electrons leaving the system in a relativistic jet, because there is no evidence for the subsequent radio emission that one would expect to see from these electrons.

While the companion IC model of γ -ray emission from microquasars is clear and compelling, the lack of consistent radio/ γ -ray timing lags call into question its completeness. There could be more than one mechanism responsible for the γ -ray emission of Cyg X-3, becoming more or less relevant in different circumstances, or the reprocessing of the γ -ray emission by effects

such as pair production could be more significant than commonly assumed. Detections of or limits to very high-energy γ -rays ($\gtrsim 100$ GeV; e.g. [Aleksić et al. 2010](#)) or neutrinos from Cyg X-3 would aid in the understanding of the processes relevant to the emission in the *Fermi* band.

The power of multiband monitoring of Cyg X-3 promises to increase significantly with the addition of γ -ray data to the set of available observations. As Cyg X-3 inevitably produces more flares, frequent observations at all wavelengths are important to build a detailed understanding of the launching and propagation of its relativistic jets. Radio observations of $\lesssim 2$ h cadence with near-continuous coverage would be ideal for establishing an unambiguous relationship between radio and γ -ray flaring. Based on the phenomenology we observe, intensive radio observations should be triggered at the ends of very soft X-ray states without waiting for the detection of γ -ray flares.

Acknowledgments

I thank S. Corbel for useful discussions. This work was done with the support of a Space Sciences Lab Summer Fellowship and partial support from Fermi and INTEGRAL Guest Investigator grants NNX08AW58G, NNX08AX91G, and NNX10AG50G. AMI is operated by the University of Cambridge and supported by STFC.

Chapter 7

Conclusion

In this final chapter, I summarize my principal results and look ahead to their future applications and extensions.

7.1 Principal Results

I have developed and described `miriad-python`, an open-source framework for the rapid development of interferometric software in the Python programming language (Chapter 2). I have found it to be an essential tool both for developing a toolkit for processing ATA data and for the more general problem of exploring techniques to be used with the data that will be produced by the next generation of radio interferometers, which in many cases simply can't be handled by existing software packages. The fast experimentation made possible by the Python language is invaluable and I expect this to continue to hold as it will take a great deal of experimentation with real data to determine the best processing techniques for handling these next-generation data sets. Although Python is not suited to high-speed number crunching, it has many advantages, and its numerical speed is sufficient surprisingly often.

I have also developed and deployed a commensal observing system for the ATA (Chapter 3). Although the particular implementation is not expected to be very portable to other observatories, the lessons learned in approach the problem of running a commensal observing campaign should be broadly applicable. It's extremely important to have a clear division of responsibilities and to work out organizational matters before sitting down to engineer the system. Stateless systems in which components broadcast telemetry information are much more stable and easy to implement than ones in which they establish stateful, bidirectional connections. Although it is tempting to make such systems very flexible, relatively constrained systems have an enormous advantage in that they generally get "on the air" much more quickly.

I investigated the "calorimeter" model of the far-infrared/radio correlation seen in normal galaxies by taking advantage of one of the technical innovations of the ATA, its continuous

broadband frequency coverage (Chapter 4). I observed M82, NGC 253, and Arp 220 with the ATA at a frequencies ranging between 1 to 7 GHz, and combined these data with measurements from the literature to separate thermal and nonthermal contributions to the spectra. I find that simple power-law models of the nonthermal spectra are insufficient given the data. I investigated a model developed by Thompson et al. (2006) that predicts the shape of these nonthermal spectra assuming that the magnetic fields in starforming galaxies are stronger than is commonly assumed. I find mixed agreement between the Thompson et al. (2006) model and the data, and strongly rule out steep cosmic-ray electron energy injection indices if the model holds.

I conducted a large commensal survey for slow Galactic radio variability called ASGAR, the ATA Survey of Galactic Radio Dynamism (Chapter 5). The full survey comprises ~ 900 hr of integration on science targets over the course of ~ 2 yr. I described the observing strategy and data processing, which requires careful u - v modeling and imaging in order to cope with the significant large-scale static radio emission present in the Galactic plane. Early results confirm the success of the data processing techniques, demonstrate high-quality wide-field images of these complex fields, and place basic limits on the presence of Galactic variables and transients that are compatible with recent work. Analysis of the full survey, which is in progress, will yield sensitive limits on rates of relatively bright Galactic transients and measure the statistical properties of radio variability of Galactic sources in detail.

I performed a detailed multiwavelength investigation of the microquasar Cyg X-3 during its minor flare of 2010 May (Chapter 6). Only recently has this object been detected in the γ -ray regime with the *Fermi* and *AGILE* satellites, adding another tool by which the nature of its accretion disk and relativistic jets can be investigated. During this minor flare event, we detect a softening and recovery of the object's X-ray emission followed by a radio flare, which is not unusual. We also, however, detect a γ -ray flare following these events ~ 1.5 days later, which is difficult to reconcile with a model in which the high-energy emission is due to inverse Compton scattering in the inner part of the binary system. Other flaring events show a reversed sequencing, which suggests that the multiple processes may in fact be responsible for the high-energy emission.

7.2 Looking Ahead

There is no doubt that the dynamic radio sky will remain a vibrant area of study for some time to come. Indeed, with powerful telescopes such as the JVLA, ALMA, and LOFAR coming online, and other facilities such as LWA, MWA, ASKAP and MeerKAT just around the corner, there is almost certain to be even more interest in this area in the near future. Although routine monitoring of variability in the radio sky is proving to be a challenge, the science case remains compelling (Chapter 1; Cordes et al. 2004; Fender & Bell 2011; Frail et al. 2012).

One of the themes of this dissertation is the need for new tools to analyze the data that

come out of these new observatories. Compared to some of the facilities mentioned above, the ATA is a relatively simple system, and yet I still found that many elements of the existing data processing infrastructure were severely limited or simply unable to meet my scientific needs. As enumerated in §2.1, there are many efforts underway to remedy this situation, though there is very little cross-pollination among projects. This state of affairs should be somewhat cautionary: future research into improved techniques should stay strongly rooted in a concrete science case, and attention should be paid to keeping the scope of projects small. That being said, it appears that one-size-fits-all reduction packages will be less and less able to meet *all* the needs of future projects, so that even relatively small campaigns may be expected to require targeted bursts of nontrivial software development.

Complete processing of the ASGARD dataset is underway and is a high priority. As shown in Chapter 5, the full survey will be a powerful probe of Galactic radio variability, and the sooner any interesting events are discovered, the better. Although there’s always room for improvement in the pipeline, the system for analyzing these data is in place. Comparison with the surveys of [Becker et al. \(2010\)](#) and [Ofek et al. \(2011\)](#) will allow an analysis of the spatial distribution and variability structure function of Galactic radio variables across a number of dimensions. The combination of our radio monitoring of the *Kepler* field with optical data from the *Kepler* mission itself will provide an interesting probe of AGN variability. Likewise, combination of our static GC and Cygnus region maps with data at other radio bands will be a rich dataset for investigating extended radio structures such as nonthermal filaments ([Yusef-Zadeh et al. 1984](#); [Law et al. 2008](#)), H II regions ([Brogan et al. 2003](#); [Nord et al. 2006](#)), and supernova remnants ([Gray 1994](#)) at a wide range of GC separations, with implications for WIMP dark matter models motivated by *Fermi* γ -ray observations ([Hooper & Goodenough 2011](#); [Linden et al. 2011](#)), the acceleration and composition of Galactic cosmic rays ([Protheroe et al. 2008](#); [Crocker et al. 2011](#)), and the energetics of the interstellar medium in the GC region ([Morris & Serabyn 1996](#); [Crocker et al. 2010](#)).

The mechanisms of broadband emission from microquasars, and their implications for the astrophysics of their disks and jets, remain active areas of research (*e.g.* [Araudo et al. 2009](#); [Romero & Vila 2008](#); [Romero et al. 2010](#); [Cerutti et al. 2011](#); [Corbel et al. 2012](#); [Zdziarski et al. 2012](#)). Radio observations of systems such as Cyg X-3 allow investigation of the states of their relativistic jets, and detection of unexpected “radio quenching” events in Galactic variability surveys may provide an early-warning system for upcoming flares. Further multiwavelength campaigns, especially ones incorporating infrared and/or (sub)millimeter observations, will be vital in developing a complete astrophysical understanding of these systems ([Zdziarski et al. 2012](#)).

It is truly an exciting time for radio astronomy. In the next five years, an extremely powerful suite of new observatories will become available that will probe a new astronomical parameter space that has currently only begun to be explored. Should the Square Kilometer Array project proceed as envisioned, a truly game-changing instrument will be in our hands not much later. Different observatories will have different strengths and weaknesses when it comes to characterizing the dynamic radio sky.

With their ability to see virtually the entire sky instantaneously and lack of moving parts, the new dipole arrays (LWA, MWA, LOFAR) are superbly suited to continuous, commensal blind monitoring of the radio sky for radio transients, and in fact there are plans to do this monitoring at all of these telescopes. Since these facilities operate at long wavelengths, the dominant variable population that they will detect will likely be steep-spectrum coherent flares from flare stars, ultracool dwarfs and possibly hot Jupiters. These telescopes are less well-suited to the discovery of incoherently-emitting transient radio sources since these generally peak sooner and more dramatically at shorter wavelengths ([van der Laan 1966](#)).

For facilities operating in this shorter-wavelength regime, rapid followup of externally-detected events will be an extremely important use case, with the most exciting application being the challenging hunt for electromagnetic counterparts to gravitational-wave (GW) events ([Nakar & Piran 2011](#); [Metzger & Berger 2012](#)). Dynamic rescheduling and efficient subarray modes will thus be key operational capabilities, with the latter being important for quick exploration of the large error boxes that will be associated with the first GW detections ([Harry & LIGO Scientific Collaboration 2010](#)).

In the near term, the search for radio transients seems to be converging on a two-pronged approach featuring analysis of large archival datasets and targeted surveys of potential hosts of known event classes ([Frail et al. 2012](#)). Both of these techniques will become less relevant in the longer term. Regarding the first technique, as data rates become ever larger, anything other than realtime analysis will become infeasible, and indeed the designs for several next-generation facilities call for raw visibility data to be discarded. Regarding the second technique, the large fields of view and sensitivities of next-generation facilities (MeerKAT, ASKAP, SKA) make targeted searches decreasingly relevant. Significant fractions of observatory time will likely be dedicated to large surveys, so efficiency will be greatly improved if these are designed to search for variability and transients commensally. That being said, in the case of the SKA, current plans do not exclude a traditional TAC-based scheduling model ([Kellermann et al. 2006](#)), so standalone transient searches may need to be proposed for. Given the instrument’s survey speed, the most efficient approach would likely be a single catch-all campaign. A strawman project might observe a large area on a hierarchy of cadences emphasizing the monthly timescales common to extragalactic incoherent emitters. The revisitation cadence and tradeoff between footprint and sensitivity (as realized by the use of subarrays) would depend on the details of the dominant transient populations, which will be much better known by the time that the SKA is being scheduled. The relevant issues are considered in much more detail by [Cordes \(2007\)](#).

