A Matched Filter Technique for Slow Radio Transient Detection and First Demonstration with the Murchison Widefield Array

by

Lu Feng

B.S., The Pennsylvania State University (2009)

Submitted to the Department of Physics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Physics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2016

Author ..... Signature redacted
Department of Physics
May 20, 2016

Certified by ..... Signature redacted
Jacqueline N. Hewitt
Professor of Physics
Thesis Supervisor

Accepted by ..... Signature redacted
Nergis Mavalvala
Professor of Physics
Associate Department Head for Education

A Matched Filter Technique for Slow Radio Transient Detection and First Demonstration with the Murchison Widefield Array

by

Lu Feng

Submitted to the Department of Physics
on May 20, 2016, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Physics

Abstract

Many astronomical sources produce transient phenomena at radio frequencies, but the transient sky at low frequencies (< 300 MHz) remain relatively unexplored. Blind surveys with new widefield radio instruments are filling this gap. Although many of these instruments are limited by the classical confusion noise, one can in principle detect transients below the classical confusion limit. This thesis develops a technique for detecting radio transients that is based on temporal matched filters applied directly to time series of images. This technique has well-defined statistical properties and is applicable to variable and transient searches for any instrument. Using the Murchison Widefield Array as an example, we demonstrate that the technique works well on real data despite the presence of classical confusion noise, sidelobe confusion noise, and other systematic errors. We search for transients lasting between 2 minutes and 3 months and improve the upper limits on the transient surface density at 182 MHz for fluxes between $\sim$ 20–200 mJy. We use this technique to characterize detectability of radio afterglows from compact binary coalescence, which are predicted electromagnetic counterparts of gravitational wave (GW) sources and the most promising progenitor of short gamma-ray bursts. While the next generation of GW detectors have come online and detected the first GW event, their ability to localize these events will remain poor during the early days of their operation. Many new widefield radio instruments will be able to cover large areas of the sky in a short amount of time. We use simulated afterglow light curves to estimate the rates of detection for different radio instruments under ideal conditions. We find that some widefield radio instruments might be able to detect radio afterglows and constrain their properties.

Thesis Supervisor: Jacqueline N. Hewitt
Title: Professor of Physics

We looked up at the stars that year

With so many brilliant dreams...

Starry Sky, Mayday

Acknowledgments

It has been a long, strange trip. But all in all, this was a really cool thesis project.

I want to thank Jackie Hewitt, my thesis adviser, for giving me this opportunity. I still remember her enthusiasm in the first email she replied to me when I asked about her research and her group at the end of my second year, and I also remember her warmth when we met for the first time. That had been uplifting for me during a challenging time during graduate school. It was not the last challenge I faced, but Jackie had been patient and supportive throughout. She has also kept me grounded in research and pushed me forward to become a better scientist.

I want to thank Erik Katsavounidis and Scott Hughes for serving on my thesis committee. Perhaps no other student has had as many thesis committee meetings as I did, but they had been incredibly understanding and supportive, taking time from their busy schedules to see this project through and providing valuable feedback that improved this thesis. Erik had also invited me to a few of his group meetings when I first started on this project, teaching me about the opportunities available in the electromagnetic follow-up of gravitational wave events. In my last semester of graduate school, I also had the honor of being Scott's teaching assistant; the few 8.901 lectures I attended were great, and his excitement for GW150914 contagious.

I am extremely grateful to Ruslan Vaulin for introducing me to this project and for guiding it all the way. I learned so much from our weekly discussions, whether it was about physics, statistics, or career paths. He had been very kind and encouraging, especially during the most difficult time I experienced in graduate school. This thesis would not be here if it were not for him, either its conception or completion. He and Tamara had also invited me to their wonderful Thanksgiving dinners over the past few years, making Thanksgiving away from home much livelier.

I want to thank Nergis Mavalvala for being my academic adviser. We met mainly at the beginning of each semester, but it was always a pleasure to chat with her. She helped me gain perspective on my research progress, and offered immense support when I was looking to switch research groups.

I want to thank Ron Remillard for working with me on the MWA data reduction process. He taught me much about problem solving while he helped me refine the CASA pipeline and correct the issues with the primary beam. The data I used would not have been reduced as quickly as they had if it were not for his persistent efforts.

I want to thank Ed Morgan for his support on the MWA computing resources and for his guidance on one of the first MWA projects I worked on. It gave me a better understanding of the MWA tiles and a way to become acquainted with the rest of the collaboration when I presented the results at the first project meeting I attended.

I want to thank Al Levine for not just his advice on MWA work but also his incredibly helpful and indispensable Part III study sessions. He also taught me a lot about radio interferometry and synthesis imaging after I joined the group.

I want to thank Chris Williams for teaching me about MWA and reducing data in CASA. I also want to thank Aaron Ewall-Wice and Abraham Neben for the informative discussions on MWA data reduction as we wrestled with calibration and imaging issues. Together with Chris and Aaron and Dave Emrich, I had a great and adventurous trip during X16 to the MWA site, where the Halleen family provided warm hospitality at the Boolardy station. I would like to acknowledge and thank the Wajarri Yamatji people, who are the traditional owners of the observatory site.

I would also like to thank the MWA collaboration as this thesis would not have been possible without them, in particular David Kaplan, Tara Murphy, Martin Bell, Natasha Hurley-Walker, Danny Jacobs, Randall Wayth, André Offringa, Miguel Morales, Patti Carroll, Ian Sullivan, Adam Beardsley, Brian Crosse, and Bob Goeke.

I want to thank my teachers at MIT: Josh Winn, who taught me 8.901 and offered insight on research and alternate career paths; Rob Simcoe, who taught me 8.902; Max Tegmark, who taught me cosmology and chaired my Part III committee; Anna Frebel and Nevin Weinberg, who served on my Part III committee.

I want to thank my peers: Luke Bouma, Keaton Burns, Ani Chiti, Tom Cooper, Fei Dai, Kat Deck, Josh Dillon, Greg Dooley, Reed Essick, Aaron Ewall-Wice, Sherry Guo, David Hernandez, Alex Ji, Ryan Lynch, Ryan McKinnon, Abraham Neben, Uchupol Ruangsri, Champ Somboonpanyakul, Hang Yu, Liang Yu, and Jeff Zheng.

I am really grateful to have met Mingming Yang and Ying Zhao and to have their continuous support. They have become my closest friends from MIT, even though we were separated by six hours of flight time more often than we were at MIT together. No matter where we go in the future, I believe we will continue to support each other.

I am also really grateful to have met Arlene Rieneke and Kim Chi Vu. They are my best friends and most loyal cheerleaders despite being on the other side of the world in the ten, eleven years we have known each other. Their steady emotional support helped me live through the toughest times in college and graduate school, and I have no idea where I would be without them.

Someone once said that sometimes a stranger at a coffee shop could influence your life more than the people you had known for thirty years. So I want to acknowledge two strangers for giving me strength where no one else could: to be honest with myself; to know what truly matters in life; to move forward and always forward, no matter how small the step and how big the obstacle; and to understand that sometimes the things you start on a whim can become more important to you than the things you do with a thousand sound reasons, and that sometimes simple curiosity can propel you further than immense passion.

Finally, I am forever grateful to my family. My sister is a longtime friend who shares many of my interests and brings in fresh perspective. My parents are the people I have known the longest, and who know me the best. Their unconditional love and support is something that no one else can compare.

Mom and Dad, this thesis is for you.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Radio Transients

Everything in the universe evolves: planets, stars, galaxies. The process is at times gradual and at times violent, and reveals itself as changes in the electromagnetic radiation, from gamma rays at the shortest wavelengths to radio waves at the longest. When these changes appear and then disappear, we call them "transients."

Perhaps the earliest known transients were the "guest stars" recorded in East Asian history since 2nd century B.C., now thought to be novae (accreting white dwarfs) and supernovae (exploding stars) [1]. With the advent of telescopes and new technology, more classes of transients have been discovered and studied; for example, radio pulsars (rapidly rotating neutron stars) [2], X-ray binaries (accreting neutron stars or black holes) [3], and gamma-ray bursts (the collapse of massive stars or the coalescence of binary neutron stars) [4]. There is much room for exploration even now, as evident from the recent discovery of fast radio bursts whose origin remains a mystery [5].

Our knowledge of the transient sky has grown tremendously because of surveys that have large and recurrent sky coverage. Dedicated projects, such as the Palomar Transient Factory [6], the Rossi X-ray Timing Explorer [7], and the Swift Gamma-Ray Burst Mission [8], have pushed the frontiers of time domain astronomy at optical wavelengths or shorter. By contrast, surveys at radio wavelengths have been limited, rarely achieving wide fields of view and high time resolution simultaneously [9].

Nevertheless, there are many known and predicted populations of radio transients, which allow us to study a variety of physical and astrophysical processes ranging from the mechanisms of particle acceleration to the cosmological star formation history.

Extrasolar planets with magnetic fields could produce bursty cyclotron radio emission similar to that of the giant planets in the Solar System [10, 11, 12, 13]. Detecting this emission would not only constitute a direct detection of the extrasolar planet, but also offer a way to measure the planetary magnetic field and the orbital parameters [14]. So far none has been detected [15, 16, 17].

Late-type stars, from M dwarfs to brown dwarfs, often generate radio flares due to their magnetic activities. This has been observed at high radio frequencies (1–10 GHz) [18, 19, 20], and contains information about the generation of stellar magnetic fields and the structure of stellar magnetospheres [21].

Black hole accretion, whether in X-ray binaries or as tidal disruption events at the centers of galaxies, also produces radio flares. These arise from jets that create shocks in the surrounding medium and are correlated with X-ray flares caused by instabilities in the accretion disk [22, 23, 24]. The correlation between radio and X-ray emissions reveals the connection between the jet and the accretion disk, the details of which we do not yet fully understand.

Supernovae and gamma-ray bursts create shocks in their surrounding medium and produce synchrotron radiation across the entire electromagnetic spectrum. Radio observations, in particular, would constrain the energetics of the explosion and the properties of the circumburst environment [25, 26]. They could also uncover a hidden population not detectable at other wavelengths, i.e. supernovae obscured by dust or orphan afterglows, and hence better constrain their rates [27, 28, 29, 30].

Finally, although this was not an exhaustive list (see [9, 31]), gravitational wave (GW) sources, such as binary neutron star (BNS) mergers, are expected to generate radio flares and afterglows [32, 33]. Electromagnetic observations of these sources will break the degeneracies in GW measurements, such as the one between source distance and inclination angle, and test the BNS progenitor model of short gamma-ray bursts. This is especially timely given the detection of the first GW event [34, 35].

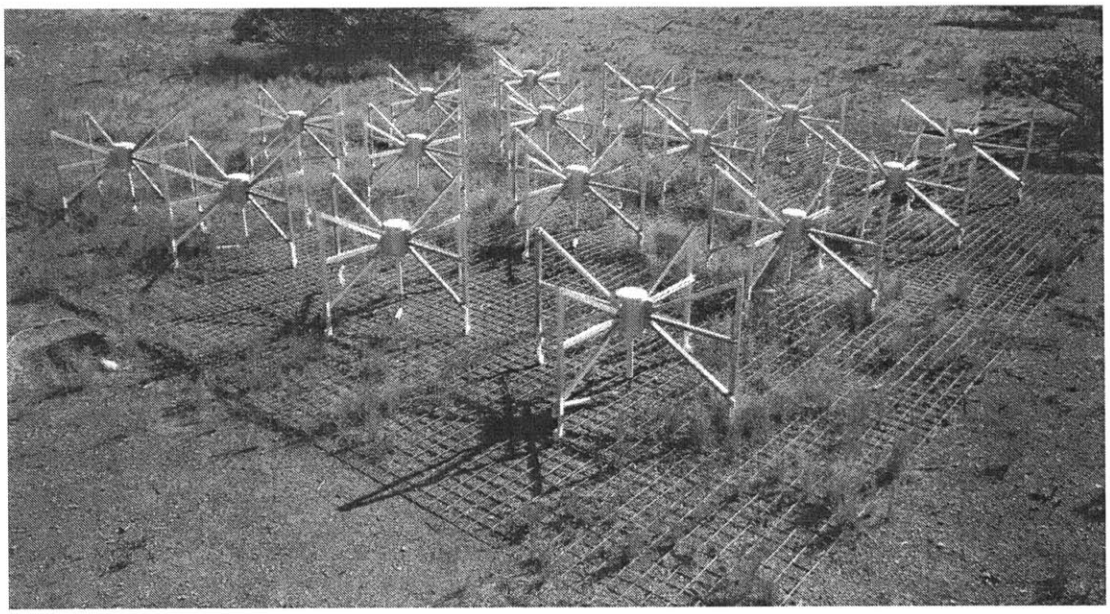

Figure 1-1: An antenna or "tile" of the Murchison Widefield Array, an aperture array operating at low radio frequencies (80–300 MHz). Each tile consists of 16 dual-polarization dipoles. The full array has 128 such tiles.

1.2 Radio Interferometry and Synthesis Imaging

The advancement of signal processing and digital electronics has introduced a new generation of widefield radio interferometers [36] that are bridging the gap between radio and the rest of the electromagnetic spectrum in the efforts to study the transient sky. Many of these interferometers are aperture arrays—arrays of connected dipoles—operating at meter wavelengths, i.e. low frequencies ($\lesssim$ 300 MHz), where there have been few blind, unbiased surveys previously. Here we describe how an aperture array operates and how one obtains astronomical images from radio interferometric measurements.

Figure 1-1 shows one antenna of such an array. Each antenna comprises several dipoles, each of which receives radiation from the sky in all directions. A beamformer combines the voltage signals from the different dipoles after applying appropriate time delays to align the phases of the signals coming from a particular direction; this forms the radiation pattern or "primary beam" of the antenna. The signal of each antenna then passes through a bandpass filter in the receiver, which determines the spectral

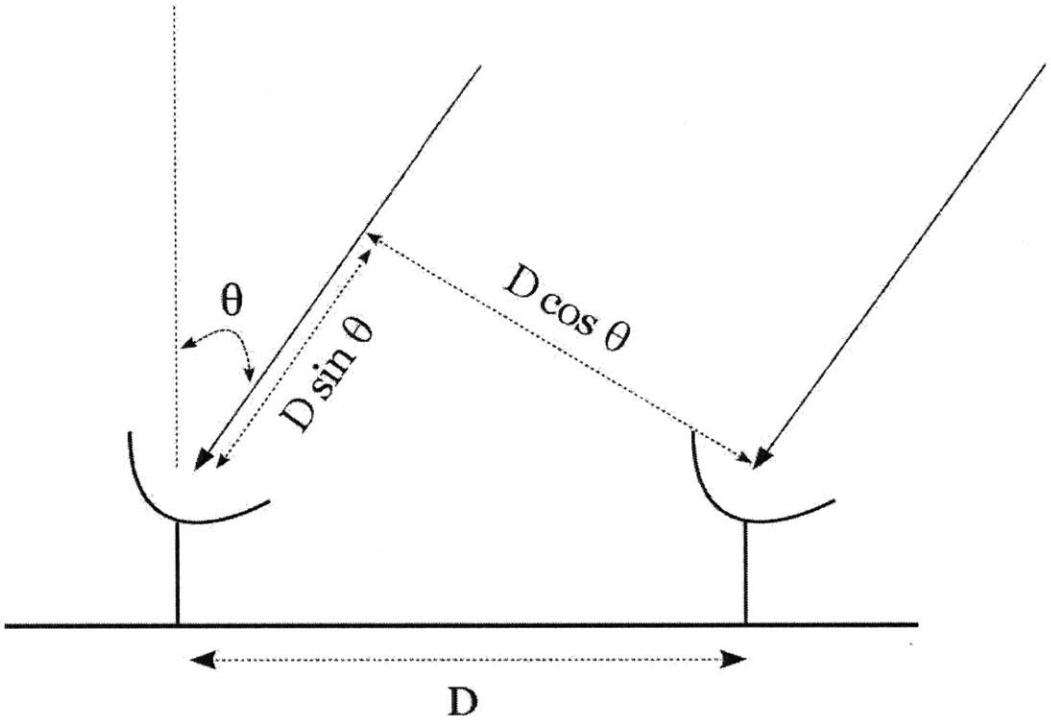

Figure 1-2: Illustration of a simple interferometer. D is the distance between two antennas, i.e. the baseline of the interferometer. θ is the angle at which the radiation is incident on each antenna.

or frequency range of the signal. The signals from pairs of antennas are digitized and correlated, representing the Fourier mode that corresponds to the antenna separation.

The correlated output, a complex quantity known as the "visibility," measures the intensity or brightness of the radiating source:

$$V_\nu(\mathbf{r}_1, \mathbf{r}_2) = \langle \mathbf{E}_\nu(\mathbf{r}_1)\mathbf{E}^*_\nu(\mathbf{r}_2)\rangle = \int A_\nu(\mathbf{s})I_\nu(\mathbf{s})e^{-2\pi i\nu\mathbf{s}\cdot(\mathbf{r}_1-\mathbf{r}_2)/c}\,\mathrm{d}\Omega \tag{1.1}$$

Here, V_ν is the visibility at frequency ν, $\mathbf{r}_1$ and $\mathbf{r}_2$ are the locations of the two antennas, $\mathbf{E}_\nu$ is the electric field of the source as measured by each antenna, A_ν is the normalized antenna pattern, I_ν is the intensity of the source, $\mathbf{s}$ is the position of the source on the celestial sphere, c is the speed of light, and Ω is the solid angle in the sky. Figure 1-2 illustrates an interferometer with two antennas, and one can see that the position of the source determines the time difference between when the same signal arrives at

one antenna and the next, i.e. the geometric delay $\tau_g = D\sin\theta/c$ where $D = |\mathbf{r}_1 - \mathbf{r}_2|$ is the length of the baseline and θ is the zenith angle of the source. This delay affects the interference fringe pattern and needs to be applied to the antenna receiving the earlier signal to ensure that the same signals are correlated.

One can rewrite Equation 1.1 in a coordinate system suitable for imaging:

$$V(u,v,w) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} A(l,m)I(l,m)e^{-2\pi i[ul+vm+w(\sqrt{1-l^2-m^2}-1)]}\frac{\mathrm{d}l\,\mathrm{d}m}{\sqrt{1-l^2-m^2}} \quad (1.2)$$

where u is the eastward component of the baseline vector, v is the northward component, w is the component toward the phase center (the center of the image), all measured in units of wavelengths, while l and m are the direction cosines measured with respect to the u- and v-axes on the celestial sphere. If (l,m) are small such that $\sqrt{1-l^2-m^2} \approx 1$, Equation 1.2 becomes a 2-dimensional Fourier transform and can, in principle, be inverted to recover $I(l,m)$. However, an interferometer has a finite and discrete number of baselines described by the sampling function $S(u,v) = \sum_{k=1}^{M}\delta(u-u_k, v-v_k)$. One can increase the number of (u,v) points sampled by using the rotation of the Earth, which changes the projection of each baseline relative to the source position and hence (u,v), but the sampling, also known as the uv-coverage, will never be complete. This impacts the inversion process.

Inverting the sampled and calibrated visibilities, one obtains the "dirty" image:

$$I^D(l,m) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} S(u,v)V(u,v)e^{2\pi i(ul+vm)}\,\mathrm{d}u\,\mathrm{d}v \quad (1.3)$$

In practice, this is usually done by interpolating the data on a rectangular grid and then running a fast Fourier transform. One can rewrite Equation 1.3 according to the convolution theorem:

$$I^D = \mathcal{F}^{-1}(SV) = \mathcal{F}^{-1}(S) * \mathcal{F}^{-1}(V) \quad (1.4)$$

where $\mathcal{F}^{-1}$ denotes the inverse Fourier transform and the notation $f*g \equiv \int_{-\infty}^{\infty} f(\tau)g(t-\tau)\,\mathrm{d}\tau$ represents the convolution of the functions f and g. The quantity $\mathcal{F}^{-1}(S)$ is

called the "synthesized beam" or "dirty beam" and represents the point spread function of the image. It is related to the angular resolution of the interferometer characterized by λ/D where λ is the wavelength of the radiation and D is the (longest) baseline.

The dirty image contains the response of the instrument characterized by $S(u, v)$. When $S(u, v)$ contains null points, one cannot divide it from the measured visibilities to remove it. Instead, deconvolution algorithms such as CLEAN have been developed [37, 38, 39, 40]. There are several variations of CLEAN, but they proceed in general as follows: determine the location and brightness of the peak intensity in the image, subtract the dirty beam with a fraction of the peak intensity from that location, record the position and the amount subtracted in the model image, repeat the previous steps until the peak is below the threshold specified by the user, convolve the model image with the CLEAN beam that is usually an elliptical Gaussian fitted to the central lobe of the dirty beam, and add the residuals of the dirty image to the CLEAN image.

The CLEAN image is the deconvolved version of $I^D(l, m) \equiv A(l, m)I(l, m)$ and thus still contains the response of the primary beam. One then divides the CLEAN image by a model of the primary beam, which is usually obtained by electromagnetic simulations of the antenna elements or empirical measurements of the antenna patterns, to recover the "true" sky intensity distribution.

The preceding discussion relies on the assumption that (l, m) are small, i.e. the field of view is narrow. This assumption obviously breaks down for the widefield radio interferometers. Nonetheless, there are algorithms to account for the widefield effects. w-projection is one [41], which interprets the factor $e^{-2\pi i w(\sqrt{1-l^2-m^2}-1)}$ as a convolution kernel that lets one project the visibility $V(u, v, w)$ to the $(u, v, w = 0)$ plane. Another is w-stacking implemented in WSCLEAN [42], which, instead of convolving the kernel, multiplies it after the inverse Fourier transform.

There is much more to synthesis imaging in radio astronomy, and the interested reader may refer to [43, 44] for detailed discussions.

1.3 The Murchison Widefield Array

The Murchison Widefield Array (MWA) is one of the new widefield radio interferometers, and began operating in 2013 [45, 46]. Located in Murchison Shire of Western Australia, where there is very little radio frequency interference, it is one of the three precursor telescopes for the Square Kilometer Array (SKA), which will be the world's largest and most sensitive radio telescope if built [47]. The MWA operates at low frequencies (80–300 MHz), complementary to the other two precursor telescopes that operate at high frequencies—the Australian SKA Pathfinder (ASKAP, also located in Murchison [48]) and MeerKAT (located in Northern Cape of South Africa [49]). The MWA, in the southern hemisphere, is also complementary to the Low Frequency Array (LOFAR [50]), which operates at similar frequencies but is located in the Netherlands.

The MWA is an aperture array consisting of 128 antennas ("tiles," see Figure 1-1) distributed across a circular area with a diameter of ~ 3 km, which gives an angular resolution of $\sim 2'$. Each antenna consists of 16 dual-polarization dipoles arranged in a 4×4 regular grid on a 5 m×5 m ground screen. Although the MWA operates at frequencies 80–300 MHz, its bandwidth is 30.72 MHz split into 1.28-MHz coarse channels, which need not cover contiguous frequency bands, and 40-kHz fine channels that set the spectral resolution. The temporal resolution is 0.5 s. The MWA has good thermal sensitivity ($\sim$ 10 mJy after 30 s of integration) and an extremely wide field of view (610 deg^2 at 150 MHz), making it an excellent survey instrument.

1.4 Image Noise

The radio transients described in Section 1.1 are classified as "slow transients." They last longer than the integration time of an image, which is typically minutes or longer. For example, radio flares of X-ray binaries might last minutes [23] or days [51], while radio afterglows could last weeks to months [32, 33]. This is in contrast to fast transients, such as pulsars or fast radio bursts, which last $\sim$ ms and require different techniques for detection. Slow transients, traditionally, have been identified as sources

that appear or disappear over time in individual images via source-finding algorithms, which search for islands of bright pixels above the background image noise and group them into sources, e.g. [52]. This technique is limited by the image root-mean-square (RMS) noise and the detection threshold of the source-finding algorithm; the noisier the image, the brighter the source would need to be for a significant detection.

Several components contribute to the image RMS noise. One component is the thermal noise $\sigma_{\rm th}$, or the random noise in the instrument, which depends on the design of the instrument but always decreases with longer integration time:

$$\sigma_{\rm th} = \left(\frac{2k_{\rm B}T_{\rm sys}}{A_{\rm eff}N_{\rm ant}\epsilon_c}\right)\frac{1}{\sqrt{N_{\rm pol}B\,t_{\rm int}}} \tag{1.5}$$

where $k_{\rm B}$ is the Boltzmann constant, $T_{\rm sys}$ is the system temperature, $A_{\rm eff}$ is the effective area of each antenna, $N_{\rm ant}$ is the number of antennas, ϵ_c is the correlator efficiency, $N_{\rm pol}$ is the number of polarizations, B is the instantaneous bandwidth, and $t_{\rm int}$ is the image integration time.

Another component is the classical confusion noise σ_c, which arises from a background of faint, unresolved sources [53]. This is a spatial noise in the image that depends on the source distribution in the sky and the instrument resolution:

$$\sigma_c^2 = \Omega_b \int_{S_{\rm min}}^{S_c} S^2 \frac{{\rm d}n}{{\rm d}S}\,{\rm d}S \tag{1.6}$$

where Ω_b is the synthesized beam, S is the source flux density, and ${\rm d}n/\,{\rm d}S$ is the differential number density of sources [54]. The lower limit of integration $S_{\rm min}$ is set by the sensitivity while the upper limit of integration S_c is the flux density of a source detected at a particular signal-to-noise ratio $q = S_c/\sigma_c$; usually S_c, referred to as the confusion limit, is determined iteratively until $q = 5$. Note that Equation 1.6 diverges for $S_{\rm min} \to 0$ if ${\rm d}n/\,{\rm d}S \propto S^{-3}$; however, measurements suggest that at low flux densities, ${\rm d}n/\,{\rm d}S \propto S^{-1.6}$ [55] while in the Euclidean universe, ${\rm d}n/\,{\rm d}S \propto S^{-2.5}$, so this is generally not an issue.

As the instrument resolution improves and more sources are resolved, the classical

confusion noise decreases. For a modest angular resolution like that of the MWA, however, the classical confusion noise will become the limiting factor in the image RMS noise because σ_c does not decrease with longer integration times as $\sigma_{\rm th}$ does. This, in turn, limits the ability of source-finding algorithms to identify faint sources in an image. The theoretical estimate of the confusion limit for the MWA is $\sim 10\,\mathrm{mJy}$ at 150 MHz [46], suggesting that the MWA images will be limited by the classical confusion noise.

However, the classical confusion noise is largely independent of time, so it is, in principle, not a limit for detecting fainter but varying flux densities [53]. A simple method for detecting flux density variations below the classical confusion noise is image subtraction. For example, one could subtract images taken at the same local sidereal time to remove both classical confusion noise and sidelobe confusion noise, the latter of which is due to the synthesized beam sidelobes of unsubtracted sources. For many surveys, however, the images are taken at different local sidereal times. In this case, image subtraction can be prone to artifacts. Even without image subtraction, CLEAN artifacts impact radio transient searches; despite the widespread use of CLEAN, its noise properties are not well-understood. For instance, [56] found that many transient candidates reported by [57] were in fact artifacts, and those that were not determined to be artifacts were reported to be detections at lower significance.

Many radio transients are expected to be faint ($\lesssim$ mJy). Thus we seek a method that not only takes advantage of the time-independent nature of the classical confusion noise to detect transients without relying on source-finding algorithms that could otherwise limit the sensitivity of the search, but also has well-defined statistical properties that takes into account the distribution of artifacts.

1.5 The Matched Filter and Radio Afterglows

The matched filter is optimal for detecting transient signals with known forms in the presence of stochastic noise. It uses the template of the known signal to search for a "match" in the noisy unknown signal by applying a linear filter that maximizes

the signal-to-noise ratio. It is widely used in engineering, e.g. radar applications [58], as well as gravitational wave astronomy [59], and we can also apply it to radio astronomy, for example in radio afterglow searches.

The coalescence of two compact objects—binary neutron stars (BNS) or a neutron star and a black hole—are predicted sources of electromagnetic (EM) and gravitational wave (GW) emission. Joint EM and GW observations of these systems are complementary as they probe different physical processes and are necessary for certain science objectives, e.g. [60, 61, 62]. For example, EM detections of GW events will measure the source redshift and break the degeneracy between the source distance and its inclination angle. This will improve estimates of astrophysical parameters such as the Hubble parameter since GW detectors have different systematic uncertainties, e.g. [63, 64, 65]. GW measurements of the inclination angles will also improve our understanding of the dynamics and energetics of the EM counterparts [66].

The next generation of GW detectors has come online or will soon do so, e.g. Advanced LIGO (aLIGO, [67]) and Advanced Virgo [68], and has made the first detection of a binary black hole event GW 150914 [34]. EM follow-up of GW events will be important for studying them. Many EM counterparts have been proposed, including kilonovae, short gamma-ray bursts (SGRB), and afterglows, e.g. [69, 70, 71, 72, 33, 73]. To date, observational evidence supporting the connection between these EM counterparts and compact binary coalescence (CBC) remains indirect or uncertain, for instance the diverse properties of SGRB host galaxies [74, 75] and one possible kilonova association with a SGRB [76, 77]. Coincident detections of these EM counterparts and GW emission will firmly establish the origin of these events. However, the sky localization of GW events will be poor during the early days of GW detector operation, ranging from 100–1000 deg^2 [78, 79]. This presents a challenge for EM follow-up as most telescopes have much smaller fields of view in comparison.

Radio interferometers such as the MWA, however, have extremely large fields of view (e.g. $\sim 600\,\mathrm{deg}^2$) that make them promising for EM follow-up of GW events. Most of these instruments operate at low frequencies ($< 500\,\mathrm{MHz}$), where the expected EM counterpart of a GW event is a SGRB afterglow. So far, there have been

no detections of SGRB afterglows at low frequencies and only 4 detections at high frequencies ($> 5\,\mathrm{GHz}$): GRB 050724, GRB 051221A, GRB 130603B, and GRB 140903A [80, 81, 82, 83]. This is not surprising given the sample of radio afterglows discussed in [84]. Few SGRBs, if any, have been observed at low frequencies. SGRBs are also intrinsically fainter than long GRBs, releasing less energy in total and occurring in a less dense medium. Radio emission from afterglows usually peak on timescales of weeks to months, if not longer, and few SGRBs have been observed on this timescale. Furthermore, SGRBs triggered by γ-rays have been cosmological ($z > 0.1$) if they have measured redshifts at all. By contrast, detectable BNS events will be nearby ($z < 0.1$). For these events, radio afterglows are still expected to be faint and long-lasting [32, 33, 85], but the detailed properties of the light curves are sensitive to many model parameters that remain uncertain.

The general shape of an afterglow light curve rises and falls on timescales of months to years as shown in Figure 1-3. At early times, an on-axis observer sees more emission than an off-axis observer because of collimated outflows and relativistic beaming. At late times, the on-axis and off-axis light curves become indistinguishable as the emission becomes isotropic. The counter-jet contributes to a late-time brightening of the light curve, an effect that is most prominent for an on-axis observer. While synchrotron emission becomes stronger as the circumburst density n increases, synchrotron self-absorption becomes even stronger at low radio frequencies. In these light curves, synchrotron self-absorption is most prominent when $n = 1\,\mathrm{cm}^{-3}$, causing a much slower rise in flux at early times.

These model afterglow light curves contain much physical information about the bursts that produced them. Hence, using them as templates for the matched filter in a search for radio afterglows will lead to a more sensitive search and better constrain the rates of these events. Extracted light curve parameters of any detected afterglow will also measure the burst properties.

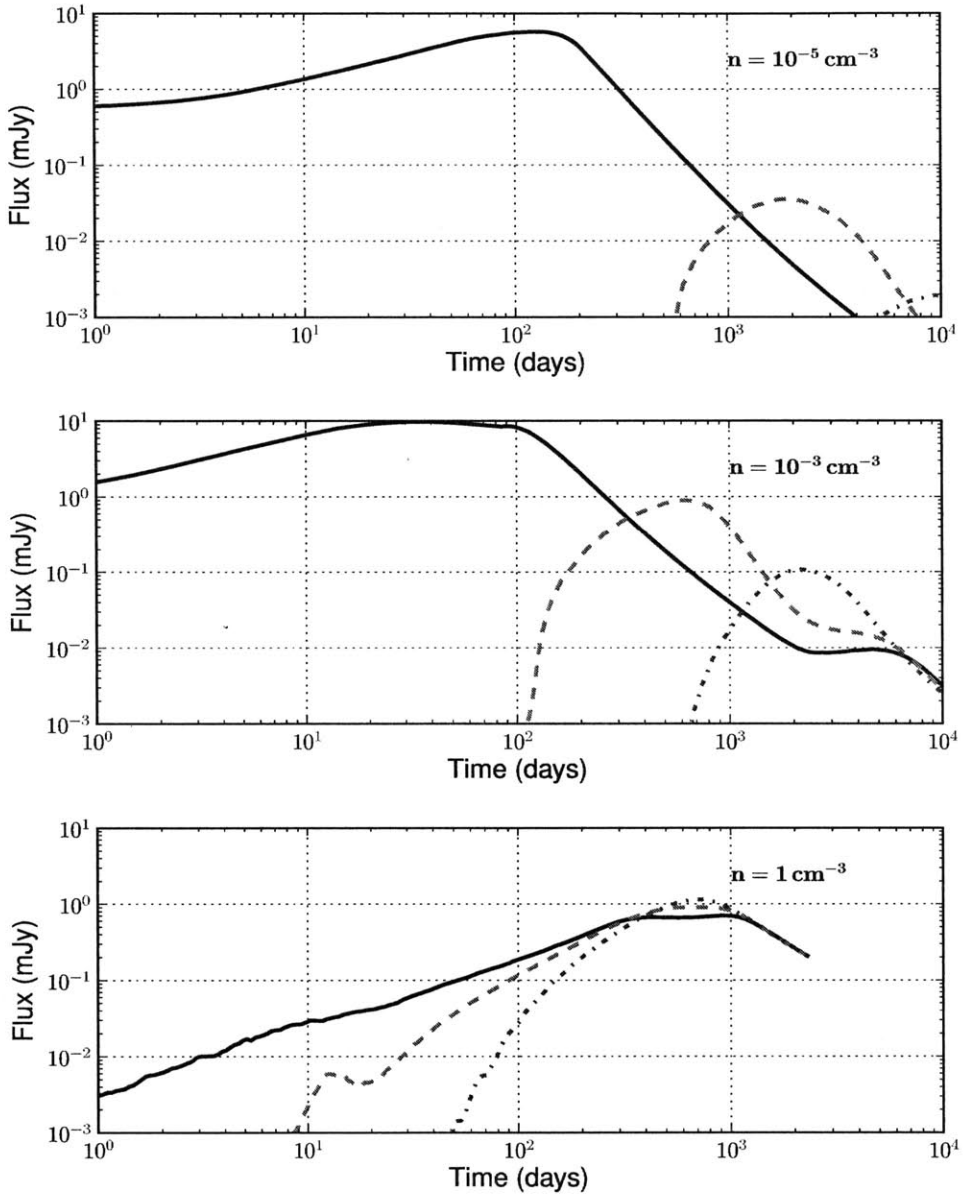

Figure 1-3: Examples of simulated afterglow light curves at 150 MHz for different values of n and $\theta_{\rm obs}$ to illustrate the characteristic features of the bursts independent of any instrument (black solid line is on-axis, red dashed line is 45°, and blue dash-dotted line is off-axis). These bursts are located at 100 Mpc with $E_{\rm iso} = 5 \times 10^{51}$ ergs and $\theta_{\rm jet} = 11.5°$. Early-time emission is brighter for on-axis observers because of collimated outflows and relativistic beaming whereas late-time emission is isotropic. The late-time bump is caused by the counter-jet. The light curves in the bottom panel are qualitatively different from the ones in the top two panels because synchrotron self-absorption is stronger at higher densities. These light curves are generated using the numerical tool BOXFIT [86].

1.6 Thesis Outline

This thesis adapts the well-known matched filter technique [58] to detect radio transients in the presence of classical confusion noise by drawing on the experience of the LIGO community, which has developed techniques to detect GW signals in the presence of non-Gaussian noise (e.g. [87, 59]). This technique operates on the image pixel level, has well-defined statistical properties, and is applicable to variable and transient searches for both confusion-limited and non-confusion-limited instruments. We apply this technique to search for slow transients in the MWA data.

In Chapter 2, we describe the mathematical framework for our radio transient detection technique, derive its statistical properties, and describe the software implementation of our technique: *Simetra*. In Chapter 3, we describe the MWA data reduction procedure. In Chapter 4, we discuss the performance of *Simetra* on real MWA data to demonstrate the potential of our technique for sensitive transient searches. In Chapter 5, we describe the transient search analysis, discuss the results of our search, and set an improved upper limit on the transient surface density. In Chapter 6, we use the matched filter technique to examine the detectability of late-time afterglows from compact binary coalescence for various widefield radio interferometers. In Chapter 7, we conclude that our technique is capable of detecting faint transients and discuss areas of improvement and application as well as future work.

Chapter 2

Matched Filter Technique for Radio Transient Detection

The presence of classical confusion noise limits the ability of source-finding algorithms to identify sources with flux densities near or below the classical confusion noise, but the source population that contributes to the classical confusion noise is independent of time unless they are genuine transient or variable events. Thus a transient detection technique that searches for flux density variations on top of a constant signal without relying on a source-finding algorithm is needed to detect transient signals fainter than the classical confusion noise. In this chapter we describe and implement a technique that identifies transients in individual image pixels despite the classical confusion noise.

2.1 Theory

Adapted from matched filter techniques, which have been used in engineering applications [58] and gravitational wave astronomy [59, 87], this technique searches for flux density variations on top of a constant signal in individual pixels without using source-finding algorithms. As it is capable of probing below the classical confusion noise, it is sensitive to faint transients in addition to bright ones. We derive a new transient detection statistic from this technique, discuss its statistical properties, and

relate it to the sensitivity of a radio transient search. Although our formalism is derived for transient detection in the image domain, it is similar to the formalism for source detection in the visibility domain [88].

To determine whether or not there is a transient signal in a particular image pixel, we compare two hypotheses: the transient is absent (the null hypothesis H_0), and the transient is present (the alternative hypothesis H_1). Usually, H_0 only includes random noise, but in this case, we add a constant background[1] to H_0 to represent the time-independent contribution from the classical confusion noise (or a steady source) because we are only interested in the change in flux density over time:

$$H_0 : x_i = c + \sigma_i \tag{2.1}$$

$$H_1 : x_i = c + Af_i + \sigma_i \tag{2.2}$$

For a fixed pixel, x_i is the measured flux density in the ith snapshot ($i = 1, 2, \ldots, N$), c is the constant background, σ_i is the RMS noise (thermal, sidelobe confusion, and other random errors), and Af_i is the transient signal, where A is the overall amplitude for a light curve template $\mathbf{f} \equiv \{f_1, f_2, \ldots, f_N\}$. Given the two hypotheses and the data $\mathbf{x} = \{x_1, x_2, \ldots, x_N\}$, we compute the ratio of the likelihood functions known as the Bayes factor or the likelihood ratio as part of hypothesis testing [89, 90]:

$$\Lambda(\mathbf{x}) = \frac{p(\mathbf{x} \mid H_1)}{p(\mathbf{x} \mid H_0)} \tag{2.3}$$

where $p(\mathbf{x} \mid H_i)$ is the probability of observing $\mathbf{x}$ given that H_i is true for $i = (0, 1)$.

