Radio Follow-up Of Gravitational Wave Events

Dougal Dobie

A thesis submitted in fulfilment of the requirements for the degree of

Doctor of Philosophy (PhD)

Supervisor:
Prof. Tara Murphy

Co-supervisors:
Prof. Richard Hunstead & Dr. Keith Bannister

The University of Sydney
Faculty of Science
School of Physics

2020

In memory of those I lost along the way.

Declaration

This is to certify that to the best of my knowledge, the content of this thesis is my own work. This thesis has not been submitted for any degree or other purposes.

I certify that the intellectual content of this thesis is the product of my own work and that all the assistance received in preparing this thesis and sources have been acknowledged.

Dougal Dobie

January 18, 2021

Abstract

I present comprehensive analysis of the role of radio observations in the field of multi-messenger astronomy. I demonstrate evidence for a turnover in the radio lightcurve of GW170817, the first detection of a neutron star merger. I propose an optimised observation strategy for follow-up of gravitational wave events with widefield radio telescopes, and apply this strategy to follow-up of a possible neutron star-black hole merger, GW190814. I discuss prospects of using Very Long Baseline Interferometry, and high-cadence observations of scintillation-induced variability, to constrain properties of the relativistic outflows produced by mergers. I quantify the capability of existing and planned radio facilities to perform gravitational wave follow-up and monitoring of detected afterglows. I conclude by summarising this work and place it in the broader context of the multi-messenger era.

This thesis was primarily written at my home and at the University of Sydney, on the land of the Darug and Gadigal people respectively.

The work presented in this thesis uses observations from the Australia Telescope Compact Array, which stands on the land of the Gomeroi people, and the Australian Square Kilometre Array Pathfinder, which stands on the land of the Wajarri Yamatji people.

Acknowledgements

None of this work would have been possible without the endless mentoring of my supervisor, Prof. Tara Murphy, who has been guiding my research for the past 5 years and opened my eyes to the politics and drama of science. I would also like to thank David Kaplan for his constant advice and feedback (particularly the conversations on Slack at 1am) and Keith Bannister for encouraging me to think outside the box and pursue what may seem like crazy ideas.

My other supervisor, Prof. Richard Hunstead, passed away 6 months before the submission of this thesis. He was a mentor to me from the very start of my University career, and was the most enthusiastic lab supervisor I've ever seen. He was always excited be shown new results, dig up archival data, and share his vast wealth of knowledge and experience. Dick, I reckon I've done a pretty good job of this thesis, but it would be even better with a bit of feedback from your green pen.

During the course of this degree I also lost a close friend of mine. James, thanks for always being interested in my work. One of the saddest parts of finishing this thesis is knowing that you're not here to demand I explain every detail of it to you.

I've been lucky to share an office with a number of excellent people, but Andrew and Mohammad and Harry were there from start to finish. Thanks for all the help, the food, the arguments and the company. I also profusely apologise to everyone in the neighbouring offices who had to listen to our nonsense.

Emil – you taught me everything I know about radio interferometry. Thanks for taking your time to answer all of my questions, troubleshoot data reduction and telling me all of the tips and trick that only a wizard like you could possibly know.

The follow-up of GW170817 was the most stressful experience of my life, so a big thankyou to Christene for riding that wave with me. The follow-up of GW190814 was less stressful, but I couldn't have done it without the help of Adam and the rest of the radio transients group.

I thank all of the ATNF staff, but particularly Robin for teaching me everything I know about the ATCA, Jamie for handling the bombardment of requests for observing time, and Phil for the many friendly chats we had in the hallways, and in the ATCA Portal.

I express the deepest gratitude to everyone outside of astronomy thatI've dragged along this long journey. To each and every one of my friends for listening to my rants, all the fun you've provided, and not getting too upset when I get flakey. To my Mum, Dad, sister, Grandma, Oma for their endless love and support, I couldn't have done it without you.

Contents

List of Figures

List of Tables

Summary of Works

This thesis includes four published works:

Chapter 2 has been published as Dobie et al., 2018, ApJ, 858, L15;

Chapter 3 has been published as Dobie et al., 2019a, PASA, 36, e019;

Chapter 4 has been published as Dobie et al., 2019b, ApJ, 887, L13;

Chapter 5 has been published as Dobie et al., 2020, MNRAS, 494, 2449.

Work completed as part of this thesis has also been included in twelve other publications. The following publications relate to observations of GW170817:

Abbott et al., 2017, ApJ, 848, L12;

Hallinan et al., 2017, Science, 358, 1579;

Kasliwal et al., 2017, Science, 358, 1559;

Andreoni et al., 2017, PASA, 34, e069;

Mooley et al., 2018a, Nature, 554, 207;

Mooley et al., 2018b, ApJ, 868, L11;

Makhathini et al., 2020, ApJ, submitted.

The following publications relate to the third LIGO/Virgo Observing Run:

Andreoni et al., 2020, ApJ, 890, 131;

Kasliwal et al., 2020 ApJ, 905, 145.

The following publications relate to radio follow-up of transients:

Ho et al., 2019, ApJ, 871, 73;

Kaplan et al., 2019, ApJ, 884, 96;

Horesh et al., 2020, ApJ, 903, 132.

An outline of the software I have developed as part of this work can be found in Appendix A, while a reproduction of my submissions to the Gamma-ray Coordinates Network and the Astronomer's Telegram can be found in Appendix B.

Chapter 1

Introduction

1.1 Preface

The dynamic nature of the sky has been known for millennia. While most early records of the variable sky relate to the motion of the Sun, the planets and the Moon, the first record of extra-solar variability can be dated to 185 AD with the discovery of a “guest star”, or supernova (Zhao et al., 2006). However, the art and oral traditions of many ancient civilisations across the world describe phenomena that are broadly consistent with stellar variability (Hamacher, 2018) and supernovae (Hamacher, 2014, and references therein). Our understanding of the Universe and our place in it has always been shaped by the variable sky, with early cosmological theories revolving around incremental advances in observations of the planets.

It was not until the discovery of SN 1572 that a true astronomical transient shaped our understanding of the Universe (Brahe, 1573). While cosmology had evolved since the time of Aristotle, his belief that the Universe beyond the planets was perfect and unchanging remained the prevailing view of the time. Naked-eye observations established that the supernova exhibited no detectable parallax or proper motion, confirming that it existed beyond the planetary sphere and hence that the stars, and the Universe, are not immutable.

Over the last few centuries, advances in telescope technology have allowed astronomers to probe the Universe. We can now perform precise observations of the most distant objects in the Universe across the entire electromagnetic spectrum, and searches in the time domain remain a vital part in answering the some of the most fundamental questions in modern astronomy. Observations of supernovae are used to measure the expansion of the Universe (Riess et al., 1998), searches for small brightness fluctuations in stars have found thousands of planets beyond our Solar System, and observations of Fast Radio Bursts have shed light upon the decades-old problem of missing baryonic matter (Macquart et al., 2020).

The 2015 detection of gravitational waves from the merger of two black holes (Abbott et al., 2016b) saw the dawn of a new era in time-domain astronomy, the multi-messenger era, where astronomers probe the Universe with electromagnetic radiation, gravitational waves, neutrinos and cosmic rays[1]. However, it was not until the first detection of a neutron star merger (Abbott et al., 2017d) that an electromagnetic counterpart to a gravitational wave event was first detected (Abbott et al., 2017f). In the coming decades, a number of revolutionary electromagnetic and gravitational wave facilities will come online, which will drastically expand our view of the Universe and help answer many of the outstanding questions in the field. For example, what is the underlying relation between neutron star mergers and gamma-ray bursts? Can we detect gravitational waves from sources other than compact object mergers? Can we use gravitational waves to measure the expansion of the Universe, and resolve the tension between current measurements? Answering these "known unknowns" will also undoubtedly lead to the discovery of "unknown unknowns" – unexpected discoveries that have not been predicted by theorists.

This thesis is focused on the contribution of radio observations to the multi-messenger parameter space. In this chapter I introduce the concept of gravitational waves and the process that led to their detection (§1.2), summarise the history of gamma-ray bursts and the science behind the afterglows they produce (§1.3), recap the discovery of GW170817 along with my contributions to the follow-up effort (§1.4) and conclude by enumerating the goals of this thesis (§1.5).

1.2 Gravitational Waves

1.2.1 General Relativity in the Dynamical Regime

The General Theory of Relativity proposes that gravity is the result of spacetime curvature caused by the presence of matter or energy. This reformulation of mechanics explained the long-known deviation of the orbit of Mercury from the laws of Newton and Kepler, as well as making a number of predictions for which observational evidence has since been found including the existence of black holes (Webster & Murdin, 1972; Event Horizon Telescope Collaboration, 2019), gravitational lensing (Walsh et al., 1979), the deflection of light by the Sun (Dyson et al., 1920) and gravitational redshift (Adams, 1925; Popper, 1954).

In the dynamical regime, General Relativity predicts that the curvature of space time can propagate as gravitational waves. Observationally, gravitational waves manifest as small distortions of the separation between two objects. This effect is incredibly small, with the change in separation of two objects on Earth caused by waves from an extragalactic merger typically a factor of $\sim 10^{-20}$. The most common source of gravitational waves in the Universe is likely the orbital decay of a binary system consisting of either neutron stars or black holes, although there may be other sources such as some supernovae (Ott, 2009; Radice et al., 2019) or single, asymmetric, neutron stars (Prix, 2009).