There’s great promise to the synoptic exploration of the dynamic radio sky. Although the planned next-generation radio observatories are ambitious engineering projects, many of the relevant risks are retired and their challenges are well-understood. It is much less clear, however, that it will be possible to analyze their data streams in a robust and automatic way. The challenge to the radio community is to tame these angry beasts and show how they can be used to produce the astrophysical breakthroughs that we know they are capable of.

Appendix A

Confused, Wide-Field Interferometric Imaging and Photometry

The ASGARD project described in Chapter 5 faces several technical challenges in its data processing pipeline due to the significant large-scale structure (LSS) present in the ASGARD fields and the wide field of view (FOV) of the Allen Telescope Array. I discuss some of the techniques used to handle wide-field interferometric imaging (§A.1), modeling and subtraction of the LSS (§A.2), and extraction of photometry from the compact sources (§A.3).

A.1 The Wide-Field Challenge

One of the first things learned by a budding interferometrist is that there is a Fourier transform (FT) relationship between what an interferometer measures (the “visibilities”) and the image of the sky:

$$\tilde{I}(l, m) = \iint V(u, v) e^{2\pi i(ul+vm)} du dv, \quad (\text{A.1})$$

where \tilde{I} is the apparent sky brightness distribution, V are the measured visibilities, u and v are the coordinates of the visibility domain, measured in wavelengths, and l and m are direction cosines relative to some “phase center” of the imaging process. (The apparent sky brightness is modified by, for instance, the mean primary beam response of the interferometer elements.) Visibilities measured by multi-element interferometers can be interpolated onto a regular grid and inverted with the fast Fourier transform (FFT; Cooley & Tukey 1965; Brigham 1988) to produce a “dirty” interferometric image. The classic references for this process are Thompson, Moran, & Swenson (2001) and Taylor, Carilli, & Perley (1999). This latter book is a conference proceedings; though its chapters will hereafter be referenced as separate publications, it is very worthwhile to obtain or peruse the entire volume¹. The best

¹The following link pulls up an ADS query listing all of the chapters in the volume in order: http://adsabs.harvard.edu/cgi-bin/nph-abs_connect?bibcode=1999ASPC..180&sort=RBIBCODE

introductory reference on Fourier transformations is [Bracewell \(2000\)](#).

Soon enough, however, the ugly truth emerges: Equation [A.1](#) is only an approximation, and a more accurate formulation is the much less elegant ([Thompson 1999](#), equation 2–21)

$$V(u, v, w) = \iint \frac{\tilde{I}(l, m)}{\sqrt{1 - l^2 - m^2}} e^{-2\pi i[ul + vm + w(\sqrt{1 - l^2 - m^2} - 1)]} dl dm. \quad (\text{A.2})$$

The w component of the visibility coordinate system has appeared, as have some inconvenient square roots, and now the apparent sky brightness \tilde{I} is unfortunately expressed as the solution of an integral equation. (The “minus one” in the parenthesized term in the phase exponent is due to the assumed use of fringe rotation in the interferometer signal path.) Equation [A.1](#) is valid when imaging small fields, but runs into problems when $|(l^2 + m^2)w| \not\ll 1$. It also breaks down when different array elements have different primary beam (PB) response patterns, when there is mixing between different polarization components, if individual array elements have varying pointing errors, and so on. There are also the well-known effects of time and bandwidth smearing ([Cotton 1999](#)), though these are not significant for standard ATA observing modes if the data are not heavily averaged and multifrequency synthesis ([Sault & Wieringa 1994](#); [Sault & Conway 1999](#)) is used. Rigorously coping with these sorts of issues requires the measurement equation formalism as developed in [Hamaker et al. \(1996\)](#), [Sault et al. \(1996\)](#), and [Hamaker & Bregman \(1996\)](#) and masterfully presented by [Smirnov \(2011a,b,c,d\)](#). The MeqTrees software package ([Noordam & Smirnov 2010](#)) implements this formalism for calibration and imaging.

Visibilities from the ATA can have w coordinates that approach $\pm 3000\lambda$, in which case Equation [A.1](#) starts becoming inaccurate for phase-center separations of $\sim 1^\circ$. In most fields, sources are easily detected out to this distance. I have furthermore found that the imaging of large-scale structures tends to be more sensitive to the w term, with effects becoming noticeable around the half-power radius, $\sim 0.5^\circ$. (It may just be that these visibilities have large amplitudes and so breakdown of Equation [A.1](#) becomes apparent more readily.) One effect of the breakdown of Equation [A.1](#) is to cause source positions to acquire an hour angle (HA) dependence, so that if wide-field image-based sky modeling is to be used, the w term must be accounted for correctly. The same logic holds if you want to perform image-domain fitting on compact sources far from the phase center. Wide-field imaging algorithms are therefore necessary to meet the science goals of ASGARD.

The simplest technique for dealing with wide-field emission is to avoid imaging altogether. In particular, unresolved sources can be dealt with in closed form, so they can be modeled in the u - v domain and subtracted from the data if necessary. This approach is obviously insufficient if the sources far from the phase center are resolved or their properties are not accurately known.

If images are necessary, there are several choices. Polyhedral or “faceted” imaging techniques can be used, in which the visibilities are rephased to multiple phase centers to build up a larger image ([Cornwell & Perley 1992](#); [Perley 1999](#)). This approach is fairly inconvenient and computationally expensive, as is the three-dimensional FT algorithm also

outlined by [Cornwell & Perley \(1992\)](#). The w -projection algorithm achieves very good results with less severe computational costs and little added user complexity ([Cornwell et al. 2008](#)). It has been generalized to the A -projection algorithm ([Bhatnagar et al. 2008](#)), which can correct for other effects (*e.g.*, pointing errors) but is correspondingly more complex to drive and does not yet appear to be widely available.

MIRIAD unfortunately does not implement any wide-field imaging algorithms. The `invert` task does support rephasing of the input data, but MIRIAD’s image-domain deconvolution algorithms are fundamentally unable to cope with wide-field emission. Imaging with CASA is about an order of magnitude slower than with MIRIAD, and conversion between MIRIAD format and the Measurement Set format is inconvenient at best, but these steps are necessary in order to correctly image the ASGAR data. (Though this is far from the largest issue that would be involved in its use, it should be noted that the u - v modeling subsystem of MeqTrees does not currently support w -projection either, although at the time of writing this feature was being implemented.)

A.2 Large-Scale Structure

The technical challenge faced by ASGAR is to find transients and perform photometry in fields confused by large-scale structure². Note that the problems presented by widefield imaging and LSS are not necessarily linked: one can certainly find and image wide fields that do not contain LSS, and large densely-packed interferometer elements can detect LSS while still having a PB response narrow enough that the w term remains unimportant. As mentioned above, however, LSS as seen by the ATA can remain significant past the half-power point, where empirically the w term is relevant.

Interferometric imaging of LSS is a notoriously thorny problem, so it’s worth considering what our scientific motivations are: perhaps we can avoid dealing with it altogether. Unfortunately, it is hard to see a way out. The primary goal of ASGAR is to search for Galactic radio transients, and it is difficult to envision any systematic way to do that in LSS-confused images. Image subtraction isn’t an option because the dirty beam is different for every observation — even if each observation of a given field had consistent HA coverage, which would be difficult enough to accomplish in practice, different sets of antennas would be available, and their weights would differ. Any variability in the apparent sky would also have nonlocal effects due to sidelobes. The same issues would affect photometry of known, compact sources: unless the image is thoroughly deconvolved, image-domain fits will be severely affected by sidelobes.

Weighting or filtering of the u - v data is not as effective as one might like. Although the most LSS power is at short u - v distances, it is present on long baselines too. Figure [A.1](#) shows an image made with short baselines filtered out and may be compared directly to Figure [5.5](#),

²The term “confused” seems less than ideal here, since it could also suggest a field with densely-packed compact sources, but I’m not able to come up with anything better.

in which LSS subtraction is used. When u - v filtering is used, the large-scale structures are removed, but low-level “fuzz” remains and degrades the quality of the deconvolution. A secondary penalty to this approach is that image S/N is reduced nontrivially because the ATA’s compact configuration leads to a large number of short baselines.

It thus seems that the LSS must be modeled. Fortunately, the LSS emission is expected to be static. If a good model image of the LSS can be generated, it can be subtracted from the u - v data, and the imaging process can then proceed as usual with a standard CLEAN-type deconvolution dealing with any differences from the model. Figure A.2 shows a typical example of the effect of this subtraction in terms of the amplitudes of observed visibilities. Figure A.3 shows the visibility amplitudes in the *Kepler* field, which is virtually unaffected by LSS, and may be compared to the top panel of Figure A.2. The subtraction approach indeed works well in practice, although there are hints that there is some apparent variability in the LSS. The most obvious examples for this are the three bright, marginally-resolved sources found at the half-power points of the ASGARD Cyg X-3 pointing. These sources are included in the Cyg X-3 LSS model. In some epochs, they subtract extremely well, but in other epochs they are noticeably under- or over-subtracted. Calculations using typical numbers (Harp et al. 2011) show that the scale of this effect is consistent with it being due to pointing errors, which would indeed have the largest effect on emission at the half-power radius. (There are epochs where one of the sources is undersubtracted but the other two are oversubtracted, so the differences cannot simply be due to a scalar gain error.) Our calibration and imaging tools are not able to account for this effect, so I haven’t investigated this possibility more deeply.

The most straightforward approach to creating LSS models is simply to generate and deconvolve deep images for each field. ASGARD uses a more complicated procedure: once compact sources are identified in the deep images, they are subtracted from each epoch’s visibilities, and a new deep image without compact sources is made. This approach is taken for two reasons. Firstly, deconvolution of the LSS is best done using maximum-entropy-type algorithms (Gull & Daniell 1978), which generally fare poorly with compact sources. Secondly, by removing compact sources from the LSS model, they are preserved in the LSS-subtracted epoch images. It would be possible to build a photometric pipeline around a paradigm in which each epoch image is (ideally) noiselike, with tiny variations comparable to the noise as the subtracted sources vary, but my intuition is that the approach taken will be more reliable and easy to implement. I haven’t investigated the full-subtraction alternative, however.