To derive an analytical result, we assume that the image noise follows a Gaussian distribution with $\mu = 0$ and $\sigma_{\text{im}} = \{\sigma_1, \sigma_2, \ldots, \sigma_N\}$, but we show later that our transient analysis does not rely on this assumption. For the MWA data, this noise is a combination of thermal noise, (residual) sidelobe confusion noise, and other random

[1] Pure random noise is a special case where the constant background is zero.

errors. The likelihood functions are thus the following:

$$p(\mathbf{x} \mid H_0) = \int \mathcal{N}_0 \exp\left(-\sum_{i=1}^{N} \frac{b_i^2(x_i - c)^2}{2\sigma_i^2}\right) p(c \mid H_0)\,\mathrm{d}c \tag{2.4}$$

$$p(\mathbf{x} \mid H_1) = \int \mathcal{N}_1 \exp\left(-\sum_{i=1}^{N} \frac{b_i^2(x_i - c - Af_i)^2}{2\sigma_i^2}\right) p(c \mid H_1) p(A \mid H_1)\,\mathrm{d}c\,\mathrm{d}A \tag{2.5}$$

$\mathcal{N}_0$ and $\mathcal{N}_1$ are the normalization factors for a multivariate normal distribution. b_i is the value of the ith primary beam for the given pixel. $p(c \mid H_0)$, $p(c \mid H_1)$, and $p(A \mid H_1)$ are the probability distributions of c and A given the respective hypotheses. We assume that c and A are independent and uniformly distributed, and we estimate them by using a least-squares approach [91]. To do that, we solve Equations 2.4 and 2.5 by approximating the integral with the value at its extremum, i.e. $p(\mathbf{x} \mid H_j) \approx \text{const} \times \exp(-\chi^2_{\min}/2)$ where $\chi^2_{\min}$ is the solution to $\nabla\chi^2 = 0$, and $\chi^2 \equiv \sum b_i^2(x_i - c)^2/\sigma_i^2$ for Equation 2.4 and $\chi^2 \equiv \sum b_i^2(x_i - c - Af_i)^2/\sigma_i^2$ for Equation 2.5.

For H_0, the solution $c = c_0$ is an estimate of the classical confusion noise or the constant background level, and it is given by the weighted average of the data:

$$c_0 = \langle \mathbf{x} \rangle \equiv \frac{\sum b_i^2 x_i/\sigma_i^2}{\sum b_i^2/\sigma_i^2} \tag{2.6}$$

For H_1, the solution $c = c_1$ is an estimate of the classical confusion noise in the presence of a transient signal; c_1 becomes c_0 in the absence of the transient signal. The other solution $A = A_1$, which is unitless, is the amplitude of the transient signal given the predefined template $\mathbf{f}$, which has units of flux density.

$$c_1 = \langle \mathbf{x} \rangle - A_1 \langle \mathbf{f} \rangle \tag{2.7}$$

$$A_1 = \frac{(\mathbf{x}, \mathbf{f} - \langle \mathbf{f} \rangle)}{(\mathbf{f} - \langle \mathbf{f} \rangle, \mathbf{f} - \langle \mathbf{f} \rangle)} \tag{2.8}$$

The notation $\langle \mathbf{j} \rangle$ represents the weighted average of $\mathbf{j}$ as in Equation 2.6, and $(\mathbf{j}, \mathbf{k}) \equiv \sum b_i^2(j_i k_i)/\sigma_i^2$ denotes the "weighted" inner product between $\mathbf{j}$ and $\mathbf{k}$. Note that an equivalent way of writing $(\mathbf{x}, \mathbf{f} - \langle \mathbf{f} \rangle)$ is $(\mathbf{x} - \langle \mathbf{x} \rangle, \mathbf{f})$, which we interpret as how

well the light curve template matches the data after the constant component has been subtracted. However, $(\mathbf{f} - \langle \mathbf{f} \rangle)$ is simpler and computationally less intensive to calculate since $\mathbf{f}$ is predefined and identical for all pixels, so we write Equation 2.8 in its given form.

Having determined the best fit values for c and A, we substitute Equations 2.6, 2.7, and 2.8 into the approximations of Equations 2.4 and 2.5, which then lets us solve for Equation 2.3. As $\Lambda(\mathbf{x})$ is essentially a ratio of exponents normalized by a constant, we rewrite Equation 2.3 in the form $\Lambda(\mathbf{x}) \approx \text{const} \times \exp\left[\rho^2/2\sigma_\rho^2\right]$ and define the new quantities ρ to be the "detection statistic" and σ_ρ to be the standard deviation of the ρ distribution:

$$\rho = (\mathbf{x}, \mathbf{f} - \langle \mathbf{f} \rangle) \tag{2.9}$$

$$\sigma_\rho = \sqrt{(\mathbf{f} - \langle \mathbf{f} \rangle, \mathbf{f} - \langle \mathbf{f} \rangle)} \tag{2.10}$$

ρ is a modified version of the matched filter [58] and determines the likelihood that $\mathbf{x}$ contains the signal $A_1\mathbf{f}$. In other words, it determines which hypothesis is favored (transient absent or present) and by how much. As ρ is a linear superposition of Gaussian random variables, its distribution is also Gaussian with width σ_ρ. When the transient signal is absent, the mean of ρ is $\mu_0 = 0$. When the transient signal is present, the mean of ρ is shifted by the signal and becomes $\mu_1 = A_1(\mathbf{f} - \langle \mathbf{f} \rangle, \mathbf{f} - \langle \mathbf{f} \rangle) = A_1\sigma_\rho^2$. This is illustrated in Figure 2-1.

ρ and σ_ρ are also related to the more familiar quantity σ_{im} through the weighted inner product: $\rho \propto A_1/\sigma_{\text{im}}^2$ and $\sigma_\rho \propto 1/\sigma_{\text{im}}$. As σ_{im} decreases, both ρ (or μ_1) and σ_ρ increase, but ρ (or μ_1) increases faster than σ_ρ, which leads to a better separation of background and signal (see Figure 2-1) and hence an improved sensitivity. Similarly, the brighter the transient (larger A_1), the better the sensitivity.

So far we have considered ρ for a single template. However, in a real search, we maximize over different start times and many templates, so the practical statistic is

$$\tilde{\rho} = \max(\rho/\sigma_\rho) \tag{2.11}$$

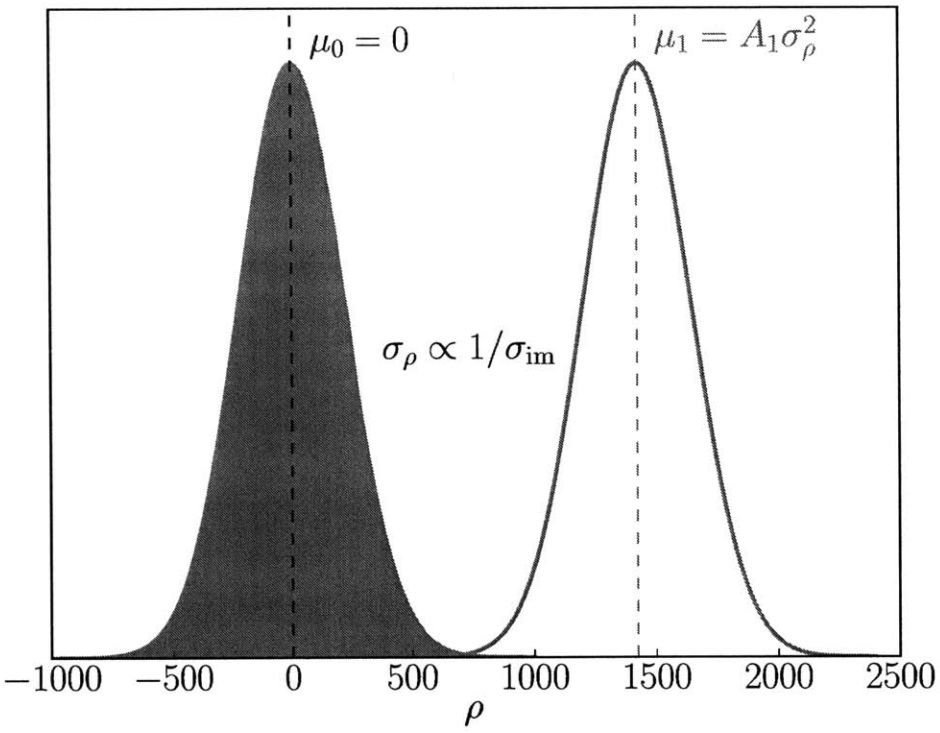

Figure 2-1: Theoretical distributions of the transient detection statistic ρ: background (gray-filled) and signal (red-line) with arbitrary image noise $\sigma_{\rm im}$ and signal amplitude A_1. As $\sigma_{\rm im}$ decreases, the width of the distributions increase according to $\sigma_\rho \propto 1/\sigma_{\rm im}$, but the mean of the signal distribution increases faster according to $\mu_1 \propto 1/\sigma_{\rm im}^2$, so the sensitivity improves as there is better separation between background and signal. Similarly, the brighter the signal (larger A_1), the better the sensitivity.

Using the ratio ρ/σ_ρ ensures that all the pixels are drawn from the same standard normal distribution. Note that $\tilde{\rho}$ is the maximum of a Gaussian random variable, so its distribution is no longer Gaussian, but this is not an issue. We will discuss this in detail in Chapter 4 where we describe how to use $\tilde{\rho}$ in a transient search and how to characterize the sensitivity and efficiency of the search.

2.2 Implementation

We describe the *Simetra* pipeline, a `Python` implementation of the transient detection technique formulated in the previous section. The pipeline determines the transient detection statistic and the parameter values of the light curve template that best match the observed light curve for every image pixel. One can also use this pipeline to inject transient light curves with known parameters into the data before running the transient search; this determines the search efficiency. The pipeline is illustrated in Figure 2-2, and the software is described in Appendix A.

The pipeline takes, as input, a list of sky images, the corresponding primary beam images, and a choice of light curve template. At the moment of writing, two templates choices, the top-hat and the power-law, are available, but other choices are easy to implement as the code is designed to be as modular as possible. The pipeline reads the input FITS images and converts the flux density and primary beam values for every pixel into the time series $\mathbf{x}$ and $\mathbf{b}$, i.e. light curves for every pixel. Out of memory consideration, the pipeline only loads a subset of the image pixels each time. The pipeline then estimates the noise σ_{im} in the sub-images by calculating the median absolute deviation and then converting that into the standard deviation to account for outliers in a robust manner.

Here the user has the option to inject transients. The user chooses the type and the number of transient light curve templates to inject, and specifies the range and the distribution of the template parameters. Given these parameters, the pipeline generates the light curves on the fly and injects them into the data before it runs the transient search. We decided to inject transients directly into the pixel light curves

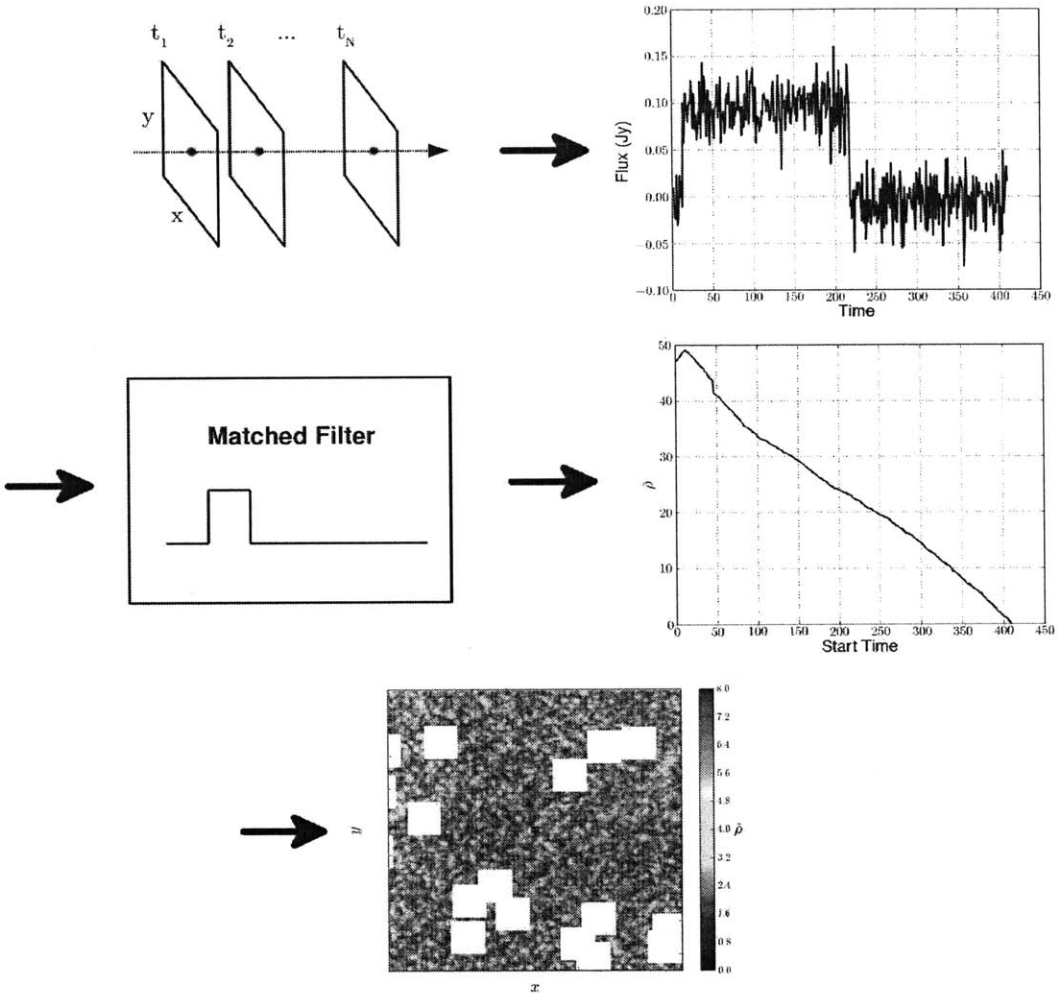

Figure 2-2: Illustration of the *Simetra* pipeline with a mix of real and synthetic data. *Simetra* aligns snapshot images in time, extracts the light curve of an individual pixel (for illustration purposes, we show the synthetic light curve of a transient source with an amplitude of 100 mJy that is later injected into real data), runs the matched filter to calculate ρ and σ_ρ, iterates over different start times of the light curve template to determine $\tilde{\rho}$, and repeats for every pixel, thus producing an "image" of $\tilde{\rho}$. Except for the injection located at the center of the image region, the rest of the pixels are real data. The white boxes are the locations of masked bright sources. The color scale of $\tilde{\rho}$ has been cut off at 8 to show the contrast of the real data. See text for details.

instead of the visibilities because of computation concerns.

The transient search is the matched filter calculation. First the pipeline generates a phase space of template parameters, which can be different from the injection parameters depending on the user's choice. Then it iterates over every set of parameters, including all possible transient start times, generates the corresponding light curve template $\mathbf{f}$, and calculates ρ and σ_ρ according to Equations 2.9 and 2.10. Gaps in the data are handled properly by sampling $\mathbf{f}$ at the existing image time stamps, and do not pose a problem. Finally the pipeline outputs a FITS table that contains the most significant (ρ, σ_ρ) and the corresponding template parameters for every pixel. As the amplitude can be determined by $A_1 = \rho/\sigma_\rho^2$, it is not stored in the output file.

There are two more steps before the user can identify transient candidates: setting the threshold for detection and characterizing the detection efficiency. As we will describe these steps in detail in Chapter 4, we only summarize them here. To set the threshold, the user designates a small part of the image as the playground region where there are assumed to be no transients, corrects the cumulative distribution of $\tilde{\rho}$ in the playground region by the trials factor, i.e. the number of synthesized beams in the search region compared to the number in the playground region, and extrapolates the tail of this corrected distribution to a tolerable probability of false alarm. The extrapolated distribution then determines the significance of any detection during the actual search. To characterize the efficiency, the user applies the threshold to the injected transients and computes the fraction that is recovered. This procedure handles non-Gaussianity in the data, such as artifacts or sidelobe confusion noise as well as the effects of maximization, thus making it a powerful technique. We demonstrate the technique in Chapter 4.

Chapter 3

Data Reduction and Empirical Primary Beam Correction

This chapter contains content adapted from the MWA memo titled An Implementation of Empirical Primary Beam Correction in MWA Imaging [92].

3.1 Data Description

The data for this thesis were taken according to the commensal MWA observing proposals[1] G0009 ("Epoch of Reionisation," EOR) and G0005 ("Search for Variable and Transient Sources in the EOR Fields with the MWA") for Semester 2013-B. These observations were done using the "point-and-drift" strategy, where the primary beam pointing (beamformer setting) changed every 20–30 minutes to track the field after it had drifted across the field of view of the instrument.

We used 1251 snapshot observations of the EOR0 field, which is centered on $(\mathrm{RA}, \mathrm{Dec}) = (0°, -27°)$, taken on 18 nights between 2013 September 2 and November 30. We included only snapshots taken when the field center was within 20° from the zenith, which corresponded to $\sim 2.5\,\mathrm{h}$ of observation each night, and excluded data from 2013 October 15 because of known ionospheric activity [93]. Each snapshot is a multi-frequency synthesis image integrated over 112 s with a bandwidth of

[1] http://mwatelescope.org/astronomers/

Table 3.1: Summary of observations used in this thesis. Each snapshot is centered on the EOR0 field, or (RA, Dec) = (0°, −27°), and integrated over 112 s at 182.4 MHz.

Date	Time Range (UT)	Number of Snapshots
2013-09-02	16:08:08—17:39:36	46
2013-09-04	16:00:16—18:32:48	74
2013-09-06	15:52:22—18:24:54	72
2013-09-09	15:40:39—18:13:03	75
2013-09-11	15:32:47—18:05:11	75
2013-09-13	15:24:55—17:57:19	72
2013-09-17	15:09:11—17:41:35	76
2013-09-19	15:01:19—17:31:43	71
2013-09-30	14:17:59—16:50:31	76
2013-10-02	15:07:03—16:42:39	47
2013-10-04	14:02:15—16:34:47	76
2013-10-08	13:46:31—16:19:03	76
2013-10-10	13:39:43—16:12:15	76
2013-10-23	12:48:39—15:21:03	76
2013-10-25	12:40:47—15:13:11	76
2013-10-29	12:25:03—14:57:27	75
2013-11-18	11:18:31—13:38:55	70
2013-11-29	11:30:41—12:56:09	42

30.72 MHz centered on 182.40 MHz. Table 3.1 lists a summary of the observations.

3.2 Preprocessing

Raw interferometric data were converted into the UVFITS format [94] by the Cotter MWA preprocessing pipeline [95]. During this process, Cotter used AOFlagger [96, 97] to flag radio-frequency interference, frequency channels affected by bandpass aliasing [98], as well as known bad tiles, which might vary from night to night depending on the state of the instrument; roughly 40–50% of the data were discarded. To decrease the file size, Cotter also averaged the data to 1-s time resolution and 80-kHz frequency resolution.

3.3 Calibration

We developed a data reduction pipeline based on the Common Astronomy Software Applications (CASA) package[2] (v4.1.0) [99] and the widefield imager WSCLEAN [42]. Appendix B outlines the pipeline in detail, while we summarize the key steps here.

We built a point-source sky model for each snapshot observation, using the 11 brightest point sources in the field after primary beam attenuation according to the MWA Commissioning Survey Catalog [100]. We generated the model visibilities and the calibration solutions in CASA, using in particular the tools `componentlist`, `ft`, `bandpass`, and `gencal`. Then we performed one iteration of self-cal with WSCLEAN.

Finally, we averaged over many observations on the same night the calibration solutions generated in the previous step to produce a single calibration solution for this night. This was done in two steps: (1) we selected a list of observations for which the AEGEAN source finder (v951) [52] detected at least 1500 sources in each 112-s snapshot as we found this to be a practical indicator of image quality; (2) we averaged the calibration amplitude and phase solutions for each tile, polarization, and frequency channel for the selected observations while ignoring the highest and lowest 10% of the data.

This procedure provided stable and smooth calibration solutions, which were not expected a priori to vary significantly over time and frequency. It also gave more accurate estimates of the source flux densities. We illustrate the improvement of the average calibration process over the individual calibration process through the calibration amplitude and phase solutions in Figure 3-1 as well as the raw light curves in Figure 3-2.

3.4 Imaging

After we applied the average calibration solutions to each snapshot, we generated multi-frequency synthesis images over 30.72 MHz bandwidth in the instrumental XX

[2]http://casa.nrao.edu/

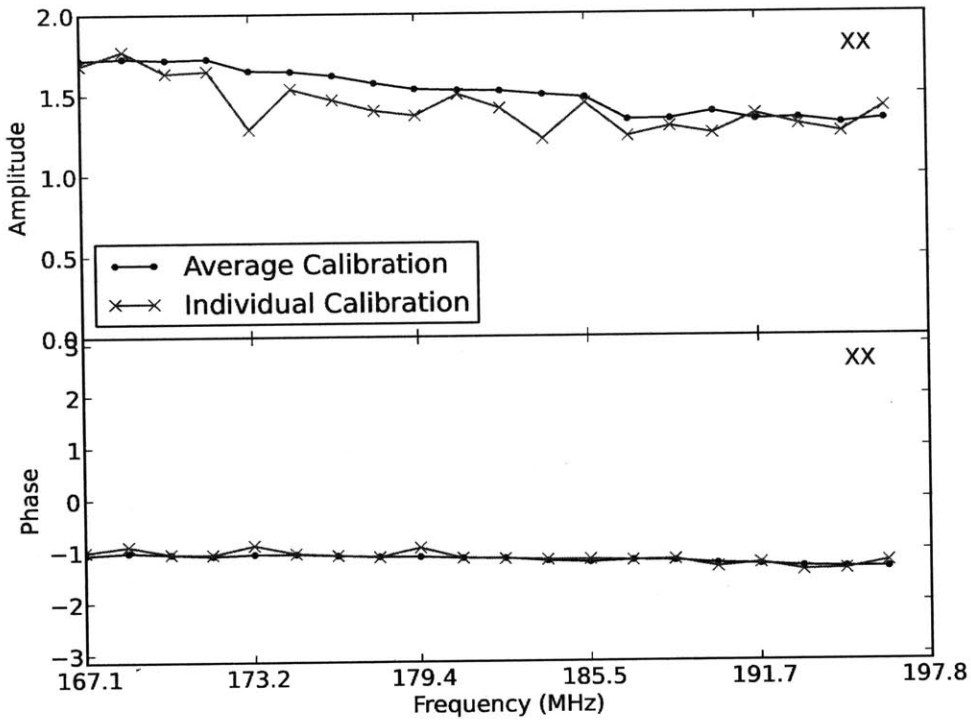

Figure 3-1: Example calibration solutions for MWA antenna Tile 25. The calibration solutions obtained from an individual snapshot show fluctuations that are not expected for the bandpass. In contrast, the solutions averaged over many snapshots are smoother and thus considered to be more reliable.

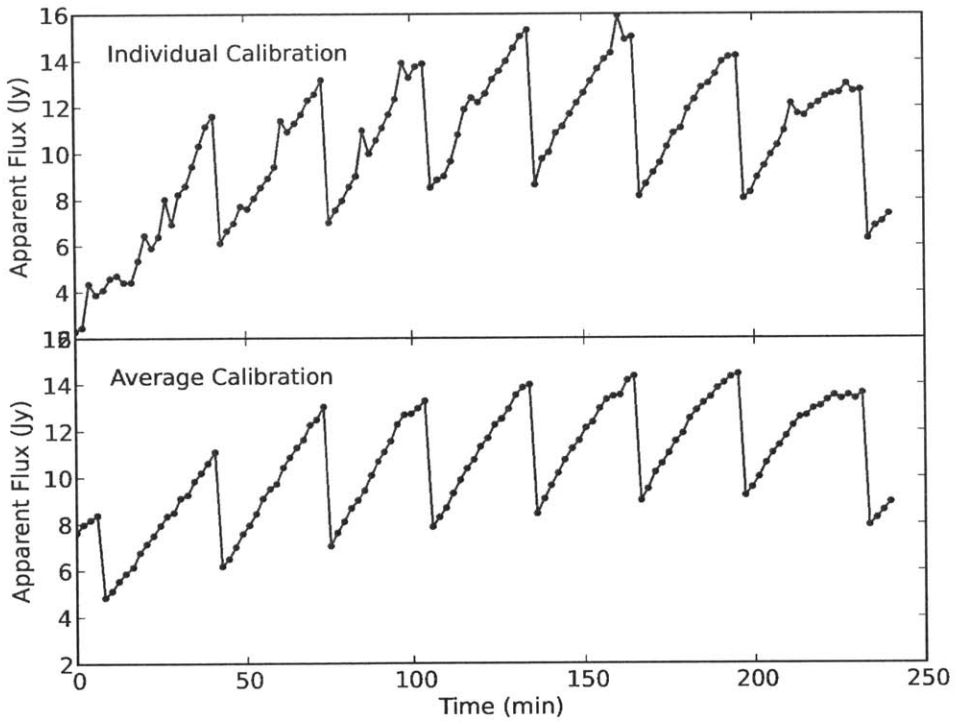

Figure 3-2: Example raw light curves prior to primary beam correction for the same source (MWACS J0025.8-2602) but different calibration methods. Each jump corresponds to a pointing change. The top panel shows fluctuations on short timescales. These fluctuations are present for all sources, suggesting that there are errors in the antenna gain calibration as the primary beam is expected to be smooth. The average calibration solution, on the other hand, removes these fluctuations.

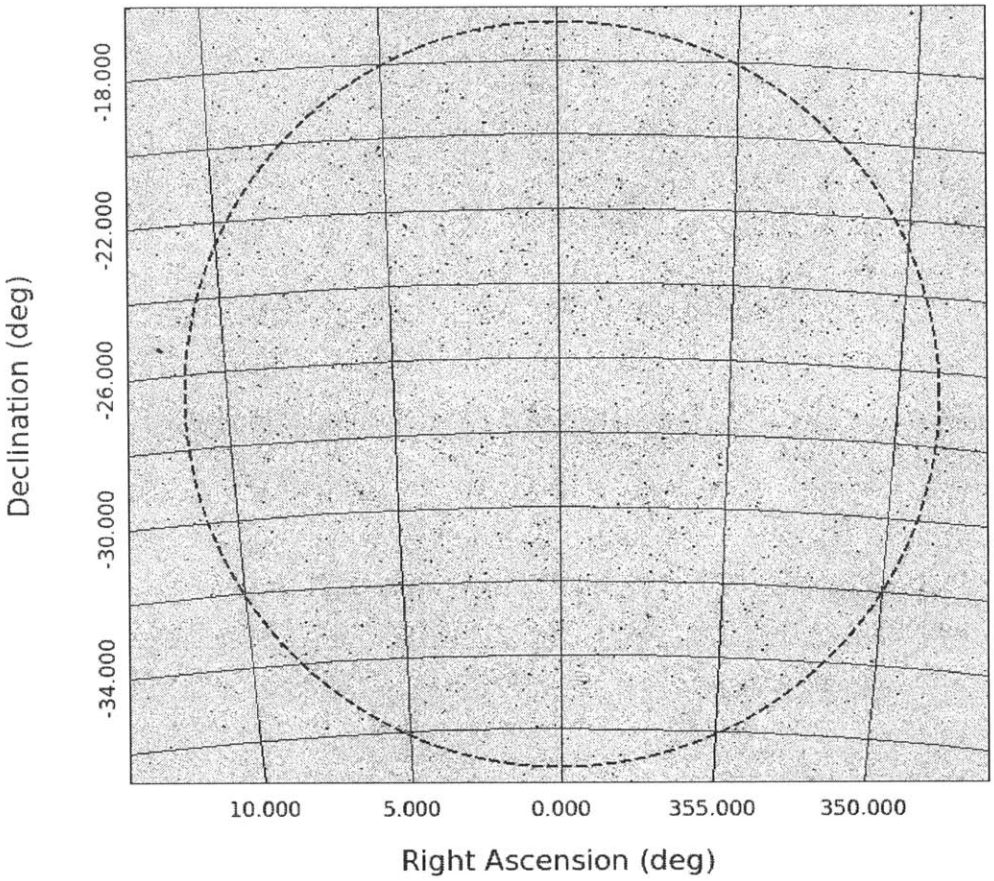

Figure 3-3: An example snapshot of the EOR0 field. It was integrated over 112 s and cleaned in the XX polarization, plotted with the J2000 coordinate grid and a squared color scale. The dashed circle has a radius of 10°, which is approximately the outer boundary of the field included in the analysis of this thesis.

and YY polarizations, setting WSCLEAN to use uniform weighting that gave a synthesized beam of $\sim 2'$, a pixel size of $0.5'$, and an image size of 4096×4096 pixels, which corresponded to a field of view of $34° \times 34°$. Figure 3-3 shows an example snapshot of the EOR0 field.

3.5 Primary Beam Correction

The primary beam model establishes the flux scale during calibration and after imaging, so there have been many efforts to measure and model the primary beam of the

MWA. For example, there is a project to map the beam pattern with an octocopter and a transmitter, while another has used ORBCOMM satellites to measure the beam pattern of an MWA tile at 137 MHz [101]. However, extrapolating this result to other MWA frequencies is not straightforward, so we need to rely on antenna modeling.

The best available model for the MWA primary beam is that developed by [102], hereafter referred to as "the Curtin beam." It improved on the previous model, which treated each antenna element as a Hertzian dipole, by incorporating mutual coupling between the elements and using an average embedded element pattern. This decreased the amount of instrumental Stokes leakage that was more prominent at frequencies $\gtrsim$ 180 MHz. However, the Curtin beam assumed that all tiles were identical and unchanging over time, whereas in reality the MWA site is not perfectly flat and different tiles have different malfunctioning dipoles or beamforming errors [103]. In fact, we found that systematic errors related to inaccurate primary beam modeling had affected the source light curves as measured from the images. As a result, we decided to measure and model the primary beam empirically to remove these errors, which more severely affected the higher frequency observations ($\gtrsim$ 180 MHz) as the Hertzian dipole approximation is less accurate than it is at $\lesssim$ 150 MHz where the MWA is designed to operate best.

To develop the procedure of empirical primary beam correction, we processed EOR0 observations taken on three nights: 2013 August 23, 26, and 27. Two of these (August 23 and 27) consisted of "high-band" observations at 182.42 MHz, while the other consisted of "low-band" observations at 154.28 MHz. We included 111 snapshots ($\sim$ 3.5 h) from each night with a maximum zenith angle of 20.8°. Using two sets of data taken at the same frequency but on different nights allowed us to confirm the presence of systematic errors, and comparing the data at different frequencies allowed us to characterize the frequency dependence of these errors.

After reducing the data from those three nights, we formed Stokes I images corrected by the Curtin beam. We identified sources in each Stokes I image with the AEGEAN source finder and extracted light curves by matching the sources within a radius of 0.1° from snapshot to snapshot. Our initial light curves revealed residual

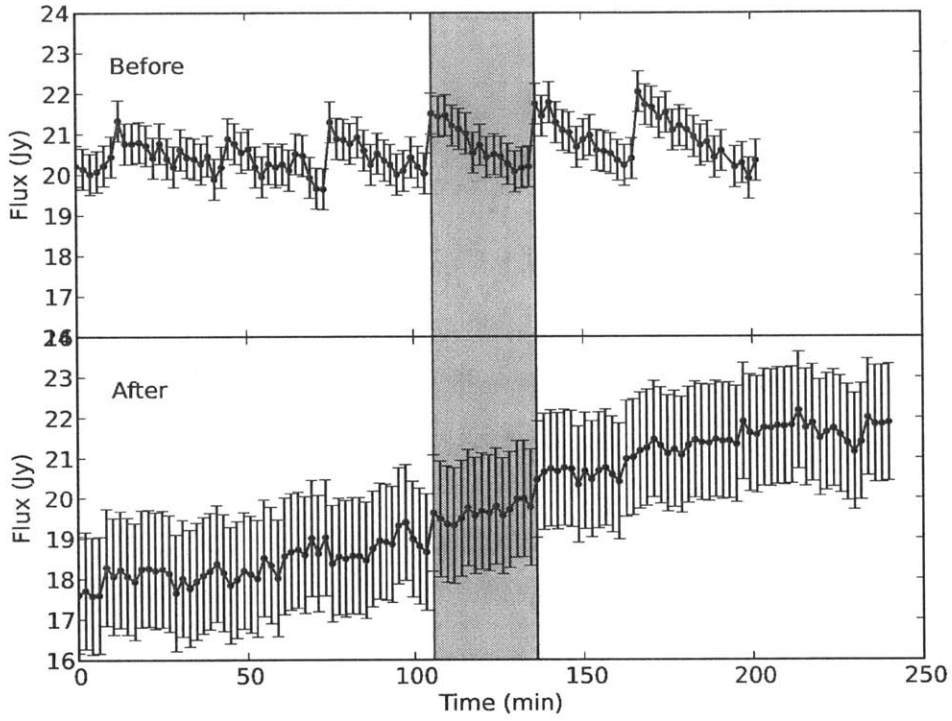

Figure 3-4: Example light curves for the same source before and after a software fix that affected the primary beam model generation. The sawtooth artifacts in the top panel were caused by a one-day date error that introduced a difference of 1° between the pointing centers of the model and the true primary beam. Each discontinuity corresponds to a pointing change; for reference, the zenith pointing is marked by the shaded area. The error bars are set to the standard deviation of the light curve.

"sawtooth" artifacts shown in Figure 3-4, leading to the identification of a software bug that caused a difference of 1° between the pointing centers of the true primary beam and the model primary beam[3]. This illustrates the importance of using accurate models of the primary beam as small errors could manifest as large effects.

Even after the software bug was fixed, the light curves continued to show systematic trends that correlated with source locations in the primary beam. To characterize these trends, we fit a line through each nightly light curve and computed the slope. These slopes showed flux density changes as large as 4% per hour. They appeared to

[3]The observation dates in the UVFITS files were off by 24 h.

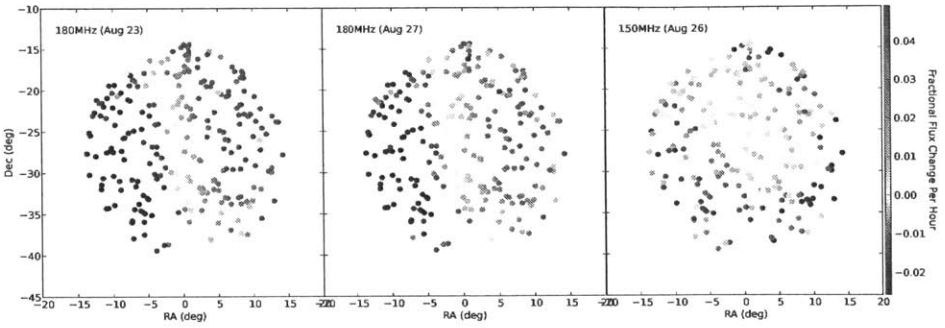

Figure 3-5: RA-dependent light curve slopes for three nights of data and two observing frequencies. The slopes could be as large as 4% flux density change per hour. These trends, independent of calibration or imaging procedures but weaker at lower frequencies, suggested there was a systematic error associated with how the model primary beam differs from the true primary beam as the sources drifted through.

depend on source RA and observing frequency as shown in Figure 3-5, even though the mean flux densities of these sources agreed relatively well with the catalog flux densities after an image gain correction as shown in Figure 3-6. We iterated our analysis pipeline in unsuccessful attempts to eliminate the light curve slopes by changing flagging schemes (80 kHz vs. 240 kHz edge channels), calibration methods (individual calibration vs. average calibration), imaging bandwidth (3.84 MHz vs. 30.72 MHz), and widefield effects (phased to the center of the EOR0 field vs. the zenith). A separate data reduction pipeline based on Fast Holographic Deconvolution (FHD) [104], which uses different calibration and imaging techniques, confirmed the light curve trends at 182 MHz; they had not processed the data at 154 MHz at the time of this particular analysis.

The spatial correlation of the light curve slopes led us to conclude that these systematic trends were most likely caused by inaccuracies in the primary beam model. The RA dependence suggested that they were related to how the primary beam changed as the sources drifted across the sky. In addition, since the theoretical primary beam model was more accurate at lower frequencies, the fact that the trends became less severe at lower frequencies also supported the hypothesis that they were related to the primary beam.

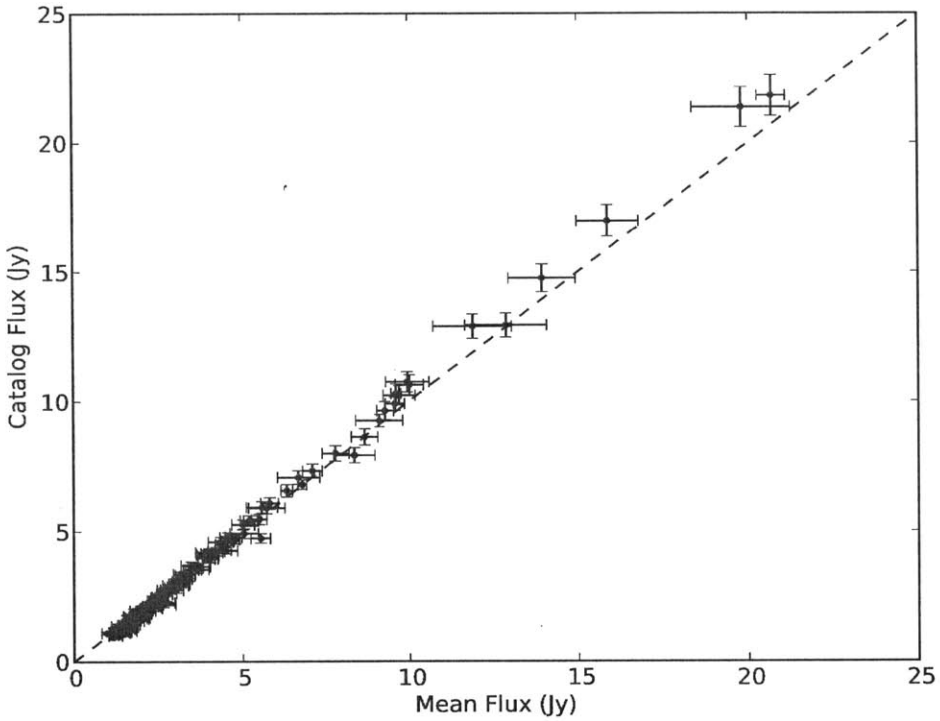

Figure 3-6: Flux density comparison between the MWA Commissioning Survey Catalog and the measured mean flux density at 182 MHz for the sources > 1 Jy after we applied the Curtin beam correction along with an image gain correction. The dashed line indicates perfect agreement. The x-error bar is the standard deviation of the measured light curve, and the y-error bar is the uncertainty reported in the catalog. Light curve slopes were present even though the measured mean flux densities of the sources agreed relatively well with the catalog flux densities.