The discovery of PSR B1913+16 enabled the first observational study of compact binary systems (Hulse & Taylor, 1975). Continued observations spanning 14 years

---

[1] Although the field began with the detection of neutrinos from SN 1987A (Bionta et al., 1987; Hirata et al., 1987)

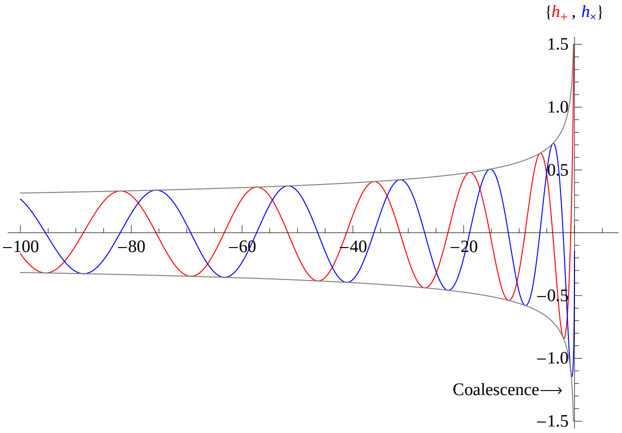

Figure 1.1: Waveform of a compact binary coalescence inspiral, with the two gravitational wave polarisations shown in red and blue. Figure reproduced from Singer (2015).

allowed precise determination of the orbital parameters of the binary, and showed that the orbit was decaying in a manner consistent with the emission of gravitational waves (Taylor & Weisberg, 1989). This discovery was the first indirect detection of gravitational waves and was later awarded the 1993 Nobel Prize in Physics.

Early attempts to detect gravitational waves involved the use of 'Weber bars' – large metal cylinders which should oscillate at their resonant frequency if a gravitational wave propagated through them (Weber, 1960). While there were claims of a detection using these methods (e.g. Weber, 1969), they were later shown to be spurious. Contemporaneously to Weber, various independent groups (Cervantes-Cota et al., 2016) investigated the possibility of detecting gravitational waves using Michelson interferometers (Michelson & Morley, 1887). Eventually this work culminated in the building of the Laser Interferometer Gravitational-Wave Observatory (LIGO; LIGO Scientific Collaboration, 2015) and Virgo (Acernese et al., 2015) gravitational wave detectors and ultimately led to the detection of gravitational waves from a binary black hole merger in 2015 (Abbott et al., 2016b). This was the first direct detection of gravitational waves and was awarded the 2017 Nobel Prize in Physics.

1.2.2 The Advanced LIGO and Virgo Detectors

The Advanced LIGO and Virgo detectors (LIGO Scientific Collaboration, 2015; Acernese et al., 2015) are Michelson interferometers (Michelson & Morley, 1887) with Fabry-Peŕot cavities along each arm (3 km for Virgo, 4 km for both LIGO detectors), that increased the effective length of the arms by a factor of 400. In the presence of gravitational waves produced by a compact binary coalescence, the distance along each arm becomes time-variable. The fractional change in distance takes the form oscillating signal that gradually increase in frequency and amplitude as seen in Figure 1.1. A detailed discussion of gravitational waveforms is beyond the scope of this thesis, but the signal depends not only on the intrinsic properties of the merger (the componenent spins and masses), but also extrinsic properties including the inclination angle, the distance and sky position (Finn & Chernoff, 1993).

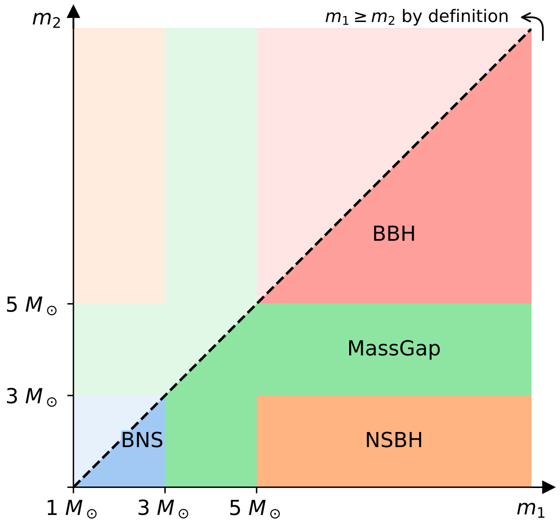

Figure 1.2: Categorisation of compact binary coalescences during the third LIGO/Virgo observing run in terms of component masses m_1 and m_2. Four mutually exclusive astrophysical categories exist – binary neutron stars (BNS), binary black holes (BBH), neutron star-black holes (NSBH) and “mass gap”. Figure reproduced from the LIGO/Virgo Public Alerts User Guide (`https://ligo.org`)

Signals are detected using four independent pipelines that carry out different forms of matched-filtering analysis to search for signals consistent with compact binary coalescences, while two further pipelines are used to search for unmodelled bursts like those predicted to be produced by supernovae. Once a candidate has been detected a rapid parameter estimate is carried out using BAYESTAR (Singer & Price, 2016), that produces a sky-map representing the posterior probability density function of the merger’s sky position. While BAYESTAR computes all merger parameters, during the third observing run (O3) only limited information was publicly released prior to formal publication[2]. Of interest are the probabilities that at least one merger component was a neutron star, the merger produced a remnant, and the probability classification. Events are classified as either binary black hole, neutron star-black hole or binary neutron star mergers based on component masses using cutoffs of $M < 3M_\odot$ for neutron stars and $M > 5M_\odot$ for black holes. Any merger with at least one component in the range $3M_\odot < M < 5M_\odot$ was classified as a “mass-gap” event (see Figure 1.2). A final category described the probability that the event was a false-positive caused by an unrelated terrestrial signal.

Alerts are distributed via the Gamma-ray Coordinates Network[3] (GCN) using the XML based VOEvent Transport protocol (Williams & Seaman, 2006). An automated preliminary notice is distributed within minutes of a merger being detected, and is followed by a human-vetted alert approximately half an hour later. Further alerts containing updated parameter estimates are distributed in the following days.

---

[2] `https://emfollow.docs.ligo.org/userguide/index.html`

[3] `https://gcn.gsfc.nasa.gov/`

All human vetted alerts are accompanied by a GCN circular that summarises the contents of the VOEvent in human-readable form, and sometimes contains further commentary, e.g. the reason for a retraction notice.

1.3 Gamma-ray Bursts, Neutron Star Mergers & Radio Afterglows

1.3.1 Detection, Classification and Origins

Recognising the detrimental effects of widespread nuclear testing, most countries signed and ratified the *Treaty Banning Nuclear Weapon Tests in the Atmosphere, in Outer Space and Under Water*, which took effect in 1963. One week after the treaty was ratified, the United States launched the Vela satellites, which were capable of detecting nuclear detonations anywhere on Earth via the associated gamma-ray emission. The satellites did not detect the signature of any nuclear weapon tests, but they did detect a number of bursts of gamma-ray emission that were not consistent with a terrestrial or solar origin (Klebesadel et al., 1973). Much like gravitational waves, the existence of gamma-ray bursts was predicted prior to their discovery – Colgate (1968) proposed that some supernovae should produce gamma-rays as relativistic shockwaves propagate outwards through the layers of ejected material.

Twenty years after the first detections, the *Burst and Transient Source Explorer* (BATSE) was launched with primary goal of detecting more bursts and discovering their origin. The large sample of bursts discovered by BATSE were distributed isotropically across the sky, suggesting an extragalactic origin (Meegan et al., 1992). It soon became clear that the distribution of burst durations is bimodal, with a divide at ~ 2 seconds (Kouveliotou et al., 1993). The short variety makes up roughly 30% of the population with a typical duration of 0.2 seconds.

The large amount of energy released by the bursts combined with their non-thermal spectra suggested that they are produced by a relativistic outflow of material (Ruderman, 1975; Paczynski, 1986) that must be collimated into a jet (Paczyński, 1993; Panaitescu & Mészáros, 1999; Rhoads, 1999; Sari et al., 1999). A multiwavelength afterglow was predicted (Paczynski & Rhoads, 1993), but follow-up observations were hindered by the localisation capability of BATSE, which detected bursts with positional uncertainties of $\sim 10\ \mathrm{deg}^2$. This was soon remedied by the launch of BeppoSAX (Boella et al., 1997), which contained a GRB monitor and multiple X-ray telescopes capable of localising X-ray sources with arcminute precision. Soon after launch, BeppoSAX detected a simultaneous burst of X-rays and gamma-rays (Costa et al., 1997). Targeted follow-up observations 8 hours later revealed a fading X-ray afterglow that declined over the following few days.

While various origins for GRBs had been proposed, it was not until the detection of GRB980425 and the subsequent detection of SN1998bw, that the origin of long GRBs as core collapse supernovae was first demonstrated observationally (Hjorth & Bloom, 2012, and references therein). The precise localisation of this burst allowed for detailed follow-up observations and confirmed the extragalactic origin of long

GRBs (Galama et al., 1998). Over the following years a larger sample of bursts and afterglows was acquired (e.g. Modjaz et al., 2016), thereby conclusively demontrating the connection between long GRBs and core collapse supernovae.

Short GRB afterglows remained elusive prior to the launch of the Swift and HETE-2 satellites, both designed for rapid localisation of GRBs via X-ray follow-up. These telescopes led a rapid improvement in our understanding of short GRBs, with the first localisations to host galaxies (Gehrels et al., 2005), the first detection of an optical afterglow (Hjorth et al., 2005; Fox et al., 2005) and the first detection of a radio afterglow (Berger et al., 2005) occurring within the span of 3 months. In the 15 years since the first detection, a large sample of short GRBs have been localised via their optical afterglows, although radio counterparts remain relatively rare (Fong et al., 2015; Lien et al., 2016).

1.3.2 Properties of Short GRBs and Their Environments

It makes little sense to discuss the population properties of neutron star mergers given the sample size of one. Recognising that at least some neutron star mergers produce short GRBs, based on the detection of a short GRB from GW170817, this section instead looks at the population of short GRBs and their environments.