In the ASGARD approach, the LSS model can be thought of as simply an aid to the deconvolution process: taking out all of the known static flux lets a CLEAN-like algorithm zero in on the easy-to-handle and possibly-variable compact sources. In fact, after imaging an epoch with the above procedure, the LSS model can be added back into the original data (convolving with the synthesized beam). To the extent that the imaging process is linear, the final image should be equivalent to what you’d get from deconvolving the unmodified data with a really accurate initial model. Doing this addresses a potential pitfall of the LSS modeling system, namely that each source is formally subject to an additive flux bias from

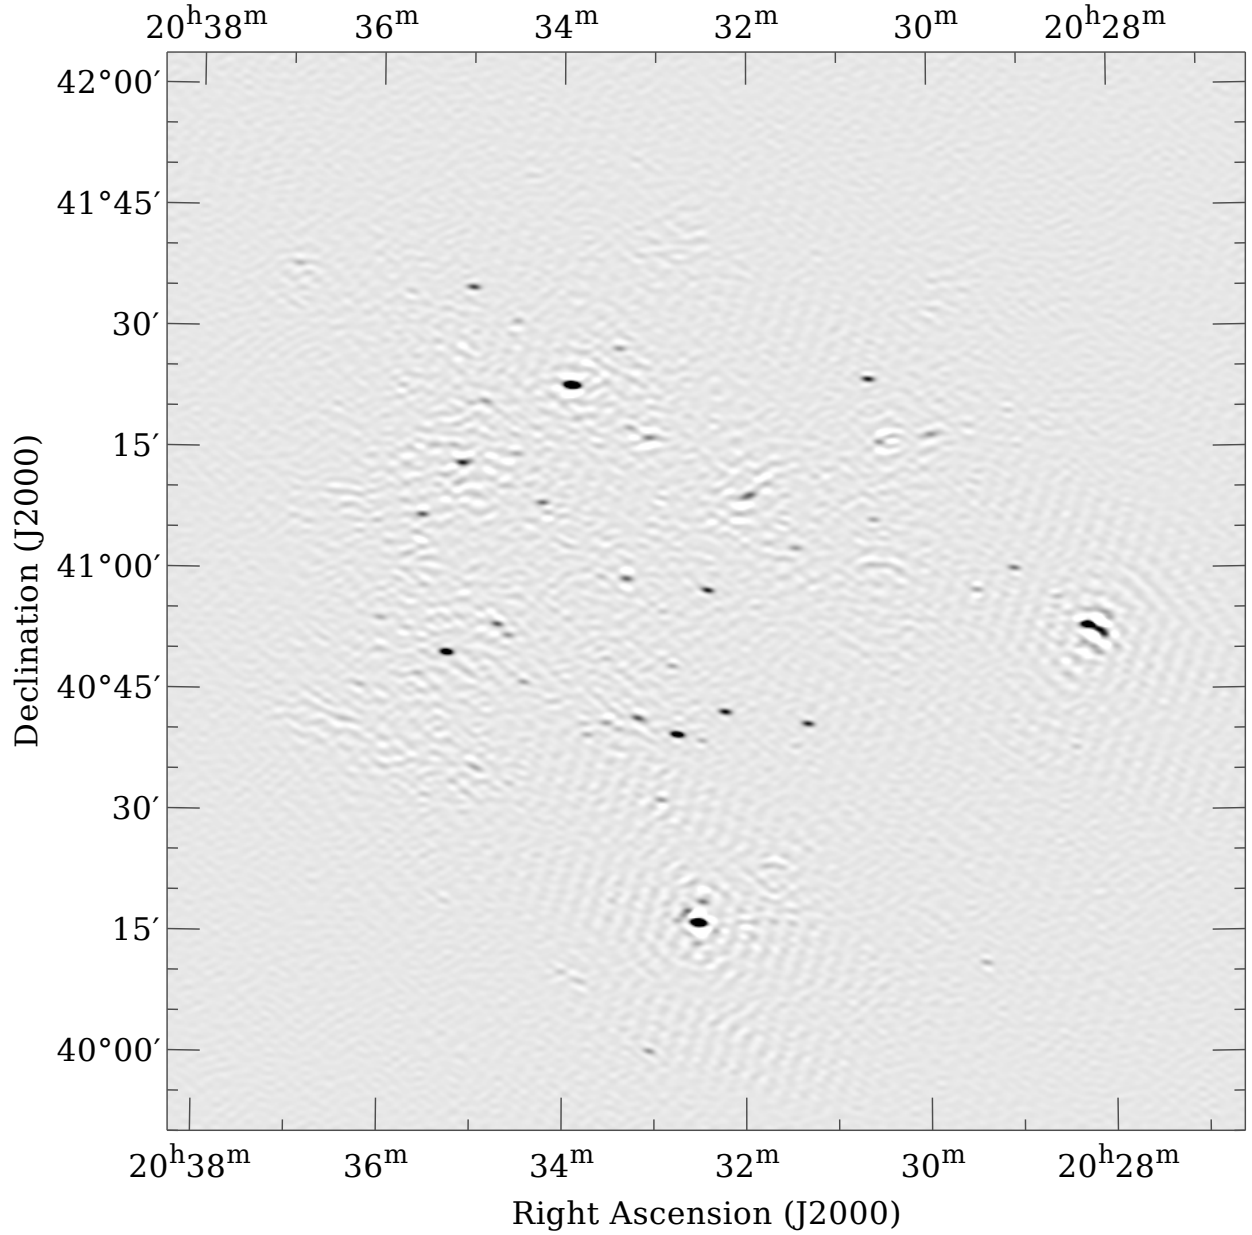


Figure A.1: Image of the Cyg X-3 pointing from 2011 Feb 01 with baselines shorter than 150 m filtered out. The grayscale is linear from -5 mJy/bm (white) to 50 mJy/bm (black). This figure may be compared directly to Figure 5.5, in which the visibilities have an LSS model subtracted rather than simply being truncated.

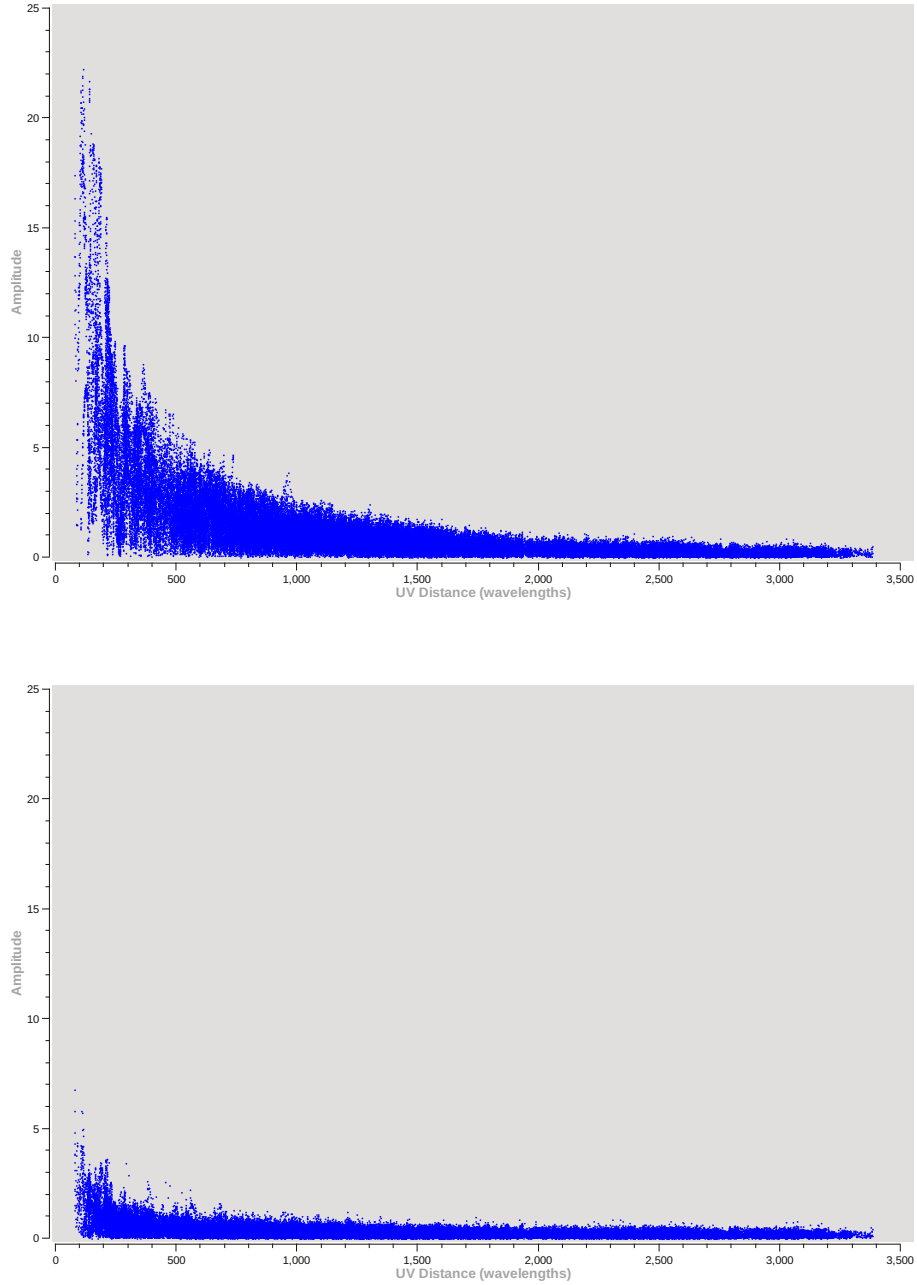


Figure A.2: Visibility amplitude distribution in the 2011 Feb 01 observations of Cyg X-3. *Upper panel:* before LSS model subtraction. *Lower panel:* after LSS model subtraction. Amplitudes are measured in Jy.

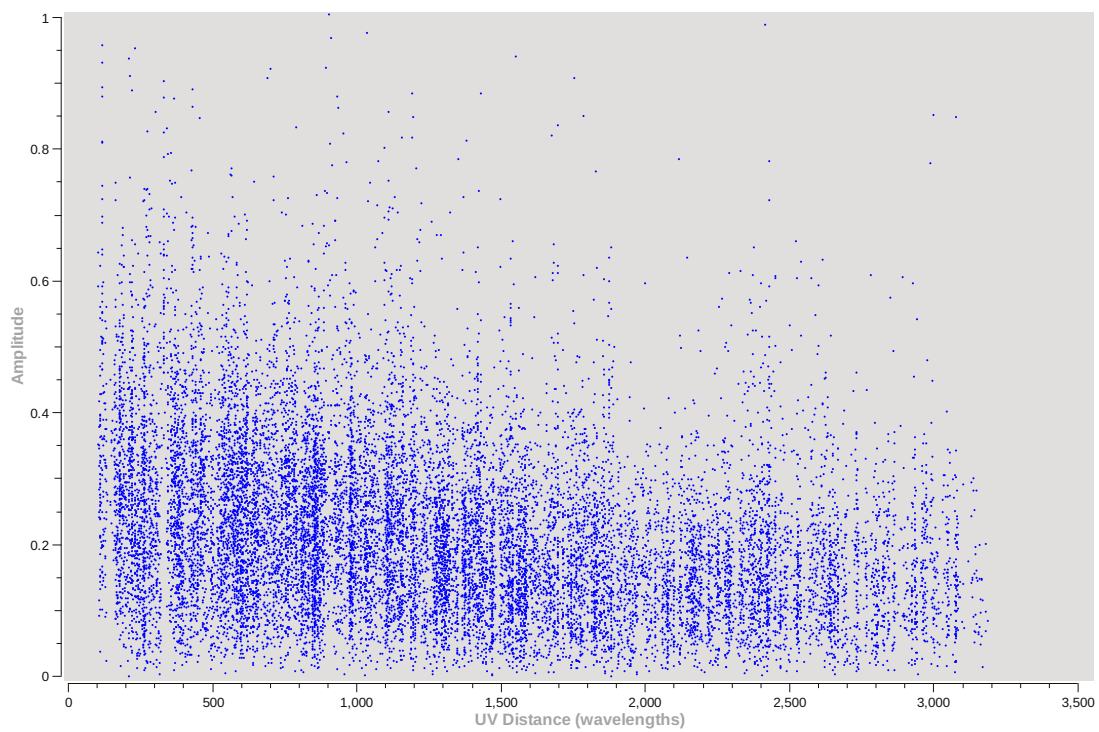


Figure A.3: Visibility amplitude distribution in the 2010 Oct 28 observations of the *Kepler* field. Amplitudes are measured in Jy. Note the differing vertical scale than is used in Figure A.2.

the LSS model, but investigations show that adding back the convolved LSS model doesn't change photometric results significantly.