3.5.1 Empirical Beam Modeling

To mitigate the systematic errors associated with primary beam modeling, we decided to measure and model the primary beam empirically. If we assume that we know the true flux densities of the sources, we can determine the empirical primary beam according to the following relationship:

$$b_{\mathrm{emp,pol}}(\alpha, \delta) = \frac{S_{\mathrm{meas,pol}}(\alpha, \delta)}{S_{\mathrm{ref,pol}}(\alpha, \delta)} \tag{3.1}$$

where $b_{\mathrm{emp,pol}}$ is the empirically measured primary beam for a particular instrumental polarization (XX or YY), $S_{\mathrm{meas,pol}}$ is the measured XX or YY flux density of a source with equatorial coordinates (α, δ) and $S_{\mathrm{ref,pol}}$ is the reference catalog ("true") flux density of the same source. Similar analyses have been done for the Very Large Array [105].

To measure the empirical beam, we used a fixed subset of sources instead of the entire ensemble detected in individual XX and YY snapshots. We did this for two reasons: (1) we wanted to ensure that the sources we used to measure the empirical beam have reliable flux density measurements, and (2) we wanted to avoid overfitting when we fit a smooth function to the measured data points.

For self-consistency, we used the MWA Commissioning Survey Catalog as the reference catalog, which we also used for calibration. This avoids issues that could arise if we used source catalogs from other instruments, which might have different angular resolutions, frequency bands, sky coverage, and so on. We assumed the sources to be unpolarized and used the same catalog flux densities for both X and Y polarizations.

When we selected the subset of sources, we filtered out the sources close to the null of the primary beam ($> 13°$ from the phase center, corresponding to $\lesssim 0.3$ of the primary beam gain). We also filtered out the sources that appeared to have unreliable flux density measurements, which we determined by comparing the catalog flux densities > 1 Jy to the mean flux densities that we measured: if the mean flux density that we measured was 3-σ away from the catalog flux density, we removed it

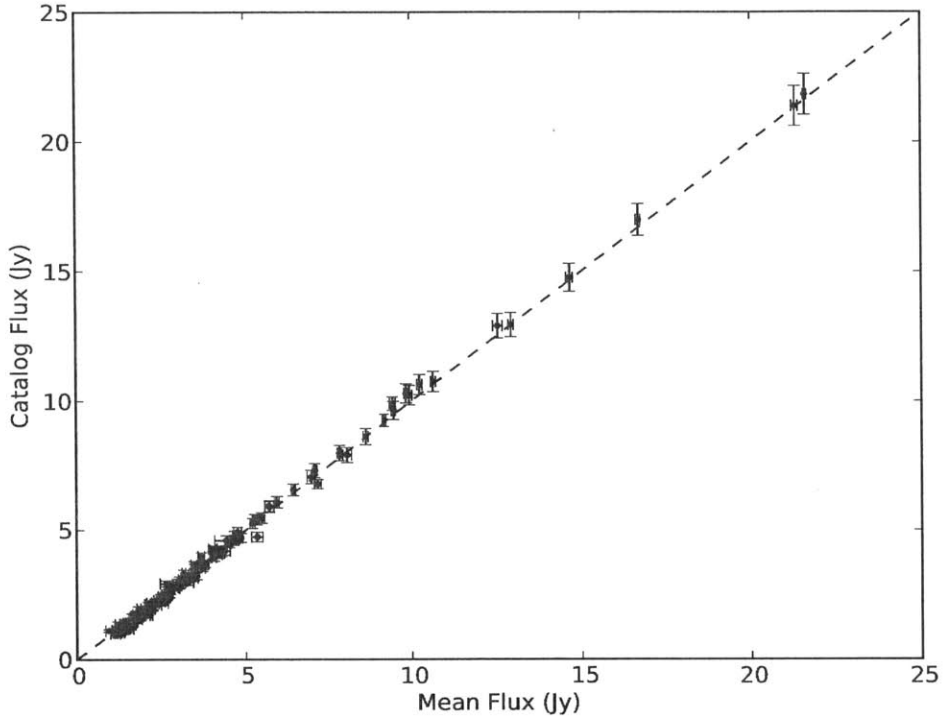

Figure 3-7: Flux density comparison between the MWA Commissioning Survey Catalog and the measured mean flux density at 182 MHz for the sources used to measure the empirical beam after the empirical fitted primary beam correction. The dashed line indicates perfect agreement. The x-error bar is the standard deviation of the measured light curve, and the y-error bar is the uncertainty reported in the catalog. The agreement between the two measurements is better than that shown in Figure 3-6, which used the Curtin beam correction.

from the final catalog that we used to measure the empirical beam. As we could not measure a mean flux density without applying the primary beam correction to the images, we did this step iteratively by first using the Curtin beam and then refining it with the empirical beam once. After this process, the reference catalog contained 245 sources. The agreement between the catalog flux densities and the measured flux densities, shown in Figure 3-7, ensured that the fitting procedure used reliable data.

For the fitting procedure, we assumed that the primary beam was smooth and fit a smoothing spline to the empirical beam measurements on a snapshot-by-snapshot basis. A spline function of degree $k \geq 0$ is a continuous function composed of one or

more polynomial functions with degree $\leq k$ [106]. Each polynomial is defined on a knot interval, where a knot marks the boundary between two polynomial functions. There may be a discontinuity in the kth-order derivative at each knot, but the lower order derivatives of the spline function are continuous. A zeroth-order spline, for example, is a piecewise constant function. We used biquartic splines $(k_x, k_y) = (4, 4)$ as they provided a better fit (lower residuals) than the default bicubic splines, whereas biquintic splines did not improve the fit significantly.

Fitting a spline function $s(x_i, y_i)$ to a set of data z_i that have measurement errors involves a trade between the smoothness of the spline and the goodness of the fit. The algorithm used by the `SmoothBivariateSpline` routine in `scipy`[4], which we used for this analysis, determines the smoothest spline given the constraint that the goodness-of-fit is less than the smoothing factor S:

$$\sum_{i=1}^{m} w_i \left[z_i - s(x_i, y_i)\right]^2 \leq S \tag{3.2}$$

where z_i is the empirical beam measurement for each source, (x_i, y_i) is the equatorial coordinate of the source, and w_i is the weight of each measurement [107]. If S is very large, the spline function becomes the least-square polynomial fit, which would be smooth but likely underfitting. If S is very small, the process becomes an interpolation, which would likely be overfitting. We chose S to be the number of sources that entered the fit (245), which was the default value, as it gave a satisfactory fit; choosing $S = 1$, for example, resulted in an overfit, though setting $S = 1000$ did not change the fit. We set w_i to be the catalog flux density for each source, because the flux density errors derived from AEGEAN appeared to be too large to provide reliable inverse variance weights and did, in fact, make the fit worse. The unreliable errors reported by AEGEAN are a known issue and will be fixed.

We also verified that the fitting function was robust. Instead of deriving an empirical primary beam based on flux density measurements in clean XX and YY images before primary beam correction, we derived an empirical "correction factor" for the

[4]http://www.scipy.org/

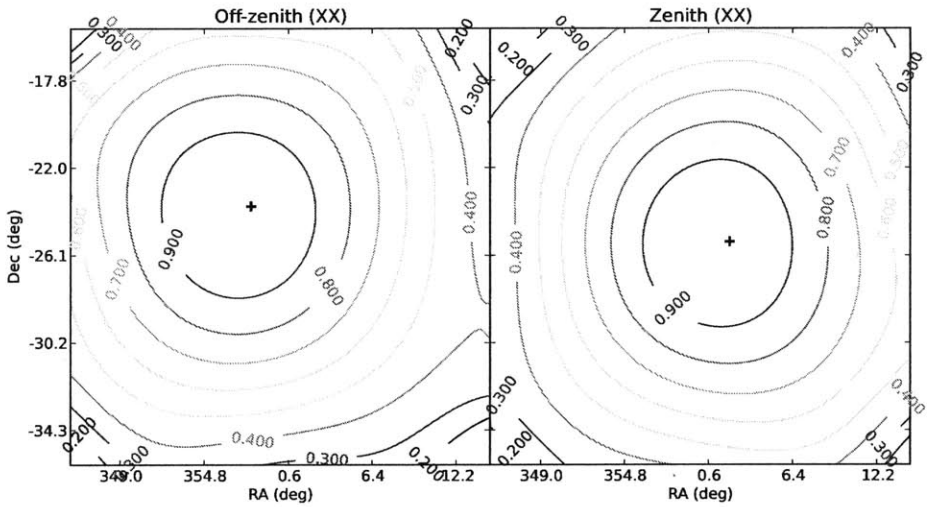

Figure 3-8: Example empirical fitted primary beam for different pointings at 182 MHz, both normalized to 1. The cross marks the pointing center. The spline fit produces a smooth beam, which is expected for the true beam. Only the sources within 13° of the phase center were considered for the fit, so we only show this region, beyond which the fit is an extrapolation and therefore unreliable.

Stokes I images corrected by the Curtin beam, using the same fitting procedure. Both procedures gave the same results, thus demonstrating that the fitting function was reliable.

3.5.2 Empirical Beam Performance

The empirical fitted beam behaves in a reasonable manner. Its maximum gain is around the pointing center, and it preserves the smoothness expected for the primary beam, both evident in Figure 3-8. Deviation from the Curtin beam shows that the empirical fitted beam has a $\sim$ 10% tilt from the phase center in the RA direction as shown in Figure 3-9, which is consistent with the light curve trends. This is also evident in the ratio between the empirical beam measurements and the Curtin beam, whereas the ratio between the data and the fit shows mostly statistical fluctuations, illustrated in Figure 3-10. Although we only show one example of the fit for an off-zenith pointing ($\sim$ 30° from zenith) and one for a zenith pointing, both XX polarization, others behave similarly.

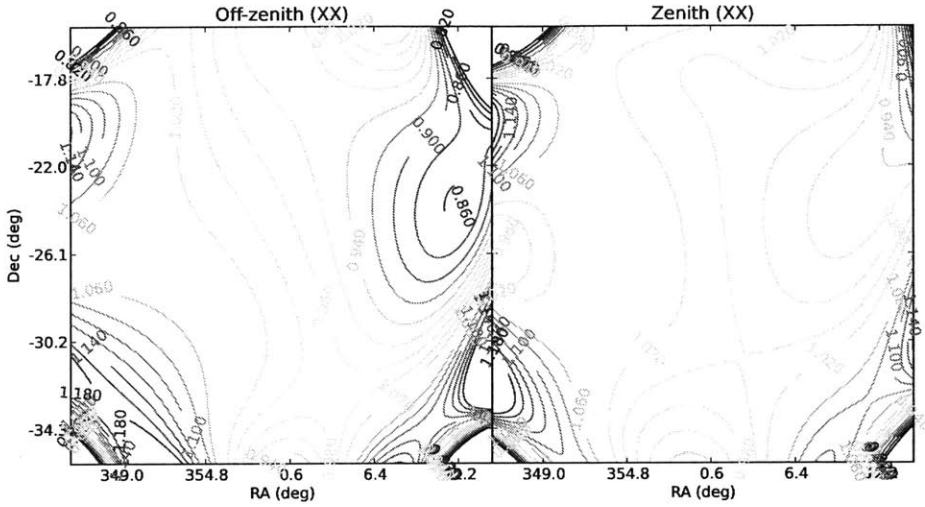

Figure 3-9: Ratio between the empirical fitted beam and the Curtin beam for the pointings in Figure 3-8, all normalized to 1. The difference between them is $\lesssim 5\%$ within $\sim 5°$ of the pointing center and increases to $\sim 10\%$ toward the edge of the beam. There is a tilt in the east-west direction that is likely related to the RA-dependent light curve slopes (see also Figure 3-10).

The empirical fitted beam also removes most of the systematic errors in the light curves. It reduces the light curve slopes from $\lesssim 4\%$ to $\lesssim 2\%$ flux density change per hour and recovers (by design) a more accurate source flux density as shown in Figure 3-11. However, there remain some residual systematic errors $\lesssim 2\%$ with a weaker RA dependence as shown in Figure 3-12, which could be due to calibration errors, imperfect fits, or systematic issues present in the reference source catalog, i.e. systematic primary beam issues could have already been present during the MWA commissioning survey.

Despite residual systematic errors, the empirical fitted beam is a big improvement over the Curtin beam. To compare more quantitatively the light curve variations from the two different primary beam corrections, we used the modulation index m as defined in [108]:

$$m = \frac{\sigma_S}{\bar{S}} \tag{3.3}$$

where σ_S is the standard deviation of the light curve and $\bar{S}$ is the mean flux density of the light curve. For constant sources, m is the inverse of the source signal-to-

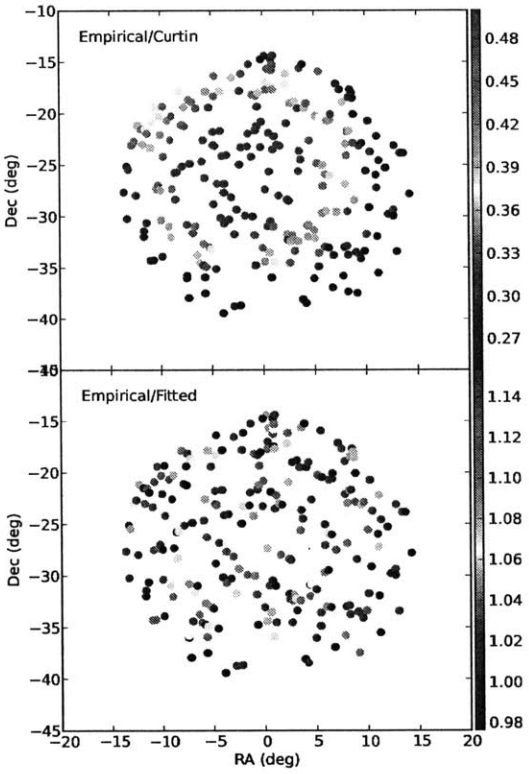

Figure 3-10: Ratio between the empirical beam measurements and the Curtin beam (top panel), as well as the ratio between the empirical beam measurements and the empirical fitted beam (bottom panel). This is the same off-zenith pointing as Figures 3-8 and 3-9. The panels have different color scales. The trend is prominent in the top panel and likely related to the light curve slopes (see Figure 3-5). This trend goes away for the fitted beam as the fit reproduces the main features in the measurements. The empirical beam measurements have *not* been normalized to 1 as they are the ratios between the measured flux densities and the catalog flux densities, so the normalization offset between the empirical beam and the Curtin beam is indicative of a problem with the calibration or the primary beam normalization.

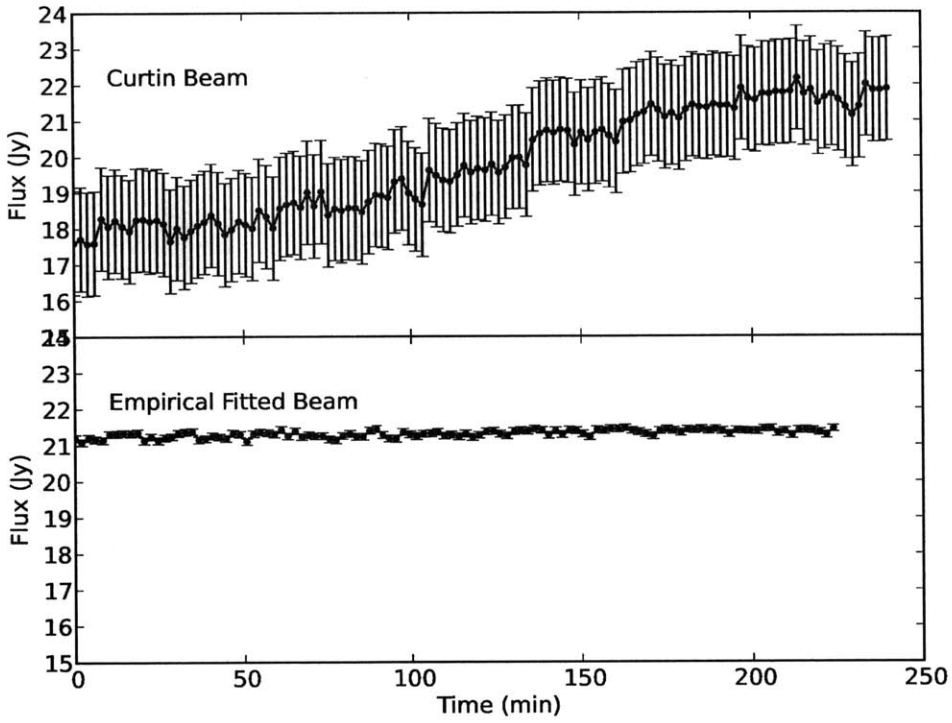

Figure 3-11: Example light curves for the same source but different primary beam corrections. The error bars are set to the standard deviation of the light curve. The empirical fitted beam removes most of the trend in the light curve and recovers a more accurate source flux density (21.74 Jy in the catalog).

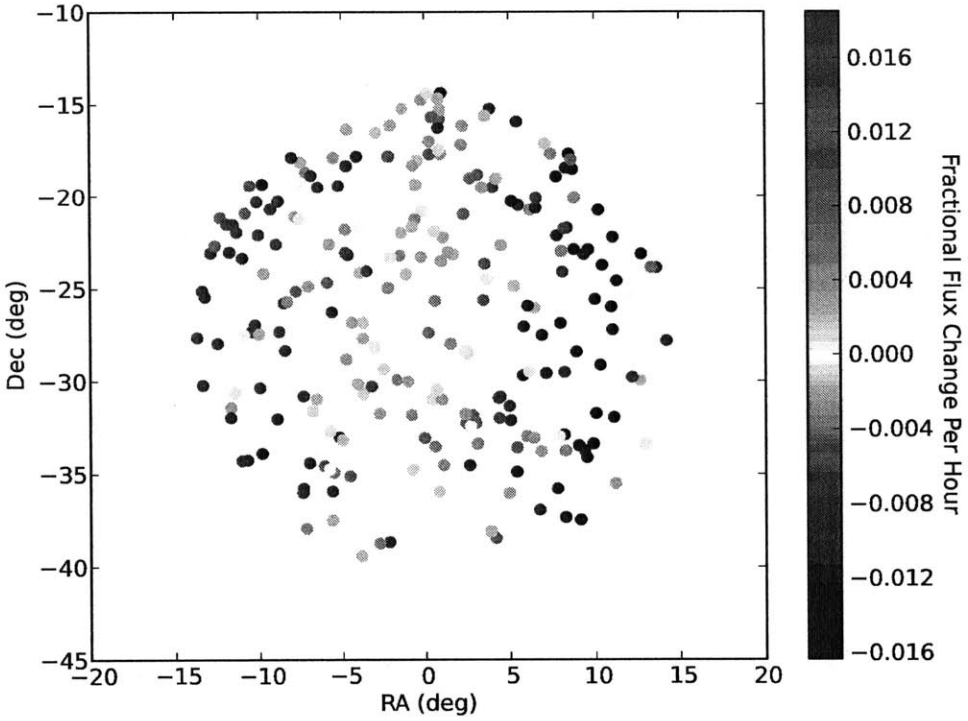

Figure 3-12: Residual light curve slopes after the empirical fitted beam correction at 182 MHz. While the fitted beam is not perfect, it is a big improvement over the Curtin beam, decreasing the slopes from $\lesssim 4\%$ per hour to $\lesssim 2\%$ per hour (compare with Figure 3-5). There could be remaining issues with the calibration, the fitting process, or the reference source catalog used to measure the empirical beam.

noise ratio (SNR) if σ_S is statistical, e.g. if $\bar{S} = 5\sigma$, then $m = 0.2$ from statistical uncertainties. Thus, m is a measure of the expected level of statistical fluctuations present in light curves of constant sources. Comparing m between the light curves produced from the empirical fitted beam and the Curtin beam, we found that the light curve variations after the empirical fitted beam correction were largely statistical, which is in contrast to the same light curves after the Curtin beam correction; this is illustrated in Figure 3-13. We computed SNR = peak flux/RMS with the values reported by AEGEAN; as mentioned previously in Section 3.5.1, the flux density errors reported by AEGEAN were too large, so we used the RMS (image noise over a 20×20 synthesized beam area) instead, which we considered to be a conservative estimate of the flux density uncertainties but more accurate than the reported flux density errors.

The differences between the Curtin beam and the empirical beam, besides theoretical modeling inaccuracies, could be due to the following factors: different tiles could have different malfunctioning dipoles or beamforming errors at any given time; the tiles are located on a hill-like surface [103], so each tile is pointed at a slightly different location in the sky; the model primary beam is generated at an instant in time (the middle of the observation) rather than integrated over 112-s (the duration of each snapshot observation). Investigating which of these has the largest effect is beyond the scope of this work. The empirical primary beam correction is sufficient for our purposes, although future improvements would include fitting for a single empirical primary beam model for each pointing as opposed to the current snapshot-by-snapshot correction, using a more accurate source catalog, and further improving on the calibration method as it is more accurate to average the calibration solutions only within the same pointing rather than over all the different pointings.

3.6 Source Subtraction

In preparing for the transient search, we also explored the option of using dirty images because clean artifacts had affected previous transient search results [56].

To reduce the amount of sidelobe confusion noise in the dirty images, we sub-

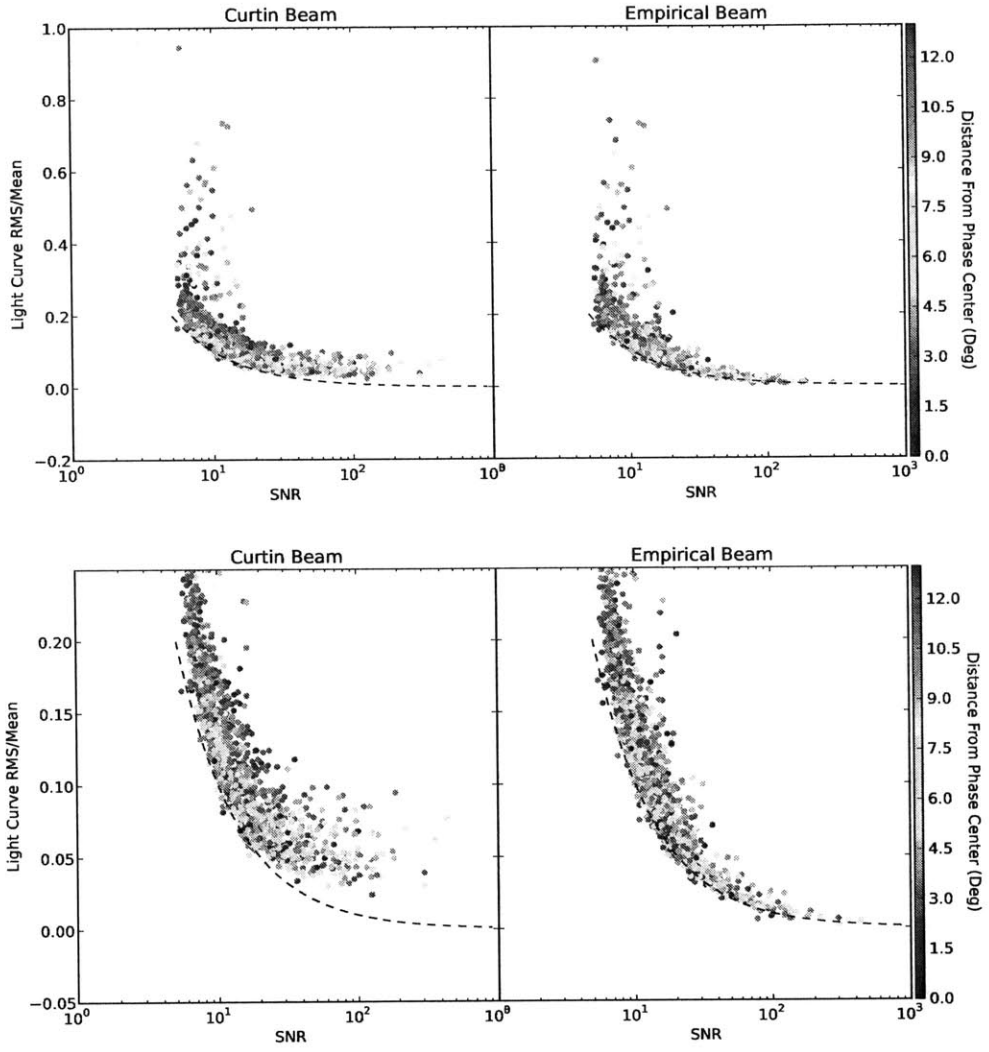

Figure 3-13: Light curve variation as a function of the mean SNR of each source. All sources, detected in 90% of the snapshots, are included, in addition to the subset used to measure the empirical primary beam; the bottom panel is a zoom-in of the top panel. The dashed line indicates the expected level of statistical fluctuations from flux density measurement uncertainties. See text for the details of the calculation. The light curve variations after the empirical beam correction are mostly statistical, in contrast to the systematic errors present in the light curves after the Curtin beam correction, showing the improvement of the empirical beam over the Curtin beam. The tail at low SNR is present for both versions of the primary beam correction (but absent for the subset of sources used to measure the empirical beam) and appears independent of the source location in the primary beam, suggesting that it is caused by systematic issues unrelated to the primary beam and requires more investigation. For example, it could be related to the source fitting algorithm.

tracted known sources from the calibrated visibilities. We subtracted the same sources from every snapshot observation to avoid introducing time-varying mistakes into the images. We used the CASA task **uvsub**, subtracting the model visibilities derived from a point source catalog of 5673 sources with flux densities > 100 mJy that we extracted by running AEGEAN on a 2.3-h deep integration image taken on 2013 September 4.

We only subtracted sources located within 13° from the phase center because the empirical primary beam was accurate to that radius, where the primary beam value was $\sim 30\%$ of its maximum gain beyond which was difficult for us to measure the primary beam empirically. Although the sidelobe noise from unsubtracted sources outside the 13° region was present in the image and the process of source subtraction was not perfect, the noise in the dirty image was drastically reduced from before and approached that of a clean image as shown in Figure 3-14.

We then passed these images and their corresponding primary beam images to *Simetra* for the transient search analysis.

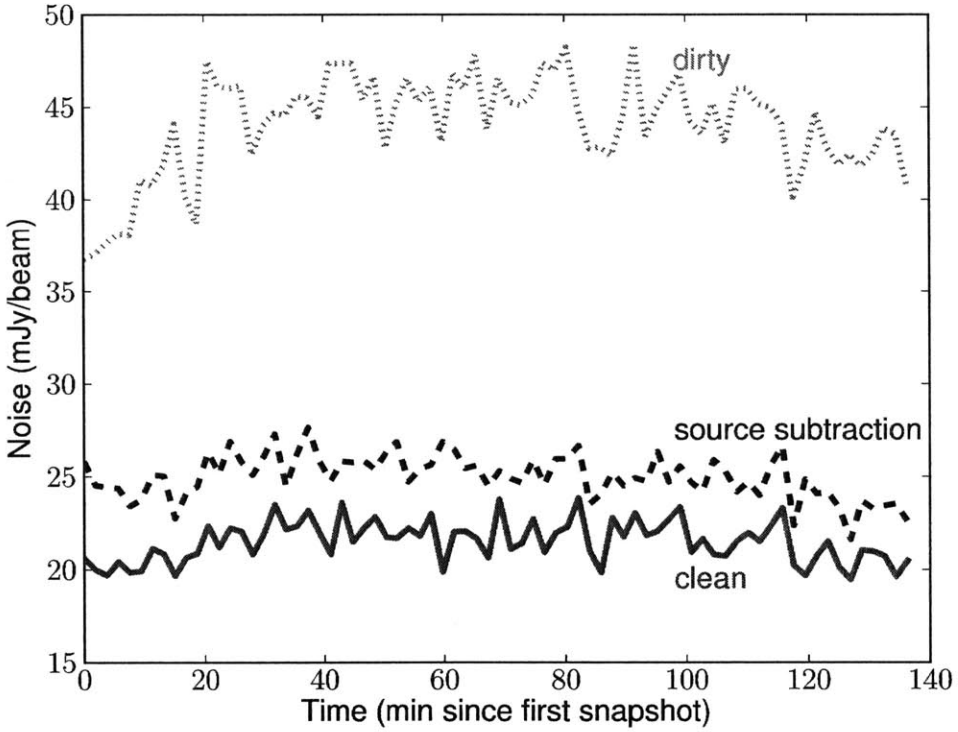

Figure 3-14: Image noise comparison: "dirty" without source subtraction (red dotted line), "dirty" with source subtraction (black dashed line), and "clean" (blue solid line). The noise is drastically reduced in the source-subtracted dirty images and approaches that of the clean images. The noise $\sigma_{\rm im}$ was estimated from the standard deviation as calculated from the median absolute deviation of the flux density values in the same central 86×86-pixel region of each 112-s snapshot. The data presented here are from 74 snapshots of EOR0, XX polarization, taken on 2013 September 4.

Chapter 4

Technique Demonstration

This chapter demonstrates how a transient search with our matched filter technique proceeds, and that it performs well on real data. First we characterize the background distribution for the transient search and the flux sensitivity. Then we estimate the efficiency at which we accurately recover injected transients to demonstrate that our technique is capable of detecting transients below the classical confusion noise.

4.1 Background Characterization

Before we identify possible transient signals according to the matched filter detection statistic, we characterize the background distribution of $\tilde{\rho} \equiv \max(\rho/\sigma_\rho)$. Using $\tilde{\rho}$ instead of ρ ensures that all pixels are drawn from the same distribution. Characterizing the background sets the threshold $\tilde{\rho}^*$ for an event to be classified as a transient candidate and determines the significance of such a detection.

We used a small area ($\sim 10\%$) labeled as the "playground region" in the images for the background characterization. In this region, we assumed that there were no transient events; existing limits on the rate of radio transients at low frequencies suggest that these events are relatively rare (e.g. [109, 110]), so our assumption should be valid. Significant transients would appear as a tail in the $\tilde{\rho}$ distribution that we would examine further.

To demonstrate our transient detection technique, we present the results for one

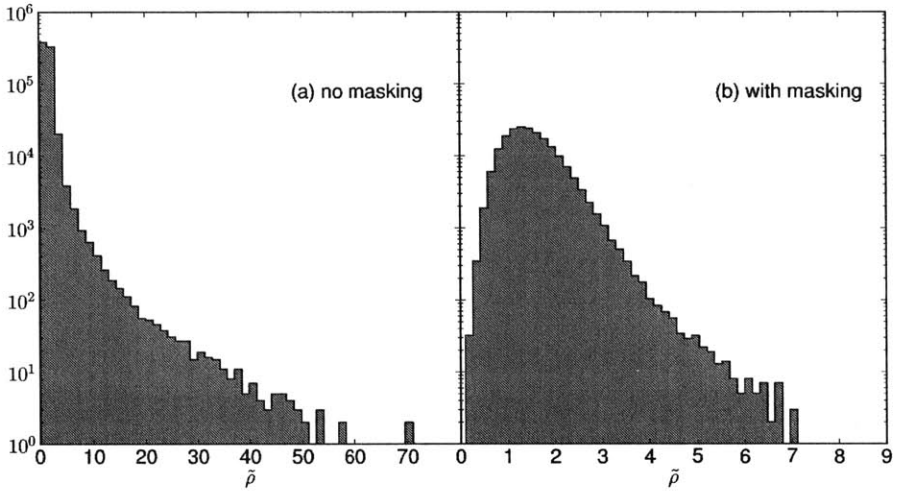

Figure 4-1: (a) Distribution of $\tilde{\rho}$ for all pixels in the playground region of the source-subtracted dirty images. (b) The same distribution of $\tilde{\rho}$ but without the "bad" pixels, i.e. the pixels in the region where the primary beam was poorly modeled and the pixels that contained the subtracted bright sources. The tail in the background distribution was removed by masking bad pixels.

light curve template: the top-hat with a duration of 15 days and a flux density of 1 Jy, searched over different start times t_0. In principle, we could treat the observation time of every 112-s snapshot as a unique t_0, but since computation time scaled with the number of search parameters, we shifted t_0 by $\sim 10\%$ of the duration where there were data, which in this case corresponded to the start of every night of the observation. The distribution of $\tilde{\rho}$ for the entire playground region is presented in Figure 4-1.

As evident in Figure 4-1a, there is a significant tail in the $\tilde{\rho}$ distribution. This is due to residual primary beam effects as the pixels in the tail were located in the regions where the primary beam was poorly modeled or that contained the subtracted bright sources. We masked the pixels beyond $\sim 10^\circ$ from the phase center where the primary beam was poorly modeled and excluded the source pixels by creating a 20×20-pixel square mask centered on the (RA, Dec) of each source. We also masked regions where the image RMS was $\gtrsim 15$ mJy/beam in a 2.3-h integration image. After masking, we removed the tail in the distribution as illustrated in Figure 4-1b. The distribution after masking is what we work with and refer to as the "background distribution."

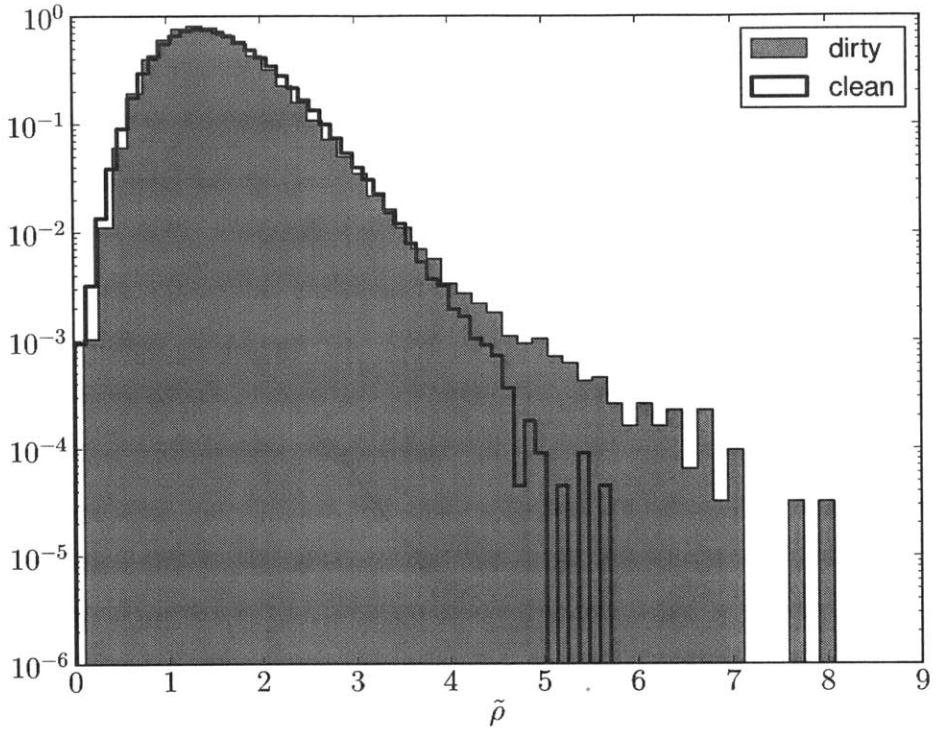

Figure 4-2: Background distribution of $\tilde{\rho}$ for clean images compared to source-subtracted dirty images. The same masking of bad pixels applies. As the clean images have lower noise, they are more sensitive to fainter transients. We report the results for the clean images in the rest of this thesis.

In Figure 4-2, we compare the background distributions for source-subtracted dirty images and clean images. Because of their lower noise properties, clean images are more sensitive to fainter transients. In the rest of this thesis, we work with the clean images with pixel masking.

The background distribution of $\tilde{\rho}$ lets us derive the probability of false alarm P_{FA} (equivalent to reliability) that sets $\tilde{\rho}^*$. Here P_{FA} is the probability that our experiment contains a false positive. Every experiment has its own observation timescale and sky coverage, which need to be factored accordingly. We started with a quantity closely related to P_{FA}: the number of background detections $N(\geq \tilde{\rho})$ above a particular threshold. We scaled $N(\geq \tilde{\rho})$ as derived from the playground region to the search region according to the number of synthesized beams for the different sky areas, thus

accounting for the trials factor correctly. Then we fitted a smooth function to the tail of $N(\geq \tilde{\rho})$ to extrapolate the background rate for larger values of $\tilde{\rho}$ where we might not have measurements. Non-Gaussianity and non-thermal components of the image noise, such as sidelobe confusion and image artifacts, are modeled in this empirical fit. We determined the fit parameters by minimizing the negative log-likelihood for a Poisson distribution, where $N(\geq \tilde{\rho})$ was treated as the Poisson mean. As illustrated in Figure 4-3, an exponential function of the form $f_N(\tilde{\rho}) = \hat{N}\exp(-\tilde{\rho}/\hat{\rho})$, where $\hat{N}$ and $\hat{\rho}$ are the fit parameters, fits the tail of $N(\geq \tilde{\rho})$ well within the errors $\sigma_N \equiv \sqrt{N}$. The probability for $N(\geq \tilde{\rho})$ to be non-zero, assuming a Poisson distribution, is $P(N > 0) = 1 - P(N = 0) = 1 - e^{-N}$. We took $P_{FA} = P(N > 0)$, and since we required $P_{FA} \ll 1$, $P_{FA} = P(N > 0) \approx N(\geq \tilde{\rho})$. Having determined $f_N(\tilde{\rho})$, we chose a tolerable value of P_{FA} (for example $P_{FA} = 10^{-3}$) and solved for $\tilde{\rho}^*$ such that $P_{FA} = f_N(\tilde{\rho}^*) = N(\geq \tilde{\rho}^*)$. For our exponential function, the threshold is of the form

$$\tilde{\rho}^* = \hat{\rho}\ (\log \hat{N} - \log P_{FA}) \tag{4.1}$$

The threshold $\tilde{\rho}^*$ depends on how much of the distribution tail is used in the fit. For Figure 4-3, we fitted the tail 500 points, which gave $\tilde{\rho}^* = 7.98$ at $P_{FA} = 10^{-3}$ and a reduced $\chi^2 = 0.25$ as computed from $\chi^2 = \sum_{i=1}^{500}[N(\geq \tilde{\rho}_i) - f_N(\tilde{\rho}_i)]^2/f_N(\tilde{\rho}_i)$ where the number of degrees of freedom is $(500-2)$ for 2 fit parameters. If we fitted 100 points, we obtained $\tilde{\rho}^* = 7.37$ at $P_{FA} = 10^{-3}$ and a reduced $\chi^2 = 0.18$; if we fitted 50 points, we obtained $\tilde{\rho}^* = 7.26$ at $P_{FA} = 10^{-3}$ and a reduced $\chi^2 = 0.31$. This dependence implies that a single threshold value is not very robust for identifying transient events near the threshold, so in the actual search, we opt for the "loudest event statistic" [111, 112], which we will describe in more detail in Chapter 5. Nevertheless, the threshold is useful for comparing how well different searches perform, for providing an estimate of the flux sensitivity, and for verifying the recovery of injected transients.