One of the most important properties of short GRBs is the opening angle of the jet, θ_j, which determines properties of the afterglow emission (see Section 1.3.3). Measurement of θ_j for a large number of bursts allows calculation of the beaming fraction, f_b, given by

$$f_b = 1 - \cos\theta_j. \tag{1.1}$$

This quantity (valid for a top-hat jet structure) determines the total energy produced by the burst and also allows calculation of the true burst rate, and the expected rate of joint detections of short GRBs and gravitational waves (Howell et al., 2019).

Fong et al. (2015) use a sample of four bursts with measurable jet opening angles to estimate a median value of $\theta_j \sim 6 \pm 1$ deg. Including 7 further bursts with measurable lower-limits the typical value increases to $\theta_j \sim 15$ deg. Using a larger sample of 14 short GRBs, Wu & MacFadyen (2019) find $\theta_j = 6.9 \pm 2.3$ deg, although this estimate is model-dependent. Both estimates are in general agreement with measurements of the opening angle of GW170817 (see Section 1.4.4). Indeed, most properties of GW170817 inferred from the synchrotron afterglow (energetics, microphysics parameters, circum-burst density, electron distribution) are consistent with the short GRB population[4] (Wu & MacFadyen, 2019). The only parameter that differs significantly is the inclination angle – short GRBs typically have $\theta_{\rm obs} \sim 5$ deg while the inclination of GW170817 was $\theta_{\rm obs} \sim 30$ deg (see Figure 5.1). This discrepancy can be attributed to selection effects, since prompt gamma-ray emission from these events is generally only detectable if $\theta_{\rm obs} \leq \theta_j$.

---

[4] The luminosity of the GRB accompanying GW170817 was 3 orders of magnitude lower than any previously detected short GRB (Abbott et al., 2017g). However, von Kienlin et al. (2019) have found 12 other bursts with similar spectral properties. Therefore the outlying luminosity of this burst is likely due to inclination angle selection effects

Short GRBs are typically detected at cosmological distances, and generally span a redshift range of $z = 0.1 - 1.3$ (Fong & Berger, 2013), although this estimate is likely dominated by sensitivity selection effects – the redshifts of bursts are generally estimated from their host galaxies and therefore most bursts with measured redshifts are subject to the detection threshold of *Swift* (Berger, 2014). However, modelling of GW170817 and the associated short GRB suggests that there should be more bursts detected at lower redshifts (Howell et al., 2019). The offset of short GRBs from the nucleus of their host galaxy is typically $\sim 5\,\mathrm{kpc}$, although many are offset by $> 10\,\mathrm{kpc}$ (Fong et al., 2015). This is significantly larger than the typical offset of long GRBs, the positions of which generally trace stellar mass within galaxies (Bloom et al., 2002). However, the offset is in good agreement with the neutron star merger offset distribution predicted by a variety of population synthesis models (Fong & Berger, 2013). In comparison, the offset of GW170817 from NGC 4993 is 2 kpc, which is closer than 90% of short GRBs (Levan et al., 2017a).

To determine whether all short GRBs originate from neutron star mergers it is necessary to compile a larger sample of both types of event. Tight constraints on the short GRB jet opening angle distribution combined will improve estimates of the total volumetric burst rate and allow comparison to the rate of neutron star mergers. Localisation of these events to host galaxies will allow comparisons of the redshift distributions and host galaxy properties as well as inferences about the typical formation channels and environments based on host galaxy offsets. Finally, radio observations of synchrotron afterglows will allow comparison of event energetics, the density of the surrounding environments and typical microphysics parameters. The dependence of the afterglow on these properties is outlined below.

1.3.3 The Radio Afterglow

While there are a variety of models for the radio afterglows of neutron star mergers, most start from the same premise – the merger causes a relativistic jet to be launched into the surrounding material, and the accelerated electrons produce non-thermal synchrotron radiation. However, there is little agreement on the details of that process and the nature of the central engine driving the jet is entirely unclear. Some models predict that the merger immediately forms a black hole, which produces the jet via accretion (Lee & Ramirez-Ruiz, 2007; Metzger, 2017a, and references therein), while others predict that a hypermassive neutron star is formed instead (Hotokezaka et al., 2013a; Rowlinson et al., 2013). In this instance the emission may be driven by the collapse of the unstable neutron star into a black hole (Murguia-Berthier et al., 2017) or by magnetar spin down (Metzger et al., 2008; Bucciantini et al., 2012). The former case would potentially explaining the delay between the arrival of gravitational waves and gamma-rays in GW170817, see Section 1.4.1.

There are at least three sources of radio emission – the forward and reverse shocks, and the dynamical ejecta. The reverse shock component is dominant at early times, peaking $\lesssim 1$ day post-merger and fading on timescales of a few weeks when the forward shock component becomes detectable (e.g. Sari & Piran, 1995; Resmi & Zhang, 2016; Lamb & Kobayashi, 2019). The comparable timescales means that resolving

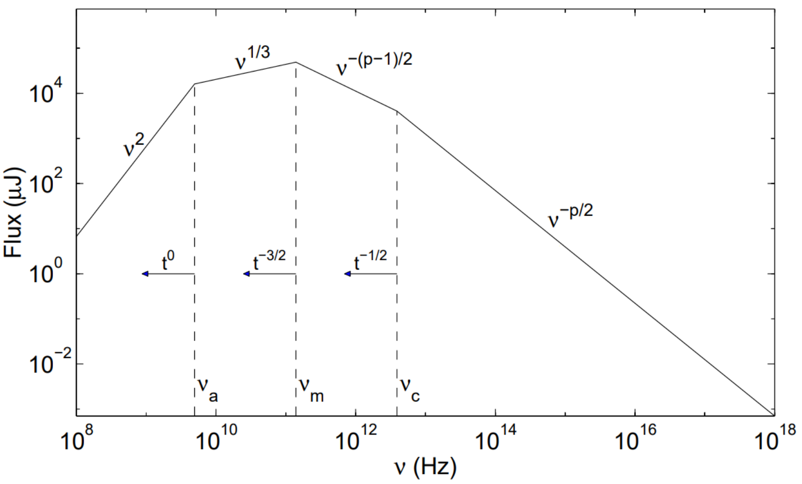

Figure 1.3: Spectrum of the relativistic shock in the canonical GRB afterglow model of Granot & Sari (2002). The spectrum consists of four segments with break frequencies ν_a, ν_m and ν_c that evolve adiabatically and scale with time as denoted by the arrows. Figure modified from (Sari et al., 1998).

one component from the other is difficult without high-cadence monitoring, and similarly, modelling of the early-time lightcurve is incomplete without considering both components. The precise nature of the forward shock remains unclear, and there exists numerous competing models that aim to describe the exact geometry of the outflow (e.g. Sari et al., 1998; Granot & Sari, 2002; Panaitescu & Kumar, 2004; Duffell & MacFadyen, 2013; Salafia et al., 2015; Kathirgamaraju et al., 2019a; Gottlieb et al., 2021). Simultaneously, each of these models is also highly dependent on the properties of the merger and the surrounding environment and therefore predictions for the lightcurve of this component vary greatly, with peak timescales generally ranging from days–years and peak flux density ranging from sub-μJy to tens of mJy. The dynamical ejecta is predicted to produce detectable radio emission on timescales of decades, with a similarly wide range of expected flux densities (Nakar & Piran, 2011; Piran et al., 2013; Hotokezaka & Piran, 2015; Margalit & Piran, 2015; Dietrich & Ujevic, 2017; Radice et al., 2018, e.g.). However, the behaviour of this component is generally independent of the forward and reverse shock components, and instead depends on the mass and velocity of the dynamical ejecta.

This thesis focuses on the forward shock component. Not only is it the only component detected in the afterglow of GW170817 to-date, but it is also the most viable component to detect in future follow-up. The detection of the reverse shock is limited by follow-up latency while the long delay in the emission from the dynamical ejecta becoming detectable means that conclusively associating it with the merger is difficult in the absence of any other electromagnetic counterpart.

The canonical GRB afterglow model of Granot & Sari (2002) is based on the self-similar solution for the expansion of a relativistic blast wave with Lorentz factor γ, described by Blandford & McKee (1976) with a variety of adjustments to account for radiative effects (Sari, 1997; Cohen et al., 1998). While this model explicitly describes a spherical blast wave, it is a valid approximation to the outflow of a jet as long as the lateral expansion of the jet remains negligible, i.e. $\gamma \gg \theta_j^{-1}$ (Piran, 1994). It is assumed that a fixed fraction, ϵ_B, of the blast waves energy is deposited into the magnetic field, and that the energy of the electrons behind the shock follows a power-law distribution of the form $N(\gamma) \propto \gamma^{-p}$ above some Lorentz factor $\gamma > \gamma_m$. The fraction of the total energy deposited into the electrons is given by ϵ_e.

The spectrum of the resulting emission is described by a series of 4 power-law segments as seen in Figure 1.3. The location of the 3 break frequencies are given in terms of the isotropic equivalent energy (in units of 10^{52} erg), E_{52}, the time post-merger (in days), t_d, the redshift of the merger, z, the energy density of the electrons, e, and the circum-merger density (in units of cm^{-3}), n_0,

$$\nu_a = 1.24 \times 10^9\,\mathrm{Hz}\left(\frac{p-1}{3p+2}\right)^{3/5}(1+z)^{-1}\bar{\epsilon}_e^{-1}\epsilon_B^{1/5}n_0^{3/5}E_{52}^{1/5} \tag{1.2}$$

$$\nu_m = 3.73 \times 10^{15}\,\mathrm{Hz}\left(\frac{p-0.67}{\sqrt{1+z}}\right)\bar{\epsilon}_e^{2}\epsilon_B^{1/2}E_{52}^{1/2}t_d^{-3/2} \tag{1.3}$$

$$\nu_c = 6.37 \times 10^{13}\,\mathrm{Hz}\left(\frac{p-0.46}{e^{1.16p}\sqrt{1+z}}\right)\epsilon_B^{-3/2}E_{52}^{-1/2}t_d^{-1/2} \tag{1.4}$$

where ν_a is the frequency below which synchrotron self-absorption is dominant, ν_m is the frequency corresponding to the synchrotron emission produced by electrons with minimum Lorentz factor γ_m, and ν_c is the synchrotron cooling frequency.