In theory, these well-developed LSS models should provide a great model with which to self-calibrate the data. I have not been able to get this to work in practice. My general experience with ATA data has been that self-calibration has had very little success in improving results from what one gets after an initial gain and bandpass calibration. My intuition is that there should be some particular combination of solution interval, data selection, solution constraints, *etc* that leads to good results, but I've experimented with this on-and-off for quite some time now and have not yet found a recipe that succeeds reliably. This holds true even for simple calibrator observations, so there's at least the hope that once a solution is found, it may work as well for the complicated ASGARD fields as for less challenging fields. (The fact that imaging works as well as it does shows that the u - v models of the ASGARD fields are pretty good.) Close examination of ATA antenna phases as a function of time reveals systematic sawtooth-like variations with timescales of \sim minutes and amplitudes increasing with east-west antenna displacement from the array center, up to scales of $\sim 20^\circ$ at times. Self-calibration should be able to solve for these phase offsets in theory, but they certainly represent a challenge since the requisite solution interval is very small and the derivatives in the phase offsets can be large.

The above discussion has ignored polarization. It is fortunate that good results can be obtained simply by assuming that all emission is completely unpolarized. There would be two substantial problems if this were not the case. Firstly, the ASGARD observing scheme did not allow for standard polarimetric calibration observations. It is possible that the proper calibration could be stitched together by making aggressive assumptions about the system stability, but that would be a risky undertaking at best. Secondly, as mentioned before, wide-field polarimetric imaging is difficult, especially when LSS and PB cross-polarization terms are present, and the available software packages are essentially unable to cope in this situation. The A -projection algorithm of [Bhatnagar et al. \(2008\)](#), for instance, substantially improves results even in a relatively unchallenging case (see their Figure 4). The challenge would be even greater for ASGARD since the polarization properties of the ATA PB, and their variation from antenna to antenna, are only incompletely characterized and appear to be difficult to calibrate ([Law & de Villiers 2010](#); [de Villiers & Law 2010](#)). That being said, even in the complex Cyg X-3 field, the XX and YY visibilities on most ASGARD baselines agree very well. Some short baselines show significant differences between the polarization products, however. It is possible that this could be indicative of large-scale polarized emission, but is more likely suggestive of cross-polarization effects in the PBs of certain antennas, which are known to be present ([Harp et al. 2011](#)). Tests indicate that cross-polarization terms are far from the limiting factor in ASGARD imaging, although it should be remembered that correct Stokes I imaging fundamentally requires correct treatment of all polarization products ([Heiles 2002](#)).

Before the imaging steps can be performed, of course, the raw data must be flagged and calibrated. Tools and procedures for doing so are described in [Williams \(2010\)](#) and

throughout this dissertation, in §§2.6.2, 2.6.3, 4.3, and 5.3.2.

A.3 Photometry

Once the visibilities have had LSS subtracted out and been imaged in a way that correctly deals with the large FOV, the problem of searching for transients and performing photometry is the same as for any other multiepoch radio survey. All of the techniques used in ATATS (Croft et al. 2010a,b, 2011), PiGSS (Bower et al. 2010, 2011) and other surveys can be applied virtually unchanged.

In ASGAR, analysis of the epoch images is divided into two steps. First, photometry of the known sources is performed with constrained image-domain fitting, and these sources are subtracted from the image. Second, `sfind` is used to search the residual image for unexpected sources. In contrast, the SLOW pipeline used in ATATS runs `sfind` on the epoch images and cross-matches the catalogs generated from each epoch to build a master catalog and photometry database. I disfavor this approach because it doesn't take advantage of prior knowledge of the locations and shapes of all of the steady sources in each field. This information should improve the fits, especially for faint sources that are only marginally detected in each epoch. (Analogous reasoning should also hold for the deconvolution process. Prior knowledge can be encoded somewhat via “CLEAN boxes,” but some information is lost and the ability to deconvolve unexpected emission from beyond the expected source population is lost.)

The actual photometric extraction of a source can be done in myriad ways. For instance, one might compute some kind of statistic on a set of source pixels (*e.g.*, take the maximum), or fit a profile. If fitting a profile, various profile functions can be used (*e.g.*, biquadratic, 2D Gaussian, *etc*), various parameters (*e.g.*, location, shape) can be left free or unconstrained, and either the peak or the integrated flux result can be used as a basis for further processing. Background terms can likewise be modeled in various ways, or ignored. In ASGAR, Gaussian profiles with fixed shapes and locations are fit, the integrated flux density is extracted, and no background is modeled in the epoch images. Bannister et al. (2011a,b) describe a particularly elaborate system that uses different techniques depending on the source properties. It's not clear how to compare the efficacy of different extraction methodologies besides deploying them and examining summary statistics of the source flux results.

Bannister et al. (2011b) and Ofek et al. (2011) use a “post-imaging calibration” step after photometric extraction. In this process, image-wide flux density scaling factors are calculated for and applied to each epoch. The motivation for such a calibration step is understandable: in most cases, the flux density scale for an epoch is set from one calibrator observation, so there's a clear opportunity for the scale to be mis-set. Likewise, since there are many sources in each field of view and most are not systematically varying, it seems that some sort of mean ratio between the epoch fluxes and fiducial values ought to give a good measure of any error in the flux scale. In practice, however, it is found in ASGAR and Bower et al.

(2011) that this technique is not as straightforwardly helpful as might be thought. Our experience is that there is usually a *subset* of sources in each epoch, especially the bright ones, that benefit from the calibration, but there is also a group of sources that clearly do not. This holds over a variety of approaches to determining each epoch’s scale factor, including simple averaging, medians, sums of fluxes, and the RANSAC method (Fischler & Bolles 1981). Barring compelling benefits, it seems better to avoid the extra data manipulation. The fundamental rationale for this extra calibration step seems very sound, though, so perhaps better-behaved images will yield better results.

Appendix B

Probability of Constancy From a Set of Measurements

In §5.4.5 we consider the probability P_c that a source is unvarying, given a set of flux measurements. In this Appendix we give the equation used to compute this value and derive a condition that indicates the presence of incompletely controlled systematic measurement errors.¹

To compute P_c we must use the full cumulative distribution function (CDF) of the χ_k^2 family of distributions. The probability of accepting the hypothesis that a given source is constant is the probability of finding a χ^2 value at least as large as the one obtained for that source. This is

$$P_c = Q\left(\frac{k}{2}, \frac{\chi^2}{2}\right), \quad (\text{B.1})$$

where Q is the complementary regularized Γ function

$$Q(s, x) = \frac{\Gamma(s, x)}{\Gamma(s, 0)}, \quad (\text{B.2})$$

and in turn $\Gamma(s, x)$ is the upper incomplete Γ function

$$\Gamma(s, x) = \int_x^\infty t^{s-1} e^{-t} dt. \quad (\text{B.3})$$

In the theoretical case of no varying sources and truly Gaussian errors, the observed P_c will be uniformly distributed between 0 and 1. If these conditions do not hold — for instance, if systematics are present — the actual distribution of P_c values can differ. Denote the CDF of the observed P_c values $F_c(p)$; that is, $F_c(p)$ is the probability of measuring $P_c < p$ for an arbitrary source in the ensemble. We claim that when systematics are controlled, genuine source variability can only lead to $F_c(p) \geq p$, ignoring variations due to finite sample

¹This appendix was subsequently published with Chapter 5 in Williams *et al.*, 2013 [ApJ 762 85](#).

size. Therefore, data that show $F_c(p) < p$ are suggestive of uncontrolled systematics. By construction,

$$\begin{aligned} F_c(0) &= 0, \\ F_c(1) &= 1, \text{ and} \\ dF_c/dp &\geq 0. \end{aligned} \tag{B.4}$$

We furthermore argue that the underlying P_c probability density function

$$f_c(p) = \frac{dF_c}{dp} \tag{B.5}$$

must be nonincreasing, because absent systematic measurement errors it is impossible for sources to appear statistically less variable (biased towards larger P_c) than ones that are in fact unvarying. Therefore

$$d^2F_c/dp^2 \leq 0, \tag{B.6}$$

and combining with Equations [B.4](#) we must in fact have $F_c(p) \geq p$ on the unit interval. The finite number of sources being measured leads to uncertainty in F_c that allows only a probabilistic statement as to whether a given observed F_c is consistent with completely controlled systematics.

Bibliography

- Abada-Simon, M., & Aubier, M. 1997, [A&AS](#), **125**, 511
- Abdo, A. A., Ackermann, M., Ajello, M., et al. 2009a, [Science](#), **326**, 840
- . 2009b, [Science](#), **325**, 848
- Abdo, A. A., Ackermann, M., Ajello, M., et al. 2010, [ApJS](#), **188**, 405
- AIPS++ Steering Committee. 1993, [AIPS++ Memo Series](#), #110
- Ait-Allal, D., Weber, R., Dumez-Viou, C., Cognard, I., & Theureau, G. 2012, [Comptes Rendus Physique](#), **13**, 80
- Alcock, C., Akerlof, C. W., Allsman, R. A., et al. 1993, [Nature](#), **365**, 621
- Aleksić, J., Antonelli, L. A., Antoranz, P., et al. 2010, [ApJ](#), **721**, 843
- Aller, M. A., Hughes, P. A., & Aller, H. D. 2010, in [Fermi Meets Jansky — AGN in Radio and Gamma-Rays](#), ed. T. Savolainen, E. Ros, R. W. Porcas, & J. A. Zensus, 65
- Anantharamaiah, K. R., Viallefond, F., Mohan, N. R., Goss, W. M., & Zhao, J. H. 2000, [ApJ](#), **537**, 613
- Appleton, P. N., Fadda, D. T., Marleau, F. R., et al. 2004, [ApJS](#), **154**, 147
- Araudo, A. T., Bosch-Ramon, V., & Romero, G. E. 2009, [A&A](#), **503**, 673
- Atwood, W. B., Abdo, A. A., Ackermann, M., et al. 2009, [ApJ](#), **697**, 1071
- Aubourg, E., Bareyre, P., Brehin, S., et al. 1993, [Nature](#), **365**, 623
- Avni, Y. 1976, [ApJ](#), **210**, 642
- Baars, J. W. M., Genzel, R., Pauliny-Toth, I. I. K., & Witzel, A. 1977, [A&A](#), **61**, 99
- Backer, D. C. 2007, [PAPER Memo Series](#), #11
- Bannister, K. W., Murphy, T., Gaensler, B. M., Hunstead, R. W., & Chatterjee, S. 2011a, [MNRAS](#), **412**, 634
- . 2011b, [MNRAS](#), **418**, 2813
- Barrett, P. E., & Bridgman, W. T. 1999, [ASP Conf. Ser.](#), **172**, 483
- Barthelmy, S. D., Barbier, L. M., Cummings, J. R., et al. 2005, [Space Sci. Rev.](#), **120**, 143
- Baumgartner, S., Kantke, T., & Dietz, U. S. 2011, “Bahnknoten München”, http://www.stadtcreation.de/munich/bahnknoten_muenchen.html
- Beck, R., Biermann, P., Emerson, D. T., & Wiełebinski, R. 1979, [A&A](#), **77**, 25
- Beck, R., & Reich, W. 1985, in [IAU Symposium #106: The Milky Way Galaxy](#), ed. H. van Woerden, R. J. Allen, & W. B. Burton, 239
- Becker, R. H., Helfand, D. J., White, R. L., & Proctor, D. D. 2010, [AJ](#), **140**, 157
- Becker, R. H., White, R. L., & Helfand, D. J. 1995, [ApJ](#), **450**, 559