The flux sensitivity of the search can be calculated according to $A^* = \tilde{\rho}^*/\sigma_\rho$ (see Chapter 2; note $\tilde{\rho}^*$ already contains a factor of σ_ρ unlike ρ). Strictly speaking, A^* is unitless, so one needs to multiply by the template $\mathbf{f}$ to convert into flux density units,

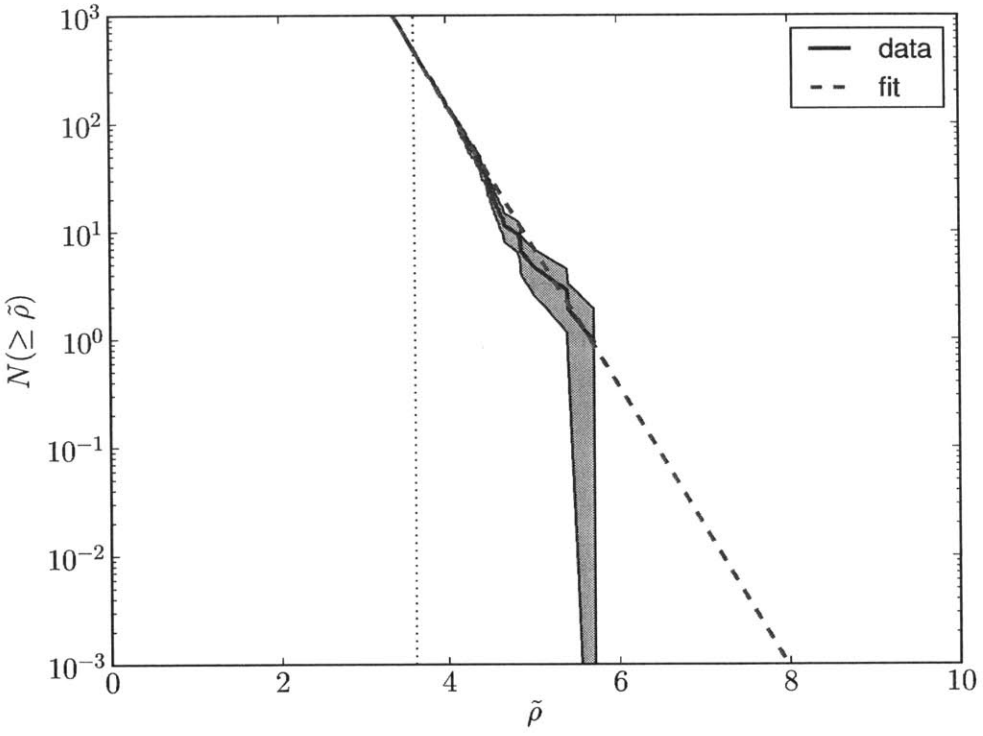

Figure 4-3: Cumulative background distribution of $\tilde{\rho}$ (black solid line). The shaded area is the error region as computed from $\sqrt{N}$. Only the tail of the distribution (500 data points to the right of the vertical dotted line) was used for fitting. The fit (blue dashed line) is an exponential function: $f_N(\tilde{\rho}) = \hat{N}\exp(-\tilde{\rho}/\hat{\rho}) = 2.38 \times 10^7 \exp(-\tilde{\rho}/0.33)$. The reduced χ^2 for the fit is 0.25; when we used instead 50 data points or 100 data points for the fitting, the reduced χ^2 became 0.31 and 0.18 respectively. The fit was then used to compute the false alarm probability P_{FA}; see Section 4.1 for the details of the calculation. For $P_{FA} = 10^{-3}$, the transient detection threshold is $\tilde{\rho}^* = 7.98$ in this case.

but we drop $\mathbf{f}$ for a simpler notation as we choose $\mathbf{f}$ to have units of 1 mJy. Because σ_ρ depends on the primary beam, as presented in Equations 2.6 and 2.10, pixels closer to the edge of the primary beam have different values of σ_ρ compared to the pixels closer to the center of the primary beam. This implies that we have non-uniform flux sensitivity across the image, where we are more sensitive to fainter transient sources toward the center of the primary beam. This is illustrated in Figure 4-4, where we also show an example detection threshold $6\sigma_{\rm im}$ for a source finder.

To report a single number for the flux sensitivity, we computed A^* separately for each pixel, then we computed the median value of A^*. Since the distribution of σ_ρ and hence A^* is skewed, the median of A^* is a better estimator of the flux sensitivity than the mean. However, we stress that this value only provides an estimate of how well the search might perform and should not be interpreted as a strict flux threshold because a single value for the flux sensitivity is a convenience and cannot capture the complexity of the data.

It is $\tilde{\rho}^*$ that matters in this technique. Since our technique uses $\tilde{\rho}^*$ and not flux density as a metric to identify transient sources, it is capable of detecting sources fainter than the median flux sensitivity at the *same significance* (reliability) but a lower efficiency (completeness), making it potentially more powerful than the techniques that apply a more stringent flux threshold. For our clean images and a 15-day top-hat template, $P_{FA} = 10^{-3}$ corresponds to $\tilde{\rho}^* = 7.98$ and a median flux sensitivity of 25.0 mJy; note that if we had fitted 50 points instead of 500, $\tilde{\rho}^* = 7.26$ corresponds to a median flux sensitivity of 22.7 mJy, which is not very different from 25.0 mJy. By contrast, the source-subtracted dirty images have a median flux sensitivity of 52.4 mJy or $\tilde{\rho}^* = 13.4$ for the same P_{FA}.

4.2 Transient Injection

Having chosen the threshold $\tilde{\rho}^*$ for transient detection, we characterize the accuracy and efficiency (or completeness) at which we recover transient events. To do so, we injected transients with known light curve parameters into the data and processed

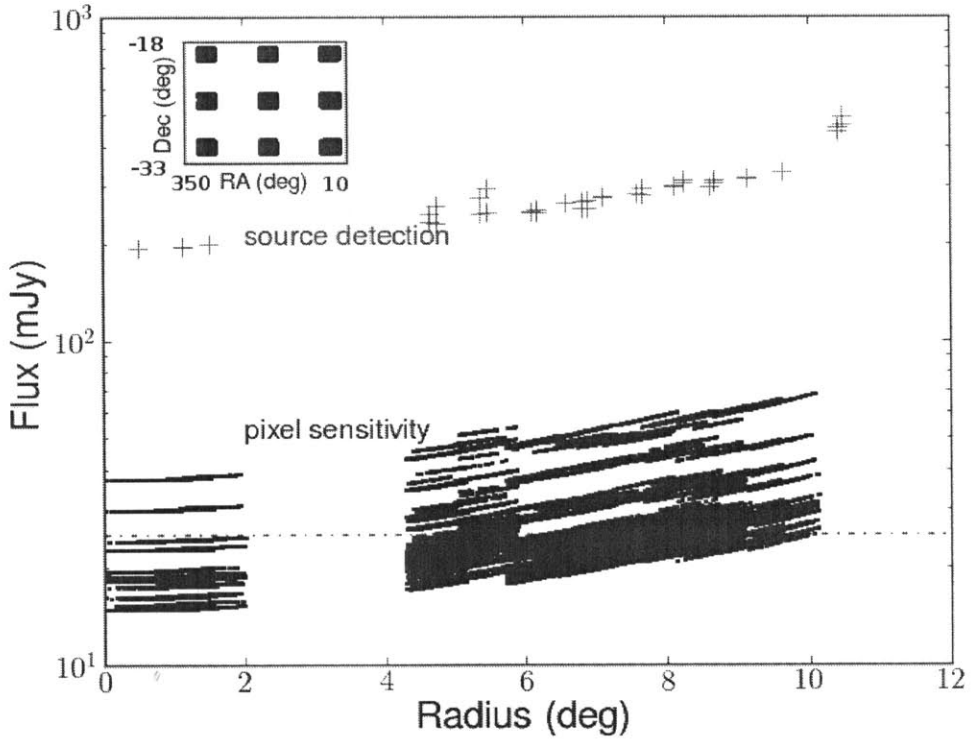

Figure 4-4: Flux sensitivity comparison between the matched filter transient detection technique (using a 15-day top-hat template) and a source detection technique (e.g. [110]). Each black dot is $A^* = (\tilde{\rho}^*/\sigma_\rho) = (7.98/\sigma_\rho)$ of a pixel and each blue cross is $6\sigma_{\rm im}$ (an example source detection threshold) at a particular radius from the EOR0 field center. The dash-dot line is the median flux sensitivity 25.0 mJy quoted in the text for the matched filter technique. The matched filter technique achieves a better sensitivity than a source-finding algorithm for the same set of images. The sensitivity decreases away from the phase center because of increased noise, but the significance remains the same. The gap between 2° and 4° is because of the discontinuous playground region sampling; see the inset where black rectangles mark the playground regions, each of which consists of several smaller 86×86-pixel patches. $\sigma_{\rm im}$ is calculated independently for each patch, so the stripe-like features in A^* correspond to the different patches where noise properties are different.

these light curves through the *Simetra* pipeline. The pixels in which injection occurred defined the "injection region," which consisted of 10^4 randomly sampled pixels across the image. The injected light curves might be different from the search template, but the better matched they are, the better the transients are recovered.

For simplicity and as a proof-of-concept, we injected transients with the same shape as our search template, i.e. top-hat with a duration of 15 days, and varied the flux densities (amplitudes) from 10 mJy to 100 mJy in increments of 10 mJy. Because of observational gaps in the data, we chose a start time for the injected transients such that the start and end times of the transients were present in the available data, though there could be gaps in the middle. Depending on how one schedules observations, there could be situations where a 2-week transient appears indistinguishable from, say, a 1-week transient because of missing data, but this type of transient will be identified as a 1-week transient regardless of the transient search technique. This impacts the estimation of "true" transient occurrence rates but not the performance of our pipeline.

As mentioned in Section 4.1, the pipeline searched over all possible start times for the best match and estimated the amplitude of the transient. Then we identified pixels with $\tilde{\rho} \geq \tilde{\rho}^*$ as pixels that contained 15-day-long transients. Figure 4-5 illustrates the identification of one such injected transient and the power of this technique. Figure 4-6 compares the recovered start times and amplitudes to the injected values for these pixels, demonstrating that the pipeline is recovering injected transients accurately. Figure 4-7 shows the efficiency at which the injected transients are recovered, and the flux sensitivity is consistent with our expectation.

Real transient light curves will differ from the top-hat shape and from each other. For example, radio supernova light curves are qualitatively different from extreme scattering events [51]. Our technique does not require an exact match between the search template and the observed light curve to detect a transient. Figure 4-8 shows the results when there is a mismatch, where the search template is 15 days and the injected light curve ranges from 1–15 days depending on the start time. Compared to a perfect match, a transient detected from a mismatch will be less significant, less

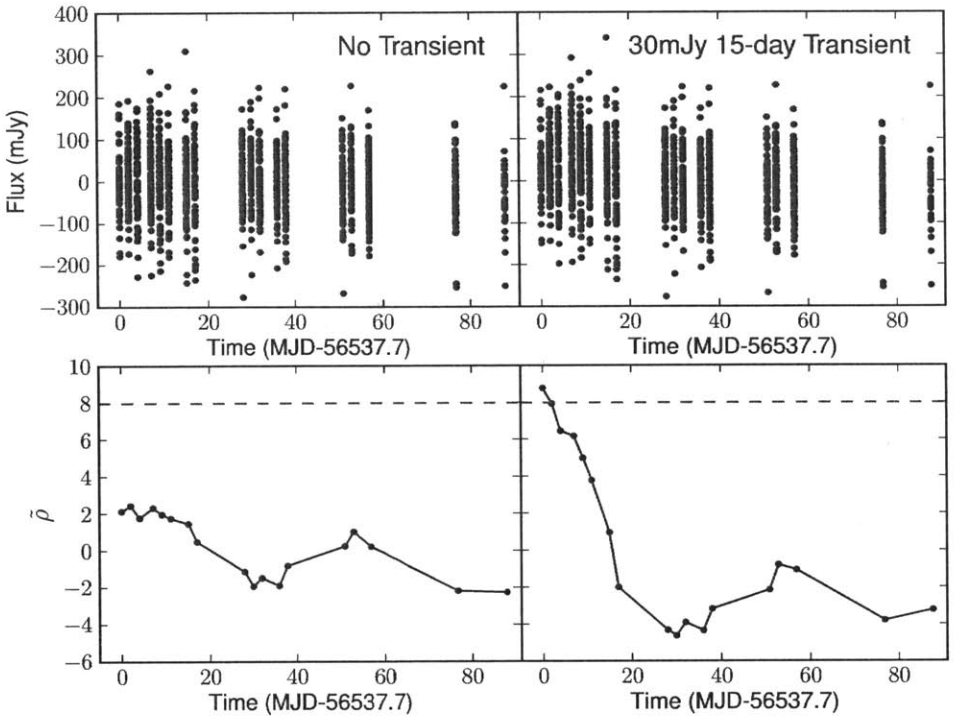

Figure 4-5: Identification of an injected transient. The top left and right panels compare the light curves of the same pixel before and after a 30-mJy, 15-day transient was injected starting on the first night of observation. The bottom left and right panels compare $\tilde{\rho}$ as a function of the start times searched for a 15-day template. The blue dashed line is $\tilde{\rho}^*$, above which the event is classified as a transient candidate. Even though the injected transient is faint and difficult to identify by eye, it is easy to identify using our technique as $\max(\tilde{\rho}) > \tilde{\rho}^*$.

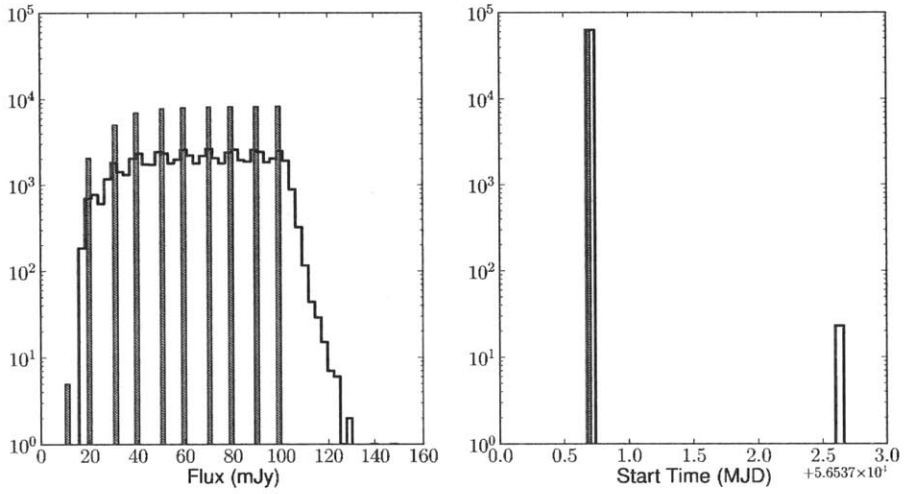

Figure 4-6: Distribution of recovered light curve parameters compared to their injected values for a top-hat template with a width of 15 days: flux density (left) and start time (right) for events with $\tilde{\rho} \geq \tilde{\rho}^* = 7.98$. The injected values are the shaded histograms. The recovered flux densities are distributed about the injected values. There are a few events with recovered start times that are ~ 2 days after the injected start time, which is the second night of observation (there are no observations between those two start times). The recovered values are in good agreement with the injected values.

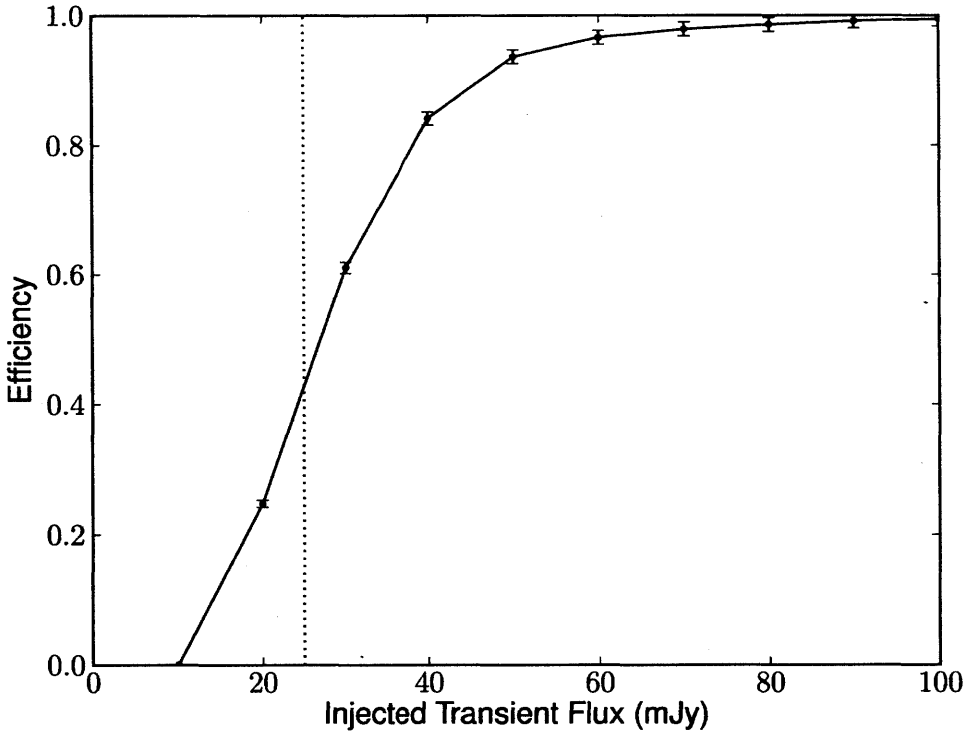

Figure 4-7: Efficiency (completeness) at which the injected 15-day transients are recovered. These are significant events that have been identified as transients according to $\tilde{\rho} \geq \tilde{\rho}^* = 7.98$ for $P_{FA} \leq 10^{-3}$. The median flux sensitivity is indicated by the vertical dotted line at 25.0 mJy. Transients below the median flux sensitivity can still be identified with the same significance but at a lower efficiency.

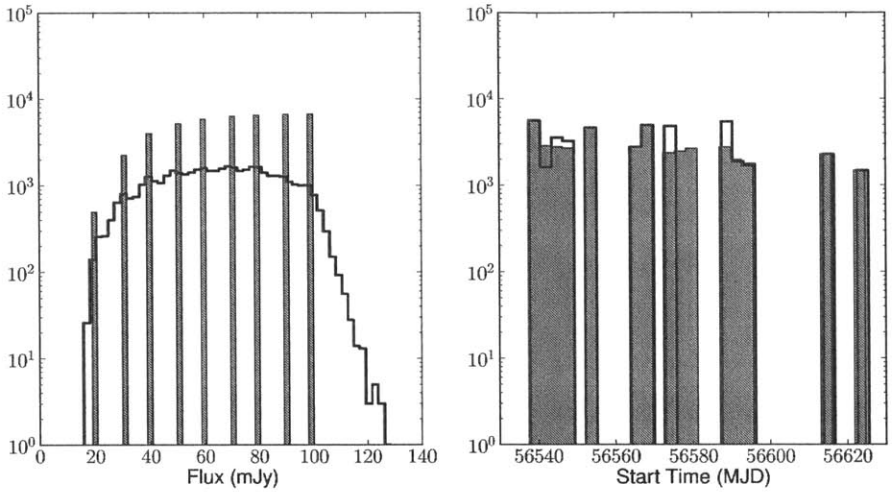

Figure 4-8: Distribution of recovered light curve parameters compared to their injected values when there is a mismatch in the template duration between injection and search: flux density (left) and start time (right) for events with $\tilde{\rho} \geq \tilde{\rho}^* = 7.98$. The injected values are the shaded histograms. The recovered flux densities are distributed about the injected values but with more spread compared to Figure 4-6. Start times are recovered relatively accurately.

accurate, and with less efficiency, and the amount by which it differs from a perfect match will depend on the nature of the mismatch. One could, however, rerun the analysis with more refined templates on the pixels with transient detections in order to characterize the transient properties more accurately.

Nonetheless, our technique treats all templates equally. There is nothing special about the top-hat template we have chosen for this analysis except that it is the simplest to implement. One could repeat the analysis with supernova light curves, for example, which will identify and constrain supernova-like transient events but not transients that are very unlike supernovae. To expand the search, one simply includes other types of templates. The sensitivity to different types of transients will be different even for the same dataset, but this is due to the nature of the data, not the technique, e.g. certain systematic errors might affect one timescale more than another.

4.3 Computation Time

The pipeline is highly parallelizable and not limited by computation as the matched filter calculation for each pixel is independent. I/O, i.e. converting the image FITS files into light curves stored as `numpy npz` files, is the bottleneck, but it only needs to be done once. Afterwards, the maximization over different start times is the slowest step and scales with the number of time samples. For $\mathcal{O}(10^3)$ time samples, the computation time per 86×86 pixels, which was the smallest image unit we processed, for the entire pipeline was $\mathcal{O}(\text{min})$. We ran this process on 1 CPU core with 3 GB of allocated RAM on a computing cluster that comprised 14 computers with 99 GB of RAM and 24 CPU cores per machine, each core operating at a speed of 60 MFLOPS. Running the pipeline for the full search of this thesis, for example, can be completed in < 1 week with 36 CPU cores.

Chapter 5

Limits on Slow Radio Transients at 182 MHz

After demonstrating that the matched filter technique worked as expected, we ran three separate blind searches on the MWA data. Each search used a different set of light curve templates and thus was sensitive to transients on a different timescale.

5.1 Defining the Searches

The timescales spanned by the data ranged from 2 min to 3 months, but we did not have uniform sensitivity over all possible timescales. Some timescales could not be probed because of gaps in the observations, while other particular timescales suffered from systematic effects, such as those due to the periodic change in the primary beam pointing.

In order to determine which searches to run and which templates to use, we did a test run on the playground region to compare the expected transient thresholds for various timescales. Figure 5-1 shows the expected thresholds from the test run; we sampled roughly logarithmically between 2 min and 3 months but also sampled every ~ 5 min between 15 and 50 min. There is a drastic increase in threshold around the 30-min timescale, which corresponds to the time between consecutive changes in the primary beam pointing. Hence we excluded that timescale from our search and

Table 5.1: Summary of our searches and their respective injection runs. $\tilde{\rho}^*$ is the threshold with $P_{FA} = 10^{-3}$. For the injection parameters, we sampled uniformly in durations and peak fluxes as well as start times (not listed) that spanned the entire 3 months of observation. For the FRED profile, τ_1 is the characteristic rise time, and τ_2 is the characteristic decay time. See Section 5.4 for a discussion on the revised values for $\tilde{\rho}^*$.

Search	Template	Duration	$\tilde{\rho}^*$ [revised]
1: minute	top-hat	4 m	7.7 [7.6]
2: hour	top-hat	1.5 h	9.7 [7.9]
3: day-to-month	top-hat	2 d, 4 d, 7 d, 9 d 11 d, 15 d, 17 d, 28 d 30 d, 32 d, 36 d, 38 d 51 d, 53 d, 57 d, 77 d, 88 d	9.1 [8.0]
Injection	**Template**	**Duration**	**Peak Flux**
1: minute	top-hat	2–12 m	< 1300 mJy
2: hour	top-hat	1–2 h	< 450 mJy
3: day-to-month	top-hat	1–90 d	< 150 mJy
	FRED	$\tau_1 = 1$–2 d, $\tau_2 = 30$–40 d	< 160 mJy

defined the following three searches: minute, hour, and day-to-months. Despite the variation in $\tilde{\rho}^*$ values for the day-to-month templates, the median flux sensitivities were about the same, so we grouped them together. The specific templates we used for the searches are listed in Table 5.1. Because of the primary beam systematics, we avoided templates with durations between 15–60 min. The 4-min template is capable of recovering transients with durations < 15 min as the difference between ρ computed from the 4-min template and ρ computed from the perfectly matched template is $< 10\%$. As each night consisted of ~ 2 hours of observation, we chose the 1.5-hour template that is halfway between 1 and 2 hours for the second search. Finally, for the last search, we considered all possible durations sampled by the data, excluding the gaps, and separated by at least 1-day.

Each search consisted of three parts: (1) setting the threshold $\tilde{\rho}^*$ to identify transient candidates, (2) characterizing the efficiency at which the transient events are identified, and (3) identifying the transient candidates.

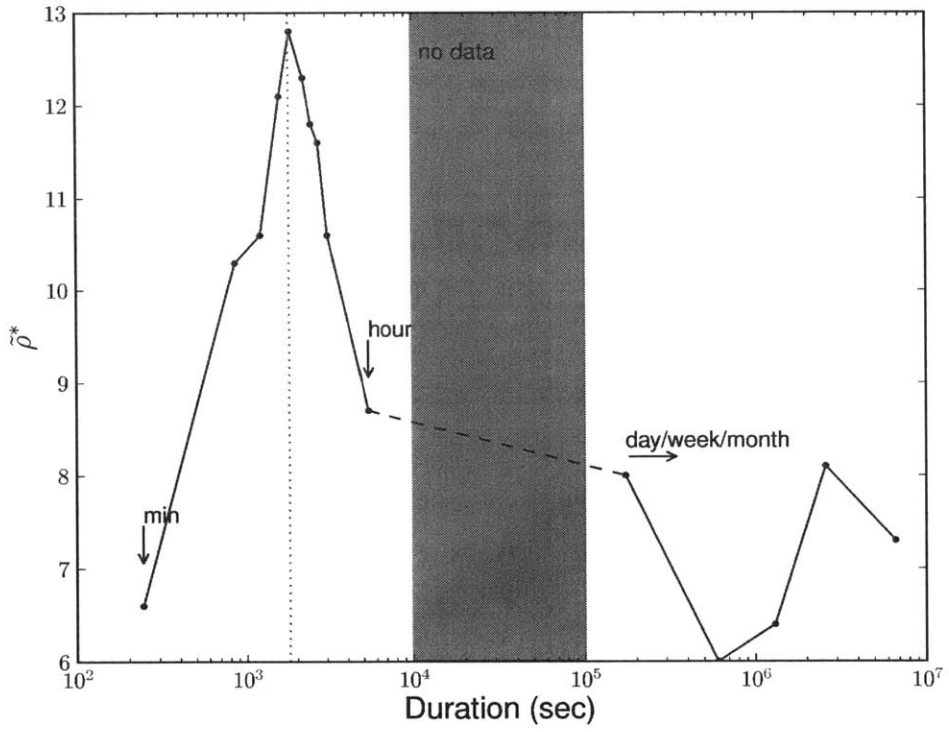

Figure 5-1: Expected thresholds for transients on different timescales, characterized by the matched filter statistic $\tilde{\rho}^*$, which depends on the image noise, the light curve template, and how the observations were taken. Smaller values of $\tilde{\rho}^*$ correspond to better sensitivities. The vertical blue dotted line marks the timescale for consecutive changes in the primary beam pointing; as we have poor sensitivity on this timescale, we excluded it from our search. The vertical grey region marks the approximate timescales that correspond to the gaps between observations where we have no data. Based on this plot, we divided our transient search into three separate ones: min, hour, and day-to-months.

5.2 Setting the Threshold

The threshold $\tilde{\rho}^*$ was set by running the pipeline on the playground region. This region was $\sim 10\%$ of the image, which we assumed contained no transients. First we divided the inner $\sim 13°$ of the image, or 3096×3096 pixels, into 86×86-pixel squares as the pipeline, at any given time, ran on one CPU core allocated 3 GB of RAM and would encounter memory issues if it processed more than $\sim 100 \times 100$ pixels each with 1251 flux density measurements. Because of primary beam systematics, we only searched for transients in the inner $\sim 10°$, or 2096×2096 pixels, of the image. Then we chose the playground region to be nine 172×172-pixel patches divided into rows of three across the image. This choice sampled uniformly across the image to capture any spatial noise variation.

We further divided the playground region into two parts, A and B, by choosing alternating 86×86-pixel squares. Playground A was used to characterize the background distribution of $\tilde{\rho}$ and to set the threshold $\tilde{\rho}^*$ by extrapolating the tail of the distribution to our choice of false alarm probability P_{FA} (the reliability of detected candidates), while playground B was used to verify that extrapolation and the trials factor normalization. The extrapolation was done by fitting an exponential function to the tail of the cumulative distribution of $\tilde{\rho}$ as described in Section 4.1. For all three of our searches, we chose $P_{FA} = 10^{-3}$ to determine $\tilde{\rho}^*$, which meant that the probability of detecting a false positive in each search (experiment) is $\leq 10^{-3}$. If the distribution is Gaussian, which it is not, this probability corresponds to a significance of 3.3-σ. Figure 5-2 shows the extrapolation and verification for the three searches. While the normalization is set to the number of synthesized beams (independent pixels) in playground B for the verification process, the threshold was determined after normalizing the background distribution to the number of synthesized beams in the search region. The values of $\tilde{\rho}^*$ for the three searches are listed in Table 5.1.

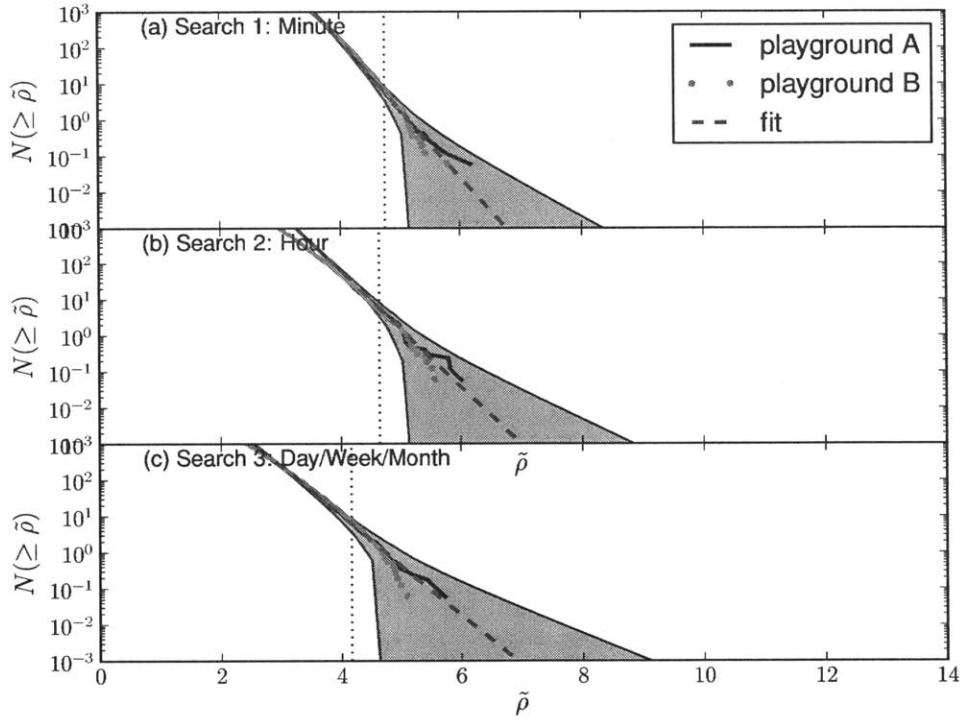

Figure 5-2: Cumulative background distribution of $\tilde{\rho}$ for the playground region for the three searches. The blue dashed line is an exponential fit to the tail of the playground A distribution, where the tail consists of 100 points to the right of the vertical dotted line. The yellow shaded region is the $\sqrt{N}$ error region for the fit, showing that the distributions from playground A and B agree within the error bars. All curves are normalized to the number of synthesized beams in playground B. The fit determines $\tilde{\rho}^*$ at a predetermined $P_{FA} = N(\geq \tilde{\rho})$ for $N(\geq \tilde{\rho}) \ll 1$ when it is normalized to the number of synthesized beams in the search region. See text for further discussion.

5.3 Characterizing the Efficiency

The efficiency (completeness) of each search characterizes the fraction of real transients successfully recovered and plays a role in the upper limit calculation of transient surface density. We determined the efficiency by running the same search on the injection region. This region consisted of 10^4 random pixels, and we injected transient light curves with known parameters into these pixels.

For all three searches, we injected transients with the top-hat profile, sampling uniformly in transient duration, start time, and flux density (amplitude). For the day-to-month search, we also injected transients with the fast-rise-exponential-decay (FRED) profile to mimic what real transients might look like, e.g. radio flares from X-ray binaries [51], sampling uniformly in characteristic rise and decay times, start time, as well as peak flux. All injection parameters are listed in Table 5.1.

After running the search on the injection region, we applied a cut on $\tilde{\rho}$, choosing only the events with $\tilde{\rho} \geq \tilde{\rho}^*$. These were the recovered transient events. By computing the ratio of the number of recovered events to the total number of injected events, we determined the efficiency as a function of flux density for each search, as shown in Figure 5-3. Since sensitivity improves with lower image noise or longer integration time, the day-to-month search is able to recover fainter transient sources at a higher efficiency than the minute or hour search.

Since we also knew the true injection parameters, we checked the accuracy at which the pipeline recovered these parameters, as shown in Figures 5-4 and 5-5. The pipeline is able to recover injected parameters fairly accurately ($\lesssim$ 10–30% 1-σ errors), even when the search is run with top-hat templates on injected transients with the FRED profile. This demonstrates that we are capable of detecting real transients in the data, if they are present, despite using simple top-hat templates for the search.

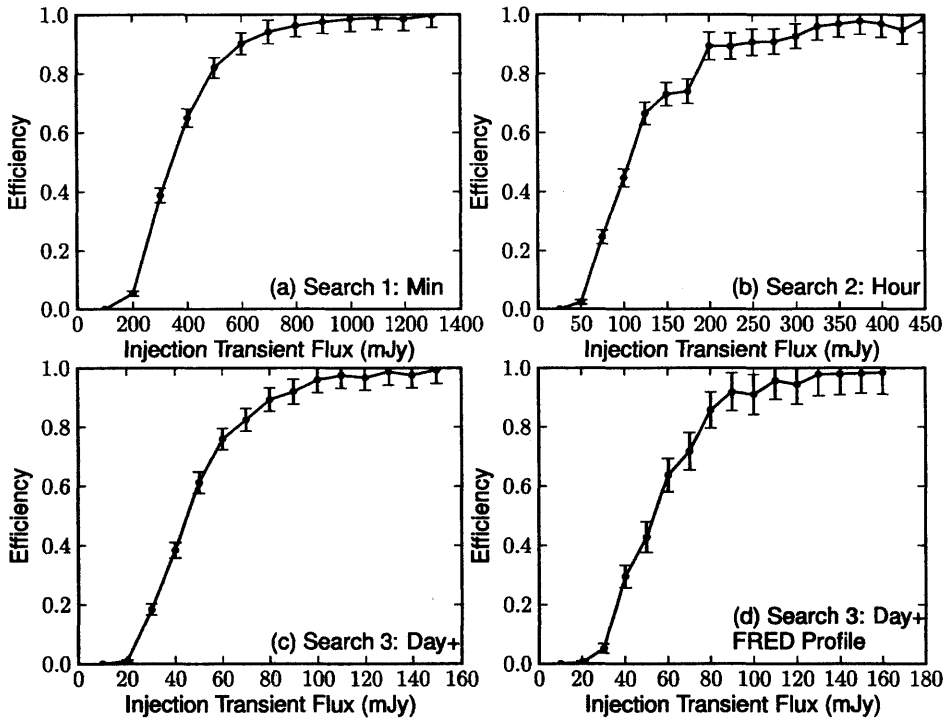

Figure 5-3: Efficiency (completeness) of recovered transients for the three searches. The efficiency increases at lower flux densities as the searched duration increases because longer durations imply longer integration times and lower image noise. Panel (d) shows the efficiency for when the transients were injected with the FRED profile, which is qualitatively different from the top-hat templates used in the search. It is not drastically different from the efficiency for when the transients were injected with the top-hat profile as in panel (c), demonstrating that the top-hat template is capable of recovering transients that are not top-hat in shape.

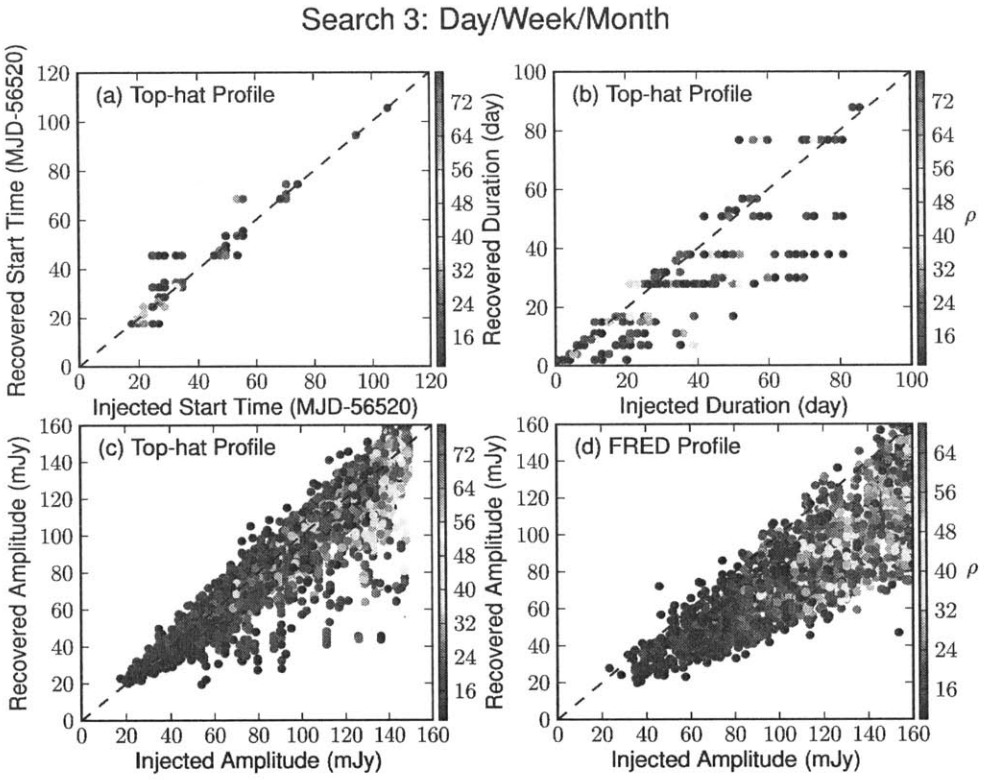

Figure 5-4: Recovery of transient parameters. The pipeline is able to recover injected transient parameters fairly accurately; see also Figure 5-5. There is a slightly bigger spread in the recovered durations, but that is because the search used a limited set of durations compared to the uniform sampling of injected durations. The flux densities recovered for the FRED profile injections are systematically lower than the injected peak fluxes, and that is a result of the qualitative difference between the injected and the searched light curve profiles; the FRED profile is sharper than the top-hat profile so the recovered flux density is smeared out. Although we only show this for the day-to-month search, the pipeline performs similarly for the other two searches.