Berger (2014) estimates typical values are given by $\nu_a = 1\,\mathrm{GHz}$, $\nu_m = 200t_d^{-1.5}\,\mathrm{GHz}$ and $\nu_c = 200t_d^{-0.5}\,\mathrm{PHz}$ based on typical short GRB parameters of $E_{\mathrm{iso}} = 10^{51}\,\mathrm{erg}$, $n = 0.1\,\mathrm{cm}^{-3}$ and $z = 0.5$ and microphysics parameters of $\epsilon_e = 0.1$, $\epsilon_B = 0.01$, $p = 2.4$. To compare expectations for afterglows associated with mergers detected by LIGO we also calculate typical break frequencies for $z = 0.05$ (corresponding to the LIGO horizon; Abbott et al., 2018) with microphysics parameters $p = 2.2$, $\epsilon_e = 0.3$ $\epsilon_e = 0.02$ based on typical values associated with known short GRBs (Wu & MacFadyen, 2019). This results in $\nu_a \sim 700\,\mathrm{MHz}$, $\nu_m \sim 160t_d^{-1.5}\,\mathrm{GHz}$ and $\nu_c = 240t_d^{-0.5}\,\mathrm{PHz}$, suggesting that by ~ 1 month post-merger radio and X-ray observations will probe the same segment of the spectrum ($\nu_m < \nu < \nu_c$). This is in agreement with observations of GW170817 (see Section 1.4.3).

If energy is continuously injected into the ejecta the lightcurve will rise, otherwise any observed increase is due to the evolution of the break frequencies. The decay of the lightcurve can then be used to characterise the geometry of the shock. In the spherical outflow regime the lightcurve will decay as $t^{-3(p-1)/4}$, while if the emission is jet-dominated it will decay as t^{-p}, assuming $\nu_m < \nu < \nu_c$ (Sari et al., 1999).

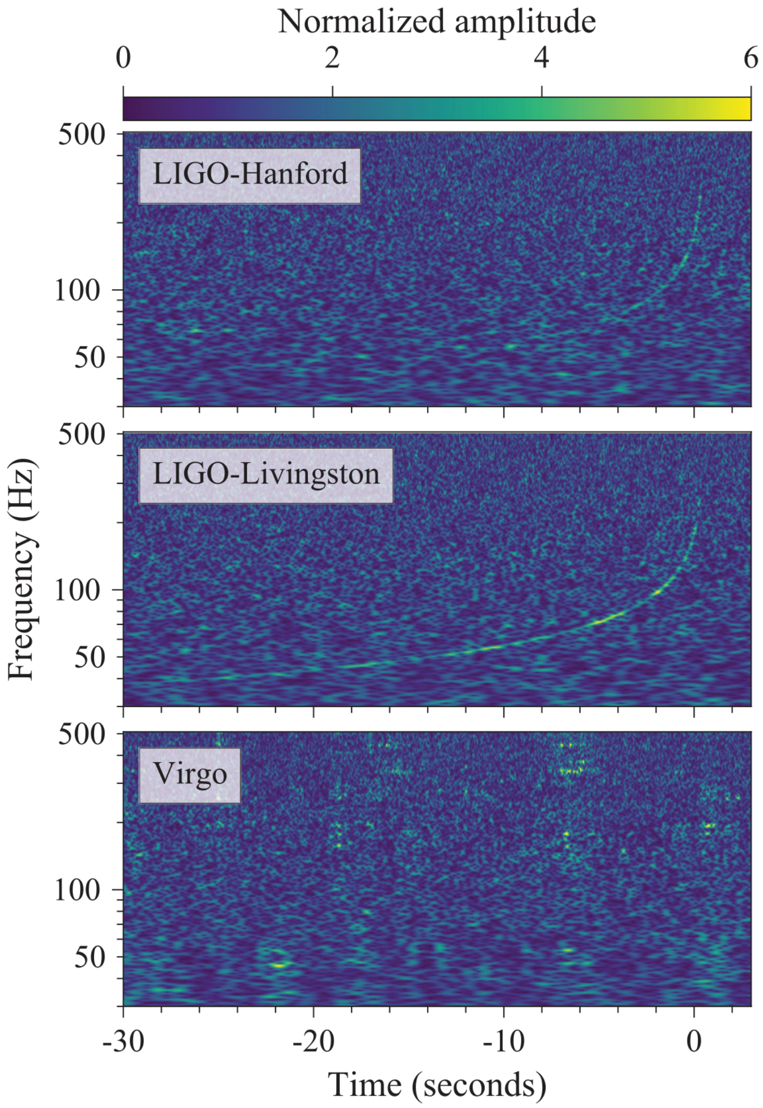

Figure 1.4: Time-frequency representations of the gravitational wave signal produced by GW170817 as observed by the LIGO-Hanford (top), LIGO-Livingston (middle), and Virgo (bottom) detectors. Times are shown relative to August 17, 2017 12:41:04 UTC. Figure reproduced from Abbott et al. (2017d)

1.4 GW170817: The First Detection of a Neutron Star Merger

1.4.1 Gravitational Waves and Gamma-rays

Abbott et al. (2017d) report the detection of a binary neutron star merger candidate on August 17, 2017 at 12:41:04 UTC in data from the LIGO Hanford detector with a signal-to-noise (SNR) ratio of 18.8. An offline analysis of data from the LIGO Livingston detector revealed the presence of an unrelated noise artefact 1.1 s pre-merger, and after modelling and subtracting the artefact, the gravitational wave was clearly visible with an SNR of 26.4. While the Virgo detector was also operating at the time, it did not detect the signal due to its lower overall sensitivity and the antenna pattern relative to the location of the merger.

The merger was initially localised to 31 deg^2, before being refined to 28 deg^2 within days (Abbott et al., 2017d). A search of galaxy catalogues found 49 galaxies within the localisation (Kasliwal et al., 2017), as seen in Figure 1.5. Further parameter estimates at later times improved the localisation to 16 deg^2 (Abbott et al., 2019a).

The Fermi Gamma-Ray Burst Monitor (GBM) detected a short GRB spatially coincident with the binary neutron star candidate 1.74 s after to the merger (Goldstein et al., 2017; Abbott et al., 2017g). This burst was also detected at low significance (Savchenko et al., 2017) by the International Gamma-Ray Astrophysics Laboratory (INTEGRAL). The probability that the gravitational wave signal and GRB are unrelated was quantified by considering both events as independent Poisson processes and comparing the rate of both types of detection. The probability that the two events were temporally coincident but unrelated is 5×10^{-6}. A comparison of the sky localisation of both detections showed that the probability of a chance spatial coincidence was 0.01, resulting in an overall probability that the events are unrelated of 5×10^{-8} (Abbott et al., 2017g).

This joint detection of light and gravitational waves from the same source has a variety of implications for astrophysics and fundamental physics as a whole. The small difference in arrival time suggests that gravitational waves propagate at the speed of light to within a factor of $\sim 10^{15}$, although this precision may be underestimated by up to two orders of magnitude if the gravitational wave signal and GRB emission were asynchronous as many models predict (e.g. Granot et al., 2017; Nakar et al., 2018; Lazzati et al., 2020). In turn, this rules out a multitude of theories regarding Dark Energy (Ezquiaga & Zumalacárregui, 2017) and gravity (Sanders, 2018; Abbott et al., 2019c). Most importantly this joint detection confirms the long-predicted connection between neutron star mergers and short GRBs (Paczynski, 1986; Eichler et al., 1989; Narayan et al., 1992), although questions remain as to whether all neutron star mergers produce short GRBs and similarly, whether all short GRBs are produced by neutron star mergers.

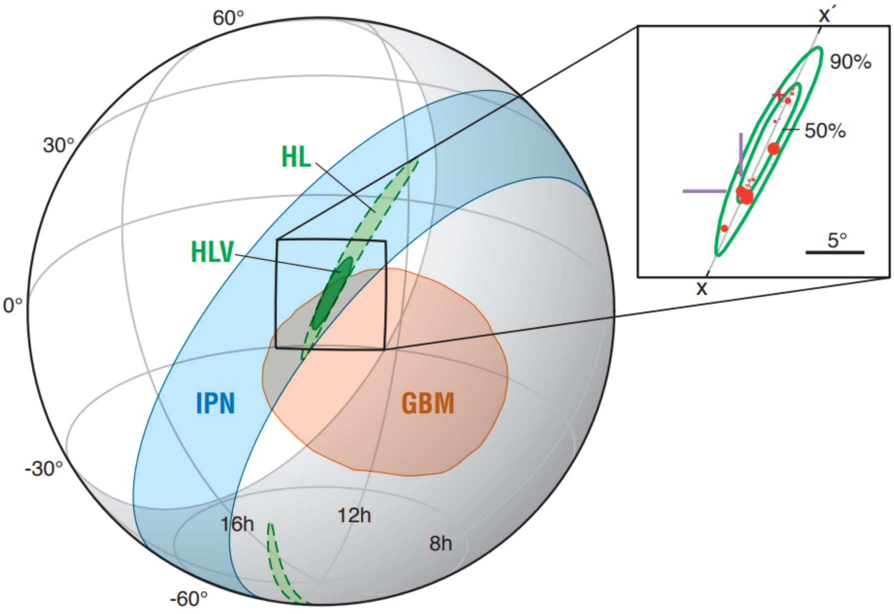

Figure 1.5: Localisation of GW170817. The LIGO and full LIGO/Virgo localisations are shown in light and dark green respectively. The localisation from Fermi is shown in orange, and triangulation of the GRB using the Interplanetary Network (Fermi and INTEGRAL) is shown in blue. The inset shows the 50% and 90% contours with crosshairs marking the location of GW170817. Red markers show the 49 candidate host galaxies within the localisation volume, with marker size proportional to stellar mass. Figure adapted from Kasliwal et al. (2017).