- Becker, R. H., White, R. L., Helfand, D. J., & Zoonematkermani, S. 1994, [ApJS](#), **91**, 347
- Bell, M. E. 2011, [PhD thesis, University of Southampton](#)
- Bell, M. E., Fender, R. P., Swinbank, J., et al. 2011, [MNRAS](#), **415**, 2
- Berger, Ball, Becker, et al. 2001, [Nature](#), **410**, 338
- Berger, E. 2002, [ApJ](#), **572**, 503
- Berger, E., Zauderer, A., Pooley, G. G., et al. 2012, [ApJ](#), **748**, 36
- Bhatnagar, S., Cornwell, T. J., Golap, K., & Uson, J. M. 2008, [A&A](#), **487**, 419
- Binder, R. V. 1996, [Software Testing, Verification, & Reliability](#), **6**, 125
- Birkinshaw, M. 1978, [MNRAS](#), **184**, 387
- Blair, S. K., & ATA Team. 2009, [BAAS](#), **41**, 670
- Blakeslee, J. P., Anderson, K. R., Meurer, G. R., Benítez, N., & Magee, D. 2003, [ASP Conf. Ser.](#), **295**, 257
- Blandford, R., & Eichler, D. 1987, [Physics Reports](#), **154**, 1
- Bloom, J. S., Giannios, D., Metzger, B. D., et al. 2011, [Science](#), **333**, 203
- Bodaghee, A., Tomsick, J. A., Rodriguez, J., & James, J. B. 2012, [ApJ](#), **744**, 108
- Bond, I. A., Udalski, A., Jaroszyński, M., et al. 2004, [ApJL](#), **606**, L155
- Bonnet-Bidaud, J. M., & Chardin, G. 1988, [Physics Reports](#), **170**, 325
- Booth, R. S., & Jonas, J. L. 2012, [African Skies](#), **16**, 101
- Bopp, B. W., & Fekel, F. 1977, [AJ](#), **82**, 490
- Boulanger, F., & Perault, M. 1988, [ApJ](#), **330**, 964
- Bower, G. C., Roberts, D. A., Yusef-Zadeh, F., et al. 2005, [ApJ](#), **633**, 218
- Bower, G. C., & Saul, D. 2011, [ApJL](#), **728**, L14
- Bower, G. C., Saul, D., Bloom, J. S., et al. 2007, [ApJ](#), **666**, 346
- Bower, G. C., Whysong, D., Blair, S., et al. 2011, [ApJ](#), **739**, 76
- Bower, G. C., Croft, S., Keating, G., et al. 2010, [ApJ](#), **725**, 1792
- Bracewell, R. N. 2000, [“The Fourier transform and its applications”](#), ISBN 978-0073039381 (Boston: McGraw-Hill)
- Breitschwerdt, D., McKenzie, J. F., & Völk, H. J. 1991, [A&A](#), **245**, 79
- Brigham, E. O. 1988, [“The fast Fourier transform and its applications”](#), ISBN 978-0133075052 (Upper Saddle River, NJ: Prentice Hall)
- Brinkmann, W., Siebert, J., Feigelson, E. D., et al. 1997, [A&A](#), **323**, 739
- Brogan, C. L., Nord, M., Kassim, N., Lazio, J., & Anantharamaiah, K. 2003, [Astronomische Nachrichten Supplement](#), **324**, 17
- Broten, N. W., Vallée, J. P., & MacLeod, J. M. 1986, [A&AS](#), **64**, 499
- Brown, R., Wild, W., & Cunningham, C. 2004, [Advances in Space Research](#), **34**, 555
- Brunthaler, A., Menten, K. M., Reid, M. J., et al. 2009, [A&A](#), **499**, L17
- Bulgarelli, A., Striani, E., Sabatini, S., et al. 2010, [The Astronomer’s Telegram](#), **#2609**
- Bunton, J. D., & Hay, S. G. 2004, [Experimental Astronomy](#), **17**, 381
- Bunton, J. D., Jackson, C. A., & Sadler, E. M. 2002, [SKA Memo Series](#), **#23**
- Burbidge, G. R. 1956, [ApJ](#), **124**, 416
- Burke-Spolaor, S., Bailes, M., Ekers, R., Macquart, J. P., & Crawford, F. 2011, [ApJ](#), **727**, 18

- Cai, X., Langtangen, H. P., & Moe, H. 2005, [J Scientific Programming](#), 13, 31
- Calabretta, M. R. 2011, in *Astrophysics Source Code Library*, [record ascl:1108.003](#)
- Camilo, F., Ransom, S. M., Halpern, J. P., et al. 2006, [Nature](#), 442, 892
- Camilo, F., Ray, P. S., Ransom, S. M., et al. 2009, [ApJ](#), 705, 1
- Carico, D. P., Keene, J., Soifer, B. T., & Neugebauer, G. 1992, [PASP](#), 104, 1086
- Carilli, C., & Rawlings, S. 2004, [New A Rev.](#), 48, 979
- Carilli, C. L., Wrobel, J. M., & Ulvestad, J. S. 1998, [AJ](#), 115, 928
- Carlstrom, J. E., & Kronberg, P. P. 1991, [ApJ](#), 366, 422
- Carns, P. H., Ligon, W. B., Ross, R. B., & Thakur, R. 2000, in [Proceedings of the 4th Annual Linux Showcase & Conference](#), 28
- Cerutti, B., Dubus, G., Malzac, J., et al. 2011, [A&A](#), 529, A120
- Cherepashchuk, A. M., & Moffat, A. F. J. 1994, [ApJ](#), 424, L53
- Cheung, A. C., Rank, D. M., Townes, C. H., Thornton, D. D., & Welch, W. J. 1969, [Nature](#), 221, 626
- Clark, B. G. 1980, [A&A](#), 89, 377
- Clark, M. A., La Plante, P., & Greenhill, L. J. 2012, [Intl J of High Performance Computing Applications](#), 1094342012444794
- Cohen, M. H. 1975, [Annals of the New York Academy of Sciences](#), 262, 428
- Cohen, M. H., Kellermann, K. I., Shaffer, D. B., et al. 1977, [Nature](#), 268, 405
- Condon, J. J. 1992, [ARA&A](#), 30, 575
- Condon, J. J., Cotton, W. D., Greisen, E. W., et al. 1998, [AJ](#), 115, 1693
- Connell, J. J. 1998, [ApJL](#), 501, L59
- Conway, J. E., Cornwell, T. J., & Wilkinson, P. N. 1990, [MNRAS](#), 246, 490
- Cooley, J. W., & Tukey, J. W. 1965, [Mathematics of Computation](#), 19, 297
- Corbel, S., & Fermi Large Area Telescope Collaboration. 2010, [The Astronomer's Telegram](#), #2611
- Corbel, S., & Hays, E. 2010, [The Astronomer's Telegram](#), #2646
- Corbel, S., Dubus, G., Tomsick, J. A., et al. 2012, [MNRAS](#), 421, 2947
- Cordes, J. 2007, [SKA Memo Series](#), #97
- Cordes, J. M., Lazio, T. J. W., & McLaughlin, M. A. 2004, [New A Rev.](#), 48, 1459
- Cornwell, T., Braun, R., & Briggs, D. S. 1999, [ASP Conf. Ser.](#), 180, 151
- Cornwell, T., & Humphreys, B. 2010, [SKA Memo Series](#), #128
- Cornwell, T. J. 1987, [A&A](#), 180, 269
- Cornwell, T. J., Golap, K., & Bhatnagar, S. 2008, [IEEE J Selected Topics in Signal Processing](#), 2, 647
- Cornwell, T. J., & Perley, R. A. 1992, [A&A](#), 261, 353
- Cotton, W. D. 1999, [ASP Conf. Ser.](#), 180, 357
- Cotton, W. D. 2008a, [Obit Development Memo Series](#), #1
- . 2008b, [PASP](#), 120, 439
- Cotton, W. D., & Perley, R. 2010, [Obit Development Memo Series](#), #21
- Cox, M. J., Eales, S. A. E., Alexander, P., & Fitt, A. J. 1988, [MNRAS](#), 235, 1227

- Crocker, R. M., Jones, D. I., Aharonian, F., et al. 2011, *MNRAS*, **413**, 763
- Crocker, R. M., Jones, D. I., Melia, F., Ott, J., & Protheroe, R. J. 2010, *Nature*, **463**, 65
- Croft, S., Bower, G. C., Keating, G., et al. 2011, *ApJ*, **731**, 34
- Croft, S., Bower, G. C., Ackermann, R., et al. 2010a, *ApJ*, **719**, 45
- . 2010b, *ApJ*, **724**, 827
- Cutri, R. M., Wright, E. L., Conrow, T., et al. 2012, “Explanatory Supplement to the WISE All-Sky Data Release Products” (IPAC/Caltech)
- Dalcin, L., Paz, R., Storti, M., & Delia, J. 2008, *J Parallel & Distributed Computing*, **68**, 655
- Davies, R. D., Walsh, D., Browne, I. W. A., Edwards, M. R., & Noble, R. G. 1976, *Nature*, **261**, 476
- de Jong, T., Klein, U., Wielebinski, R., & Wunderlich, E. 1985, *A&A*, **147**, L6
- de Villiers, M., & Law, C. J. 2010, *ATA Memo Series*, #88
- de Vries, W. H., Becker, R. H., White, R. L., & Helfand, D. J. 2004, *AJ*, **127**, 2565
- DeBoer, D. R., Gough, R. G., Bunton, J. D., et al. 2009, *IEEE Proc.*, **97**, 1507
- Dennett-Thorpe, J., & de Bruyn, A. G. 2002, *Nature*, **415**, 57
- Devereux, N. A., & Eales, S. A. 1989, *ApJ*, **340**, 708
- Downey, E. C. 2011, in *Astrophysics Source Code Library*, record ascl:1112.013
- Drake, A. J., Djorgovski, S. G., Mahabal, A., et al. 2009, *ApJ*, **696**, 870
- Drake, S. A., Simon, T., & Linsky, J. L. 1989, *ApJS*, **71**, 905
- Dubus, G., Cerutti, B., & Henri, G. 2010, *MNRAS*, **404**, L55
- Dunne, L., & Eales, S. A. 2001, *MNRAS*, **327**, 697
- Duric, N., Bourneuf, E., & Gregory, P. C. 1988, *AJ*, **96**, 81
- Eales, S. A., Wynn-Williams, C. G., & Duncan, W. D. 1989, *ApJ*, **339**, 859
- Eker, Z., Ak, N. F., Bilir, S., et al. 2008, *MNRAS*, **389**, 1722
- Ellingson, S. W., Clarke, T. E., Cohen, A., et al. 2009, *IEEE Proc.*, **97**, 1421
- Ellingson, S. W., & Hampson, G. A. 2002, *IEEE Transactions on Antennas & Propagation*, **50**, 25
- Emmanoulopoulos, D., McHardy, I. M., & Uttley, P. 2010, *MNRAS*, **404**, 931
- Ewen, H. I., & Purcell, E. M. 1951, *Nature*, **168**, 356
- Fanaroff, B. L., & Riley, J. M. 1974, *MNRAS*, **167**, 31P
- Favata, F., Micela, G., & Sciortino, S. 1995, *A&A*, **298**, 482
- Feldman, S. I. 1979, *Software: Practice and Experience*, **9**, 255
- Fender, R. P., & Bell, M. E. 2011, *BASI*, **39**, 315
- Fender, R. P., Belloni, T. M., & Gallo, E. 2004, *MNRAS*, **355**, 1105
- Fender, R. P., Burnell, S. J. B., Waltman, E. B., et al. 1997, *MNRAS*, **288**, 849
- Fender, R. P., Pooley, G. G., Brocksopp, C., & Newell, S. J. 1997, *MNRAS*, **290**, L65
- Fiedler, R., Dennison, B., Johnston, K. J., Waltman, E. B., & Simon, R. S. 1994, *ApJ*, **430**, 581
- Fiedler, R., Pauls, T., Johnston, K. J., & Dennison, B. 1994, *ApJ*, **430**, 595
- Fiedler, R. L., Dennison, B., Johnston, K. J., & Hewish, A. 1987, *Nature*, **326**, 675
- Fischler, M. A., & Bolles, R. C. 1981, *Communications of the ACM*, **24**, 381