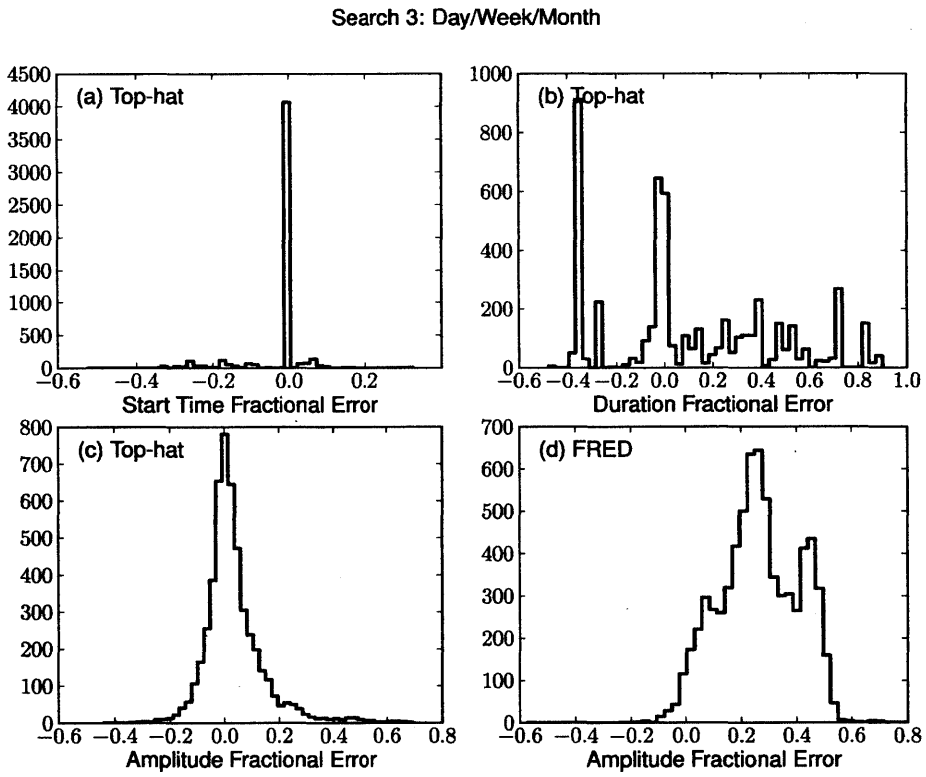

Figure 5-5: Accuracy of recovered transient parameters, corresponding to the panels in Figure 5-4. The accuracy of each light curve parameter p is quantified by its fractional error: $(p_{\rm inj} - p_{\rm rec})/p_{\rm inj}$ except for the start time, which is characterized by the start time difference relative to the injected duration: $(t_{0,\rm inj} - t_{0,\rm rec})/(t_{\rm dur,inj})$. The means of the distributions are (a) -1.9%, (b) 10.2%, (c) 3.3%, (d) 26.3%; and the standard deviations are (a) 7.8%, (b) 35.5%, (c) 11.6%, (d) 14.3%.

5.4 Identifying the Candidates

Finally, we ran the pipeline on the search region, which contained $\sim 3 \times 10^6$ image pixels or $\sim 2 \times 10^5$ synthesized beams, where one synthesized beam contains ~ 16 image pixels. We applied the cut $\tilde{\rho} \geq \tilde{\rho}^*$ on individual pixels to identify transient candidates. If multiple pixels passed the cut, we grouped the adjacent ones together and considered them as one candidate because the image is oversampled. Since we computed $\tilde{\rho}$ for every pixel, we can produce a "transient" sky map, as shown in Figure 5-6, where each pixel contains the corresponding $\tilde{\rho}$ instead of flux density. This map visualizes the variability on a particular timescale across the image, and a transient candidate would stand out as a source.

When we first ran the search, we found 4 candidates: 3 for the hour search and 1 for the day-to-month search. However, upon closer inspection, they appeared to be sidelobe artifacts because they were located near two bright sources (within $\sim 5'$–$10'$ of a 6 Jy source and a 11 Jy source). The $\tilde{\rho}$ distribution for the search region also deviated from the background expectation as shown in the top panels of Figures 5-8, 5-9, and 5-10. This led us to examine the source flux density distributions of the search region and the playground region, which we had used to tune and verify the fit parameters for the tail of the background distribution. As shown in Figure 5-7, the playground region did not contain any of the brightest sources in the field. This suggests that sampling uniformly across the field to select the playground region as we had done is not sufficient. Perhaps a better strategy is to define the region according to the source locations such that the source flux density distributions in both regions are similar, e.g. draw randomly from the locations of the bright sources and define boxed regions around them.

Instead of changing the playground region, however, we considered another approach to make our search more robust. In our initial search, we used the same source mask for all of the sources regardless of the brightness. However, brighter sources have larger sidelobes. One can consider the source mask as an "auxillary channel" for the search, where events that occur within a certain area of a (bright) source are vetoed.

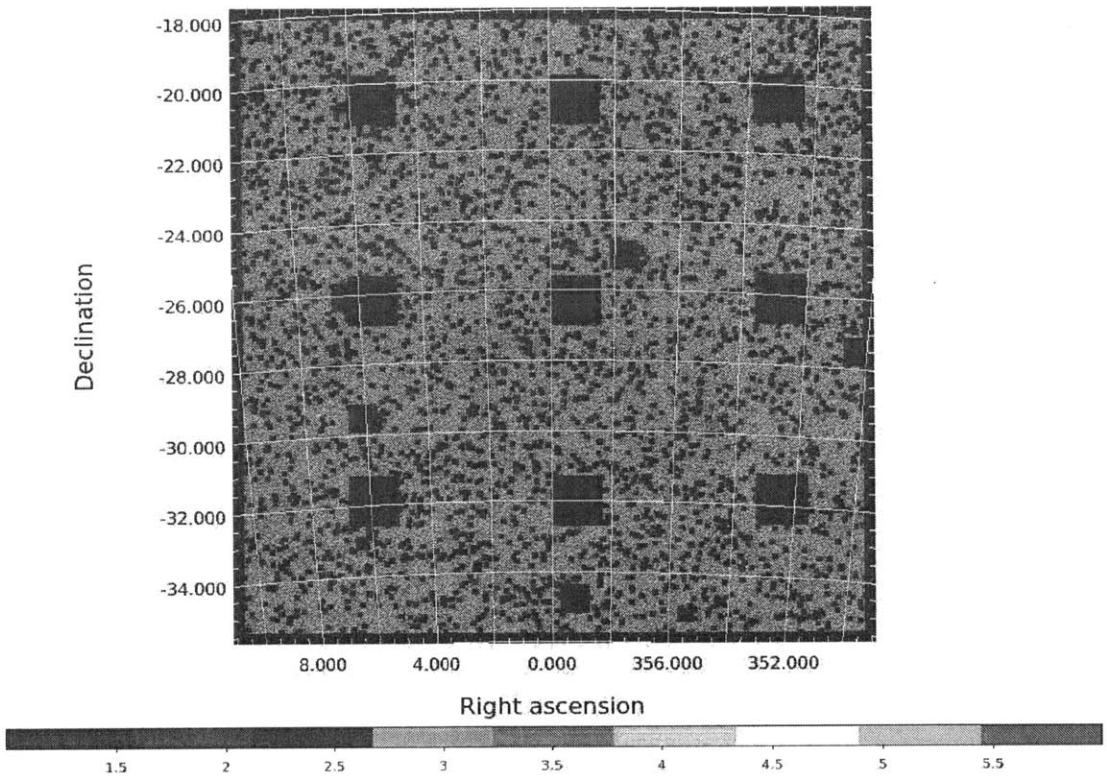

Figure 5-6: Example map of the transient sky, where the value of each pixel is $\tilde{\rho}$ instead of flux density. This is for the minute search, so it shows the variability across the image on the minute timescale. The largest nine black patches make up the playground region and are not part of the map. The other smaller black squares are masked pixels, regions excluded from the search because they surround radio sources with flux densities $> 100\,\mathrm{mJy}$ or have RMS $> 15\,\mathrm{mJy/beam}$.

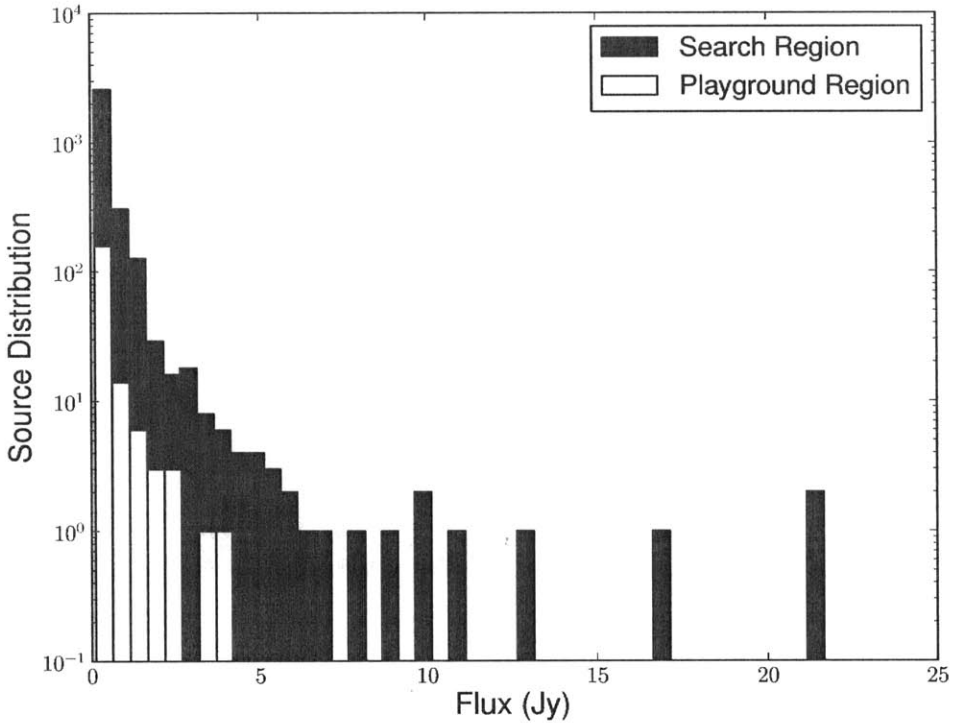

Figure 5-7: Source flux density distributions for the search and the playground regions. The playground region, which we naively chose to sample the image uniformly, did not contain the brightest sources in the field. As a result, the sidelobe noise properties near the bright sources were not well-captured by the playground region.

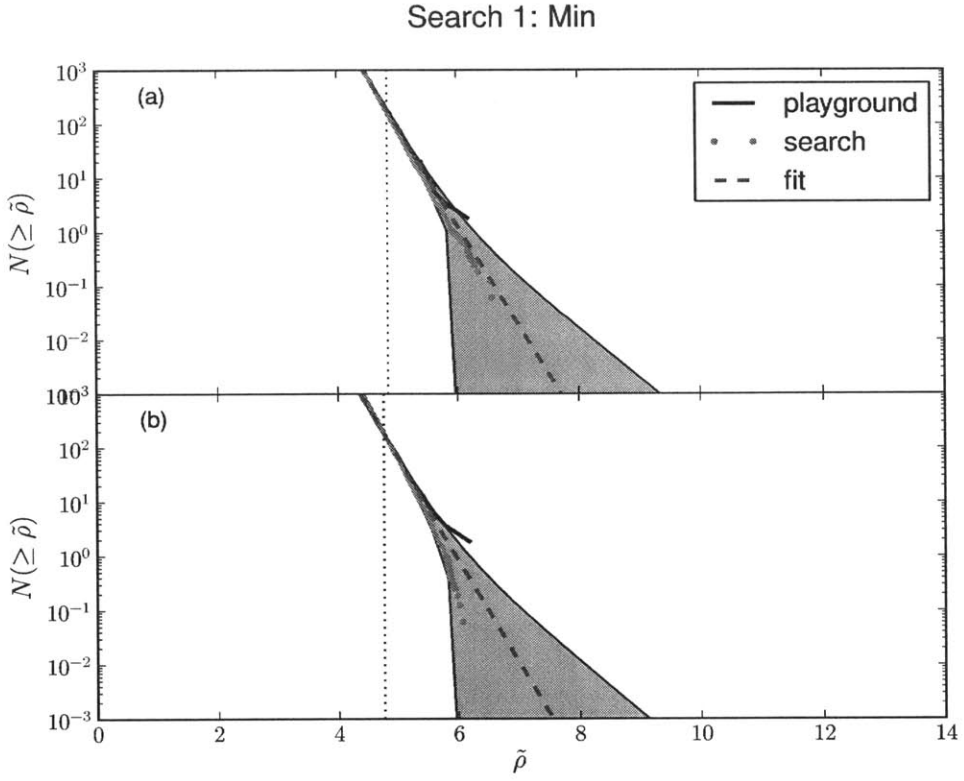

Figure 5-8: Cumulative distribution of $\tilde{\rho}$ for the minute search. This is similar to Figure 5-2, but instead of comparing two playground regions, we compare the search and the playground regions. The tail to the right of the vertical dotted line was used for the fit. Panel (a) shows the distribution when the source masks are the same size, while panel (b) shows the distribution when the source masks are proportional to the source brightness. The search is consistent with the background expectation and returned no transient candidates; the loudest event has $\tilde{\rho} = 6.1$ in (b). See text for more discussion.

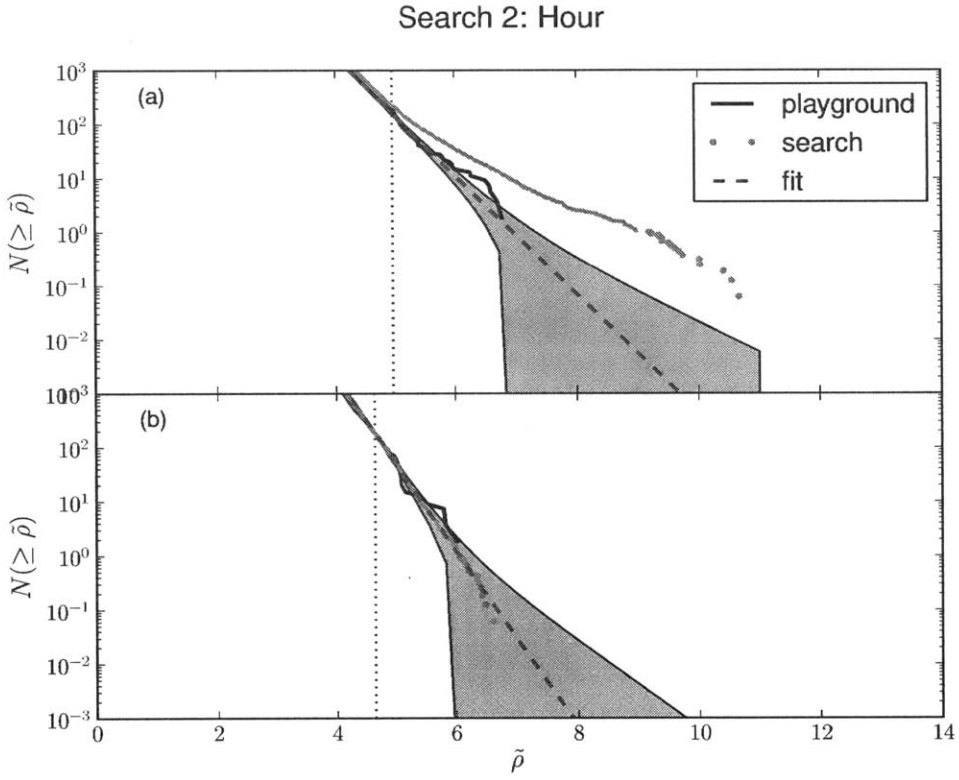

Figure 5-9: Cumulative distribution of $\tilde{\rho}$ for the hour search, similar to Figure 5-8. Panel (a) shows the distribution when the source masks are the same size, while panel (b) shows the distribution when the source masks are proportional to the source brightness. In panel (b), the search is consistent with the background expectation and returned no transient candidates; the loudest event has $\tilde{\rho} = 6.6$. See text for more discussion.

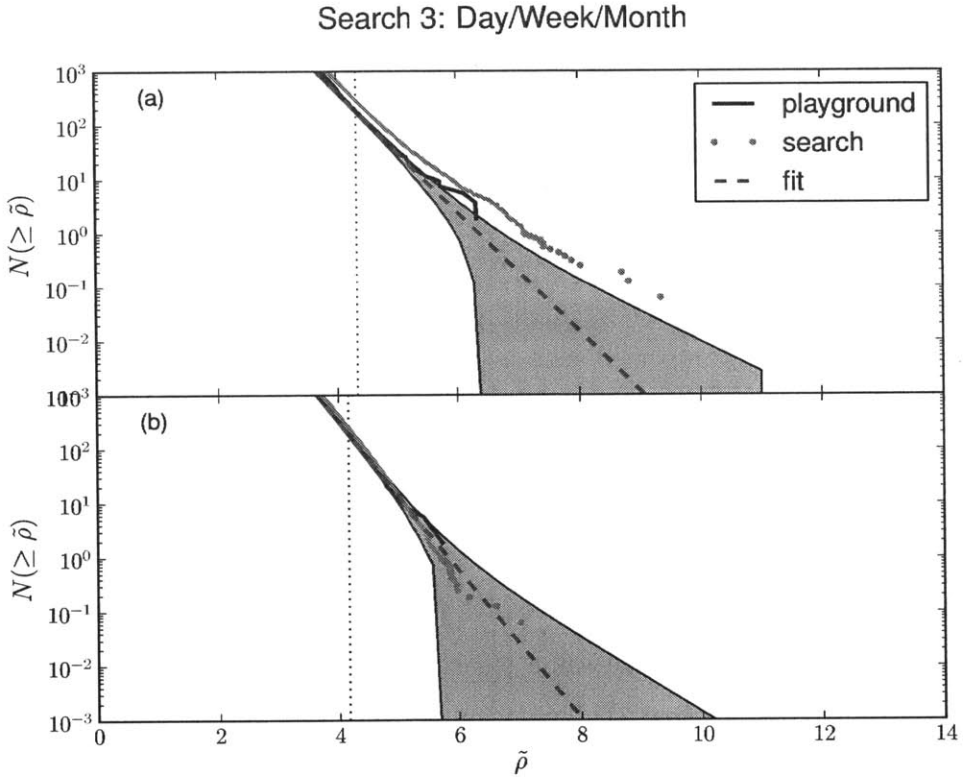

Figure 5-10: Cumulative distribution of $\tilde{\rho}$ for the hour search, similar to Figure 5-8. Panel (a) shows the distribution when the source masks are the same size, while panel (b) shows the distribution when the source masks are proportional to the source brightness. In panel (b), the search is consistent with the background expectation and returned no transient candidates; the loudest event has $\tilde{\rho} = 7.0$. See text for more discussion.

Table 5.2: Summary of the loudest events.

Search	$\tilde{\rho}$	RA (deg)	Dec (deg)	Amp (mJy)	Duration	Start Time (MJD)
min	6.1	355.16	-18.34	220.8	4 m	56539.72
hour	6.6	353.60	-32.41	58.3	1.5 h	56614.52
month	7.0	351.64	-25.86	18.7	30 d	56537.67

To define the veto area and refine the background characterization, we performed an empirical fit to determine the size of the source mask as a function of source flux density. We picked 6 sources with flux densities above 5 Jy, measured their sidelobe contamination areas in the unmasked transient sky map for the hour search, and fitted a straight line through the two points with the steepest slope to obtain the most conservative relationship between the size of the source mask and the source flux density:

$$\Delta n_{\rm pix} = 6.6 \times S + 0.5 \tag{5.1}$$

where $\Delta n_{\rm pix}$ is the number of pixels to mask on each side of the source and S is the value of the source flux density. The fit is illustrated in Figure 5-11. If the fit returned $\Delta n_{\rm pix} < 10$, however, we set $\Delta n_{\rm pix} = 10$ to give the 20×20-pixel mask region we had used before. This is a conservative value to account for the size of the synthesized beam. We also manually flagged 31 double sources that were misidentified as single sources by AEGEAN.

We reran the searches with the new source masks. We obtained a revised threshold from the playground region for each of the three searches as listed in Table 5.1, and computed new efficiency curves, which are the ones illustrated in Figure 5-3. The revised thresholds are lower, corresponding to more sensitive searches. We found no candidates. The properties of the loudest events are listed in Table 5.2. The $\tilde{\rho}$ distributions from the search region now agree very well with the background expectation as shown in the bottom panels of Figures 5-8, 5-9, and 5-10.

Finally, we mention that another possible way to determine the background distribution is to shuffle the images in time and then use the entire image instead of defining a playground region. This avoids the need to extrapolate the tail of the $\tilde{\rho}$

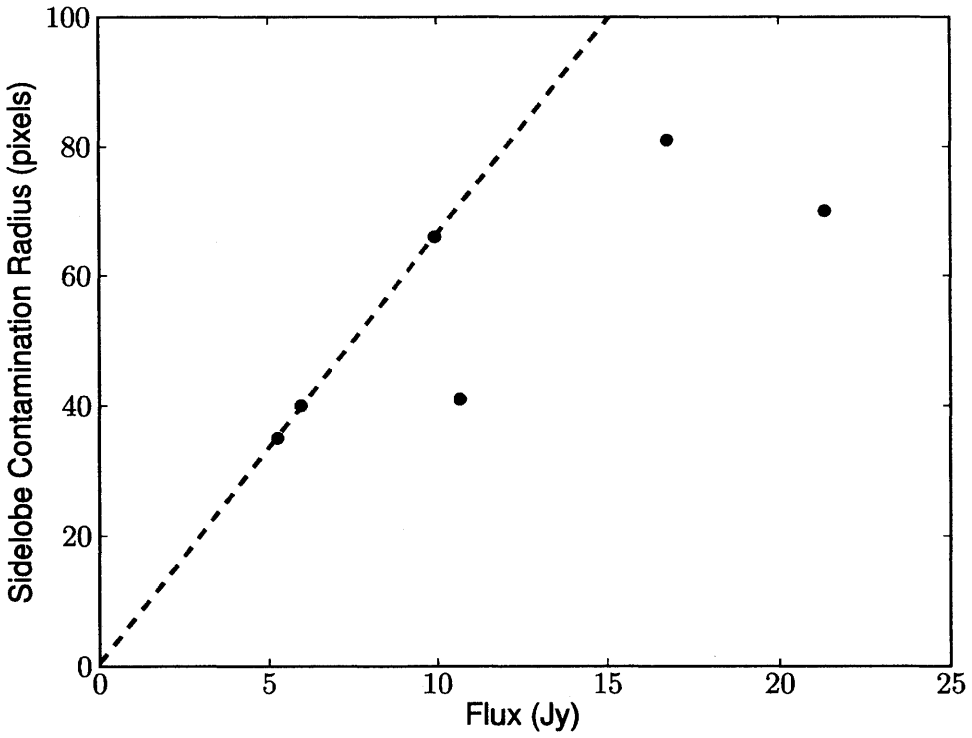

Figure 5-11: An empirical line fit through the two points with the steepest slope (the third that lies on the line is coincidental) to determine the most conservative size of the masked region according to source brightness. See Equation 5.1 for the values of the linear fit.

distribution, but we caution that random shuffling could break any temporal correlations in the systematic errors and change the noise distribution. We did not do this in our analysis. Future work is necessary to determine the timescale on which to shuffle the images that would preserve the true noise distribution.

5.5 Limits

As we did not detect any transient candidates, we placed an upper limit on the transient surface density. We based our upper limit calculation on the "loudest event statistic" as derived by [111, 112], which meant that we used the largest observed $\tilde{\rho}$ instead of the search threshold $\tilde{\rho}^*$ to determine our search efficiency. This formalism does not rely on the threshold or the extrapolation described in Section 4.1, making it more robust and stringent, and is very similar to the two-epoch equivalent snapshot rate introduced by [57], but in addition it takes into account the search efficiency.

The probability that we detect no events above $\tilde{\rho}$, assuming that the astrophysical transient events are described by a Poisson distribution, is

$$P(\tilde{\rho}) = e^{-\mu\epsilon(\tilde{\rho})} \tag{5.2}$$

where $\mu = \Sigma\Omega N_e$ is the Poisson mean, Σ is the transient surface density, Ω is the area of each searched image, N_e is the number of epochs or independent time samples, $\epsilon(\tilde{\rho})$ is the search efficiency as a function of flux density evaluated at $\tilde{\rho}$. The upper limit on Σ at a particular confidence level p is then determined by $P(\tilde{\rho}_m) = 1 - p$ or

$$\Sigma_p = -\frac{\ln(1-p)}{\Omega N_e \epsilon(\tilde{\rho}_m)} \tag{5.3}$$

where $\tilde{\rho}_m = \max(\tilde{\rho})$ is the loudest event statistic.

We computed the upper limit at 95% confidence level separately for each of our three searches as they probed different timescales that corresponded to different astrophysical sources or processes. We determined Ω by multiplying the number of pixels searched and the area of each pixel, which gave us $\Omega = 186\,\mathrm{deg}^2$. We determined

N_e by dividing the whole observation period, i.e. the time between the first and the last snapshot, by the transient duration, or the timescale of the search; if there were gaps in the observation that were as long as the transient duration, we subtracted the number of gaps from N_e. For our three searches, we have 625, 28, and 3 epochs respectively.

Figure 5-12 shows our results and compares them to the published results on the transient surface density between 150 and 330 MHz. Our technique allows us to explore a larger phase space more efficiently. Despite using images each with a 2-min integration time, we achieved sensitivities equivalent to longer integration times for the longer duration transient searches. Although our results do not set more stringent limits at the same flux densities compared to [110], also an MWA result, their analysis covered a bigger sky area (452 $\deg^2$) and a longer observation period ($\sim$ 80 hr integration time spanning 1 yr). If we naively scaled Ω and N_e to match theirs, our limits would be comparable or better, but the increased trials factor and the gaps in the data might affect $\tilde{\rho}$ and the limits in a non-linear manner. Nonetheless, the limits on these timescales will improve simply by adding more data, pushing toward lower and lower transient surface densities at the same flux sensitivities. Pushing toward fainter flux densities would require better calibration techniques or primary beam modeling to decrease the image noise. Even with our current data, we report improved limits at flux densities between $\sim$ 20–200 mJy for hour- and month-long transients.

Our limits are also consistent with the reported detections of radio transients. The transient reported by [109] was much fainter than the sensitivity we could achieve with our data even though it occurred on a timescale that we probed ($\sim$ day); if we assume a typical spectral index of -0.7, the source they detected at 2.1 mJy at 325 MHz would be 3.2 mJy at 182 MHz, which is an order of magnitude fainter than our best flux sensitivity at $\sim$ 20 mJy. While our limit for the day-to-month search appears to overlap with the transient detection reported by [113], our results are still consistent with a non-detection. Their transient lasted about $\sim$ 6 months, which is longer than the total observation time of our data. Furthermore, their transient was detected

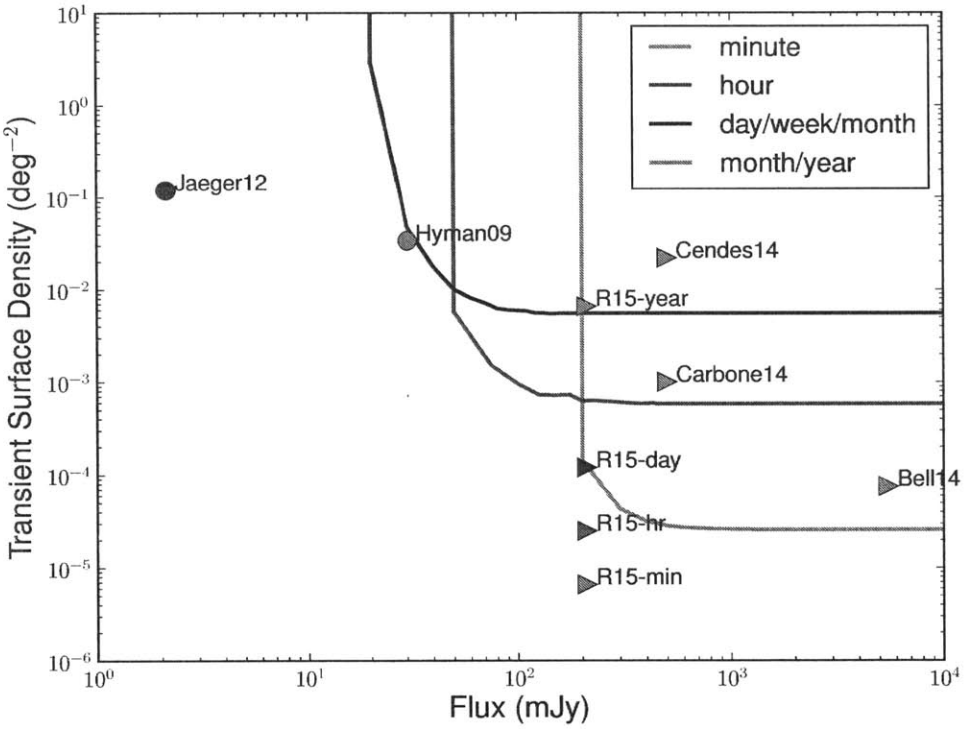

Figure 5-12: Upper limits on the transient surface density at 95% confidence level from our analysis (solid lines) and other published results (symbols) between 150 and 330 MHz. R15 denotes Rowlinson15 [110]. Jaeger12 [109] and Hyman09 [113] reported detections; the rest reported upper limits [108, 114, 115, 110]. Timescale is color-coded.

near the Galactic Center, where it is plausible that the transient population and hence transient rate might be different from the extragalactic transient population, which we observed. If the transient population is similar, however, with more data, we should also begin to detect such transients.

Chapter 6

Properties of Late-Time Radio Afterglows from Compact Binary Coalescence

This chapter is adapted from the preprint titled Detectability of Late-Time Radio Afterglows from Compact Binary Coalescence [116].

6.1 Introduction

As mentioned in Chapter 1, joint electromagnetic (EM) and gravitational wave (GW) observations of compact binary coalescence (CBC), e.g. binary neutron stars (BNS), are complementary and necessary. However, the large errors, ranging from 100–1000 deg^2 [78, 79], in the sky localization of the GW events during the early days of GW detector operation present a challenge for EM follow-up as most telescopes of much smaller fields of view in comparison.

Many new widefield radio instruments are currently operating or will soon begin operating: the Long Wavelength Array[1] Station 1 (LWA1, [117]), the Low-Frequency Array[2] (LOFAR, [50]), the Murchison Widefield Array[3] (MWA, [45, 46]), the Cana-

[1] http://lwa.phys.unm.edu/
[2] http://www.lofar.org/
[3] http://mwatelescope.org/

dian Hydrogen Intensity Mapping Experiment[4] (CHIME, [118]), and the Australian Square Kilometer Array Pathfinder[5] (ASKAP, [119, 48]). More instruments are planned for the future, such as the Hydrogen Epoch of Reionization Array[6] (HERA) and the Square Kilometer Array[7] (SKA). The wide fields of view of these instruments (30–600 $\deg^2$) make them promising for EM follow-up of GW events.

Most of these instruments operate at low frequencies (< 500 MHz), where the expected EM counterpart of a GW event is a short gamma-ray burst (SGRB) afterglow. Radio emission from afterglows usually peak on timescales of weeks to months, if not longer, and are expected to be faint [32, 33, 85]. Although there has been no detection of SGRB afterglows at low frequencies so far, this work will show that, within a plausible range of afterglow model parameters, there is a spread in the distributions of peak fluxes and durations of these afterglows, suggesting that a subset of afterglows could be detectable by the widefield radio instruments. However, these results are sensitive to many model parameters that are still uncertain.

Distinguishing faint SGRB afterglows from other slow transients such as radio supernovae could be an additional challenge. However, the radio transient sky at low frequencies is not well-understood. Many radio transients are expected to exist [9], but few have been detected so far [56]. While previous transient surveys at low frequencies were limited by sensitivity [120, 108, 114, 115, 110] or field of view [109], future surveys with the widefield radio instruments will be able to characterize the rate of background transients for EM follow-up. If SGRBs are indeed associated with CBC, these instruments can also search for on-axis and off-axis afterglows to constrain the CBC rate, which is uncertain by three orders of magnitude [121].

Previous radio searches for orphan afterglows have yielded null results [122, 28] but at relatively limited sensitivity (6 mJy). As this work shows, if the widefield radio instruments achieve their theoretical thermal sensitivities, they will be suitable for transient surveys and follow-up observations. This work is complementary to

[4]http://chime.phas.ubc.ca/
[5]http://www.atnf.csiro.au/projects/askap/
[6]http://reionization.org/
[7]http://www.skatelescope.org/

the studies of radio emission from subrelativistic outflows of CBC [73] and the detectability of radio afterglows from long GRBs at high frequencies [123, 124] and high redshifts [125]. We explore the properties of simulated SGRB afterglow light curves at radio frequencies for a range of source and observer parameters in Section 6.2, and characterize the detectability of these events for ideal radio instruments in Section 6.3. Then we estimate the rates of detection for these instruments and different survey methods in Section 6.4. We compare our results to recent work by others in Section 6.5. While the radio observations of GRB 130427A show that there is bright ($\sim$ mJy) radio emission due to the reverse shock at early times [126], this emission component is not included in our work, which only considers late-time (> 1 d) afterglow emission from the forward shock, but it will be subject to future studies.

6.2 Light Curve Properties

The afterglow emission of a SGRB is synchrotron radiation produced when the relativistic ejecta creates a shock in the surrounding medium (see [127, 128] for recent reviews). The shape of the light curve depends on the properties of the burst, the microphysics of synchrotron radiation, and the parameters specifying an observer [129, 130]. Observationally, SGRBs have isotropic energies $10^{49} \lesssim E_{\rm iso} \lesssim 10^{51}$ ergs (see [127, 77] and references therein). Their jet opening angles are difficult to measure and thus have large uncertainties, but a few jet break measurements suggest $\theta_{\rm jet} \sim 10^{\circ}$ [131, 81, 132, 82]. Their circumburst environments generally have low inferred densities $10^{-5} \lesssim n \lesssim 1\,{\rm cm}^{-3}$ [81, 133, 134, 132, 82], consistent with the expectations for BNS mergers [135]. The results from these observations motivate the parameter space that we explore in this work.

A group has developed a numerical tool BOXFIT [86] that generates afterglow light curves quickly for arbitrary burst and observer parameters, and it has allowed us to improve on the previous estimates of SGRB afterglow properties derived from the analytical approximations [32, 33]. BOXFIT calculates the fluid state of the shock by interpolating the results of two-dimensional relativistic hydrodynamics jet simulations

Table 6.1: Parameters used to generate afterglow lightcurves.

Observer Parameters	Burst Parameters[a]	Microphysics[b]
$\nu_{\rm obs} = 60, 150, 600, 1430\,{\rm MHz}$	$\theta_{\rm jet} = 11.5^\circ$ (0.2 rad)	$\xi_N = 1.0$
$\theta_{\rm obs} = 0^\circ$–$90^\circ$	$E_{\rm jet} = 10^{48}, 10^{50}\,{\rm ergs}$	$p = 2.5$
$t_{\rm obs} = 0.1$–$10^6\,{\rm d}$	$E_{\rm iso} = 5 \times 10^{49}, 5 \times 10^{51}\,{\rm ergs}$	$\epsilon_e = 0.1$
$d_{\rm ref} = 10^{27}\,{\rm cm}$ ($z = 0$)	$n = 10^{-5}, 10^{-3}, 1.0\,{\rm cm}^{-3}$	$\epsilon_B = 0.1$

[a]Jet energy $E_{\rm jet}$ is related to isotropic energy $E_{\rm iso}$ through $E_{\rm iso} = 2E_{\rm jet}/\theta_{\rm jet}^2$ for $\theta_{\rm jet} \ll 1$.
[b]ξ_N is the fraction of accelerated electrons. p is the power-law slope of the electron energy distribution. ϵ_e is the fraction of internal energy in the electrons. ϵ_B is the fraction of internal energy in the magnetic field.

after applying the analytical Blandford–McKee solutions [136] to the ultra-relativistic phase of the shock expansion. Then it calculates the light curve by solving the linear radiative transfer equations for synchrotron radiation.

Using this tool, we generated light curves of SGRB afterglows to study their properties and detectability at radio frequencies. We specified BOXFIT to use the Blandford–McKee solutions for $200 > \gamma > 25$ where γ is the Lorentz factor of the fluid directly behind the shock front. We also fixed the parameters describing the microphysics of synchrotron radiation to their characteristic values, e.g. [137, 73]. All simulation parameters are listed in Table 6.1 and are consistent with observations. We explored a range of energies $E_{\rm iso}$ and circumburst densities n corresponding to known constraints and expectations. For each combination of $E_{\rm iso}$ and n, we generated light curves at 4 observer frequencies $\nu_{\rm obs}$ sampling the range covered by widefield radio instruments and at 11 observer angles $\theta_{\rm obs}$ spaced linearly between 0° (on-axis) and 90° (off-axis). Each light curve consisted of 350 time samples spaced logarithmically between 0.1 and 10^6 d, capturing the evolution of the afterglow from early to late times. The bursts were located at $d_{\rm ref} = 10^{27}\,{\rm cm}$ (324 Mpc), a distance comparable to the average aLIGO BNS range at design sensitivity [78] but was otherwise an arbitrary choice. The light curves were generated in the source frame.

To capture the properties of an ensemble of afterglow light curves, we generated a sample of bursts that is uniformly distributed in energy ($5 \times 10^{49} \leq E_{\rm iso} \leq 5 \times 10^{51}\,{\rm ergs}$), jet orientation ($-1 \leq \cos\theta_{\rm obs} \leq 1$), and comoving volume ($z \leq 1$). The choice of $z = 1$ provides a flux-limited sample ($\sim \mu$Jy) for the radio instruments.

While some of these bursts might not be detectable by the GW detectors, they might be detectable by radio instruments in a blind survey and will be of interest.

Instead of rerunning BOXFIT with new parameters, we used the analytical energy-flux scaling relation derived by [138] to determine the peak fluxes of the light curves over a continuous range of energies: $E'_{\rm iso} = \kappa E_{\rm iso}$ and $f'_{\rm peak} = \kappa f_{\rm peak}$ where κ is a scaling parameter for a fixed density, distance, and observer angle. Then we scaled these fluxes according to their luminosity distances [139]: $f_{\rm peak}(d_{\rm L}) = (1+z) f'_{\rm peak}(d_{\rm ref}/d_{\rm L})^2$. We also scaled the durations of these light curves according to $t'_{\rm dur} = \kappa^{1/3} t_{\rm dur}$ [138], where we defined $t_{\rm dur}$ to be the time during which flux $> 0.5 f_{\rm peak}$.