1.4.2 Discovery of an Electromagnetic Counterpart

Approximately 11 hours post-merger the detection of a candidate optical counterpart was announced by the One-Meter, Two-Hemisphere (1M2H) team (Coulter et al., 2017), associated with NGC 4993, one of the top-ranked host galaxy candidates. This counterpart was also independently detected by five other teams in observations carried out prior to the initial announcement (Arcavi et al., 2017a; Lipunov et al., 2017; Soares-Santos et al., 2017; Tanvir et al., 2017; Valenti et al., 2017). Over the next few days comprehensive follow-up with ultraviolet, optical and infra-red (UVOIR) facilities was carried out, revealing a rapid dimming of the initially blue emission, and a brightening of redder emission. Spectroscopy of the source showed no significant emission or absorption lines, something rarely observed in the thousands of optical transients that have been detected. The rapidly peaking blue emission followed by a slower rising red component, along with the featureless spectra, was consistent with models for a kilonova produced by a neutron star merger (Hotokezaka et al., 2015; Kasen et al., 2015; Metzger et al., 2015a; Barnes et al., 2016). Late-time high-resolution spectroscopy found the signatures of elements produced by r-process nucleosynthesis, thereby confirming the nature of the optical emission (Chornock et al., 2017; Kasliwal et al., 2017; Pian et al., 2017). This discovery simultaneously confirmed previous predictions that neutron star mergers are the origin of many of the heavy elements in Universe (Lattimer & Schramm, 1974; Symbalisty & Schramm, 1982; Freiburghaus et al., 1999).

While deep observations with a variety of X-ray telescopes began almost immediately after the discovery of the optical counterpart, it was not until 9 days post-merger that X-ray emission was detected (Troja et al., 2017). This detection was confirmed in a second epoch of observations 15 days post-merger (Haggard et al., 2017; Troja et al., 2017), just prior to the location of GW170817 being unobservable with leading X-ray facilities due to solar proximity.

Two independent observations with the Karl G. Jansky Very Large Array (VLA) 16 days post-merger claimed a low-significance detection of radio emission consistent with the coordinates of the merger, which was confirmed by an observation with the Australia Telescope Compact Array (ATCA) two days later (Hallinan et al., 2017). Further observations in the following weeks showed that the radio emission had continued to rise (Alexander et al., 2017a; Hallinan et al., 2017). We discuss the evolution of the radio lightcurve and it's implications in Section 1.4.3.

A full description of the discovery of the electromagnetic counterpart and the initial follow-up can be found in Abbott et al. (2017f) and references therein.

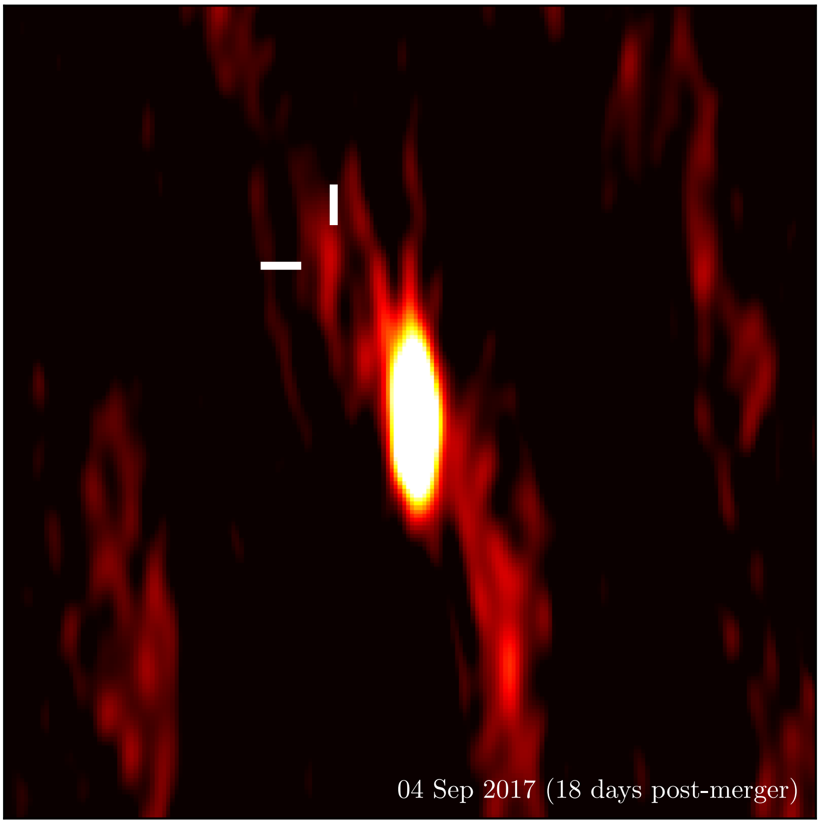

Figure 1.6: ATCA imaging of GW170817 up to one year post-merger. The position of the merger is denoted by white crosshairs, and the central radio source is host galaxy emission. Please use Adobe Acrobat to view the animated graphic.

1.4.3 Observations of the Radio Afterglow

Continued radio observations with the ATCA, VLA and Giant Metrewave Radio Telescope (GMRT) showed that the radio emission rose according to a spectral and temporal power law, $S_\nu(t) \propto \nu^{-0.6} t^{0.8}$ (Mooley et al., 2018a). These observations suggested the emission was not produced by a standard relativistic jet viewed off-axis, but instead originated from a mildly-relativistic, quasi-spherical outflow.

In December 2017 the Chandra X-ray telescope emerged from sun-block and was able to continue observations. Contemporaneous X-ray, optical and radio observations showed that the spectrum was well-fit by a single power law, suggesting that the synchrotron cooling break, ν_c, had not yet transitioned into the X-ray band (Troja et al., 2018; Margutti et al., 2018b). The long spectral baseline of these observations also refined estimates of the spectral index to $\alpha = 0.58$ and constrained the index of the electron energy distribution to $p = 2.17 \pm 0.01$ (Margutti et al., 2018b).

The next breakthrough came from radio observations carried out as part of this thesis, which provided strong evidence that the non-thermal lightcurve had peaked at approximately 150 days post-merger and had begun to decline (Dobie et al., 2018a). Chapter 2 details this discovery and it's implications for afterglow modelling.

Multiple works detail the gradual refinement of merger parameters based on continued monitoring (Alexander et al., 2018; Mooley et al., 2018d; Piro et al., 2019). However, the similarities in the lightcurves predicted by competing models for the outflow geometry necessitate some other form of measurement to distinguish between them (Nakar & Piran, 2018). Sari (1999) showed that the relativistic jet associated with GRBs should produce linearly polarised radio emission, and that the emission centroid should also exhibit proper motion. Both the polarisation fraction and the amplitude of proper motion should be maximum at the time of the lightcurve peak.

Observations between 75 and 230 days post-merger with the Very Long Baseline Array (VLBA) suggested that the centroid of the radio emission had shifted significantly (Mooley et al., 2018b). The apparent velocity of this shift was calculated to be $\beta_{\rm app} = 4.1$, measured in terms of the speed of light, c. These findings were later confirmed with an independent observation using global VLBI from 207 days post-merger (Ghirlanda et al., 2019). The observation of this motion suggests that the observed late-time emission originated from the jet, and enables a direct measurement of the associated Lorentz factor ($\Gamma = \beta_{\rm app}$).

Corsi et al. (2018b) carried out a deep search for linear polarisation and constrained the polarisation fraction to $< 12\%$ at 244 days post-merger. This result suggests that a significant component of the magnetic field of the jet is oriented perpendicular to the shock front. Specifically, the models of Gill & Granot (2018) suggest $b = \langle B_{\rm sn}^2 \rangle / \langle B_{\rm sp}^2 \rangle \gtrsim 0.75$, where $B_{\rm sn}$ and $B_{\rm sp}$ are the components of the magnetic field in the direction of the shock normal and in the plane of the shock respectively. More recent modelling suggests $0.66 < b < 1.49$ (Gill & Granot, 2020).