- Fitt, A. J., Alexander, P., & Cox, M. J. 1988, [MNRAS](#), 233, 907
- Fragos, T., Kalogera, V., Willems, B., et al. 2009, [ApJL](#), 702, L143
- Frail, D. A., Kulkarni, S. R., Ofek, E. O., Bower, G. C., & Nakar, E. 2012, [ApJ](#), 747, 70
- Frail, D. A., Kulkarni, S. R., Sari, R., et al. 2001, [ApJL](#), 562, L55
- Fuentes, E., Miller, C. J., & Gasson, D. 2007, [ASP Conf. Ser.](#), 376, 579
- Gaensler, B. M., & Hunstead, R. W. 2000, [PASA](#), 17, 72
- Gal-Yam, A., Ofek, E. O., Poznanski, D., et al. 2006, [ApJ](#), 639, 331
- Gallo, E. 2010, [Lecture Notes in Physics](#), 794, 85
- Gallo, E., Miller, B. P., & Fender, R. 2012, [MNRAS](#), 423, 590
- Gardner, J. P., Mather, J. C., Clampin, M., et al. 2006, [Space Science Reviews](#), 123, 485
- Geldzahler, B. J., & Witzel, A. 1981, [AJ](#), 86, 1306
- Geldzahler, B. J., Johnston, K. J., Spencer, J. H., et al. 1983, [ApJL](#), 273, L65
- Genzel, R., & Townes, C. H. 1987, [ARA&A](#), 25, 377
- Ghosh, T., & Rau, A. P. 1992, [A&A](#), 264, 203
- Giacconi, R., Gorenstein, P., Gursky, H., & Waters, J. R. 1967, [ApJ](#), 148, L119
- Giovanelli, R., Haynes, M. P., Kent, B. R., et al. 2005, [AJ](#), 130, 2598
- Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, [ApJ](#), 622, 759
- Gray, A. D. 1994, [MNRAS](#), 270, 861
- Greenfield, P., & White, R. L. 2000, [ASP Conf. Ser.](#), 216, 59
- Gregory, P. C., & Kronberg, P. P. 1972, [Nature](#), 239, 440
- Gregory, P. C., Scott, W. K., Douglas, K., & Condon, J. J. 1996, [ApJS](#), 103, 427
- Gregory, P. C., & Taylor, A. R. 1986, [AJ](#), 92, 371
- Greisen, E. W. 2002, [Astrophysics and Space Science Library](#), 285, 109
- Griessmeier, J. M., Zarka, P., & Spreeuw, H. 2007, [A&A](#), 475, 359
- Güdel, M. 2002, [ARA&A](#), 40, 217
- Gull, S. F., & Daniell, G. J. 1978, [Nature](#), 272, 686
- Gutierrez-Kraybill, C., Keating, G. K., MacMahon, D., et al. 2010, [Proc. SPIE](#), 7740, 77400Z
- Hall, D. S. 1976, [Astrophysics and Space Science Library](#), 60, 287
- Hallinan, G., Bourke, S., Lane, C., et al. 2007, [ApJL](#), 663, L25
- Hamaker, J. P., & Bregman, J. D. 1996, [A&AS](#), 117, 161
- Hamaker, J. P., Bregman, J. D., & Sault, R. J. 1996, [A&AS](#), 117, 137
- Hancock, P. P., Gaensler, B. M., & Murphy, T. 2011, [ApJL](#), 735, L35
- Harp, G. R., Ackermann, R. F., Nadler, Z. J., et al. 2011, [IEEE Transactions on Antennas & Propagation](#), 59, 2004
- Harris, D. E., & Miley, G. K. 1978, [A&AS](#), 34, 117
- Harry, G. M., & LIGO Scientific Collaboration. 2010, [Classical and Quantum Gravity](#), 27, 084006
- Harwit, M. 1981, [Physics Today](#), 34, 172
- Harwit, M., & Pacini, F. 1975, [ApJ](#), 200, L127
- Hassan, A. H., Fluke, C. J., & Barnes, D. G. 2011, [New A](#), 16, 100
- Heiles, C. 2002, [ASP Conf. Ser.](#), 278, 131

- Helou, G., Soifer, B. T., & Rowan-Robinson, M. 1985, [ApJL](#), **298**, L7
- Hewish, A., Bell, S. J., Pilkington, J. D. H., Scott, P. F., & Collins, R. A. 1968, [Nature](#), **217**, 709
- Hildebrand, R. H., Whitcomb, S. E., Winston, R., et al. 1977, [ApJ](#), **216**, 698
- Hjellming, R. M., & Han, X. 1995, in “[X-Ray Binaries](#)”, ed. W. H. G. Lewin, J. van Paradijs, & E. P. J. van der Heuvel, ISBN 978-0521599344 (Cambridge: Cambridge University Press), 308
- Hjellming, R. M., & Narayan, R. 1986, [ApJ](#), **310**, 768
- Ho, P. T. P., Moran, J. M., & Lo, K. Y. 2004, [ApJL](#), **616**, L1
- Högbom, J. A. 1974, [A&AS](#), **15**, 417
- Hooper, D., & Goodenough, L. 2011, [Physics Letters B](#), **697**, 412
- Hopkins, A. M., Miller, C. J., Connolly, A. J., et al. 2002, [AJ](#), **1086**
- Hull, C. L. H., Bower, G. C., Croft, S., et al. 2010, [PASP](#), **122**, 1510
- Hummel, E., Davies, R. D., Pedlar, A., Wolstencroft, R. D., & van der Hulst, J. M. 1988, [A&A](#), **199**, 91
- Hummel, E., Smith, P., & van der Hulst, J. M. 1984, [A&A](#), **137**, 138
- Hunter, J. D. 2007, [Computing in Science and Engineering](#), **9**, 90
- Hyman, S. D., Lazio, T. J. W., Kassim, N. E., & Bartleson, A. L. 2002, [AJ](#), **123**, 1497
- Hyman, S. D., Lazio, T. J. W., Kassim, N. E., Nord, M. E., & Neureuther, J. L. 2003, [Astronomische Nachrichten Supplement](#), **324**, 79
- Hyman, S. D., Lazio, T. J. W., Kassim, N. E., et al. 2005, [Nature](#), **434**, 50
- Hyman, S. D., Wijnands, R., Lazio, T. J. W., et al. 2009, [ApJ](#), **696**, 280
- Imanishi, M., Nakanishi, K., Tamura, Y., Oi, N., & Kohno, K. 2007, [AJ](#), **134**, 2366
- Jackson, P. D., Kundu, M. R., & White, S. M. 1989, [A&A](#), **210**, 284
- Jansky, K. G. 1932, [Proc. of the Institute of Radio Engineering](#), **20**, 1920
- Jewitt, D., Luu, J., & Marsden, B. G. 1992, [IAU Circulars](#), #5611
- Johnston, S., Feain, I. J., & Gupta, N. 2009, [ASP Conf. Ser.](#), **407**, 446
- Jonas, J. L. 2009, [IEEE Proc.](#), **97**, 1522
- Jura, M., Hobbs, R. W., & Maran, S. P. 1978, [AJ](#), **83**, 153
- Kaiser, N., Aussel, H., Burke, B. E., et al. 2002, [Proc. SPIE](#), **4836**, 154
- Kassim, N. E., Lazio, T. J. W., Ray, P. S., et al. 2004, [Planetary & Space Science](#), **52**, 1343
- Kawai, N., Morii, M., Sugimori, K., et al. 2010, [The Astronomer’s Telegram](#), #2635
- Kazemi, S., Yatawatta, S., Zaroubi, S., et al. 2011, [MNRAS](#), **414**, 1656
- Keane, E. F., Kramer, M., Lyne, A. G., Stappers, B. W., & McLaughlin, M. A. 2011, [MNRAS](#), **415**, 3065
- Keating, G. K., Barott, W. C., & Wright, M. 2010, [Proc. SPIE](#), **7740**, 774016
- Keller, S. C., Schmidt, B. P., Bessell, M. S., et al. 2007, [PASA](#), **24**, 1
- Kellermann, K. I., Cornwell, T., DeBoer, D., et al. 2006, [SKA Memo Series](#), #84
- Kennicutt, R. C. 1998, [ApJ](#), **498**, 541
- Kettenis, M., van Langevelde, H. J., Reynolds, C., & Cotton, B. 2006, [ASP Conf. Ser.](#), **351**, 497