As evident from the distributions of $f_{\rm peak}$ and $t_{\rm dur}$ shown in Figure 6-1, most bursts are faint ($\lesssim \mu$Jy) and long-lasting ($\gtrsim$ yr), confirming the results of previous studies. However, there is a spread, implying that there might be bursts detectable with the current widefield radio instruments. The spread is dependent on the model parameters, which are fairly uncertain. This is also evident in the cumulative distributions of $f_{\rm peak}$ for the same sample shown in Figure 6-2.

Figure 6-1 also shows that the bursts become brighter and longer-lasting when their energies increase, and they become fainter and longer-lasting when they are more off-axis. This is consistent with the results from [85], who considered only the lowest energy bursts. The trend along observer angle is absent for $n = 1\,{\rm cm}^{-3}$ because emission is isotropic by the time the system becomes optically thin. The two outlier points (red squares) just happen to be nearby samples. These properties suggest that detectable bursts will be more on-axis, have higher energies, and occur in higher density environments. Observing these bursts will take ~ 1 yr, which can be undertaken by a realistic survey.

A theoretical study of late-time afterglow light curves by [140] argues that these afterglows could be a factor of a few brighter than previously expected if the bulk of the shock-accelerated electrons are non-relativistic, which could improve the detectability of these events, but this effect is not included in our work.

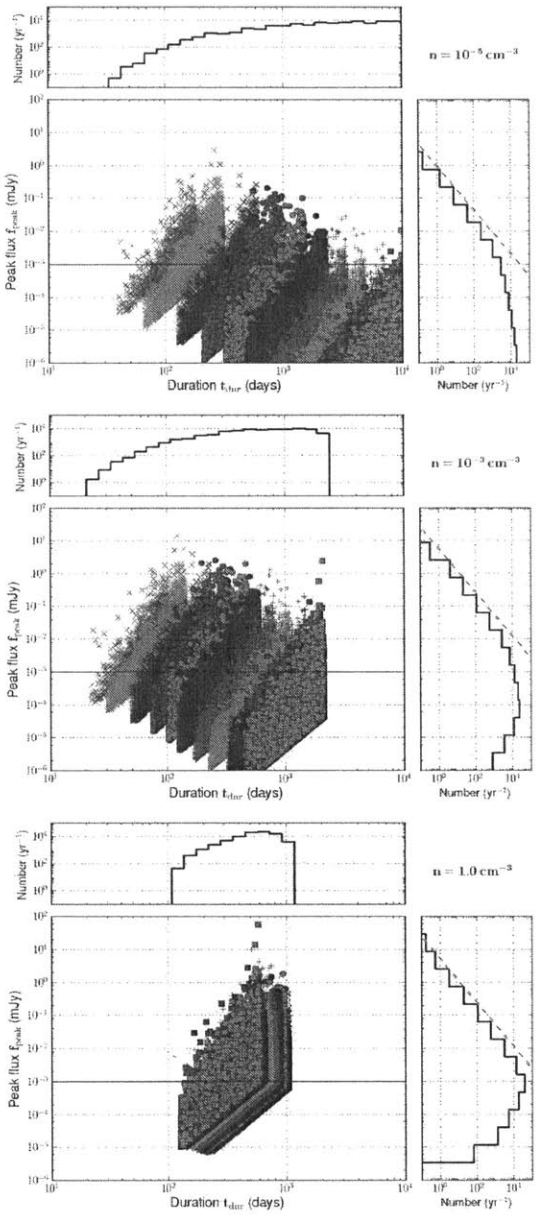

Figure 6-1: Distributions of $f_{\rm peak}$ and $t_{\rm dur}$ for the simulated afterglow light curves in a volume-limited sample at 150 MHz and for different values of n. All relevant plots in this work are normalized to the realistic CBC rate of $1\,{\rm Mpc^{-3}\,Myr^{-1}}$ for 2π sky area and will not be mentioned separately. The sample is uniformly distributed in $E_{\rm iso}$, jet orientation, and volume. The distinct clusters in the scatter plot are caused by binning in $\theta_{\rm obs}$. The vertical edges are caused by the energy cutoffs at 5×10^{49} and 5×10^{51} ergs. The diagonal edges are caused by the distance cutoff, which is chosen to be $z=1$ in order to give a flux-limited sample ($\sim\mu$Jy). The dashed line in the histogram panels on the right is a reference line with slope $-3/2$ for $N\propto f^{-3/2}$.

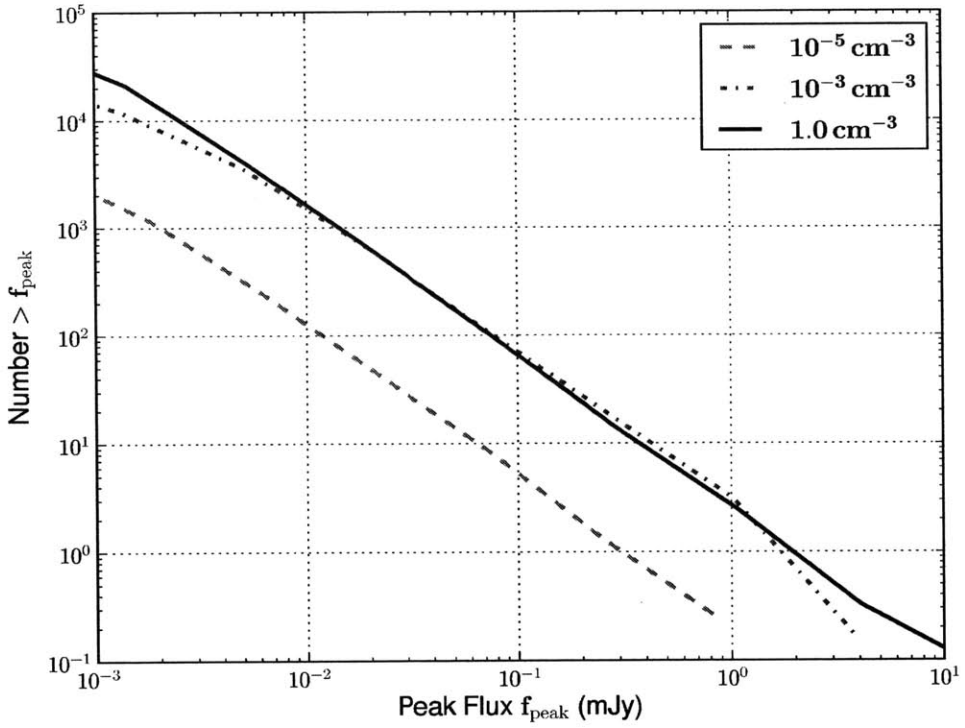

Figure 6-2: Cumulative distributions of $f_{\rm peak}$ for the simulated afterglow light curves at 150 MHz. The sample is the same as that of Figure 6-1 but plotted for $f_{\rm peak} > 1\,\mu$Jy. This shows the population of afterglows that could be detected as sources given a particular instrument flux sensitivity at 150 MHz.

Table 6.2: System parameters for present and future widefield radio instruments. $\nu_{\rm obs}$ is the observer frequency at which we generated the afterglow light curves. $\sigma_{\rm 1h}$ is the thermal sensitivity after 1 hour of integration, calculated from the published values according to Equation 6.1. $\Omega_{\rm FOV}$ is the instantaneous field of view. LWA1 and LOFAR can form multiple beams simultaneously, increasing their sky coverage by decreasing their bandwidth [117, 50]; we chose the larger $\Omega_{\rm FOV}$ for them. CHIME and HERA are drift-scan telescopes unlike the other instruments, so we calculated the average $\sigma_{\rm im}$ of a 1-day drift-scan image (see Appendix C) and listed the effective field of view for this image.

Instrument	Frequency Range (MHz)	Bandwidth (MHz)	$\nu_{\rm obs}$ (MHz)	$\sigma_{\rm 1h}$ (mJy)	$\Omega_{\rm FOV}$ (deg^2)
LWA1	10–88	16	60	16.8	4×61
LOFAR Low	10–80	3.66	60	17.5	48×74.99
LOFAR High	110–240	3.66	150	0.877	48×11.35
MWA	80–300	30.72	150	0.913	610
CHIME Pathfinder	400–800	400	600	0.240	20626
CHIME	400–800	400	600	0.036	20626
ASKAP	700–1800	300	1430	0.029	30
HERA	50–225	100	150	0.017	2712
SKA1 Low	50–350	250	150	0.002	27

6.3 Ideal Detection Metric

The detectability of radio afterglows depends not only on the intrinsic properties of SGRBs, as shown in Section 6.2, but also on the sensitivity of the radio instrument. The simplest characterization of the sensitivity of a radio interferometer is the thermal noise $\sigma_{\rm im}$ of an image, or point source sensitivity:

$$\sigma_{\rm im} = \left(\frac{2k_{\rm B}T_{\rm sys}}{A_{\rm eff}N_{\rm ant}\epsilon_c} \right) \frac{1}{\sqrt{N_{\rm pol}B\,t_{\rm int}}} \tag{6.1}$$

where $k_{\rm B}$ is the Boltzmann constant, $T_{\rm sys}$ is the system temperature, $A_{\rm eff}$ is the effective area of each antenna, $N_{\rm ant}$ is the number of antennas, ϵ_c is the correlator efficiency, $N_{\rm pol}$ is the number of polarizations, B is the instantaneous bandwidth, and $t_{\rm int}$ is the image integration time. While the radio instruments selected for this work have good thermal sensitivities as listed in Table 6.2, many of them are affected by the classical confusion noise arising from the background of unresolved sources [53].

As afterglows are generally fainter than the classical confusion noise of many widefield radio instruments, we applied the matched filter technique developed in

Chapter 2 to derive a simple metric that characterizes the detectability of afterglows. This technique is apt for afterglow searches because there are model light curve templates, which depend on the flux or amplitude A, the start time t_0, the jet energy $E_{\rm jet}$, the jet orientation $\theta_{\rm obs}$, and the circumburst density n. A real search would need to prepare a bank of templates $f(t_0, E_{\rm jet}, \theta_{\rm obs}, n)$ that samples this parameter space, and then evaluate the matched filter statistic $\rho(x, t_0, E_{\rm jet}, \theta_{\rm obs}, n)$ for every template. The highest value of $\rho(x) = \max \rho(x, t_0, E_{\rm jet}, \theta_{\rm obs}, n)$ constitutes the detection statistic.

For the purpose of establishing a criteria for afterglow detectability, however, it is sufficient to consider a simplified version of the matched filter analysis that nevertheless captures the key aspects of radio transient detection. The main simplification comes from the characterization of the afterglow light curves only by their peak flux $f_{\rm peak}$ and duration $t_{\rm dur}$. The corresponding template $f(t)$ is the top-hat profile with $f_i = f_{\rm ref}$ when the transient is on and zero otherwise. The template reference flux $f_{\rm ref}$ is $f_{\rm peak}$ of the source at a reference distance $d_{\rm ref}$.

Out of convenience, we chose the length of a survey epoch to be equal to $t_{\rm dur}$, so that the transient is on in only one epoch. Each epoch might contain more than one snapshot image with noise $\sigma_{\rm im}$, so the relevant noise for each epoch is $\sigma_{\rm e} = \sigma_{\rm im}/\sqrt{(\text{number of snapshots in each epoch})}$. If the number of epochs is $N_{\rm e}$, the average flux for the template is $\langle f \rangle = f_{\rm ref}/N_{\rm e}$ when we neglect the effects of the primary beam. The variance of ρ in the absence or presence of a signal becomes

$$\sigma_0 = \sigma_1 = \sqrt{(f - \langle f \rangle, f - \langle f \rangle)} = \frac{f_{\rm ref}}{\sigma_{\rm e}} \sqrt{\frac{N_{\rm e} - 1}{N_{\rm e}}} \tag{6.2}$$

while the mean of ρ in the presence of a signal becomes

$$\mu_1 = A(f - \langle f \rangle, f - \langle f \rangle) = A \left(\frac{f_{\rm ref}}{\sigma_{\rm e}} \right)^2 \frac{N_{\rm e} - 1}{N_{\rm e}} \tag{6.3}$$

To achieve a high efficiency ($\geq 97\%$) in the signal detection at 5σ false alarm probability, we imposed the following condition: $\mu_1 \geq 7\sigma_0$, which is satisfied when

$$f_{\rm peak} \geq f^* \equiv A f_{\rm ref} = 7 \sqrt{\frac{N_{\rm e}}{N_{\rm e} - 1}} \, \sigma_{\rm e} \tag{6.4}$$

For surveys that last much longer than $t_{\rm dur}$, $N_{\rm e} \to \infty$ and the flux threshold approaches what would be achievable in the absence of the classical confusion noise: $f^* \to 7\sigma_{\rm e}$. The afterglows tend to be long-lasting, so in practice the observations will span at best a few transient durations. We chose $N_{\rm e} = 2$, allowing for the reference epoch to be as long as the transient itself. In addition, we required an upper limit of 3 yr on $t_{\rm dur}$ independent of $f_{\rm peak}$. In the absence of archival radio data that can serve as references, it seems impractical to detect afterglows significantly longer than that. Combining the two thresholds, we defined the following detectability criteria:

$$f_{\rm peak} \geq 7\sqrt{2}\,\sigma_{\rm e} \tag{6.5}$$

$$t_{\rm dur} \leq 3{\rm yr} \tag{6.6}$$

This criteria is ideal, as it assumes that one is able to achieve the thermal noise sensitivity after the classical confusion noise is perfectly subtracted using the matched filter technique. However, this can be difficult to achieve in a real search. Calibration errors will limit the dynamic range of the images and, together with sidelobe confusion noise, will decrease the sensitivity. Errors in the primary beam model, as we saw in Chapters 4 and 5, will also limit the sensitivity. However, we can still use the ideal criteria to learn what we might be able to achieve in the best case scenario, what types of instruments are better for afterglow searches, and what kind of afterglow properties these instruments are better at constraining.

6.4 Rate Estimation

Having defined a flux threshold and a duration threshold in Section 6.3 to characterize the detectability of radio afterglows, we estimated the number of SGRB afterglows that we expected an ideal instrument to detect given the intrinsic rate of these events as well as the sensitivity, field of view, and survey strategy of the instrument.

The association between SGRBs and BNS coalescence is promising but far from conclusive. Nonetheless, the intrinsic rate of SGRBs as derived from SGRB observa-

tions is consistent with the rates of BNS coalescence as derived from binary pulsar observations and population synthesis. Hence, we assumed that the rate of SGRBs was equal to that of BNS coalescence and used the realistic value derived by [121].

Given the intrinsic rate of SGRB afterglows $R_{\rm BNS}$ (number per volume, per year), we calculated the rate of afterglow detections $R_{\rm det}$ (number per year) expected for a radio instrument by determining the volume that the instrument can observe. This volume depends on the sensitivity or flux threshold f^* of the instrument and the sky area $\Omega_{\rm sur}$ covered by the survey.

f^* sets the maximum luminosity distance $d_{\rm L}$ to which the instrument can observe a source with a fixed luminosity. We converted $d_{\rm L}$ to the horizon distance $d_{\rm H}$ that we defined to be the comoving distance corresponding to $d_{\rm L}$ at redshift z', both of which are unknown and need to be computed:

$$d_{\rm H} = \frac{d_{\rm L}(z')}{1+z'} \tag{6.7a}$$

$$\text{where } \frac{d_{\rm L}^2(z')}{1+z'} = d_{\rm ref}^2 \left(\frac{f_{\rm ref}}{f^*}\right) \tag{6.7b}$$

$d_{\rm ref}$ is the reference distance at which the afterglow light curves were generated, and $f_{\rm ref}$ is the peak flux of the afterglow at $d_{\rm ref}$. The factor of $(1+z')$ in Equation 6.7b is the k-correction term [139]. We assumed the *Planck* 2013 cosmology [141] in our analysis.

By definition, the same instrument will have a range of $d_{\rm H}$ corresponding to different light curves with different $f_{\rm ref}$. To calculate $d_{\rm H}$ for each instrument, we substituted the corresponding f^* into Equation 6.7 and numerically solved for z' (hence $d_{\rm L}$ and $d_{\rm H}$) for the light curves we generated in Section 6.2. Figures 6-3 and 6-4 show example horizon distances for the MWA, where f^* was calculated for a 1-year all-sky survey (see Section 6.4.1). The shapes of the $d_{\rm H}$ curves trace the variations of $f_{\rm ref}$ as a function of $\theta_{\rm obs}$ and $E_{\rm iso}$ while the normalization is set by f^* of the instrument. In other words, another instrument operating at the same frequency but with a different sensitivity will have $d_{\rm H}$ curves of approximately the same shape (up to cosmological corrections) but a different amplitude.

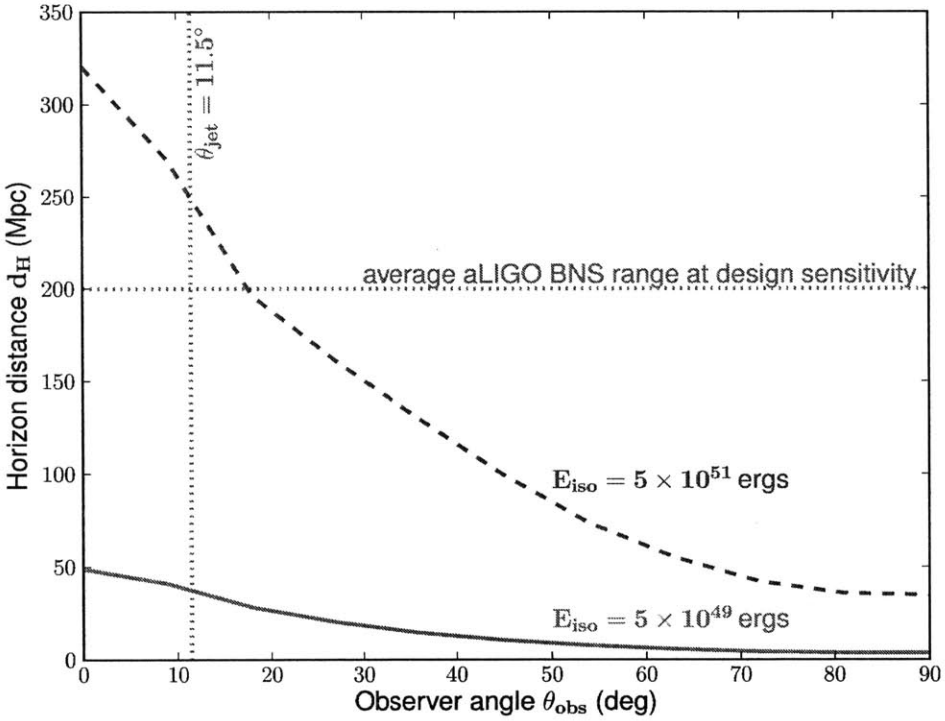

Figure 6-3: MWA horizon distances for afterglows with different $E_{\rm iso}$ but in the same environment ($n = 10^{-3}\,{\rm cm}^{-3}$). $\nu_{\rm obs} = 150\,{\rm MHz}$ and $f^* = 0.8\,{\rm mJy}$. The vertical dotted line marks the jet opening angle. Afterglows that are more energetic or more on-axis are brighter and therefore detectable to larger distances. For comparison, the average range of BNS coalescence for aLIGO at design sensitivity is 200 Mpc, as illustrated by the horizontal dotted line [78].

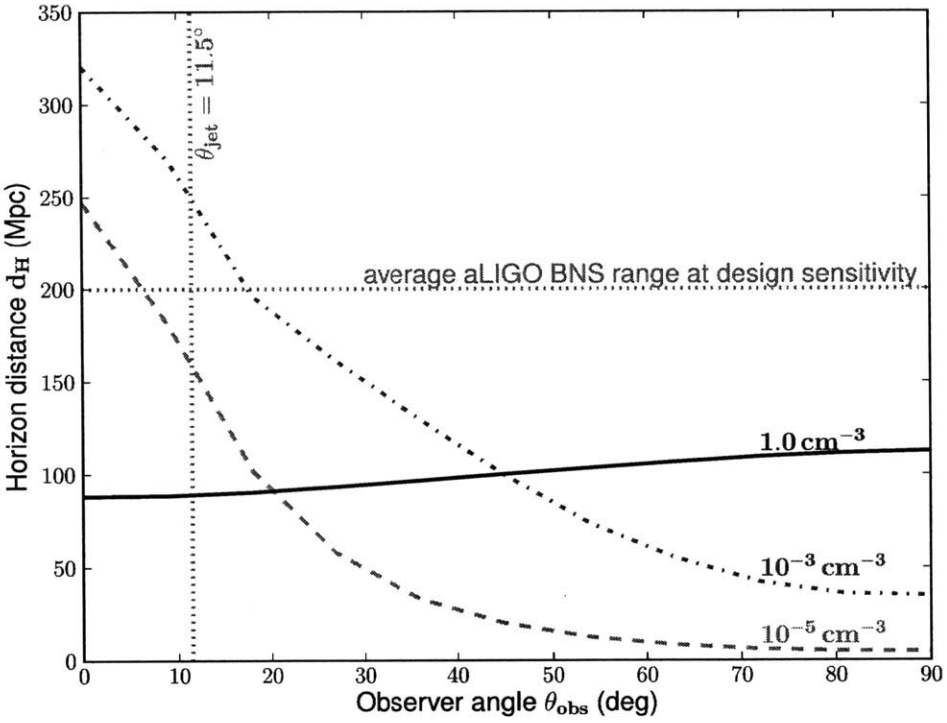

Figure 6-4: MWA horizon distances for afterglows in different environments but with the same $E_{\rm iso} = 5 \times 10^{51}$ ergs. $\nu_{\rm obs} = 150\,{\rm MHz}$ and $f^* = 0.8\,{\rm mJy}$. The vertical dotted line marks the jet opening angle. Afterglows become brighter as n increases until synchrotron self-absorption becomes dominant. The $d_{\rm H}$ curve for $n = 1\,{\rm cm}^{-3}$ is almost independent of $\theta_{\rm obs}$ because emission is isotropic by the time the system becomes optically thin. Contribution from the counter-jet makes the off-axis afterglows slightly brighter than the on-axis ones for $n = 1\,{\rm cm}^{-3}$.

Since $d_{\rm H}$ depends on $\theta_{\rm obs}$ and $E_{\rm iso}$, we integrated over $\theta_{\rm obs}$ and $E_{\rm iso}$ when we calculated $R_{\rm det}$:

$$R_{\rm det} = \left[\frac{R_{\rm BNS}}{4\pi(E_2 - E_1)}\right] \times \int_{E_1}^{E_2} \int_{4\pi} \left[\Omega_{\rm sur} \frac{d_{\rm H}^3(\theta_{\rm obs}, E_{\rm iso})}{3}\right] {\rm d}\Omega_{\rm obs}\, {\rm d}E_{\rm iso} \qquad (6.8)$$

This equation assumes that the bursts are uniformly distributed in energy ($5 \times 10^{49} \leq E_{\rm iso} \leq 5 \times 10^{51}$ ergs) and jet orientation ($-1 \leq \cos\theta_{\rm obs} \leq 1$). It also treats each n separately, where $n = 10^{-5}\,{\rm cm}^{-3}$ represents the intergalactic medium (outside the host galaxy) and $n = 1\,{\rm cm}^{-3}$ represents the interstellar medium (inside the host galaxy). If SGRBs occur equally likely in the different environments, $R_{\rm det}$ would be the average of the separate values. To integrate over $E_{\rm iso}$, we used the analytical energy-flux scaling relation for $f_{\rm ref}$ derived by [138] (see also Section 6.2) when we calculated $d_{\rm H}$. If $d_{\rm H}$ is independent of $\theta_{\rm obs}$ and $E_{\rm iso}$, Equation 6.8 reduces to $R_{\rm det} = R_{\rm BNS}(\Omega_{\rm sur} d_{\rm H}^3/3)$.

During the calculation of $R_{\rm det}$, we imposed a cut on the afterglow duration according to Equation 6.6. An afterglow that lasts longer than the survey or the availability of archival data will not be identified as a transient event even if it is bright. Figure 6-5 shows the cumulative distribution of peak fluxes for afterglow light curves with the constraint that $t_{\rm dur} \leq 3\,{\rm yr}$. This particular choice of $t_{\rm dur}$ manages to capture the majority of the detectable population without requiring a survey to last an impractical length of time. At 150 MHz, most of these afterglows last $\sim 1\,{\rm yr}$, as evident from Figure 6-6, suggesting that a survey should cover a time range that is at least as long. At higher frequencies, the durations are shorter ($\gtrsim 3\,{\rm months}$).

We now present the results for $R_{\rm det}$ for three types of observations: blind surveys, SGRB follow-up observations, and GW follow-up observations.

6.4.1 Blind Survey

There are two possible strategies for blind surveys: "narrow and deep" or "shallow and wide," where we require the total time allocated for the survey to last much longer than the time needed to reach a good sensitivity. Equations 6.5 and 6.7 show that $d_{\rm H} \propto \sigma_{\rm e}^{-1/2}$. Combining this relation with $R_{\rm det} \propto \Omega_{\rm sur} d_{\rm H}^3$ from Equation 6.8 and

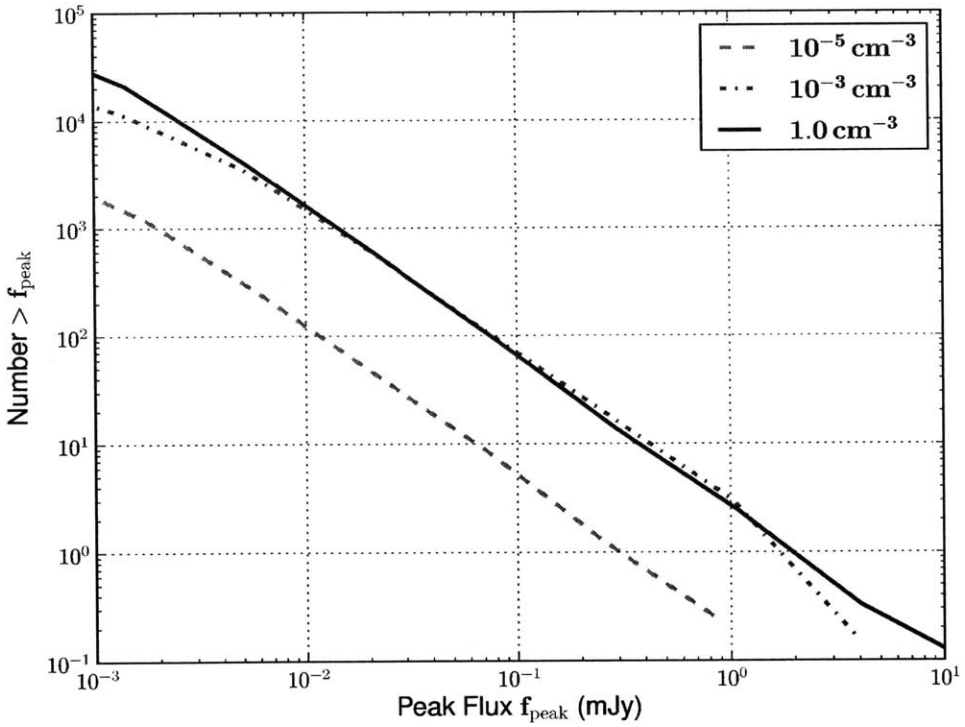

Figure 6-5: Cumulative distributions of f_{peak} for the simulated afterglow light curves at 150 MHz. This is similar to Figure 6-2 but plotted for afterglows with $t_{\mathrm{dur}} \leq 3\,\mathrm{yr}$. While certain afterglows could be bright enough to be detected as sources, they might last longer than the survey and therefore not be identified as transients.

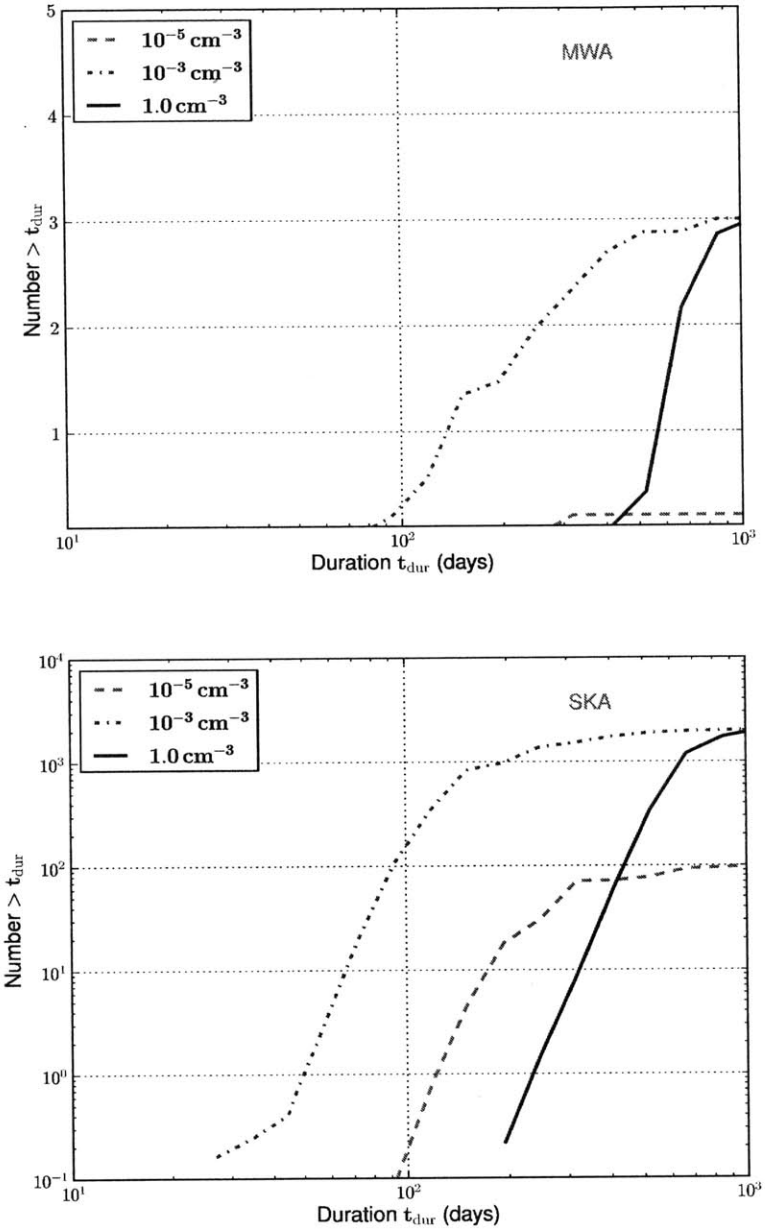

Figure 6-6: Cumulative distributions of $t_{\rm dur}$ for the simulated afterglow light curves at 150 MHz. This is plotted for afterglows with $t_{\rm dur} \leq 3$ years and $f_{\rm peak} \geq f^*$ for MWA ($f^* = 0.8$ mJy; top) and SKA ($f^* = 0.008$ mJy; bottom). As most of these afterglows last ~ 1 yr, a survey to detect these events should revisit the same area of the sky on a similar timescale. At higher frequencies, the durations are shorter ($\gtrsim 3$ months).

$\sigma_e \propto t_{int}^{-1/2}$ from Equation 6.1, we get the following dependence:

$$R_{det} \propto \Omega_{sur} t_{int}^{3/4} . \tag{6.9}$$

This shows that R_{det} increases faster with Ω_{sur} than it does with t_{int}, arguing in favor of a "shallow and wide" approach. In other words, for a fixed survey length $\tau_{sur} \gg t_{int}$, the survey should maximize its sky coverage over the time τ_{sur} rather than performing a deep integration on a small patch of sky if the goal of the survey is to increase R_{det}.

An all-sky survey covers the maximal area that any instrument can observe. The amount of accessible sky varies with the location on Earth while the Galactic plane obscures extragalactic observations. To account for this effect, we chose $\Omega_{sur} = 2\pi$ for every radio instrument that we considered in our analysis except for HERA. HERA is a drift-scan telescope with limited sky coverage; for HERA, we used $\Omega_{sur} = 2712.4\,\mathrm{deg}^2$ (see Appendix C). As each instrument has an instantaneous field of view $\Omega_{FOV} \leq \Omega_{sur}$, it needs N_p separate pointings, each with integration time t_{int}, to cover the total survey area $\Omega_{sur} = N_p \Omega_{FOV}$. We considered a survey length of $\tau_{sur} = 1\,\mathrm{yr}$, assuming 100% duty cycle. Thus σ_e (hence f^*) is set by $t_{int} = \tau_{sur}/2N_p$, where the factor of 2 comes from our requirement that the survey is split into two epochs, one of which serves as the reference epoch although both epochs will be used in the transient search. Results for selected instruments are shown in Table 6.3. These results are computed for $\Omega_{sur} = 2\pi$ sky area and depend on f^* (i.e. t_{int}), so one should use the dependence in Equation 6.9 to obtain the rates for a different sky coverage or flux sensitivity.

This is an order of magnitude estimate and a comparison between ideal instrument performances. Radio instruments at very low frequencies ($< 80\,\mathrm{MHz}$) will not be sensitive to afterglows because of strong synchrotron self-absorption and modest instrumental sensitivities. Instruments at higher frequencies, such as CHIME and ASKAP, perform better in part because they have large fields of view and in part because the afterglow emission is brighter and shorter. Afterglows that occur in denser mediums are more likely to be detectable.

Actual numbers will depend on the survey details and the achieved sensitivi-

Table 6.3: Ideal rates of afterglow detection for blind surveys. These rates are computed for ideal instrument performance and rounded to one significant digit. We used the realistic rates of BNS coalescence, where the expected rate of detection for aLIGO is $\sim 40\,\mathrm{yr}^{-1}$ [121]. Only the afterglows with $f_{\mathrm{peak}} \geq f^*$ and $t_{\mathrm{dur}} \leq 3\,\mathrm{yr}$ are included (see Section 6.3), where f^* is the ideal flux threshold for each instrument. ν_{obs} is the observer frequency at which we generated the afterglow light curves. The sky coverage is 4π for LIGO and 2π for all radio instruments except HERA, for which it is $2712.4\,\mathrm{deg}^2$. Instruments like CHIME and ASKAP are more suitable for afterglow searches, while afterglows that occur in denser mediums are more likely to be detectable. See text for a discussion of the realistic MWA rates.

Instrument	ν_{obs} (MHz)	f^* (mJy)	$n = 10^{-5}\,\mathrm{cm}^{-3}$ (outside host galaxy)	$n = 10^{-3}\,\mathrm{cm}^{-3}$	$n = 1.0\,\mathrm{cm}^{-3}$ (inside host galaxy)	
LWA1	60	23.3	0.001	0.007	6×10^{-4}	yr^{-1}
LOFAR Low	60	6.42	0.007	0.05	0.004	yr^{-1}
LOFAR High	150	0.81	0.2	3	3	yr^{-1}
MWA	150	0.80	0.2	3	3	yr^{-1}
MWA (realistic)	150	4.78	0.01	0.2	0.2	yr^{-1}
CHIME Pathfinder	600	0.176	2	50	200	yr^{-1}
CHIME Full	600	0.026	30	650	3000	yr^{-1}
ASKAP	1430	0.114	3	80	800	yr^{-1}
HERA	150	0.012	9	100	200	yr^{-1}
SKA1 Low	150	0.008	100	2000	2000	yr^{-1}

ties. As we have worked closely with the MWA, we contrasted the ideal performance of the MWA with a more realistic performance to illustrate how the rates might change. In the ideal case, $f^* = 0.8\,\mathrm{mJy}$ for the MWA. This is determined in the following manner, according to what we described earlier. We considered a 1-yr survey over 2π sky area. For an MWA field of view of $\sim 600\,\mathrm{deg}^2$, this gives $20000/600 \sim 33$ pointings. We required each pointing to contain 2 epochs, where each epoch will contain many snapshot images, so the total time integrated over each pointing per epoch is $(1\,\mathrm{yr})/(2 \times 33) \sim 130\,\mathrm{hr}$. Scaling 0.9 mJy from 1 hr of integration to 130 hr gives the epoch noise $\sigma_e = 0.08\,\mathrm{mJy}$ used to set $f^* = 7\sqrt{2}\sigma_e = 0.8\,\mathrm{mJy}$. However, our analysis in Chapter 3 suggested that the measured MWA image noise is closer to 30 mJy after 2 m of integration rather than the theoretical expectation of 10 mJy in 30 s [46]. Using the measured value to compute σ_e and f^* in the same manner, we obtain $f^* = 5\,\mathrm{mJy}$, which lowers the rates by about a factor of 10. Nonetheless, this could improve in the future with better calibration techniques and the planned upgrades for the MWA. We did not perform the same analysis with the other instruments as uncertainties in the sidelobe confusion noise, the primary beam,

and calibration are difficult to estimate theoretically, and some of these instruments have yet to be built.

In summary, instruments such as LWA1, LOFAR, and MWA are less ideal for afterglow detection compared to CHIME and ASKAP, which operate at higher frequencies and have better sensitivities that give them the potential to constrain rate predictions for CBC. CHIME, in particular, conducts an all-sky survey on a daily basis by construction, unlike the other instruments that might have to divide their duty cycle among different observing modes, making it an excellent instrument to search for afterglow-like transients.

6.4.2 Gamma-Ray Burst Follow-up

SGRB follow-up observations will be sensitive only to the population of on-axis afterglows. To calculate $R_{\rm det}$ for on-axis afterglows, we used the constraint $0 \leq \theta_{\rm obs} \leq \theta_{\rm jet}$ instead of $0 \leq \theta_{\rm obs} \leq \pi/2$ when we integrated Equation 6.8. $\theta_{\rm jet}$ is 11.5° for our simulated lightcurves (Table 6.1), consistent with the observed values of $\theta_{\rm jet}$ although the uncertainty is quite large (see [128] and references therein). A larger $\theta_{\rm jet}$ implies a larger fraction of on-axis afterglows.

Instead of estimating the number of SGRBs that the radio instruments can detect given a SGRB trigger from a γ-ray telescope, such as *Swift* or *Fermi*, we determined the fraction of on-axis bursts present in the blind surveys as shown in Table 6.4. These are bursts that, in principle, could have γ-ray counterparts that might trigger a γ-ray telescope. While we accounted for the instantaneous fields of view of *Swift* and *Fermi* [142], we did not consider other factors, such as the systematic uncertainties of γ-ray detectors and selection effects, that could lower the fraction of bursts detectable by γ-ray telescopes. The fraction of on-axis afterglows increases with decreasing circumburst density because as the synchrotron emission becomes weaker, the radio instruments become less sensitive to off-axis afterglows and detect only the population that is more on-axis.

As some afterglows detectable in a blind radio survey are on-axis and could have γ-ray counterparts, coincident detections at different wavelengths could increase the

Table 6.4: Fraction of on-axis radio afterglows in a blind survey. These could have γ-ray counterparts that trigger a γ-ray telescope. We accounted for the fields of view of *Swift* (1.4 sr) and *Fermi* (9.5 sr), which cover 0.867 of the entire sky without overlapping regions [142], but not any systematics specific to γ-ray detection.