By ~ 750 days post-merger the radio afterglow had faded beyond detection, with an observation with the VLA at $\Delta T = 767\,$d finding only marginal evidence for con-

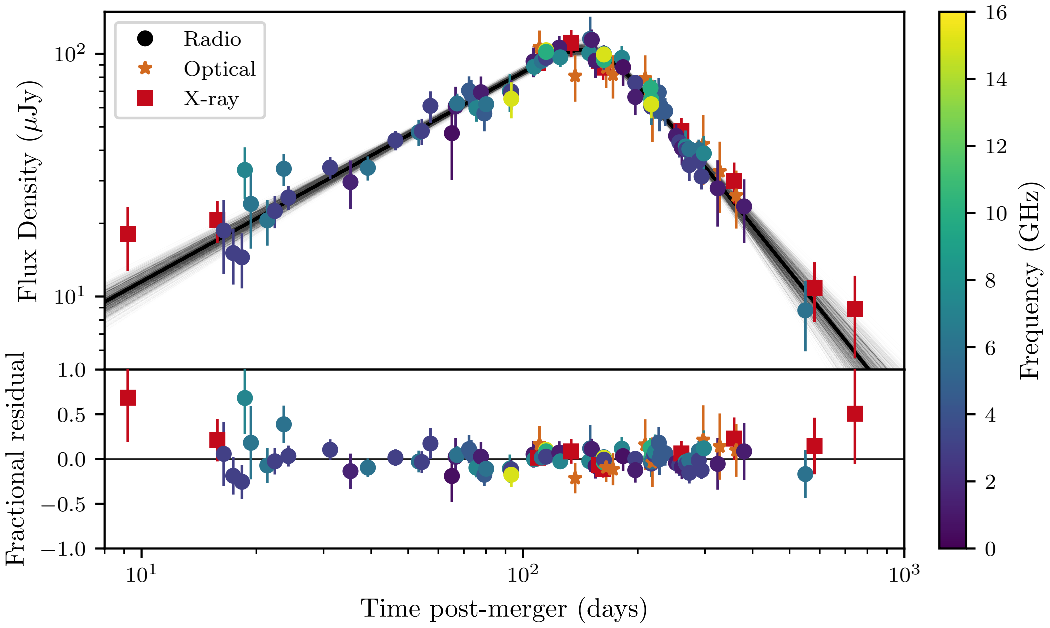

Figure 1.7: Detections of non-thermal emission from GW170817 with radio (circles), optical (stars) and X-ray (squares) facilities compiled by (Makhathini et al., 2020). The lightcurve has been scaled to 7.25 GHz, corresponding to the central frequency of our ATCA observations, based on a spectral index of $\alpha = -0.585$, while the original observing frequency is denoted by the colourbar. A smoothed broken power law has been fit to the data, with fractional residuals shown in the bottom panel.

tinued emission (Makhathini et al., 2020) and an ATCA observation at $\Delta T = 990\,\mathrm{d}$ finding no emission at 2.1 GHz (Troja et al., 2020). However, the X-ray emission remains detectable at $\Delta T = 940\,\mathrm{d}$ (the latest observation at the time of writing) although it continues to rapidly fade consistent with the afterglow of a jet (Troja et al., 2020; Hajela et al., 2020). These late-time X-ray observations appear to marginally deviate from the simple power-law decline, perhaps suggesting more complex temporal evolution or the emergence of emission from the dynamical ejecta. However, this deviation is within observational uncertainties and continued monitoring is necessary to determine whether it is statistically significant.

1.4.4 Modelling the Non-thermal Afterglow

The exact geometry of the merger, and the origins of the observed emission, were a source of strong contention for some time. While it was generally agreed that the ‘top-hat’ jet structure (Gill, 1965) commonly used to describe previous short GRB afterglows was inconsistent with the slow rise of the radio afterglow (Kasliwal et al., 2017; Margutti et al., 2018b), there was little agreement as to the outflows true nature. A wide range of models have been proposed, but they generally fall into one of two categories, with the emission being produced by either a quasi-spherical outflow or an off-axis jet.

A quasi-spherical outflow could be produced in a variety of ways. Salafia et al. (2018) suggest that the merger may not have launched a jet at all. Instead, the final stages of the merger may have resulted in the formation of a turbulent magnetic dynamo, which dissipates via magnetic reconnection, thereby accelerating electrons to relativistic speeds and launching a 'fireball' (Zrake & MacFadyen, 2013; Giacomazzo et al., 2015). While the early rise of the lightcurve is inconsistent with emission from the dynamical ejecta scenario proposed by Nakar & Piran (2011), a small fraction of the ejecta would be accelerated to mildly relativistic speeds. This 'fast tail' is sufficient to enhance the the radiation produced by the interaction between the outflow and the circum-merger medium and produce the observed emission (Hotokezaka et al., 2018). However, this scenario is inconsistent with the rapid peak and decline of the observed lightcurve and can now be ruled out.

Even if a jet is initially launched by the merger, it may not break out of the surrounding ejecta (comparable to a standard short GRB) and instead be choked, producing an energetic cocoon (Nakar & Piran, 2017; Gottlieb et al., 2018b,a). While there is a variety of jet types, all will produce a cocoon, regardless of whether or not they successfully break out (Nakar & Piran, 2018). The vast majority of jet models that aim to explain the emission from GW170817 are 'structured jets' - i.e. an axisymmetric relativistic outflow with some angular or radial dependence. Some assume that the jet has a particular angular structure, such as power-law or Gaussian profiles (e.g. Hotokezaka et al., 2019; Troja et al., 2020; Nathanail et al., 2020), while others are more physically motivated (e.g. Kathirgamaraju et al., 2018; Xie et al., 2018) or motivated by the observed lightcurve alone (Takahashi & Ioka, 2020). However, Nakar & Piran (2020) note that the majority of these fits ignore the observed centroid motion and consider only the lightcurve, which cannot uniquely determine the inclination angle, $\theta_{\rm obs}$. The result is large uncertainties in the value of $\theta_{\rm obs}$ inferred by each model, and inconsistencies between models.

The observed centroid motion indicates that the radio emission originates from some form of (at least) mildly relativistic ejecta, but at the time of writing there is no clear evidence to preference one particular model or geometry over another. While the leading models discussed above are all objectively different, they predict qualitatively similar lightcurves that are indistinguishable with the observational data at hand. With the benefit of hindsight, higher cadence monitoring between 10 and 100 days may have been able to determine the presence more complex temporal evolution predicted by some models (e.g. Lazzati et al., 2018; Makhathini et al., 2020). However, this measurement would still likely have been limited by the sensitivity of the observations. Similarly, observations thousands of days post-merger could in principle reveal a flattening decline predicted by other models (Troja et al., 2020), but are impossible due to the limited sensitivity of X-ray and radio facilities. While population studies of future events may reveal the standard mechanisms behind non-thermal emission from neutron star mergers, it is likely that the precise structure of the jet launched by GW170817 in particular will forever elude us.

1.5 Thesis Goals

The 2015 discovery of gravitational waves from a binary black hole merger provided a new avenue for astronomers to probe the Universe. The fields of gravitational waves and multi-messenger astronomy have evolved rapidly since then, with the detection of gravitational waves and light from a neutron star merger 2 years later, and a veritable menagerie of compact object coalescences detected in 2019-20. Planned upgrades to existing gravitational wave detectors, and proposed future detectors, will enable the detection of more mergers at even larger distances, and with improved localisations that will be conducive to follow-up with electromagnetic facilities. Simultaneously, radio astronomy is undergoing a revolution as we move towards the era of the Square Kilometre Array (SKA). SKA precursors like the Australian Square Kilometre Array Pathfinder (ASKAP) and MeerKAT are finally achieving design sensitivity and are beginning to carry out unprecedented surveys of the sky. These next-generation telescopes will enable astronomers to undertake most sensitive surveys for radio transients to-date, including searches for radio counterparts to gravitational wave events.

The work in this thesis spans a critical period in the development of the fields of multi-messenger and radio-transient astronomy. It begins with the detection of the first neutron star merger, GW170817, where I played a significant role in the detection of radio emission and the long-term monitoring of its evolution. I developed an algorithm to optimise follow-up observations with ASKAP, and then carried out ASKAP follow-up of the first neutron star-black hole merger, GW190814. This thesis also comprehensively outlines future prospects for the field and quantifies expectations for future detections of radio counterparts, the use of VLBI observations to directly measure outflow structure and how high-cadence radio monitoring can be used to constrain the size of outflows.

In Chapter 2 I demonstrate evidence for a turnover in the radio lightcurve of GW170817 based on observations with the ATCA and VLA

In Chapter 3 I outline the capabilities of ASKAP in performing follow-up of gravitational wave events, and describe an optimised observing strategy to do so.

Chapter 4 applies the results of the previous chapter to follow-up of GW190814, where I undertook the most sensitive widefield search for radio transients to-date, and placed constraints the merger properties.

In Chapter 5 I discuss prospects for constraining the geometry of merger outflows using Very Long Baseline Interferometry to detect relativistic expansion and motion of the outflow and high-cadence monitoring of scintillation to infer source size.

Chapter 6 outlines the various contributions radio observations can make to the gravitational wave follow-up effort and quantifies prospects for detecting afterglows with existing and planned facilities.

To conclude, Chapter 7 summarises the work of this thesis and places it in the broader context of multi-messenger astronomy.

Chapter 2

A Turnover in the Radio Light Curve of GW170817

Dougal Dobie, David L. Kaplan, Tara Murphy, Emil Lenc,
Kunal P. Mooley, Christene Lynch, Alessandra Corsi,
Dale Frail, Mansi Kasliwal, Gregg Hallinan

This chapter is reproduced from Dobie et al. (2018a), which was published in The Astrophysical Journal as "A Turnover in the Radio Light Curve of GW170817". My contributions to the project include carrying out all observations with the ATCA under project code CX391, production all figures and tables, all lightcurve analysis and writing the majority of the text. Emil Lenc contributed to the description of the ATCA data reduction. Kunal Mooley wrote the description of the VLA data reduction. Kunal Mooley, David Kaplan and Gregg Hallinan contributed to Section 2.3.3. Figures have been reformatted in a style consistent with the rest of this thesis.

We present 2–9 GHz radio observations of GW170817 covering the period 125–200 days post-merger, taken with the Australia Telescope Compact Array and the Karl G. Jansky Very Large Array. Our observations demonstrate that the radio afterglow peaked at 149 ± 2 days post-merger and is now declining in flux density. We see no evidence for evolution in the radio-only spectral index, which remains consistent with optically-thin synchrotron emission connecting the radio, optical, and X-ray regimes. The peak implies a total energy in the synchrotron-emitting component of a few $\times 10^{50}$ erg. The temporal decay rate is most consistent with mildly- or non-relativistic material and we do not see evidence for a very energetic off-axis jet, but we cannot distinguish between a lower-energy jet and more isotropic emission.