- Kida, S., Niinuma, K., Suzuki, S., et al. 2008, [New A](#), **13**, 519
- Klein, U., Wielebinski, R., & Morsi, H. W. 1988, [A&A](#), **190**, 41
- Kleinman, S. J., Gunn, J. E., Boroski, B., et al. 2008, [Proc. SPIE](#), **7016**, 70160B
- Klöckner, A., Pinto, N., Lee, Y., et al. 2012, [Parallel Computing](#), **38**, 157
- Koch, D. G., Borucki, W. J., Basri, G., et al. 2010, [ApJL](#), **713**, L79
- Koljonen, K. I. I., Hannikainen, D. C., McCollough, M. L., Pooley, G. G., & Trushkin, S. A. 2010, [MNRAS](#), **406**, 307
- Körding, E., Colbert, E., & Falcke, H. 2005, [A&A](#), **436**, 427
- Kotani, T., Kawai, N., Matsuoka, M., & Brinkmann, W. 1996, [PASJ](#), **48**, 619
- Kotani, T., Kawai, N., Aoki, T., et al. 1994, [PASJ](#), **46**, L147
- Kotani, T., Miyazaki, A., Tsuboi, M., et al. 2011, [The Astronomer's Telegram](#), #3130
- Kraus, A., Krichbaum, T. P., Wegner, R., et al. 2003, [A&A](#), **401**, 161
- Kronberg, P. P., Emerson, D. T., Klein, U., & Wielebinski, R. 1979, [ApJ](#), **230**, L149
- Kuniyoshi, M., Matsumura, N., Takefuji, K., et al. 2007, [PASP](#), **119**, 122
- Lammers, U., Beijersbergen, M., Thomas, M., & Vacanti, G. 2002, [ASP Conf. Ser.](#), **281**, 269
- Lattner, C., & Adve, V. 2004, in [CGO 2004: IEEE/ACM International Symposium on Code Generation and Optimization](#), 75
- Law, C. J., & Bower, G. C. 2012, [ApJ](#), **749**, 143
- Law, C. J., Brentjens, M. A., & Novak, G. 2011a, [ApJ](#), **731**, 36
- Law, C. J., & de Villiers, M. 2010, [ATA Memo Series](#), #87
- Law, C. J., Jones, G., Backer, D. C., et al. 2011b, [ApJ](#), **742**, 12
- Law, C. J., Yusef-Zadeh, F., & Cotton, W. D. 2008, [ApJS](#), **177**, 515
- Law, C. J., Gaensler, B. M., Bower, G. C., et al. 2011, [ApJ](#), **728**, 57
- Law, N. M., Kulkarni, S. R., Dekany, R. G., et al. 2009, [PASP](#), **121**, 1395
- Lazio, J., Bloom, J. S., Bower, G. C., et al. 2009, in “[Astro2010: The Astronomy and Astrophysics Decadal Survey](#)”, science white paper #176
- Lazio, T. J. W., Fey, A. L., Dennison, B., et al. 2000, [ApJ](#), **534**, 706
- Lebrun, F., Leray, J. P., Lavocat, P., et al. 2003, [A&A](#), **411**, L141
- Leeuw, L. L., & Robson, E. I. 2009, [AJ](#), **137**, 517
- Levine, A. M., Bradt, H., Cui, W., et al. 1996, [ApJL](#), **469**, L33
- Levinson, A., Ofek, E. O., Waxman, E., & Gal-Yam, A. 2002, [ApJ](#), **576**, 923
- Linden, T., Hooper, D., & Yusef-Zadeh, F. 2011, [ApJ](#), **741**, 95
- Ling, Z., Zhang, S. N., & Tang, S. 2009, [ApJ](#), **695**, 1111
- Lisenfeld, U., Voelk, H. J., & Xu, C. 1996, [A&A](#), **306**, 677
- Lister, M. L., Aller, M., Aller, H., et al. 2011, [ApJ](#), **742**, 27
- Longair, M. S. 1994, “[High energy astrophysics. Volume 2. Stars, the Galaxy and the interstellar medium](#)”, ISBN 978-0521434393 (Cambridge: Cambridge University Press)
- Lonsdale, C. J., Cappallo, R. J., Morales, M. F., et al. 2009, [IEEE Proc.](#), **97**, 1497
- Lorimer, D. R., Bailes, M., McLaughlin, M. A., Narkevic, D. J., & Crawford, F. 2007, [Science](#), **318**, 777
- Lovell, J. E. J., Rickett, B. J., J.-P. Macquart, et al. 2008, [ApJ](#), **689**, 108

- Lundgren, S. C., Cordes, J. M., Ulmer, M., et al. 1995, [ApJ](#), 453, 433
- Lutovinov, A., Revnivtsev, M., Gilfanov, M., & Sunyaev, R. 2007, in [The Obscured Universe. Proceedings of the VI INTEGRAL Workshop](#), ed. S. Grebenev, R. Sunyaev, & C. Winkler, 241
- Machado, P., Vincenzi, A., & Maldonado, J. C. 2010, [Lecture Notes in Computer Science](#), 6153, 1
- Magee, D. K., Bouwens, R. J., & Illingworth, G. D. 2007, [ASP Conf. Ser.](#), 376, 261
- Magro, A., Karastergiou, A., Salvini, S., et al. 2011, [MNRAS](#), 417, 2642
- Marscher, A. P., & Brown, R. L. 1975, [ApJ](#), 200, 719
- Marsh, K. A., & Helou, G. 1995, [ApJ](#), 445, 599
- Matsumura, N., Daishido, T., Kuniyoshi, M., et al. 2007, [AJ](#), 133, 1441
- Matsumura, N., Niinuma, K., Kuniyoshi, M., et al. 2009, [AJ](#), 138, 787
- Mayor, M., & Queloz, D. 1995, [Nature](#), 378, 355
- McDonald, A. R., Muxlow, T. W. B., Wills, K. A., Pedlar, A., & Beswick, R. J. 2002, [MNRAS](#), 334, 912
- McLaughlin, M. A., Lyne, A. G., Lorimer, D. R., et al. 2006, [Nature](#), 439, 817
- McMullin, J. P., Golap, K., & Myers, S. T. 2004, [ASP Conf. Ser.](#), 314, 468
- McMullin, J. P., Waters, B., Schiebel, D., Young, W., & Golap, K. 2007, [ASP Conf. Ser.](#), 376, 127
- Mehring, D. M., & Plante, R. 2004, [ASP. Conf. Ser.](#), 314, 42
- Metzger, B. D., & Berger, E. 2012, [ApJ](#), 746, 48
- Migliari, S., Fender, R., & Méndez, M. 2002, [Science](#), 297, 1673
- Migliari, S., & Fender, R. P. 2006, [MNRAS](#), 366, 79
- Mirabel, I. F., & Rodríguez, L. F. 1998, [Nature](#), 392, 673
- . 1999, [ARA&A](#), 37, 409
- Morris, M., & Serabyn, E. 1996, [ARA&A](#), 34, 645
- Murphy, T., Chatterjee, S., Kaplan, D. L., et al. 2012, PASA, in press, [arxiv:1207.1528](#)
- Muxlow, T. W. B., Beswick, R. J., Garrington, S. T., et al. 2010, [MNRAS](#), 404, L109
- Myers, C. R., Gutenkunst, R. N., & Sethna, J. P. 2007, [Computing in Science & Engineering](#), 9, 34
- Nakar, E., & Piran, T. 2011, [Nature](#), 478, 82
- Narayan, R. 1992, [Royal Society of London Philosophical Transactions Series A](#), 341, 151
- Niinuma, K., Asuma, K., Kuniyoshi, M., et al. 2007, [ApJL](#), 657, L37
- Niklas, S., Klein, U., & Wielebinski, R. 1997, [A&A](#), 322, 19
- Nishiyama, S., Hatano, H., Tamura, M., et al. 2010, [ApJL](#), 722, L23
- Noordam, J. E., & Smirnov, O. M. 2010, [A&A](#), 524, A61
- Nord, M. E., Henning, P. A., Rand, R. J., Lazio, T. J. W., & Kassim, N. E. 2006, [AJ](#), 132, 242
- Norris, R. P. 1988, [MNRAS](#), 230, 345
- O’Dea, C. P., & Owen, F. N. 1985, [AJ](#), 90, 927
- Ofek, E. O., & Frail, D. A. 2011, [ApJ](#), 737, 45

- Ofek, E. O., Frail, D. A., Breslauer, B., et al. 2011, [ApJ](#), 740, 65
- Offringa, A. R., van de Gronde, J. J., & Roerdink, J. B. T. M. 2012, [A&A](#), 539, A95
- O'Mahony, S. 2003, [Research Policy](#), 32, 1179
- Orellana, M., Bordas, P., Bosch-Ramon, V., Romero, G. E., & Paredes, J. M. 2007, [A&A](#), 476, 9
- Osten, R. A. 2008, in [Bursts, Pulses and Flickering: wide-field monitoring of the dynamic radio sky](#), [Proceedings of Science article PoS\(Dynamic2007\)005](#)
- Ott, J., Weiss, A., Henkel, C., & Walter, F. 2005, [ApJ](#), 629, 767
- Owen, F. N., & Gibson, D. M. 1978, [AJ](#), 83, 1488
- Padovani, P., Giommi, P., Landt, H., & Perlman, E. S. 2007, [ApJ](#), 662, 182
- Paladino, R., Murgia, M., Helfer, T. T., et al. 2006, [A&A](#), 456, 847
- Parker, E. N. 1966, [ApJ](#), 145, 811
- Parsons, A. R., & Backer, D. C. 2009, [AJ](#), 138, 219
- Parsons, A. R., Backer, D. C., Foster, G. S., et al. 2010, [AJ](#), 139, 1468
- Pen, U. L., & King, L. 2012, [MNRAS](#), 421, L132
- Peng, R., Zhou, S., Whiteoak, J. B., Lo, K. Y., & Sutton, E. C. 1996, [ApJ](#), 470, 821
- Penzias, A. A., & Wilson, R. W. 1965, [ApJ](#), 142, 419
- Pérez, F., & Granger, B. E. 2007, [Computing in Science & Engineering](#), 9, 21
- Pérez, F., Granger, B. E., & Hunter, J. D. 2011, [Computing in Science & Engineering](#), 13, 13
- Perley, R. A. 1999, [ASP Conf. Ser.](#), 180, 383
- Perley, R. A., Chandler, C. J., Butler, B. J., & Wrobel, J. M. 2011, [ApJL](#), 739, L1
- Perlmutter, S., Aldering, G., Goldhaber, G., et al. 1999, [ApJ](#), 517, 565
- Petersen, K. 2011, [Information and Software Technology](#), 53, 317
- Peterson, P. 2009, [International J of Computational Science & Engineering](#), 4, 296
- Pooley, G. G., & Fender, R. P. 1997, [MNRAS](#), 292, 925
- Pound, M. W., & Teuben, P. 2012, [arXiv:1202.1030](#)
- Prat, L., Rodriguez, J., & Pooley, G. G. 2010, [ApJ](#), 717, 1222
- Predehl, P., Burwitz, V., Paerels, F., & Trümper, J. 2000, [A&A](#), 357, L25
- Price, R., & Duric, N. 1992, [ApJ](#), 401, 81
- Protheroe, R. J., Ott, J., Ekers, R. D., Jones, D. I., & Crocker, R. M. 2008, [MNRAS](#), 390, 683
- Purcell, C. R., Hoare, M. G., & Diamond, P. 2008, [ASP Conf. Ser.](#), 387, 389
- Radford, S. J. E., Delannoy, J., Downes, D., et al. 1991, in [IAU Symposium #146: Dynamics of Galaxies and Their Molecular Cloud Distributions](#), ed. F. Combes & F. Casoli, 303
- Rau, U., & Cornwell, T. J. 2011, [A&A](#), 532, A71
- Raymond, E. 1999, [Knowledge, Technology & Policy](#), 12, 23
- Reber, G. 1940, [ApJ](#), 91, 621
- Rees, M. J. 1988, [Nature](#), 333, 523
- Rengelink, R. B., Tang, Y., de Bruyn, A. G., et al. 1997, [A&AS](#), 124, 259
- Ricci, R., Prandoni, I., Gruppioni, C., Sault, R. J., & de Zotti, G. 2006, [A&A](#), 445, 465
- Richards, M. T., Waltman, E. B., Ghigo, F. D., & Richards, D. S. 2003, [ApJS](#), 147, 337