Instrument	$10^{-5}\,\mathrm{cm}^{-3}$	$10^{-3}\,\mathrm{cm}^{-3}$	$1.0\,\mathrm{cm}^{-3}$
LWA1	0.58	0.22	0.002
LOFAR Low	0.58	0.22	0.002
LOFAR High	0.60	0.37	0.01
MWA	0.60	0.37	0.01
CHIME Path	0.68	0.49	0.04
CHIME	0.66	0.46	0.04
ASKAP	0.72	0.55	0.08
HERA	0.57	0.32	0.01
SKA1 Low	0.57	0.31	0.01

significance of weak signals. While *Swift* detections are well-localized, *Fermi* detections often have large localization uncertainties (10–100 deg^2). Consequently, many *Fermi* detections do not have follow-up observations at other wavelengths (cf. [143]). Current widefield radio instruments could develop SGRB follow-up strategies with *Fermi* in conjunction with GW follow-up strategies. In particular, instruments such as CHIME and ASKAP might detect a non-zero number of afterglows in a SGRB follow-up. While the measured redshifts of some SGRBs suggest a typical $z \sim 0.5$, most SGRBs lack redshift measurements and therefore could be closer; non-detections could constrain or confirm certain parameters of the bursts, e.g. opening angle, circumburst density, distance (lower limit). Even if no SGRB afterglows are detectable by the widefield radio instruments, understanding the background rate of afterglow-like transients in such a follow-up is valuable for the GW follow-up observations, for which the events are close-by.

6.4.3 Gravitational Wave Follow-up

GW events lie within the horizon distance of the GW detector. If this distance is less than d_{H} of the radio instrument performing the follow-up observation, all GW events will be detectable, but this is more often not the case.

To estimate the expected number of afterglow detections given the detection of a

GW event by aLIGO or a similar ground-based detector, we assumed that a GW signal from an optimally located and oriented binary could be detected up to 450 Mpc, the design sensitivity of aLIGO for binary neutron stars [121]. In general, the distance at which the GW signal from a CBC is detectable depends on the location of the binary on the sky, the inclination, and the polarization angle. The horizon distance for the radio signal depends on the jet energy, the jet orientation, and the circumburst density. While the current models of jet formation predict that the SGRB jet is likely to be aligned with the inclination of the binary (binary disk is face-on when the jet is on-axis [70, 144]), the question is far from settled observationally. Consequently, we explored two cases: (i) the jet is aligned with the binary inclination, (ii) the jet is uncorrelated with the binary inclination. For both cases, we computed the fraction of detected GW events that are also detectable by a radio instrument, averaging over all intrinsic and extrinsic parameters of the binary and the SGRB. See Appendix D for the calculation.

The calculation of $R_{\rm det}$ for GW follow-up is otherwise similar to that of a blind survey. Unlike a blind survey, $\tau_{\rm sur}$ is divided by the expected number of aLIGO events or the number of pointings needed to cover the entire sky, whichever is smaller. For each event, we used the expected aLIGO localization error as $\Omega_{\rm sur}$, choosing $100\,{\rm deg}^2$ as our value [78]. Most of the radio instruments, however, cover this error box in one pointing. If aLIGO detects many events such that the sky surface density is high, the observation strategy for a radio instrument then becomes indistinguishable from that of a blind survey. Table 6.5 presents estimates of the fraction of aLIGO events that various radio instruments can detect. Following up all the aLIGO events requires at least two telescopes, one in the northern and one in the southern hemisphere. Future radio instruments that might be operating at the same time as aLIGO will likely be able to detect most or all of the aLIGO events, except when bursts occur at the lowest densities.

Table 6.5: Average fraction of BNS events detectable by both aLIGO and radio follow-up observations. The fraction is normalized to the 2π sky area accessible to a radio telescope except for HERA. The horizon distance that we used for aLIGO is 450 Mpc, the value for an optimally located and oriented BNS system for aLIGO at design sensitivity. Results for two cases are listed $\star/\star$: (the jet is aligned with the binary inclination) / (the jet is uncorrelated with the binary inclination). HERA is suboptimal for aLIGO follow-up observations because it is a drift-scan telescope that cannot cover the whole sky, unlike CHIME, which sees the whole northern hemisphere. These results are optimistic as we assumed that the instruments would dedicate 100% of their time to the follow-up observations.

Instrument	$10^{-5}\,\mathrm{cm}^{-3}$	$10^{-3}\,\mathrm{cm}^{-3}$	$1.0\,\mathrm{cm}^{-3}$
LWA1	0.001 / 0.005	0.01 / 0.03	0.01 / 0.01
LOFAR Low	0.002 / 0.01	0.01 / 0.05	0.01 / 0.01
LOFAR High	0.02 / 0.05	0.19 / 0.27	0.58 / 0.58
MWA	0.02 / 0.05	0.18 / 0.27	0.57 / 0.57
CHIME Path	0.05 / 0.10	0.30 / 0.36	0.96 / 0.96
CHIME	0.12 / 0.15	0.55 / 0.54	1.0 / 1.0
ASKAP	0.10 / 0.12	0.43 / 0.45	1.0 / 1.0
HERA	0.02 / 0.02	0.09 / 0.09	0.13 / 0.13
SKA1 Low	0.16 / 0.17	0.68 / 0.67	1.0 / 1.0

6.5 Discussion

EM follow-up of GW candidates is important to advance the study of their progenitors. However, the next generation of GW detectors will have large localization errors during the early days of their operation, presenting a challenge for EM follow-up efforts. As shown in this work, certain widefield radio instruments have the capability to perform follow-up observations of SGRB afterglows, a possible EM counterpart of GW events.

While previous studies of SGRB radio afterglows argued that these events are too faint and long-lasting to be detectable by present and planned instruments [32, 33, 85], this work showed a spread in the distributions of peak fluxes and durations of these afterglow light curves generated from the numerical tool BOXFIT developed by [86] within a plausible range of model parameters (Section 6.2). These distributions are consistent with previous estimates, showing that most afterglows are faint ($\lesssim \mu$Jy) and long-lasting ($\gtrsim$ yr). However, there is a tail of bright ($\sim$ mJy) and short ($\sim$ yr) afterglows that could be detectable by current and future radio instruments. This tail,

however, is sensitive to the model parameters, such as $E_{\rm iso}$ and $\theta_{\rm jet}$, many of which are currently uncertain. The results in [85], who considered SGRB radio afterglows as triggers to GW searches, are more pessimistic than ours because they explored the low energy and high density ends of the plausible afterglow parameter space. Future studies exploring the dependence of the properties of radio afterglows on the various model parameters, such as a wider range of $\theta_{\rm jet}$, are needed. These late-time afterglows could also be a factor of a few brighter than previously expected [140], but this effect is not included in our study.

To characterize the detectability of these afterglows, we derived a criteria on peak flux and duration based on a simple matched filter technique for radio instruments in the presence of thermal noise and constant noise from background source confusion (Section 6.3). The actual sensitivies of radio instruments will be limited by other sources of error, such as calibration errors, primary beam errors, sidelobe confusion, etc. These are specific to the analyses performed with each instrument and are beyond the scope of this work. Nonetheless, our criteria provides an order of magnitude estimate for SGRB afterglows that could be detectable by these instruments. The actual rates measured by these instruments will be lower because of instrumental systematics. False positives from intrinsic variability of other sources, such as active galactic nuclei (AGN) variability, might make the interpretation of detected transient signals difficult but could be distinguished using counterparts at other wavelengths, e.g. the optical counterpart of the AGN. The counterparts of the afterglows at other wavelengths will no longer be visible on the timescale of the radio afterglows, so their nature can only be inferred through the process of elimination, i.e. that they are not other variable sources with persistent counterparts at other wavelengths. Identifying them as CBC requires that they are associated with GW emission, otherwise the nature of detected afterglow-like transients can only be inferred based on circumstantial evidence.

Converting the detectability criteria into a horizon distance, we estimated the rates of SGRB afterglow detection expected for various radio instruments performing three types of surveys under ideal conditions: blind surveys, SGRB follow-up observations,

and GW follow-up observations (Section 6.4). Given the context of EM follow-up of GW events, we assumed the intrinsic rate of the progenitors of SGRBs to be equal to that of BNS coalescence as summarized in [121]. Blind all-sky surveys with instruments such as CHIME and ASKAP will be able to constrain the rate predictions for BNS coalescence. They will also be able to characterize the background radio transients for future follow-up observations and perform independent studies of afterglows. A large fraction of afterglows in these blind surveys will also be on-axis bursts, suggesting that many detectable radio afterglows could have γ-ray counterparts that might trigger γ-ray telescopes. Coincident detections at different wavelengths could increase the significance of weak signals. Furthermore, *Fermi* detections of SGRBs could have large localization errors not unlike those of aLIGO. Strategies on SGRB follow-up observations with current radio instruments could thus be similar to the strategies on GW follow-up observations with future instruments that will likely have the ability to detect most or all of the aLIGO events.

The results of this work are consistent with the known limits placed on the surface density of radio transients (see [31, 110] for a summary). At the very most, CHIME or SKA1 could detect 5 SGRB afterglows $\mathrm{deg}^{-2}\ \mathrm{yr}^{-1}$ (upper limit on the rate of BNS coalescence) on the μJy level, but no radio surveys have reached that sensitivity yet. Furthermore, the upper limit rate of BNS coalescence is very unlikely.

This work is also complementary to other work on the detectability of long GRB afterglows. [123] and [124] considered the detectability of radio afterglows from on-axis and orphan long GRBs respectively over a wide range of frequencies. Their results are more pessimistic than ours for the following reasons: the rate of long GRBs that they used is roughly a factor of 10 lower than the realistic CBC rate; the circumburst densities that they explored are much higher than the densities we explored (1–$30\,\mathrm{cm}^{-3}$ compared to 10^{-5}–$1\,\mathrm{cm}^{-3}$), which is appropriate for long GRBs but synchrotron self-absorption is stronger at higher densities; the microphysics parameters ϵ_e and ϵ_B that they used have lower values and would thus reduce the radio flux. This shows the sensitivity of the results on the choice of model parameters, which are motivated by and consistent with observations but remain highly uncer-

tain. Consequently, orphan afterglow searches with radio instruments might also be able to constrain some of these parameters, such as $\theta_{\rm jet}$. Furthermore, long GRBs could be a background to future GW follow-ups for the pessimistic and realistic rate predictions of BNS coalescence. Well-sampled radio light curves with afterglow modeling or observations at other wavelengths would be necessary to distinguish the two populations.

Chapter 7

Conclusion

In this thesis we have developed a transient detection technique based on matched filters to search for transients in the presence of classical confusion noise. It searches for the light curve template that best matches the flux variation above a constant signal in an individual pixel. The criterion for identifying transient candidates is set by the transient detection statistic ρ, which follows a well-defined distribution characterized by σ_ρ. The empirical background distribution of $\tilde{\rho} \equiv \max(\rho/\sigma_\rho)$ determines the probability of false alarm P_{FA} (reliability), which establishes the threshold $\tilde{\rho}^*$ above which an event is classified as a transient candidate.

For every pixel, $\tilde{\rho}^*$ can be converted to a flux sensitivity according to $A^* = \tilde{\rho}^*/\sigma_\rho$. As different pixels have different noise properties, A^* varies across the image, but the significance of the detection remains the same. The median flux sensitivity, as computed from all the pixels, provides an estimate of the flux sensitivity of the search but is not a strict threshold below which nothing is detectable. The efficiency (completeness) is determined by recovering injected transients with fluxes and light curve parameters drawn from known or expected astrophysical distributions.

We have applied our technique for the first time to real data and demonstrated that our technique performs well despite the presence of residual sidelobe confusion noise and calibration errors. We have performed an example search, using the MWA, for transients that have light curves resembling top-hats with a duration of 15 days. For this type of transients, our technique is capable of detecting transients with fluxes

$\sim 25.0\,\mathrm{mJy}$ at $P_{FA} \leq 10^{-3}$ for the experiment. As our technique identifies transient candidates by applying a cut on $\tilde{\rho}$ and not flux, it remains sensitive to fainter transient sources at the same significance level but a lower efficiency. This is in contrast to flux-limited source detection techniques, such as the $\sim 210\,\mathrm{mJy}$ limit from one of the most constraining transient searches to date at 182 MHz [110].

Our technique also accurately recovers the injected transient light curve parameters. However, as mismatch between the template and the data increases, the accuracy and the significance decrease; one needs to expand the search to include additional templates. The search, however, will be sensitive to different templates in different ways, depending on the properties of the data, e.g. there might be more systematic effects on certain timescales, thus reducing the sensitivity to transients on those timescales. Nonetheless, our technique characterizes the complexity of the search and works for any kind of template.

The ability to detect fainter transients in the presence of classical confusion noise increases the transient parameter space that a particular instrument can explore. As calibration techniques and primary beam modeling improve, one could push the limits of an instrument even further to study astrophysical transient sources that might be missed by source-finding algorithms. Our technique is also applicable to non-confusion-limited instruments and provides a way to study fainter source variations.

We used this technique to search for transients in 3 months of MWA data. We ran three separate blind searches, using top-hat templates, to probe transients on different timescales: minute, hour, and day-to-month. For each search, we first characterized the background distribution of $\tilde{\rho}$, which allowed us to set the threshold $\tilde{\rho}^{\star}$ above which events were considered to be transient candidates. This threshold corresponded to a false alarm probability of 10^{-3}, i.e. the probability that a candidate is a false positive (reliability) in the entire search is $< 10^{-3}$. Then we characterized the efficiency of each search, or completeness, by running transient injections. The injections also demonstrated that we were able to recover transient properties accurately, even if the light curve profile of the injected transient differed qualitatively from the light curve template used in the search.

After processing the data, we initially found 4 candidates out of two searches. However, closer inspection suggested that they were more likely to be sidelobe artifacts rather than astrophysical sources. They were within $5'$–$10'$ of a bright source ($> 5\,\mathrm{Jy}$) in an image that had a synthesized beam of $\sim 2'$. This was caused by inadequate source masking, which we had introduced to flag known sources with fluxes $> 100\,\mathrm{mJy}$ because they were affected by primary beam systematics. In our initial search, the masking size was identical for every source, but brighter sources have larger sidelobes. We determined an empirical relation that scaled the masking size to the source brightness and reran the search with the new source masks. This produced excellent agreement in the $\tilde{\rho}$ distribution between the search region and the background expectation. We found no transient candidates.

Given that we did not detect any transient candidates, we set an upper limit on the transient surface density for each of our searches. We took into account the search efficiency in our upper limit calculation, thus we were able to push to fainter fluxes than would otherwise be available. We reported improved limits at fluxes between $\sim$ 20–200 mJy for hour- and month-long transients. This is consistent with reported transient detections in the literature, and it will easily improve with more data.

In the last part of this thesis, we used the matched filter technique we developed for radio transient detection to characterize the detectability of late-time radio afterglows from compact binary coalescence. These are predicted EM counterparts of GW sources, the detection of which could break degeneracies in GW measurements and test progenitor models for SGRBs. However, they are faint at low radio frequencies, where the instruments have extremely wide fields of view suitable for GW follow-up observations. The matched filter technique we developed is applicable to an EM follow-up search as it is capable of probing below the classical confusion noise, which affects many widefield radio instruments, and it can also make use of the simulated light curve templates of SGRBs. We found that the detectability of afterglows depended on the properties of the burst: the ones that are more energetic, more on-axis, and occurring in more dense medium are more detectable. We also concluded that instruments such as CHIME and ASKAP would be able to detect many afterglows

and therefore constrain the rate of these events, which is still quite uncertain.

The matched filter technique is powerful, and we have successfully adapted it for slow radio transient detection. Future transient searches will be able to probe a larger phase space more efficiently and with well-characterized statistics. There is still much to explore in the transient sky. Along with it is the great synergy between time domain astronomy and the new field of gravitational wave astronomy. More discoveries will occur as we continue to expand our knowledge about the ever-changing universe.

Appendix A

Simetra

Simetra is a transient search pipeline based on matched filter and written in `Python`, available on GitHub[1]. Its main script is `artemis.py` but it includes custom dependencies `simetra.py`, `injection.py`, and `mfilter.py`. It uses secondary scripts `join_tables.py`,
`join_tables_vmask.py`, `fit_false_alarm.py`, and an optional `mpix2tb.py`, which will all be described in detail below. It also requires `Astropy`, `numpy`, `scipy`, and `matplotlib`. Here we outline the basic steps and options on how to run *Simetra*.

1. `--template`: Choose light curve template.
   Simetra is based on matched filter, so the user needs to choose a light curve template before running a search. At the time of this writing, the valid options are the top-hat template and the power-law template. The user can define additional templates in `simetra.py` and add the option to `artemis.py`. By default, *Simetra* will iterate over all possible template parameters, e.g. all possible durations for the top-hat template, but the user can choose to run *Simetra* on a single template (`--single`) and also specify which template to use (`--which_template`).

2. `--section`: Define image regions: background, injection, playground, search.
   Due to memory limitations, *Simetra* cannot load the pixel light curves of the

[1]https://github.com/lufeng5001/simetra

entire image at once, e.g. $4096 \times 4096 \times 1251$ flux measurements. Instead, it processes sub-images and joins them together at the end. The user defines each sub-image by specifying the bottom-left and top-right pixel coordinates; this thesis defines each sub-image as 86×86 contiguous pixels.

3. `--ioconversion`: Convert FITS images to time series.

   Simetra reads the sub-images (`--images`, `--beams`) from both sky and primary beam images (FITS format) and returns the data as a set of time series for every pixel (`Python npz` format). The time series data are stored on-disk instead of in-memory because of I/O limitations that come with loading the FITS files repeatedly. If no primary beam files are specified, the primary beam correction is set to 1 for all pixels. In addition to the pixel light curves and the primary beam values, the `npz` file contains the noise for each sub-image. Unless the image noise is specified in the FITS header under the keyword `NOISE`, *Simetra* estimates the image noise by computing the median absolute deviation (MAD) of each sub-image and converting this quantity into the standard deviation according to $\sigma_{\rm im} = 1.4826 \times \mathrm{MAD}$. Note: `--images` must be specified if the user is running `ioconversion`, but it is unnecessary after the `npz` files are created; `artemis.py` will look for the `npz` files (in the working directory) created during `ioconversion` so do not rename them.

4. (optional) `--injection`, `--injnpz`: Inject transient light curves into the data. For an injection run, the user first needs to create an injection file (`Python npz` format), which can be done by running `injection.py`. This file contains the light curve template type, the amplitude(s), the start time(s), and the parameters characterizing the template, e.g. duration(s) for the top-hat. The user may specify the injection parameters directly or draw them randomly from a particular distribution. If there are multiple templates in this file, the sub-images will be split into subsections to accommodate all of the injection templates. *Simetra* generates an injected light curve for every pixel on the fly according to the parameters in the injection file and stores these parameters in an `istore`

file (FITS table), which lets the user easily match the injection parameters for each pixel to its transient search results.

5. Matched filter calculation: Calculate $\rho_{\max}$ and σ_ρ for every pixel.

   This is the heart of *Simetra*. The search template can be different from the injection template. If the search will iterate over multiple templates, the user can specify the increment at which to probe the template phase space (`--dtshift`), e.g. durations in steps of 10 min instead of 2 min; the finest resolution gives the best accuracy but is the most computationally intensive. The start time is incremented by 10% of the duration or to the start of the next observation window, whichever is smaller. *Simetra* generates the phase space of the input light curve template type, e.g. a list of durations for the top-hat, as well as a list of start times, and iterates over them. For each template and pixel, *Simetra* calculates ρ and σ_ρ and saves the template with the most significant ρ/σ_ρ to a FITS table, which also includes the start time and other template parameters.

6. (optional) Create a FITS file to mask certain pixels.

   This step is useful if there are particular pixels that the user would like to include or exclude from the transient search. The user needs to create a FITS image with the same dimensions as the full image processed through *Simetra*, and set the included pixels to 0 and the excluded pixels to 1. Then the user converts this FITS file into a FITS tables by running `mpix2tb.py`, which matches each table row to the corresponding pixel output by *Simetra*; the user must give the list of sub-image pixels in the same order as they have specified it for *Simetra*.

7. Join sub-images.

   As mentioned before, *Simetra* processes sub-images. It joins them together at the end via `join_tables.py`. After this, the user may also choose to mask certain pixels by doing the previous step and then running `join_tables_vmask.py`. This thesis uses the latter ability to flag pixels where the primary beam is poorly modeled and where there are bright (> 100 mJy) sources; in addition, it also uses this ability to define the injection and playground regions for easier book-

keeping. The same scripts work on the `istore` files.

8. Fit for the probability of false alarm P_{FA}.

   The script `fit_false_alarm.py` fits a function to the tail of the cumulative background distribution of $\rho_{\max}$ and extrapolates the threshold $\tilde{\rho}^{\star}$ from the fitted function for a given value of P_{FA}. The user should run this script on the *Simetra* output for the playground region. The user needs to choose a value for P_{FA} (the default 10^{-3} corresponds to a significance of 3.3σ), set the number of tail elements to include in the fit (recommended values are 100 to 500), and specify the number of synthesized beams in the search region so that the code correctly accounts for the trials factor; if each synthesized beam contains more than 1 image pixel, the script allows for this correction as well. At the time of this writing, the script only fits for an exponential function, but the user should choose a more appropriate function if the exponential fit is inadequate; note, however, that the threshold calculation may change if the fitted function form is changed (see Equation 4.1).

9. Post-processing: Characterize the efficiency and identify transient candidates.

   This is not officially part of *Simetra* because users can define their own analysis using the transient detection statistic $\tilde{\rho}$, but the user may follow the procedure outlined in Chapters 4 and 5 to do a traditional transient search.

Appendix B

CASApipe

The image reduction pipeline is a shell script that outlines the work flow and calls custom programs based on CASA[1] (v4.1.0) [99], `Python`, or other publicly available software developed for radio astronomy, e.g. WSCLEAN [42] and AEGEAN (v951) [52]. The flow is controlled by a configuration file that the user can modify. In this configuration file, the user can choose which parts of the pipeline to run, e.g. calibrate but do not image, and also set options for various CASA or WSCLEAN tasks, e.g. the CASA task `bandpass`. This design makes it easy for the same program to be adapted for different analysis goals without the need to change the underlying code.

Below, we outline the pipeline, marking the options that the user can change with [square brackets] and listing the values that we have used for this thesis. Note that all CASA task options are set by the user in the configuration file, and the options that we have left as the CASA default will not be listed explicitly. Each numbered step can be turned on or off, depending on what the user wants to do, but the lettered steps run together as a block. While Step 0 is not part of the pipeline, it provides the starting point and is included in the outline. The code is available on GitHub[2].

0. Cotter [95]: Convert the raw interferometric data from the MWA format to the more standard UVFITS format phased to the pointing center.

[1] http://casa.nrao.edu/
[2] https://github.com/lufeng5001/casapipe

(a) Flag bad data: RFI (AOFlagger [96, 97]), first 4 s, 80-kHz edge channels, center DC channels.

(b) Average data: 0.5-s to 1-s time integration, 40-kHz to 80-kHz frequency resolution.

1. CASA `importuvfits`: Convert the UVFITS file into a measurement set.

2. Update the measurement set.

   (a) `fixmwams`: Add MWA-specific keywords.

   (b) Update the `ANTENNA` table with the correct antenna locations, which are not parsed correctly by `importuvfits` for unknown reasons.

3. `chgcentre`: Rotate the phase center[3] from the pointing center to [the EOR0 field center].

4. CASA `flagdata`: Flag additional edge channels (80 kHz to [240 kHz]) affected by aliasing.

5. `delaycal`: Build a point source calibration model.

   (a) Apply the (analytical) primary beam model to [the MWA Commissioning Survey Catalog] to get the apparent fluxes.

   (b) Build a CASA component list with the [11] brightest sources with XX and YY fluxes converted to Stokes I and Q.

   (c) CASA `ft`: Transform the component list into model visibilities.

6. CASA `bandpass`: Generate a bandpass calibration solution.

7. CASA `applycal`: Apply the calibration solution.

8. Perform [1] round of self-calibration.

   (a) WSCLEAN: Generate model visibilities based on clean components for each polarization.

[3]FHD [104] uses the same UVFITS files with the pointing center, so we have an extra step.

(b) CASA `bandpass`: Generate the selfcal solution.

(c) CASA `applycal`: Apply the selfcal solution.

9. WSCLEAN: Image [4096×4096] pixels in XX and YY polarizations with [uniform] weighting, [0.5′] pixel size, [10^4] iterations.

10. Apply the [Curtin] primary beam correction.

    (a) `make_beam.py`: Generate the XX and YY primary beam models for the middle of the observation.

    (b) Generate Stokes I and Q images.

11. Run the AEGEAN source finder.

    (a) Determine the image RMS with AEGEAN.

    (b) Mask the pixels with RMS > [20 Jy] to prevent AEGEAN from finding false sources toward the null of the primary beam.

    (c) Rerun AEGEAN with the default settings to generate a source list.

To apply the average calibration solution, replace Steps 5–8 with Step 7. Steps 10 and 11 can be turned off if we are only interested in the empirical beam and not the Curtin beam.

Appendix C

Thermal Sensitivities of CHIME and HERA

CHIME is a drift-scan telescope with cylindrical reflectors oriented in the north-south direction without any moving parts. Its instantaneous field of view can be approximated as a 2.5° narrow band spanning the whole sky from north to south. As the Earth rotates, the telescope observes effectively half of the sky every day. The integration time of a source is a function of its declination. As a result, the instrument sensitivity will vary with source declination. In order to account for this, we computed the integration time as a function of declination $t_{\text{int}}(\delta)$.

The instantaneous field of view of CHIME is modeled by two intersecting planes with the angle between their normal directions $\Delta \approx 2.5°$ defining the aperture in the east-west direction. The slice of the celestial sphere between the two planes defines the instantaneous field of view of the telescope. Figure C-1 shows the visualization of the field of view of CHIME, which we assumed to be located at a latitude $\phi = 45°$.

The integration time during a single pass of the source across the field of view is proportional to the angle β subtended by the arc between the planes, designated as $\beta = \angle CBD$ on Figure C-1. When expressed as a fraction of the full 24-h day (corresponding to a single complete revolution), the integration time is given by

$$t_{\text{int}} = \frac{\beta}{2\pi} \text{ day} \tag{C.1}$$

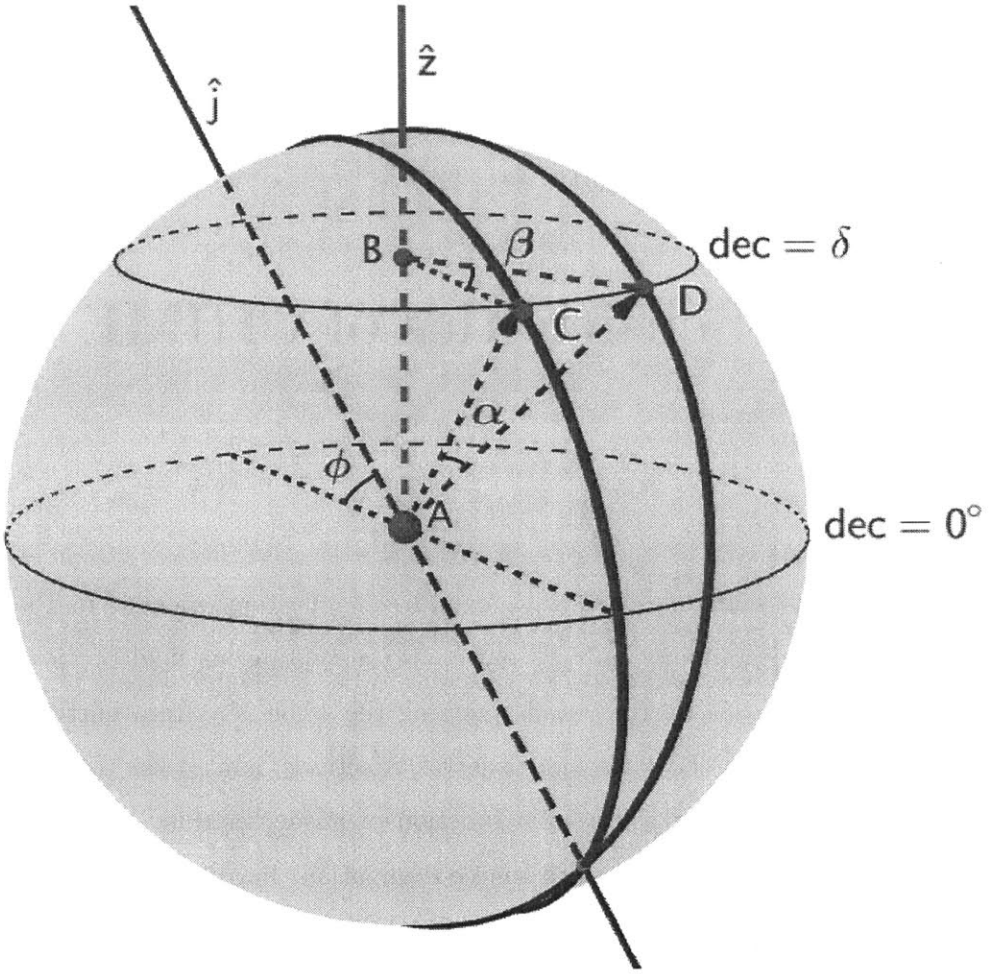

Figure C-1: Visualization of CHIME observations. The CHIME field of view is defined by two planes intersecting at an angle $\Delta = 2.5°$, where the line of intersection, shown as axis $\hat{j}$, is oriented in the north-south direction and makes an angle $\phi = 45°$ with the z-axis. The intersections of these planes with the celestial sphere, indicated by the thick solid arcs, defines the instantaneous field of view. As the Earth rotates, sources come in and out of the field of view, and the typical daily trajectory of a source at Dec $= \delta$ is shown as the circle through points C and D. This source is observable only when it is located on the short arc between points C and D, and the ratio of the angle subtended by this arc, $\beta = \angle CBD$, to 2π gives the fraction of the day during which this source is observable. The total integration time thus increases with δ in the northern hemisphere; the sources in the southern hemisphere are either observable for only a very short time or inaccessible. The sources with $\delta > 45°$ are observed twice a day, and the sources with $\delta \gtrsim 88.75°$ are always in the field of view.

Using the fact that triangles CBD and CAD share the same side, CD, one can express $\beta = \angle CBD$ in terms of the angle $\alpha = \angle CAD$ and the source declination δ:

$$\cos\beta = 1 - \frac{1 - \cos\alpha}{\cos^2\delta} \tag{C.2}$$

In deriving this expression, we used $BD = BC$, $AC = AD$, and $BC = AC\cos\delta$. Next, we found $\cos\alpha$ by computing the scalar product between the two unit vectors $\mathbf{a}_1 = \mathbf{AC}/|\mathbf{AC}|$ and $\mathbf{a}_2 = \mathbf{AD}/|\mathbf{AD}|$. These vectors lie in the planes defining the field of view. The angle between the planes is Δ. Vectors $\mathbf{a}_1$ and $\mathbf{a}_2$ are identical up to a rotation by Δ around the axis defined as the line of intersection of the planes, axis $\hat{j}$ in Figure C-1. Thus

$$\cos\alpha = (\mathbf{a}_1, \mathbf{a}_2) = \sin^2(\phi - \delta) + \cos\Delta\cos^2(\phi - \delta) \tag{C.3}$$

Using Equations C.3 and C.2 in C.1, we arrived at the final expression for the integration time in a single transit:

$$t_{\rm int}(\delta) = \frac{1}{2\pi}\arccos\left[\frac{-\sin^2\delta + \sin^2(\phi - \delta) + \cos\Delta\cos^2(\phi - \delta)}{\cos^2\delta}\right] \tag{C.4}$$

The problem of calculating the integration time for a cylindrical telescope with an arbitrary location and orientation was treated in [145]. Our derivation is different from the approach used in [145]. As a result, our formula for $t_{\rm int}$ given by Equation C.4 does not match in its functional form the formula given in [145]. However, we verified numerically that both formulas, ours and [145], lead to identical results.

The sources with $\delta \leq \phi$ are observed once a day. The sources with $\delta > \phi$ transit twice a day through the field of view of the telescope. The second pass occurs on the opposite side of the northern hemisphere as shown in Figure C-1, so its integration time is not the same as the one for the first transit. The integration time for the second pass can be computed by replacing $\phi - \delta$ with $\phi + \delta$ in Equation C.4. The total integration time for such a source is the sum of the two integration times. Finally, the sources that are very near the north pole ($\delta \gtrsim 88.75°$) are always in the

field of view of the telescope.

Substituting Equation C.4 into Equation 6.1, we can now compute the thermal noise as a function of δ:

$$\sigma_{\rm CHIME}(\delta) = \left(\frac{2k_{\rm B}T_{\rm sys}}{A_{\rm eff}N_{\rm ant}\epsilon_c}\right)\frac{1}{\sqrt{N_{\rm pol}B\,t_{\rm int}(\delta)}}, \tag{C.5}$$

where we set $T_{\rm sys} = 100\,{\rm K}$, $\epsilon_c = 1.0$, $N_{\rm pol} = 2.0$, $B = 400\,{\rm MHz}$, and $A_{\rm eff}N_{\rm ant} = 1500\,{\rm m}^2$ for the CHIME pathfinder and $A_{\rm eff}N_{\rm ant} = 10000\,{\rm m}^2$ for the full CHIME. The latitude is set to $\phi = 45°$ and the opening angle $\Delta = 2.5°$. When averaged over all declinations in the northern hemisphere $0° \leq \delta \leq 90°$, the average thermal noise in a 1-day observation is

$$\langle\sigma_{\rm CHIME\,path}\rangle = 0.240{\rm mJy} \tag{C.6}$$

$$\langle\sigma_{\rm CHIME}\rangle = 0.036{\rm mJy} \tag{C.7}$$

for the CHIME pathfinder and the full CHIME respectively.

When estimating the expected number of detections for CHIME, we needed to average the horizon distance in Equation 6.8 over all declinations. Since $d_{\rm H} \propto \sigma_{\rm im}^{-1/2}$, one should average the noise taken in the power of $-3/2$, $\left\langle\sigma_{\rm CHIME}^{-3/2}\right\rangle$, which differs from the average noise taken in the same power by a factor of 1.19,

$$\left\langle\sigma_{\rm CHIME}^{-3/2}\right\rangle = 1.19\langle\sigma_{\rm CHIME}\rangle^{-3/2} \tag{C.8}$$

In calculating the expected number of detections for CHIME, we used Equation C.6 for the CHIME pathfinder and Equation C.7 for the full CHIME in Equation 6.8, and then corrected the results by multiply them by the factor of 1.19.

HERA is a drift-scan telescope with a 8.7° primary beam full-width at half-maximum. As it will be located in South Africa, we assumed the telescope to be at the latitude $\phi = -30°$. Setting $\Delta = 8.7°$, we estimated $t_{\rm int}$ for a source at zenith ($\delta = \phi$) using Equation C.4. We neglected the effects related to the circular shape of the field of view, which should be small. Because the area surveyed by the telescope

is relatively narrow in the north-south direction, we can neglect the dependence of $t_{\rm int}$ on δ. The resulting $t_{\rm int}$ for the sources observed by HERA in a single drift-scan observation is $t_{\rm int} = 0.67\,{\rm h}$. We computed the survey area accessible to the telescope by approximating it as a strip of sky centered on $\delta = -30°$ with width 8.7°: $\Omega_{\rm HERA} = 2712.4\,{\rm deg}^2$. Lastly, when computing the thermal noise for HERA, we set $T_{\rm sys} = 351\,{\rm K}$, $\epsilon_c = 1.0$, $N_{\rm pol} = 2.0$, $B = 100\,{\rm MHz}$, and $A_{\rm total} = N_{\rm ant}\pi r_{\rm ant}^2 = 84204\ {\rm m}^2$, where the antenna radius is $r_{\rm ant} = 7\,{\rm m}$. Using all of this in Equation C.5, we estimated the thermal noise for HERA in a 1-day drift scan observation to be

$$\sigma_{\rm HERA} = 0.017{\rm mJy} \tag{C.9}$$

Appendix D

Fraction of GW Events Detectable in Radio Follow-up Observations

The necessary condition for both the GW and the afterglow radio emission to be detectable is for the CBC to be within the detectability range of both instruments. The reach of the GW detectors can be expressed in terms of the horizon distance $D^{\mathrm{GW}}_{\mathrm{horizon}}$, i.e. the distance at which the signal from an optimally located and oriented CBC produces an event with the matched filter signal-to-noise ratio (SNR) of 8 in the detector. In the nominal regime, this condition corresponds to a 5σ detection. The SNR of a GW signal from the CBC is proportional to its amplitude, which in turn is inversely proportional to the physical distance to the CBC and depends on the location and the orientation of the binary relative to the detector. All of these parameters can be combined into a single factor called the effective distance, D_{eff}:

$$D_{\mathrm{eff}} = d\left[F_{+}^{2}(\theta,\phi,\psi)\left(\frac{1+\cos^{2}\iota}{2}\right)^{2} + F_{\times}^{2}(\theta,\phi,\psi)\cos^{2}\iota\right]^{-1/2} \tag{D.1}$$

where d is the physical distance to the CBC, the polar angles (θ,ϕ) define the position of the source on the sky in the detector coordinate system (for a detector with orthogonal arms, the x- and y-axes are aligned with the arms of the detector and the z-axis is perpendicular to the plane of the detector), $0 \leq \psi \leq 2\pi$ is the angle describing the polarization of the GW event, the inclination angle ι is the angle between the normal

vector to the binary disk and the vector pointing to the detector (the disk is face-on when $\iota = 0°$ and edge-on when $\iota = 90°$), and $F_+(\theta, \phi, \psi)$ and $F_\times(\theta, \phi, \psi)$ are the GW detector antenna beam patterns given by

$$F_+(\theta, \phi, \psi) = -\frac{1}{2}(1 + \cos^2\theta)\cos 2\phi \cos 2\psi - \cos\theta \sin 2\phi \sin 2\psi \tag{D.2}$$

$$F_\times(\theta, \phi, \psi) = +\frac{1}{2}(1 + \cos^2\theta)\cos 2\phi \sin 2\psi - \cos\theta \sin 2\phi \cos 2\psi \tag{D.3}$$

For an optimally located and oriented binary, $\iota = 0°$ and $F_+^2 + F_\times^2 = 1$, in which case $D_{\rm eff} = d$. In all other cases, $D_{\rm eff} > d$. Given that SNR $\propto 1/D_{\rm eff}$, the effective distance determines the strength of the GW signal from a CBC with an arbitrary location and orientation relative to the optimally located and oriented CBC. Applying the condition used to define $D^{\rm GW}_{\rm horizon}$ to a generic CBC

$$D_{\rm eff} \leq D^{\rm GW}_{\rm horizon} \tag{D.4}$$

and solving for d, we find the reach of a GW detector as a function of the CBC location and orientation:

$$d_{\rm GW}(\theta, \phi, \psi, \iota) = D^{\rm GW}_{\rm horizon} \times \left[F_+^2(\theta, \phi, \psi)\left(\frac{1 + \cos^2\iota}{2}\right)^2 + F_\times^2(\theta, \phi, \psi)\cos^2\iota \right]^{1/2} \tag{D.5}$$

From the definition of the effective distance, $d_{\rm GW}$ is always less than or equal to $D^{\rm GW}_{\rm horizon}$. Another often quoted measure of sensitivity of GW detectors is the inspiral range distance $d_{\rm range}$ defined by $d^3_{\rm range} = \langle d^3_{\rm GW} \rangle$, where the averaging is done over all possible locations and orientations. The inspiral range distance is related to the horizon distance by $d_{\rm range} = D^{\rm GW}_{\rm horizon}/2.26$. The inspiral range is used to compute the average sensitive volume of the GW detector and the expected rate of detections.