2.1 Introduction

The neutron star merger GW170817 was detected via the concurrent observation of gravitational waves (Abbott et al., 2017d) and a short gamma-ray burst (GRB; Abbott et al., 2017g; Goldstein et al., 2017). The merger was localized to its host galaxy, NGC 4993, by the detection of an optical transient (Arcavi et al., 2017a; Abbott et al., 2017f; Coulter et al., 2017; Lipunov et al., 2017; Soares-Santos et al., 2017; Tanvir et al., 2017; Valenti et al., 2017) and subsequent ultraviolet, optical and infrared observations found evidence of kilonova emission from the source (Arcavi et al., 2017a; Cowperthwaite et al., 2017; Drout et al., 2017; Evans et al., 2017; Kasliwal et al., 2017). X-ray observations found no evidence of emission until nine days post-merger (Haggard et al., 2017; Evans et al., 2017; Margutti et al., 2017; Troja et al., 2017), suggesting that this event differs significantly from previously-observed GRBs.

Radio emission from GW170817 was first detected 16 days post-merger (Hallinan et al., 2017). Follow-up observations over the next 100 days (Alexander et al., 2017a; Margutti et al., 2018b; Mooley et al., 2018a; Troja et al., 2018) revealed a gradually rising light curve. The observed radio emission follows a power-law with temporal index $\delta = 0.78 \pm 0.05$ and spectral index $\alpha = -0.61 \pm 0.05$, where $S_\nu(t, \nu) \propto t^\delta \nu^\alpha$ (Mooley et al., 2018a). The observed radio spectral energy distribution agrees with the spectral index connecting contemporaneous radio, optical, and X-ray measurements, implying a common source for the observed synchrotron emission (Levan et al., 2017b; Mooley et al., 2017; Margutti et al., 2018b; Mooley et al., 2018a; Troja & Piro, 2018).

The late turn-on of the X-ray and radio emission from GW170817 is not consistent with emission produced via an on-axis relativistic jet (Alexander et al., 2017a; Haggard et al., 2017; Hallinan et al., 2017; Margutti et al., 2017; Troja et al., 2017). Moreover, the gradual rise of the radio light curve rules out prompt gamma-ray emission originating from a jet with a “top-hat" azimuthal density profile observed off-axis, which would have produced a much steeper peak and decline than observed (Granot et al., 2002; Nakar et al., 2002). Instead, the light curve is consistent with mildly relativistic quasi-spherical outflow called a “cocoon" (Hallinan et al., 2017; Gottlieb et al., 2018a; Mooley et al., 2018a; Nakar & Piran, 2018) which may have some contribution from an embedded relativistic jet observed off-axis (some versions of which are also referred to as a “structured jet"; Lazzati et al., 2018; D'Avanzo et al., 2018; Margutti et al., 2018b; Resmi et al., 2018).

Based on the data available in the literature to date, it is not possible to establish whether or not a successful jet is present within the cocoon, as these scenarios exhibit similar behavior in the early stages of the afterglow evolution; or to determine the energy of the cocoon itself (see Figure 5 of Margutti et al., 2018b). The timescale of the peak flux density and the rate of decline afterwards can constrain the total energy of the outflow and the properties of a successful jet (if present). If the jet did not successfully break out of the cocoon (referred to as a “choked” jet) the observed emission is dominated by the quasi-spherical outflow (cocoon or dynamical ejecta; Gottlieb et al., 2018a) and the light curve will continue to rise; if the jet is successful

Table 2.1: New radio observations of GW170817

UT date	ΔT (d)	Telescope	ν (GHz)	Bandwidth (GHz)	Beam Size (arcsec)	S_ν (μJy)
2017 Dec 20.83	125.30	ATCA[a]	5.5	2.048	5.8×1.5	82.0 ± 9.3
			9.0	2.048	3.6×1.0	63.7 ± 8.2
2018 Jan 13.79	149.26	ATCA[a]	5.5	2.048	5.4×1.5	98.9 ± 8.5
			9.0	2.048	3.3×1.0	52.7 ± 6.5
2018 Feb 01.74	168.21	ATCA[b]	5.5	2.048	⋯	⋯ [c]
			9.0	2.048	⋯	⋯ [c]
2018 Feb 15.17	181.64	ATCA[d]	5.5	2.048	4.4×1.1	89.6 ± 13.3
			9.0	2.048	2.6×0.7	57.0 ± 10.9
2018 Mar 02.32	196.79	VLA[e]	2.5	1	1.3×0.5	91.0 ± 9.1
			3.5	1	1.3×0.5	66.9 ± 6.1

[a] With the 6C configuration (maximum baselines of 6 km) and program CX391 (PI: T. Murphy).
[b] With the 750A configuration (maximum baseline of 3.75 km) and program CX394 (PI: E. Troja).
[c] Insufficient data quality
[d] With the 750B configuration (maximum baseline of 4.5 km) and program CX394 (PI: E. Troja).
[e] With the A configuration (maximum baseline of 27 km) under a Director Discretionary Time program (VLA/17B-397; PI: K. Mooley).

(structured jet; Margutti et al., 2018b; Nakar & Piran, 2018) the light curve peaks sooner and declines more rapidly. In either case identifying when and how the light curve peaks also allows calorimetry of the cocoon emission (much as was done by Frail et al. 2000; Berger et al. 2004 for long GRBs).

To date, X-ray observations provide conflicting evidence as to whether the afterglow has peaked. *XMM-Newton* observations 135 days post-merger suggest the afterglow may have flattened (D'Avanzo et al., 2018), but *Chandra* observations show a continued rise or slow turnover at about 150 days post-merger (Haggard et al., 2018; Troja & Piro, 2018; Margutti et al., 2018b). A decreasing X-ray brightness would imply that either the synchrotron cooling frequency has shifted below the X-ray band (expected on timescales of 100–1000 days post-merger) and the spectrum of the source has evolved, or the light curve of the source from the radio to X-rays has peaked, but current data are not definitive that *any* change in the X-ray light curve has occurred.

We present further radio observations of GW170817 using the the Australia Telescope Compact Array (ATCA) and the Karl G. Jansky Very Large Array (VLA), covering the period 125–200 days post-merger. These observations demonstrate (Figure 2.1) that the radio afterglow has peaked at 149 ± 2 days post-merger and is now declining in flux density.

2.2 Observations and Data Reduction

2.2.1 ATCA

We observed GW170817 on 2017 December 20 and 2018 January 13 UT with the ATCA (PI: Murphy). Further observations of GW170817 with the ATCA were obtained on 2018 February 01 and 15 UT (PI: Troja); see Table 2.1 for details. The February 01 observation only had 4 out of 6 antennas available and after removing short baselines due to the compact configuration, the data quality was insufficient to make a meaningful measurement and the observation was discarded. We determined the flux scale and bandpass response for all epochs using the ATCA primary calibrator PKS B1934–638. Observations of PKS B1245–197 were used to calibrate the complex gains during the December and January observing epochs, while PKS B1244–255 was used in the February observation. All observations used two bands of 2048 MHz centered at 5.5 and 9.0 GHz.

We reduced the visibility data using standard MIRIAD (Sault et al., 1995) routines. The calibrated visibility data were split into the 5.5 and 9.0 GHz bands, averaged to 32 MHz channels, and imported into DIFMAP (Shepherd, 1997). Bright field sources were modeled separately for each band using the visibility data and a combination of point-source and Gaussian components with power-law spectra. After subtracting the modeled field sources from the visibility data, GW170817 dominates the residual image. Restored naturally-weighted images for each band were generated by convolving the restoring beam and modeled components, adding the residual map and averaging to form a wide-band image. Image-based Gaussian fitting with unconstrained flux density and source position was performed in the region near GW170817. The resulting source position agrees with the position of GW170817 observed by the *Hubble Space Telescope* (HST, Adams et al., 2017).

To examine the stability of the absolute flux calibration from epoch to epoch we measured the flux density of the phase calibrator (PKS B1245−197) and a compact reference source in the GW170817 field (RA= $13^{\mathrm{h}}09^{\mathrm{m}}53\overset{\mathrm{s}}{.}91$, Dec= $-23^{\circ}21'34\overset{''}{.}5$, $1.9'$ from GW170817) in each epoch and frequency band of the ATCA data. We do not use the host galaxy NGC 4993 as it is extended. We find that the mean and standard deviation of the phase calibrator flux density is 2.193 ± 0.013 Jy and 1.449 ± 0.021 Jy at 5.5 GHz and 9 GHz, respectively. This compares to within 0.1% with the values reported by the ATNF Calibrator Database[1]. The reference source is three orders of magnitude fainter than the phase calibrator but is a factor of at least three brighter than GW170817 and is within the same field, so it should provide an accurate indication of the flux density scale within the target field itself. The source is also visible regardless of which phase calibrator is used and so provides an independent test of flux scale stability. Across all epochs, we find that the mean flux density and standard deviation of the reference source flux density is $452 \pm 16\,\mu$Jy and $301 \pm 18\,\mu$Jy at 5.5 GHz, and 9.0 GHz, respectively. This suggests that our field flux density measurements are stable to within 2.9% and 5.4% at 5.5 GHz and 9.0 GHz, respectively, where those additional uncertainties when added