- Rickett, B. J. 1990, [ARA&A](#), **28**, 561
- Rickett, B. J., Coles, W. A., & Markkanen, J. 2000, [ApJ](#), **533**, 304
- Rieke, G. H., Harper, D. A., Low, F. J., & Armstrong, K. R. 1973, [ApJ](#), **183**, L67
- Riess, A. G., Filippenko, A. V., Challis, P., et al. 1998, [AJ](#), **116**, 1009
- Rigopoulou, D., Lawrence, A., & Rowan-Robinson, M. 1996, [MNRAS](#), **278**, 1049
- Rodríguez-Rico, C. A., Goss, W. M., Viallefond, F., et al. 2005, [ApJ](#), **633**, 198
- Rodríguez-Rico, C. A., Viallefond, F., Zhao, J. H., Goss, W. M., & Anantharamaiah, K. R. 2004, [ApJ](#), **616**, 783
- Romero, G. E., Torres, D. F., Kaufman Bernadó, M. M., & Mirabel, I. F. 2003, [A&A](#), **410**, L1
- Romero, G. E., Valle, M. V. D., & Orellana, M. 2010, [A&A](#), **518**, A12
- Romero, G. E., & Vila, G. S. 2008, [A&A](#), **485**, 623
- Ryle, M., & Vonberg, D. D. 1946, [Nature](#), **158**, 339
- Sánchez-Sutil, J. R., Martí, J., Combi, J. A., et al. 2008, [A&A](#), **479**, 523
- Sanner, M. F. 1999, [Journal of Molecular Graphics and Modelling](#), **17**, 57
- Sault, R. J., & Conway, J. E. 1999, [ASP Conf. Ser.](#), **180**, 419
- Sault, R. J., Hamaker, J. P., & Bregman, J. D. 1996, [A&AS](#), **117**, 149
- Sault, R. J., Teuben, P. J., & Wright, M. C. H. 1995, [ASP Conf. Ser.](#), **77**, 433
- Sault, R. J., & Wieringa, M. H. 1994, [A&AS](#), **108**, 585
- Scheifler, R. W., & Gettys, J. 1986, [ACM Transactions on Graphics](#), **5**, 79
- Schwab, F. R. 1984, [AJ](#), **89**, 1076
- Schwager, M. 2008, “Schweizer Organisation des Bahnnetzes als Vorbild”, <http://www.scritti.de/text/bahn2000.html>
- Scoville, N. Z., Sargent, A. I., Sanders, D. B., & Soifer, B. T. 1991, [ApJ](#), **366**, L5
- Scoville, N. Z., Soifer, B. T., Neugebauer, G., et al. 1985, [ApJ](#), **289**, 129
- Scoville, N. Z., Yun, M. S., & Bryant, P. M. 1997, [ApJ](#), **484**, 702
- Seaquist, E. R., Bell, M. B., & Bignell, R. C. 1985, [ApJ](#), **294**, 546
- Seaquist, E. R., Carlstrom, J. E., Bryant, P. M., & Bell, M. B. 1996, [ApJ](#), **465**, 691
- Setia Gunawan, D. Y. A., de Bruyn, A. G., van der Hucht, K. A., & Williams, P. M. 2003, [ApJS](#), **149**, 123
- Shull, F., Basili, V., Boehm, B., et al. 2002, in [Proceedings of the Eighth IEEE Symposium on Software Metrics](#), 249
- Simonetti, J. H., Cordes, J. M., & Heeschen, D. S. 1985, [ApJ](#), **296**, 46
- Singh, N. S., Naik, S., Paul, B., et al. 2002, [A&A](#), **392**, 161
- Smirnov, O. M. 2011a, [A&A](#), **527**, A106
- . 2011b, [A&A](#), **527**, A107
- . 2011c, [A&A](#), **527**, A108
- . 2011d, [A&A](#), **531**, A159
- Sopp, H. M., & Alexander, P. 1991, [MNRAS](#), **251**, 112
- Spangler, S., Fanti, R., Gregorini, L., & Padrielli, L. 1989, [A&A](#), **209**, 315
- Stark, M. J., & Saia, M. 2003, [ApJL](#), **587**, L101
- Steer, D. G., Dewdney, P. E., & Ito, M. R. 1984, [A&A](#), **137**, 159

- Straten, W. V., & Bailes, M. 2011, [PASA](#), **28**, 1
- Strickland, D. K., Heckman, T. M., Colbert, E. J. M., Hoopes, C. G., & Weaver, K. A. 2004, [ApJ](#), **606**, 829
- Takano, S., Hofner, P., Winnewisser, G., Nakai, N., & Kawaguchi, K. 2005, [PASJ](#), **57**, 549
- Tavani, M., Bulgarelli, A., Piano, G., et al. 2009, [Nature](#), **462**, 620
- Taylor, A. R., Goss, W. M., Coleman, P. H., Leeuwen, J. V., & Wallace, B. J. 1996, [ApJS](#), **107**, 239
- Taylor, A. R., & Gregory, P. C. 1983, [AJ](#), **88**, 1784
- Taylor, G. B., Carilli, C. L., & Perley, R. A., eds. 1999, “[Synthesis Imaging in Radio Astronomy II](#)” (ASP Conf. Ser. vol. 180), ISBN 978-1583815168 (San Francisco: Astronomical Society of the Pacific)
- Teuben, P. 2011, [ASP Conf. Ser.](#), **442**, 533
- Thakur, R., Gropp, W., & Lusk, E. 1999, in [Proceedings of Frontiers ’99: Seventh Symposium on the Frontiers of Massively Parallel Computation](#), 182
- Thompson, A. R. 1999, [ASP Conf. Ser.](#), **180**, 11
- Thompson, A. R., Moran, J. M., & Swenson, G. W. 2001, “[Interferometry and Synthesis in Radio Astronomy](#)”, 2nd edn., ISBN 978-0471254928 (New York: Wiley-VCH)
- Thompson, T. A., Quataert, E., & Murray, N. 2009, [MNRAS](#), **397**, 1410
- Thompson, T. A., Quataert, E., Waxman, E., Murray, N., & Martin, C. L. 2006, [ApJ](#), **645**, 186
- Thyagarajan, N., Helfand, D. J., White, R. L., & Becker, R. H. 2011, [ApJ](#), **742**, 49
- . 2012, [ApJ](#), **746**, 114
- Tomsick, J. A., Tramacere, A., Ferrigno, C., et al. 2010, [The Astronomer’s Telegram](#), #2644
- Trushkin, S. A. 2000, [Astronomical and Astrophysical Transactions](#), **19**, 525
- Trushkin, S. A., Bursov, N. N., Nizhelskij, N. A., Majorova, E. K., & Voitsik, P. A. 2006, in [VI Microquasar Workshop: Microquasars and Beyond](#), ed. T. Belloni, [Proceedings of Science article PoS\(MQW6\)015](#)
- Tsutsumi, T., Gregory, P. C., Duric, N., & Taylor, A. R. 1995, [AJ](#), **110**, 238
- Tudose, V., Fender, R. P., Garrett, M. A., et al. 2007, [MNRAS](#), **375**, L11
- Tyson, J. A. 2002, [Proc. SPIE](#), **4836**, 10
- van der Kruit, P. C. 1971, [A&A](#), **15**, 110
- van der Laan, H. 1966, [Nature](#), **211**, 1131
- van Kerkwijk, M. H., Geballe, T. R., King, D. L., van der Klis, M., & van Paradijs, J. 1996, [A&A](#), **314**, 521
- Verheijen, M. A. W., Oosterloo, T. A., van Cappellen, W. A., et al. 2008, [AIP Conf. Ser.](#), **1035**, 265
- Vila, G. S., & Romero, G. E. 2010, [MNRAS](#), **403**, 1457
- Völk, H. J. 1989, [A&A](#), **218**, 67
- Walker, D. W. 1994, [Parallel Computing](#), **20**, 657
- Walker, M. A. 2006, [ASP Conf. Ser.](#), **365**, 299
- Walterbos, R. A. M., & Schwing, P. B. W. 1987, [A&A](#), **180**, 27

- Waltman, E. B., Fiedler, R. L., Johnston, K. L., & Ghigo, F. D. 1994, [AJ](#), **108**, 179
- Waltman, E. B., Ghigo, F. D., Johnston, K. J., et al. 1995, [AJ](#), **110**, 290
- Wayth, R. B., Briske, W. F., Deller, A. T., et al. 2011, [ApJ](#), **735**, 97
- Weaver, H., Williams, D. R. W., Dieter, N. H., & Lum, W. T. 1965, [Nature](#), **208**, 29
- Weiler, K. W., Panagia, N., Montes, M. J., & Sramek, R. A. 2002, [ARA&A](#), **40**, 387
- Weinreb, S., Barrett, A. H., Meeks, M. L., & Henry, J. C. 1963, [Nature](#), **200**, 829
- Weiss, A., Kovács, A., Güsten, R., et al. 2008, [A&A](#), **490**, 77
- Welch, J., Backer, D., Blitz, L., et al. 2009, [IEEE Proc.](#), **97**, 1438
- Welch, W. J., Thornton, D. D., Plambeck, R. L., et al. 1996, [PASP](#), **108**, 93
- Wendker, H. J., Higgs, L. A., & Landecker, T. L. 1991, [A&A](#), **241**, 551
- White, R. L., Becker, R. H., & Helfand, D. J. 2005, [AJ](#), **130**, 586
- Whiting, M. T. 2012, MNRAS, in press, [arXiv:1201.2710](#)
- Williams, P. 2010, in [RFI Mitigation Workshop](#), Proceedings of Science article [PoS\(RFI2010\)004](#)
- Wilms, J., Pottschmidt, K., Pooley, G. G., et al. 2007, [ApJL](#), **663**, L97
- Wilson, R. W., Jefferts, K. B., & Penzias, A. A. 1970, [ApJL](#), **161**, L43
- Winkler, C., Courvoisier, T., Cocco, G. D., et al. 2003, [A&A](#), **411**, L1
- Woody, D. P., Scott, S. L., Scoville, N. Z., et al. 1989, [ApJ](#), **337**, L41
- Woody, D. P., Beasley, A. J., Bolatto, A. D., et al. 2004, [Proc. SPIE](#), **5498**, 30
- Wright, A., & Otrupcek, R., eds. 1990, “[Parkes Catalog](#)” (Australia Telescope National Facility)
- Wright, E. L., Eisenhardt, P. R. M., Mainzer, A. K., et al. 2010, [AJ](#), **140**, 1868
- Wunderlich, E., & Klein, U. 1988, [A&A](#), **206**, 47
- Yun, M. S., Reddy, N. A., & Condon, J. J. 2001, [ApJ](#), **554**, 803
- Yusef-Zadeh, F., Morris, M., & Chance, D. 1984, [Nature](#), **310**, 557
- Zdziarski, A. A., Sikora, M., Dubus, G., et al. 2012, [MNRAS](#), **421**, 2956
- Zhao, J. H., Anantharamaiah, K. R., Goss, W. M., & Viallefond, F. 1996, [ApJ](#), **472**, 54
- Zhao, J.-H., Roberts, D. A., Goss, W. M., et al. 1992, [Science](#), **255**, 1538
- Zwart, J. T. L., Barker, R. W., Biddulph, P., et al. 2008, [MNRAS](#), **391**, 1545