For a given circumburst density, the reach of a radio telescope is characterized by $d_{\rm H}(\theta_{\rm obs}, E_{\rm iso})$, which depends on the orientation of the jet and the energy. For a CBC to be detected by a GW detector and a radio telescope, it must be within the reach of both instruments. $E_{\rm iso}$ should not be correlated with any of the extrinsic parameters

of the binary $(\theta, \phi, \psi, \iota)$. Depending on the formation mechanism, the jet orientation $\theta_{\rm obs}$ may or may not be correlated with the inclination of the binary ι. In order to encompass all possibilities, we considered two limiting cases: (i) $\theta_{\rm obs} = \iota$, (ii) $\theta_{\rm obs}$ and ι are completely uncorrelated. Averaging over all parameters describing the CBC and the SGRB jet, we found that the average fraction of the GW events that could also be detected by a radio telescope in case (i) is

$$\left\langle \frac{N_{\rm radio}}{N_{\rm GW}} \right\rangle_{\rm (i)} = \left\langle \frac{d_{\rm H}^3(\iota, E_{\rm iso})}{d_{\rm GW}^3(\theta, \phi, \psi, \iota)} \Theta\left(1 - \frac{d_{\rm H}(\iota, E_{\rm iso})}{d_{\rm GW}(\theta, \phi, \psi, \iota)}\right) + \Theta\left(\frac{d_{\rm H}(\iota, E_{\rm iso})}{d_{\rm GW}(\theta, \phi, \psi, \iota)} - 1\right) \right\rangle_{(E_{\rm iso}, \theta, \phi, \psi, \iota)} \tag{D.6}$$

and in case (ii) is

$$\left\langle \frac{N_{\rm radio}}{N_{\rm GW}} \right\rangle_{\rm (ii)} = \left\langle \frac{d_{\rm H}^3(\theta_{\rm obs}, E_{\rm iso})}{d_{\rm GW}^3(\theta, \phi, \psi, \iota)} \Theta\left(1 - \frac{d_{\rm H}(\theta_{\rm obs}, E_{\rm iso})}{d_{\rm GW}(\theta, \phi, \psi, \iota)}\right) + \Theta\left(\frac{d_{\rm H}(\theta_{\rm obs}, E_{\rm iso})}{d_{\rm GW}(\theta, \phi, \psi, \iota)} - 1\right) \right\rangle_{(\theta_{\rm obs}, E_{\rm iso}, \theta, \phi, \psi, \iota)} \tag{D.7}$$

Note that in case (i), we explicitly imposed the condition $\theta_{\rm obs} = \iota$ and performed the averaging over a reduced set of parameters. The Heaviside step functions in Equations D.6 and D.7 impose the condition for the ratio of the reach of the radio telescope and the GW detector not to exceed one. The CBC with $d_{\rm GW} < d < d_{\rm H}$ can be detected only as a radio afterglow. Thus, in Equations D.6 and D.7, we set $d_{\rm H}/d_{\rm GW} = 1$ whenever $d_{\rm H}$ exceeds $d_{\rm GW}$.

Setting $D_{\rm horizon}^{\rm GW} = 450\,{\rm Mpc}$, which is the design sensitivity of aLIGO for BNS, in Equation D.5 and substituting the result into Equations D.6 and D.7, we computed the expected fraction of aLIGO events detectable by various radio telescopes.

For each telescope that we considered, we allowed a full year of observation that is split evenly between aLIGO events. The total number of independent observations required to follow up all accessible GW events is determined by two factors: the rate of GW events and the number of pointings necessary to cover the region of the sky localization uncertainty of GW detectors. When computing the total number of

observations, we accounted for the fact that the radio telescopes can typically access only half of the sky, which reduces the number of the GW events that can be followed up by half. We took a typical uncertainty in the localization of a CBC signal with the LIGO-Virgo network to be 100 deg^2. If a telescope required multiple pointings to cover such a region, we increased the number of observations for that telescope accordingly. Most of the low frequency radio telescopes have a sufficiently large field of view to cover the localization region of a GW source with a single pointing. Knowing the approximate locations of GW sources allows one to reduce the total area of the sky that needs to be observed. This gives advantage to the targeted follow-up observations over a blind, all-sky survey. However, if the density of GW events is high, a widefield telescope might end up covering the entire accessible sky in the process of the follow-up observations. It would be equivalent to performing the blind, all-sky survey. Thus, we set the total number of observations to be either the number of pointings required to follow up all GW events or the total of number of pointings necessary to cover 2π of the sky, whichever is smaller. The total observation time of 1 yr is divided evenly between the follow-up observations. In calculating σ_{e}, we set t_{int} to be half of the time allocated for each follow-up observation. As before, we allocated half of the single observation time for reference imaging, which is required to detect transients. We calculated d_{H} following the same procedure as in the case of a blind survey in Section 6.4.1. Evaluating the multi-dimensional integrals in Equations D.6 and D.7 numerically, we computed the average fraction of aLIGO events detectable in the radio follow-up observations for each telescope and three different circumburst densities: $n = 10^{-5}, 10^{-3}, 1.0\ \mathrm{cm}^{-3}$. The results are shown in Table 6.5.

Bibliography

[1] F. R. Stephenson and D. A. Green. A Catalogue of "Guest Stars" Recorded in East Asian History from Earliest Times to A.D. 1600. *Journal for the History of Astronomy*, 40:31–54, January 2009.

[2] A. Hewish, S. J. Bell, J. D. H. Pilkington, P. F. Scott, and R. A. Collins. Observation of a Rapidly Pulsating Radio Source. *Nature*, 217:709–713, February 1968.

[3] R. Giacconi, H. Gursky, F. R. Paolini, and B. B. Rossi. Evidence for x Rays From Sources Outside the Solar System. *Physical Review Letters*, 9:439–443, December 1962.

[4] R. W. Klebesadel, I. B. Strong, and R. A. Olson. Observations of Gamma-Ray Bursts of Cosmic Origin. *ApJL*, 182:L85, June 1973.

[5] D. Thornton, B. Stappers, M. Bailes, and et al. A Population of Fast Radio Bursts at Cosmological Distances. *Science*, 341:53–56, July 2013.

[6] N. M. Law, S. R. Kulkarni, R. G. Dekany, and et al. The Palomar Transient Factory: System Overview, Performance, and First Results. *PASP*, 121:1395–1408, December 2009.

[7] H. V. Bradt, R. E. Rothschild, and J. H. Swank. X-ray timing explorer mission. *AAPS*, 97:355–360, January 1993.

[8] N. Gehrels, G. Chincarini, P. Giommi, and et al. The Swift Gamma-Ray Burst Mission. *ApJ*, 611:1005–1020, August 2004.

[9] J. M. Cordes, T. J. W. Lazio, and M. A. McLaughlin. The dynamic radio sky. *New Astronomy Reviews*, 48:1459–1472, December 2004.

[10] W. M. Farrell, M. D. Desch, and P. Zarka. On the possibility of coherent cyclotron emission from extrasolar planets. *JGR*, 104:14025–14032, June 1999.

[11] P. Zarka, R. A. Treumann, B. P. Ryabov, and V. B. Ryabov. Magnetically-Driven Planetary Radio Emissions and Application to Extrasolar Planets. *APSS*, 277:293–300, June 2001.

[12] T. J. Lazio, W., W. M. Farrell, J. Dietrick, E. Greenlees, E. Hogan, C. Jones, and L. A. Hennig. The Radiometric Bode's Law and Extrasolar Planets. *ApJ*, 612:511–518, September 2004.

[13] J.-M. Grießmeier, P. Zarka, and H. Spreeuw. Predicting low-frequency radio fluxes of known extrasolar planets. *AAP*, 475:359–368, November 2007.

[14] S. L. G. Hess and P. Zarka. Modeling the radio signature of the orbital parameters, rotation, and magnetic field of exoplanets. *AAP*, 531:A29, July 2011.

[15] T. S. Bastian, G. A. Dulk, and Y. Leblanc. A Search for Radio Emission from Extrasolar Planets. *ApJ*, 545:1058–1063, December 2000.

[16] S. J. George and I. R. Stevens. Giant Metrewave Radio Telescope low-frequency observations of extrasolar planetary systems. *MNRAS*, 382:455–460, November 2007.

[17] T. Murphy, M. E. Bell, D. L. Kaplan, and et al. Limits on low-frequency radio emission from southern exoplanets with the Murchison Widefield Array. *MNRAS*, 446:2560–2565, January 2015.

[18] T. S. Bastian. Radio emission from flare stars. *SoPh*, 130:265–294, December 1990.

[19] E. Berger. Radio Observations of a Large Sample of Late M, L, and T Dwarfs: The Distribution of Magnetic Field Strengths. *ApJ*, 648:629–636, September 2006.

[20] G. Hallinan, A. Antonova, J. G. Doyle, S. Bourke, C. Lane, and A. Golden. Confirmation of the Electron Cyclotron Maser Instability as the Dominant Source of Radio Emission from Very Low Mass Stars and Brown Dwarfs. *ApJ*, 684:644–653, September 2008.

[21] M. Güdel. Stellar Radio Astronomy: Probing Stellar Atmospheres from Protostars to Giants. *ARA&A*, 40:217–261, 2002.

[22] E. B. Waltman, F. D. Ghigo, K. J. Johnston, R. S. Foster, R. L. Fiedler, and J. H. Spencer. The Evolution of Outbursts in Cygnus X-3 at 2.25 and 8.3 GHz. *AJ*, 110:290, July 1995.

[23] J. Wilms, K. Pottschmidt, G. G. Pooley, S. Markoff, M. A. Nowak, I. Kreykenbohm, and R. E. Rothschild. Correlated Radio-X-Ray Variability of Galactic Black Holes: A Radio-X-Ray Flare in Cygnus X-1. *ApJL*, 663:L97–L100, July 2007.

[24] D. Giannios and B. D. Metzger. Radio transients from stellar tidal disruption by massive black holes. *MNRAS*, 416:2102–2107, September 2011.

[25] B. D. Metzger, D. Giannios, and P. Mimica. Afterglow model for the radio emission from the jetted tidal disruption candidate Swift J1644+57. *MNRAS*, 420:3528–3537, March 2012.

[26] E. Berger, A. Zauderer, G. G. Pooley, A. M. Soderberg, R. Sari, A. Brunthaler, and M. F. Bietenholz. Radio Monitoring of the Tidal Disruption Event Swift J164449.3+573451. I. Jet Energetics and the Pristine Parsec-scale Environment of a Supermassive Black Hole. *ApJ*, 748:36, March 2012.

[27] K. W. Weiler, N. Panagia, M. J. Montes, and R. A. Sramek. Radio Emission from Supernovae and Gamma-Ray Bursters. *ARA&A*, 40:387–438, 2002.

[28] A. Gal-Yam, E. O. Ofek, D. Poznanski, A. Levinson, E. Waxman, D. A. Frail, A. M. Soderberg, E. Nakar, W. Li, and A. V. Filippenko. Radio and Optical Follow-up Observations of a Uniform Radio Transient Search: Implications for Gamma-Ray Bursts and Supernovae. *ApJ*, 639:331–339, March 2006.

[29] A. M. Soderberg, S. Chakraborti, G. Pignata, and et al. A relativistic type Ibc supernova without a detected γ-ray burst. *Nature*, 463:513–515, January 2010.

[30] A. Kamble, A. Soderberg, E. Berger, A. Zauderer, S. Chakraborti, and P. Williams. Radio Supernovae in the Local Universe. *ArXiv e-prints*, January 2014.

[31] T. Murphy, S. Chatterjee, D. L. Kaplan, J. Banyer, M. E. Bell, H. E. Bignall, G. C. Bower, R. A. Cameron, D. M. Coward, J. M. Cordes, and et al. VAST: An ASKAP Survey for Variables and Slow Transients. *PASA*, 30:6, February 2013.

[32] E. Nakar and T. Piran. Detectable radio flares following gravitational waves from mergers of binary neutron stars. *Nature*, 478:82–84, October 2011.

[33] B. D. Metzger and E. Berger. What is the Most Promising Electromagnetic Counterpart of a Neutron Star Binary Merger? *ApJ*, 746:48, February 2012.

[34] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, and et al. Observation of Gravitational Waves from a Binary Black Hole Merger. *Physical Review Letters*, 116(6):061102, February 2016.

[35] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, and et al. Localization and broadband follow-up of the gravitational-wave transient GW150914. *ArXiv e-prints*, February 2016.

[36] M. A. Garrett. Radio Astronomy Transformed: Aperture Arrays - Past, Present & Future. In *From Antikythera to the Square Kilometre Array: Lessons from the Ancients*, page 41, 2012.

[37] J. A. Högbom. Aperture Synthesis with a Non-Regular Distribution of Interferometer Baselines. *AAPS*, 15:417, June 1974.

[38] U. J. Schwarz. Mathematical-statistical Description of the Iterative Beam Removing Technique (Method CLEAN). *AAP*, 65:345, April 1978.

[39] B. G. Clark. An efficient implementation of the algorithm 'CLEAN'. *AAP*, 89:377, September 1980.

[40] F. R. Schwab. Relaxing the isoplanatism assumption in self-calibration; applications to low-frequency radio interferometry. *AJ*, 89:1076–1081, July 1984.

[41] T. J. Cornwell, K. Golap, and S. Bhatnagar. The Noncoplanar Baselines Effect in Radio Interferometry: The W-Projection Algorithm. *IEEE Journal of Selected Topics in Signal Processing*, 2:647–657, November 2008.

[42] A. R. Offringa, B. McKinley, N. Hurley-Walker, and et al. WSCLEAN: an implementation of a fast, generic wide-field imager for radio astronomy. *MNRAS*, 444:606–619, October 2014.

[43] G. B. Taylor, C. L. Carilli, and R. A. Perley, editors. *Synthesis Imaging in Radio Astronomy II*, volume 180 of *Astronomical Society of the Pacific Conference Series*, 1999.

[44] A. R. Thompson, J. M. Moran, and G. W. Swenson. *Interferometry and Synthesis in Radio Astronomy, John Wiley & Sons, 2007.* 2007.

[45] C. J. Lonsdale, R. J. Cappallo, M. F. Morales, and et al. The Murchison Widefield Array: Design Overview. *IEEE Proceedings*, 97:1497–1506, August 2009.

[46] S. J. Tingay, R. Goeke, J. D. Bowman, D. Emrich, S. M. Ord, D. A. Mitchell, M. F. Morales, T. Booler, B. Crosse, R. B. Wayth, and et al. The Murchison Widefield Array: The Square Kilometre Array Precursor at Low Radio Frequencies. *PASA*, 30:7, January 2013.

[47] P. E. Dewdney, P. J. Hall, R. T. Schilizzi, and T. J. L. W. Lazio. The Square Kilometre Array. *IEEE Proceedings*, 97:1482–1496, August 2009.

[48] D. R. DeBoer, R. G. Gough, J. D. Bunton, and et al. Australian SKA Pathfinder: A High-Dynamic Range Wide-Field of View Survey Telescope. *IEEE Proceedings*, 97:1507–1521, August 2009.

[49] T. B. Gibbon, E. K. Rotich Kipnoo, R. R. G. Gamatham, A. W. R. Leitch, R. Siebrits, R. Julie, S. Malan, W. Rust, F. Kapp, T. L. Venkatasubramani, B. Wallace, A. Peens-Hough, and P. Herselman. Fiber-to-the-telescope: MeerKAT, the South African precursor to Square Kilometre Telescope Array. *Journal of Astronomical Telescopes, Instruments, and Systems*, 1(2):028001, April 2015.

[50] M. P. van Haarlem, M. W. Wise, A. W. Gunst, G. Heald, J. P. McKean, J. W. T. Hessels, A. G. de Bruyn, R. Nijboer, J. Swinbank, R. Fallows, and et al. LOFAR: The LOw-Frequency ARray. *A&A*, 556:A2, August 2013.

[51] Kitty Ka Yi Lo. *Understanding and Classifying the Time-Variable Sky.* PhD thesis, University of Sydney, School of Physics, 3 2014.

[52] P. J. Hancock, T. Murphy, B. M. Gaensler, A. Hopkins, and J. R. Curran. Compact continuum source finding for next generation radio surveys. *MNRAS*, 422:1812–1824, May 2012.

[53] J. J. Condon. Confusion and Flux-Density Error Distributions. *ApJ*, 188:279–286, March 1974.

[54] N. Thyagarajan, N. Udaya Shankar, R. Subrahmanyan, and et al. A Study of Fundamental Limitations to Statistical Detection of Redshifted H I from the Epoch of Reionization. *ApJ*, 776:6, October 2013.

[55] M. H. Wieringa. PhD thesis, Rijksuniversiteit Leiden, (1991), 1991.

[56] D. A. Frail, S. R. Kulkarni, E. O. Ofek, G. C. Bower, and E. Nakar. A Revised View of the Transient Radio Sky. *ApJ*, 747:70, March 2012.

[57] G. C. Bower, D. Saul, J. S. Bloom, A. Bolatto, A. V. Filippenko, R. J. Foley, and D. Perley. Submillijansky Transients in Archival Radio Observations. *ApJ*, 666:346–360, September 2007.

[58] C. W. Helmstrom. *Statistical Theory of Signal Detection, 2nd edition.* Pergamon Press, London, 1968.

[59] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. E. Creighton. FINDCHIRP: An algorithm for detection of gravitational waves from inspiraling compact binaries. *Phys. Rev. D*, 85(12):122006, June 2012.

[60] J. S. Bloom, D. E. Holz, S. A. Hughes, K. Menou, A. Adams, S. F. Anderson, A. Becker, G. C. Bower, N. Brandt, B. Cobb, and et al. Astro2010 Decadal Survey Whitepaper: Coordinated Science in the Gravitational and Electromagnetic Skies. *arXiv: 0902.1527*, February 2009.

[61] E. S. Phinney. Finding and Using Electromagnetic Counterparts of Gravitational Wave Sources. In *astro2010: The Astronomy and Astrophysics Decadal Survey*, volume 2010 of *Astronomy*, page 235, 2009.

[62] LIGO Scientific Collaboration, Virgo Collaboration, J. Abadie, B. P. Abbott, R. Abbott, T. D. Abbott, M. Abernathy, T. Accadia, F. Acernese, C. Adams, and et al. Implementation and testing of the first prompt search for gravitational wave transients with electromagnetic counterparts. *A&A*, 539:A124, April 2012.

[63] B. F. Schutz. Determining the Hubble constant from gravitational wave observations. *Nature*, 323:310, September 1986.

[64] N. Dalal, D. E. Holz, S. A. Hughes, and B. Jain. Short GRB and binary black hole standard sirens as a probe of dark energy. *Phys. Rev. D*, 74(6):063006, September 2006.

[65] S. Nissanke, D. E. Holz, S. A. Hughes, N. Dalal, and J. L. Sievers. Exploring Short Gamma-ray Bursts as Gravitational-wave Standard Sirens. *ApJ*, 725:496–514, December 2010.

[66] K. G. Arun, H. Tagoshi, A. Pai, and C. K. Mishra. Synergy of short gamma ray burst and gravitational wave observations: Constraining the inclination angle of the binary and possible implications for off-axis gamma ray bursts. *Phys. Rev. D*, 90(2):024060, July 2014.

[67] G. M. Harry and LIGO Scientific Collaboration. Advanced LIGO: the next generation of gravitational wave detectors. *Classical and Quantum Gravity*, 27(8):084006, April 2010.

[68] The Virgo Collaboration. Advanced Virgo Baseline Design, 2009. https://pub3.ego-gw.it/itf/tds/file.php?callFile=VIR-0027A-09.pdf.

[69] D. Eichler, M. Livio, T. Piran, and D. N. Schramm. Nucleosynthesis, neutrino bursts and gamma-rays from coalescing neutron stars. *Nature*, 340:126–128, July 1989.

[70] R. Narayan, B. Paczynski, and T. Piran. Gamma-ray bursts as the death throes of massive binary stars. *ApJ*, 395:L83–L86, August 1992.

[71] L.-X. Li and B. Paczyński. Transient Events from Neutron Star Mergers. *ApJ*, 507:L59–L62, November 1998.

[72] B. D. Metzger, G. Martínez-Pinedo, S. Darbha, E. Quataert, A. Arcones, D. Kasen, R. Thomas, P. Nugent, I. V. Panov, and N. T. Zinner. Electromagnetic counterparts of compact object mergers powered by the radioactive decay of r-process nuclei. *MNRAS*, 406:2650–2662, August 2010.

[73] T. Piran, E. Nakar, and S. Rosswog. The electromagnetic signals of compact binary mergers. *MNRAS*, 430:2121–2136, April 2013.

[74] E. Berger. The Host Galaxies of Short-Duration Gamma-Ray Bursts: Luminosities, Metallicities, and Star-Formation Rates. *ApJ*, 690:231–237, January 2009.

[75] W. Fong, E. Berger, and D. B. Fox. Hubble Space Telescope Observations of Short Gamma-Ray Burst Host Galaxies: Morphologies, Offsets, and Local Environments. *ApJ*, 708:9–25, January 2010.

[76] N. R. Tanvir, A. J. Levan, A. S. Fruchter, J. Hjorth, R. A. Hounsell, K. Wiersema, and R. L. Tunnicliffe. A 'kilonova' associated with the short-duration γ-ray burst GRB130603B. *Nature*, 500:547–549, August 2013.

[77] E. Berger, W. Fong, and R. Chornock. An r-process Kilonova Associated with the Short-hard GRB 130603B. *ApJ*, 774:L23, September 2013.

[78] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, and et al. Prospects for Observing and Localizing Gravitational-Wave Transients with Advanced LIGO and Advanced Virgo. *Living Reviews in Relativity*, 19, February 2016.

[79] L. P. Singer, L. R. Price, B. Farr, A. L. Urban, C. Pankow, S. Vitale, J. Veitch, W. M. Farr, C. Hanna, K. Cannon, T. Downes, P. Graff, C.-J. Haster, I. Mandel, T. Sidery, and A. Vecchio. The First Two Years of Electromagnetic Follow-up with Advanced LIGO and Virgo. *ApJ*, 795:105, November 2014.

[80] E. Berger, P. A. Price, S. B. Cenko, A. Gal-Yam, A. M. Soderberg, M. Kasliwal, D. C. Leonard, P. B. Cameron, D. A. Frail, S. R. Kulkarni, and et al. The afterglow and elliptical host galaxy of the short γ-ray burst GRB 050724. *Nature*, 438:988–990, December 2005.

[81] A. M. Soderberg, E. Berger, M. Kasliwal, D. A. Frail, P. A. Price, B. P. Schmidt, S. R. Kulkarni, D. B. Fox, S. B. Cenko, A. Gal-Yam, E. Nakar, and K. C. Roth. The Afterglow, Energetics, and Host Galaxy of the Short-Hard Gamma-Ray Burst 051221a. *ApJ*, 650:261–271, October 2006.

[82] W. Fong, E. Berger, B. D. Metzger, R. Margutti, R. Chornock, G. Migliori, R. J. Foley, B. A. Zauderer, R. Lunnan, T. Laskar, S. J. Desch, K. J. Meech, S. Sonnett, C. Dickey, A. Hedlund, and P. Harding. Short GRB 130603B: Discovery of a Jet Break in the Optical and Radio Afterglows, and a Mysterious Late-time X-Ray Excess. *ApJ*, 780:118, January 2014.

[83] W. Fong, E. Berger, R. Margutti, and B. A. Zauderer. A Decade of Short-duration Gamma-Ray Burst Broadband Afterglows: Energetics, Circumburst Densities, and Jet Opening Angles. *ApJ*, 815:102, December 2015.

[84] P. Chandra and D. A. Frail. A Radio-selected Sample of Gamma-Ray Burst Afterglows. *ApJ*, 746:156, February 2012.

[85] L. Z. Kelley, I. Mandel, and E. Ramirez-Ruiz. Electromagnetic transients as triggers in searches for gravitational waves from compact binary mergers. *Phys. Rev. D*, 87(12):123004, June 2013.

[86] H. van Eerten, A. van der Horst, and A. MacFadyen. Gamma-Ray Burst Afterglow Broadband Fitting Based Directly on Hydrodynamics Simulations. *ApJ*, 749:44, April 2012.

[87] Rahul Biswas, Patrick R. Brady, Jordi Burguet-Castell, Kipp Cannon, Jessica Clayton, Alexander Dietz, Nickolas Fotopoulos, Lisa M. Goggin, Drew Keppel, Chris Pankow, Larry R. Price, and Ruslan Vaulin. Likelihood-ratio ranking of gravitational-wave candidates in a non-gaussian background. *Phys. Rev. D*, 85:122008, Jun 2012.

[88] C. M. Trott, R. B. Wayth, J.-P. R. Macquart, and S. J. Tingay. Source Detection in Interferometric Visibility Data. I. Fundamental Estimation Limits. *ApJ*, 731:81, April 2011.

[89] Robert E. Kass and Adrian E. Raftery. Bayes factors. *Journal of the American Statistical Association*, 90(430):773–795, 1995.

[90] J. Neyman and E. S. Pearson. On the problem of the most efficient tests of statistical hypotheses. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 231:289–337, 1933.

[91] D. S. Sivia and J. Skilling. *Data Analysis: A Bayesian Tutorial, Second Edition.* Oxford University Press, Oxford, 2006.

[92] L. Feng, R. Remillard, J. N. Hewitt, and R. Vaulin. An Implementation of Empirical Primary Beam Correction in MWA Imaging. *MWA Memo*, March 2015.

[93] S. T. Loi, T. Murphy, I. H. Cairns, and et al. Real-time imaging of density ducts between the plasmasphere and ionosphere. *GRL*, 42:3707–3714, May 2015.

[94] E. W. Greisen. AIPS FITS File Format. *AIPS Memo 117*, August 2012.

[95] A. R. Offringa, R. B. Wayth, N. Hurley-Walker, and et al. The Low-Frequency Environment of the Murchison Widefield Array: Radio-Frequency Interference Analysis and Mitigation. *PASA*, 32:8, March 2015.

[96] A. R. Offringa, A. G. de Bruyn, M. Biehl, S. Zaroubi, G. Bernardi, and V. N. Pandey. Post-correlation radio frequency interference classification methods. *MNRAS*, 405:155–167, June 2010.

[97] A. R. Offringa, J. J. van de Gronde, and J. B. T. M. Roerdink. A morphological algorithm for improving radio-frequency interference detection. *AAP*, 539:A95, March 2012.

[98] A. M. Levine. The Effects of the Two-Stage PFB Architecture on the Response of the MWA. *MWA Memo*, June 2012.

[99] J. P. McMullin, B. Waters, D. Schiebel, W. Young, and K. Golap. CASA Architecture and Applications. In R. A. Shaw, F. Hill, and D. J. Bell, editors, *Astronomical Data Analysis Software and Systems XVI*, volume 376 of *Astronomical Society of the Pacific Conference Series*, page 127, October 2007.

[100] N. Hurley-Walker, J. Morgan, R. B. Wayth, and et al. The Murchison Widefield Array Commissioning Survey: A Low-Frequency Catalogue of 14 110 Compact Radio Sources over 6 100 Square Degrees. *PASA*, 31:45, November 2014.

[101] A. R. Neben, R. F. Bradley, J. N. Hewitt, and et al. Measuring phased-array antenna beampatterns with high dynamic range for the Murchison Widefield Array using 137 MHz ORBCOMM satellites. *Radio Science*, 50:614–629, July 2015.

[102] A. Sutinjo, J. O'Sullivan, E. Lenc, R. B. Wayth, S. Padhi, P. Hall, and S. J. Tingay. Understanding instrumental Stokes leakage in Murchison Widefield Array polarimetry. *Radio Science*, 50:52–65, January 2015.

[103] A. R. Neben, J. N. Hewitt, R. F. Bradley, and et al. Beamforming Errors in Murchison Widefield Array Phased Array Antennas and their effects on Epoch of Reionization Science. *ArXiv e-prints*, February 2016.

[104] I. S. Sullivan, M. F. Morales, and B. J. et al. Hazelton. Fast Holographic Deconvolution: A New Technique for Precision Radio Interferometry. *ApJ*, 759:17, November 2012.

[105] N. Thyagarajan, D. J. Helfand, R. L. White, and R. H. Becker. Variable and Transient Radio Sources in the FIRST Survey. *ApJ*, 742:49, November 2011.

[106] Paul Dierckx. *Curve and Surface Fitting with Splines.* Oxford University Press, Inc., New York, NY, USA, 1993.

[107] P. Dierckx. An algorithm for surface-fitting with spline functions. *IMA Journal of Numerical Analysis*, 1(3):267–283, 1981.

[108] M. E. Bell, T. Murphy, D. L. Kaplan, P. Hancock, B. M. Gaensler, J. Banyer, K. Bannister, C. Trott, N. Hurley-Walker, R. B. Wayth, and et al. A survey for transients and variables with the Murchison Widefield Array 32-tile prototype at 154 MHz. *MNRAS*, 438:352–367, February 2014.

[109] T. R. Jaeger, S. D. Hyman, N. E. Kassim, and T. J. W. Lazio. Discovery of a Meter-wavelength Radio Transient in the SWIRE Deep Field: 1046+59. *AJ*, 143:96, April 2012.

[110] A. Rowlinson, M. E. Bell, T. Murphy, C. M. Trott, N. Hurley-Walker, S. Johnston, S. J. Tingay, D. L. Kaplan, D. Carbone, P. J. Hancock, L. Feng, A. R. Offringa, G. Bernardi, J. D. Bowman, F. Briggs, R. J. Cappallo, A. A. Deshpande, B. M. Gaensler, L. J. Greenhill, B. J. Hazelton, M. Johnston-Hollitt, C. J. Lonsdale, S. R. McWhirter, D. A. Mitchell, M. F. Morales, E. Morgan, D. Oberoi, S. M. Ord, T. Prabu, N. U. Shankar, K. S. Srivani, R. Subrahmanyan, R. B. Wayth, R. L. Webster, A. Williams, and C. L. Williams. Limits on Fast Radio Bursts and other transient sources at 182 MHz using the Murchison Widefield Array. *MNRAS*, March 2016.

[111] P. R. Brady, J. D. E. Creighton, and A. G. Wiseman. Upper limits on gravitational-wave signals based on loudest events. *Classical and Quantum Gravity*, 21:S1775–S1781, October 2004.

[112] R. Biswas, P. R. Brady, J. D. E. Creighton, and S. Fairhurst. The loudest event statistic: general formulation, properties and applications. *Classical and Quantum Gravity*, 26(17):175009, September 2009.

[113] S. D. Hyman, R. Wijnands, T. J. W. Lazio, S. Pal, R. Starling, N. E. Kassim, and P. S. Ray. GCRT J1742-3001: A New Radio Transient Toward the Galactic Center. *ApJ*, 696:280–286, May 2009.

[114] D. Carbone, A. J. van der Horst, R. A. M. J. Wijers, and et al. New methods to constrain the radio transient rate: results from a survey of four fields with LOFAR. *ArXiv e-prints*, November 2014.

[115] Y. Cendes, R. A. M. J. Wijers, J. D. Swinbank, and et al. LOFAR Observations of Swift J1644+57 and Implications for Short-Duration Transients. *ArXiv e-prints*, December 2014.

[116] L. Feng, R. Vaulin, and J. N. Hewitt. Detectability of Late-Time Radio Afterglows from Compact Binary Coalescence. *ArXiv e-prints*, May 2014.

[117] S. W. Ellingson, G. B. Taylor, J. Craig, and et al. The LWA1 Radio Telescope. *IEEE Transactions on Antennas and Propagation*, 61:2540–2549, May 2013.

[118] K. Bandura, G. E. Addison, M. Amiri, and et al. Canadian Hydrogen Intensity Mapping Experiment (CHIME) pathfinder. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 9145 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 22, July 2014.

[119] S. Johnston, M. Bailes, N. Bartel, and et al. Science with the Australian Square Kilometre Array Pathfinder. *PASA*, 24:174–188, December 2007.

[120] T. J. W. Lazio, T. E. Clarke, W. M. Lane, C. Gross, N. E. Kassim, P. S. Ray, D. Wood, J. A. York, A. Kerkhoff, B. Hicks, E. Polisensky, K. Stewart, N. Paravastu Dalal, A. S. Cohen, and W. C. Erickson. Surveying the Dynamic Radio Sky with the Long Wavelength Demonstrator Array. *AJ*, 140:1995–2006, December 2010.

[121] J. Abadie, B. P. Abbott, R. Abbott, M. Abernathy, T. Accadia, F. Acernese, C. Adams, R. Adhikari, P. Ajith, B. Allen, and et al. TOPICAL REVIEW: Predictions for the rates of compact binary coalescences observable by ground-based gravitational-wave detectors. *Classical and Quantum Gravity*, 27(17):173001, September 2010.

[122] A. Levinson, E. O. Ofek, E. Waxman, and A. Gal-Yam. Orphan Gamma-Ray Burst Radio Afterglows: Candidates and Constraints on Beaming. *ApJ*, 576:923–931, September 2002.

[123] G. Ghirlanda, R. Salvaterra, D. Burlon, S. Campana, A. Melandri, M. G. Bernardini, S. Covino, P. D'Avanzo, V. D'Elia, G. Ghisellini, L. Nava, I. Prandoni, L. Sironi, G. Tagliaferri, S. D. Vergani, and A. Wolter. Radio afterglows of a complete sample of bright Swift GRBs: predictions from present days to the SKA era. *MNRAS*, 435:2543–2551, November 2013.

[124] G. Ghirlanda, D. Burlon, G. Ghisellini, R. Salvaterra, M. G. Bernardini, S. Campana, S. Covino, P. D'Avanzo, V. D'Elia, A. Melandri, T. Murphy, L. Nava, S. D. Vergani, and G. Tagliaferri. GRB Orphan Afterglows in Present and Future Radio Transient Surveys. *PASA*, 31:e022, May 2014.

[125] Z.-B. Zhang, S.-W. Kong, Y.-F. Huang, D. Li, and L.-B. Li. Detecting radio afterglows of gamma-ray bursts with FAST. *Research in Astronomy and Astrophysics*, 15:237–251, February 2015.

[126] G. E. Anderson, A. J. van der Horst, T. D. Staley, R. P. Fender, R. A. M. J. Wijers, A. M. M. Scaife, C. Rumsey, D. J. Titterington, A. Rowlinson, and R. D. E. Saunders. Probing the bright radio flare and afterglow of GRB 130427A with the Arcminute Microkelvin Imager. *MNRAS*, 440:2059–2065, May 2014.

[127] E. Nakar. Short-hard gamma-ray bursts. *Physics Reports*, 442:166–236, April 2007.

[128] E. Berger. Short-Duration Gamma-Ray Bursts. *ARA&A*, 52:43–105, August 2014.

[129] R. Sari, T. Piran, and R. Narayan. Spectra and Light Curves of Gamma-Ray Burst Afterglows. *ApJ*, 497:L17, April 1998.

[130] J. Granot and R. Sari. The Shape of Spectral Breaks in Gamma-Ray Burst Afterglows. *ApJ*, 568:820–829, April 2002.

[131] D. N. Burrows, D. Grupe, M. Capalbi, A. Panaitescu, S. K. Patel, C. Kouveliotou, B. Zhang, P. Mészáros, G. Chincarini, N. Gehrels, and R. A. M. Wijers. Jet Breaks in Short Gamma-Ray Bursts. II. The Collimated Afterglow of GRB 051221A. *ApJ*, 653:468–473, December 2006.

[132] W. Fong, E. Berger, R. Margutti, B. A. Zauderer, E. Troja, I. Czekala, R. Chornock, N. Gehrels, T. Sakamoto, D. B. Fox, and P. Podsiadlowski. A Jet Break in the X-Ray Light Curve of Short GRB 111020A: Implications for Energetics and Rates. *ApJ*, 756:189, September 2012.

[133] A. Panaitescu. The energetics and environment of the short-GRB afterglows 050709 and 050724. *MNRAS*, 367:L42–L46, March 2006.

[134] W. Fong, E. Berger, R. Chornock, N. R. Tanvir, A. J. Levan, A. S. Fruchter, J. F. Graham, A. Cucchiara, and D. B. Fox. The Optical Afterglow and z = 0.92 Early-type Host Galaxy of the Short GRB 100117A. *ApJ*, 730:26, March 2011.

[135] R. Perna and K. Belczynski. Short Gamma-Ray Bursts and Mergers of Compact Objects: Observational Constraints. *ApJ*, 570:252–263, May 2002.

[136] R. D. Blandford and C. F. McKee. Fluid dynamics of relativistic blast waves. *Physics of Fluids*, 19:1130–1138, August 1976.

[137] H. J. van Eerten and A. I. MacFadyen. Synthetic Off-axis Light Curves for Low-energy Gamma-Ray Bursts. *ApJ*, 733:L37, June 2011.

[138] H. J. van Eerten and A. I. MacFadyen. Gamma-Ray Burst Afterglow Scaling Relations for the Full Blast Wave Evolution. *ApJ*, 747:L30, March 2012.

[139] D. W. Hogg. Distance measures in cosmology. *arXiv: astro-ph/9905116*, May 1999.

[140] L. Sironi and D. Giannios. A Late-time Flattening of Light Curves in Gamma-Ray Burst Afterglows. *ApJ*, 778:107, December 2013.

[141] Planck Collaboration, P. A. R. Ade, N. Aghanim, C. Armitage-Caplan, M. Arnaud, M. Ashdown, F. Atrio-Barandela, J. Aumont, C. Baccigalupi, A. J. Banday, and et al. Planck 2013 results. XVI. Cosmological parameters. *AAP*, 571:A16, November 2014.

[142] K. Siellez, M. Boër, and B. Gendre. Simultaneous event detection rates by electromagnetic and gravitational wave detectors in the advanced era of LIGO and Virgo. *MNRAS*, 437:649–655, January 2014.

[143] L. P. Singer, S. B. Cenko, M. M. Kasliwal, D. A. Perley, E. O. Ofek, D. A. Brown, P. E. Nugent, S. R. Kulkarni, A. Corsi, D. A. Frail, and et al. Discovery and Redshift of an Optical Afterglow in 71 deg^2: iPTF13bxl and GRB 130702A. *ApJ*, 776:L34, October 2013.

[144] C. S. Kochanek and T. Piran. Gravitational Waves and gamma -Ray Bursts. *ApJ*, 417:L17, November 1993.

[145] M. Moniez. Positioning and orienting a static cylindrical radio-reflector for wide field surveys. *arXiv: 1208.6427*, August 2012.