---

[1] `http://www.narrabri.atnf.csiro.au/calibrators/`

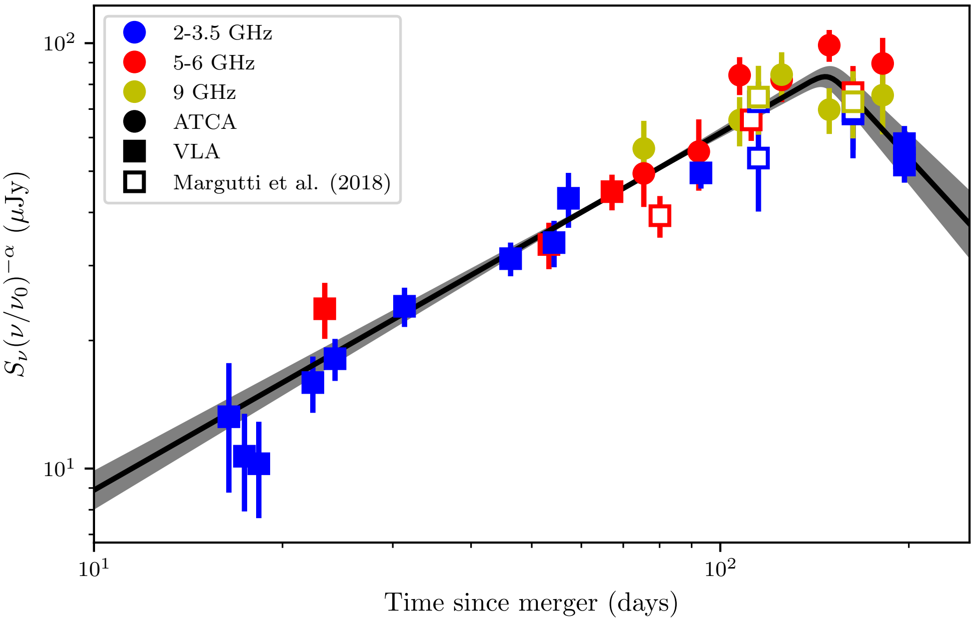

Figure 2.1: Light curve of GW170817 from ATCA (circles) and VLA (squares) observations grouped by frequency band, with 2–3.5 GHz (blue), 5–6 GHz (red), and 9 GHz (yellow). The flux densities have been adjusted to 5.5 GHz assuming a spectral index of $\alpha = -0.57 \pm 0.04$ (§ 2.3.1). Open squares denote observations from Margutti et al. (2018b), while filled symbols denote observations from this paper or other observations by our group (Hallinan et al., 2017; Mooley et al., 2018a). Our best-fit smoothed broken power-law with temporal index on the rise $\delta_1 = 0.84 \pm 0.05$, temporal index on the decay $\delta_2 = -1.6 \pm 0.2$ and peak time $t_{\rm peak} = 149 \pm 2$ days is shown in black, with uncertainties shaded.

in quadrature to the measurement uncertainties give reduced $\chi^2 = 1$ for the reference source. For GW170817 itself we measured the noise in the vicinity of the source to account for additional contributions from unmodeled sidelobes from the host galaxy NGC 4993 and included the additional uncertainties discussed above.

2.2.2 VLA

VLA observations of the GW170817 field were carried out on 2018 March 02. The Wideband Interferometric Digital Architecture (WIDAR) correlator was used at S band (2–4 GHz) to maximize sensitivity. We used J1248–1959 as the phase calibrator and 3C286 as the flux density and bandpass calibrator. The data were calibrated and flagged for RFI using the NRAO CASA (McMullin et al., 2007) pipeline. We then split and imaged the target data using the CASA tasks `split` and `clean`. We made final images by splitting the bandpass into 2 subbands of 1 GHz each.

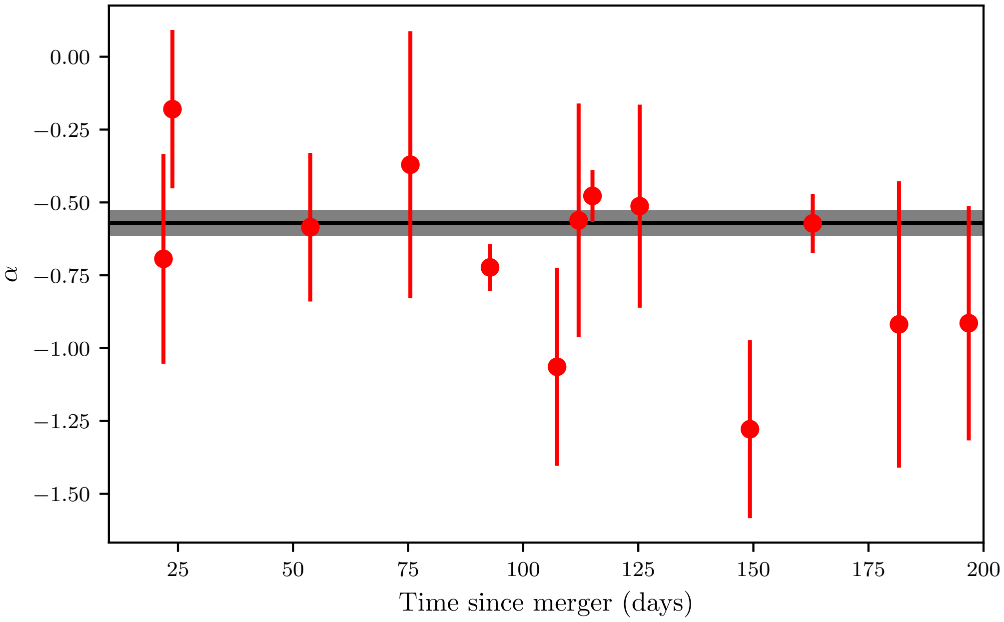

Figure 2.2: Spectral index (α) of contemporaneous radio observations from Hallinan et al. (2017), Mooley et al. (2018a), Margutti et al. (2018b) and this work. The best-fit spectral index from the first 120 days of the radio light curve, $\alpha = -0.57 \pm 0.04$, is shown in black, with uncertainties shaded.

2.3 Results and Discussion

2.3.1 Spectral Analysis

We first re-visit the spectral behavior of the radio emission. As in Mooley et al. (2018a) we fit a power-law of the form $S_\nu \propto \nu^\alpha t^\delta$ to the first 120 days of the radio light curve (before any sign of a turnover) and find a spectral index $\alpha = -0.57 \pm 0.04$ and temporal index $\delta = 0.84 \pm 0.05$. This is consistent with Mooley et al. (2018a) and with Margutti et al. (2018b), who find a joint radio-to-X-ray spectral index $\alpha = -0.585 \pm 0.005$ at 110 days and $\alpha = -0.584 \pm 0.006$ at 160 days post-merger.

We examined the variability of the spectral behavior using all quasi-simultaneous radio observations. We identified data-sets with more than one observation within ± 1 day and fit for a spectral index. These values are shown in Figure 2.2. We find the data largely consistent with a constant spectral index, with $\chi^2 = 15.9$ for 12 degrees-of-freedom. There appears to be no evidence for significant change in the spectrum of the source, consistent with previous radio, X-ray and HST observations (D'Avanzo et al., 2018; Lyman et al., 2018; Mooley et al., 2018a; Margutti et al., 2018b; Resmi et al., 2018).

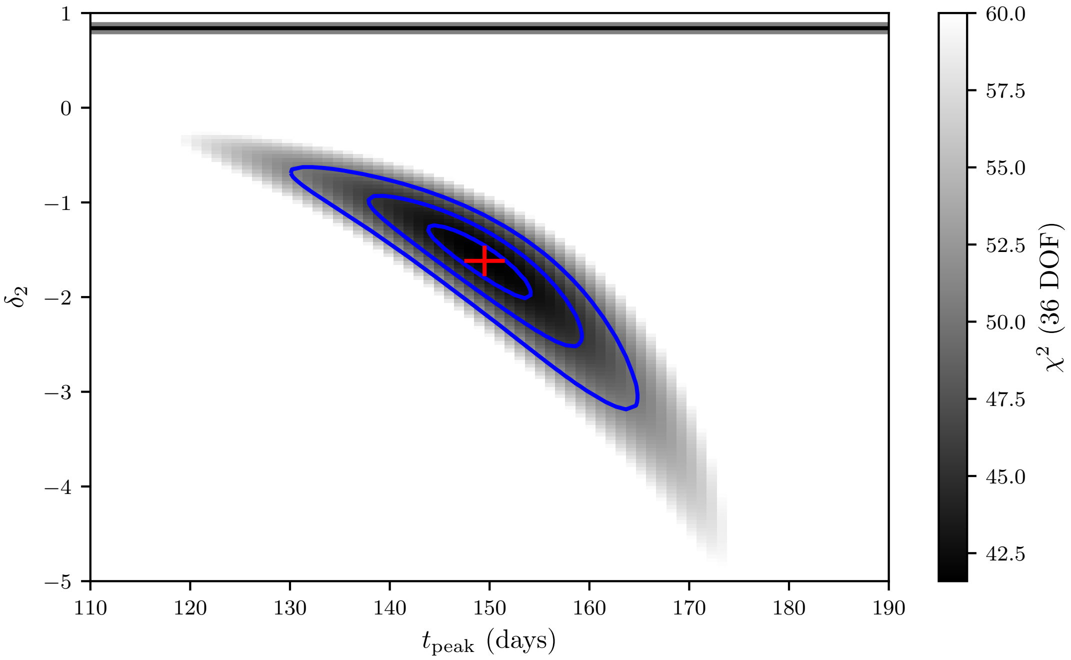

Figure 2.3: Two-dimensional joint probability distribution of δ_2 and $t_{\rm peak}$, assuming $\alpha = -0.57 \pm 0.04$ and $\delta_1 = 0.84$. The background gray-scale is the χ^2 for 35 degrees-of-freedom, with 1-, 2-, and 3-σ joint confidence contours are shown in blue. The best-fit value of $\delta_2 = -1.6 \pm 0.2$ and $t_{\rm peak} = 149 \pm 2$ days is shown in red. The temporal index of the light curve as it rises, $\delta_1 = 0.84 \pm 0.05$, is indicated by the black line with uncertainties shaded.

2.3.2 Light Curve Analysis

Figure 2.1 shows the light curve of GW170817 over the 2–9 GHz frequency range from the observations in Table 2.1 and the literature (Hallinan et al., 2017; Mooley et al., 2018a; Margutti et al., 2018b), scaling the flux density for each observation to 5.5 GHz based on the spectral index of $\alpha = -0.57 \pm 0.04$ calculated above. Assuming the light curve initially rises with a temporal index of $\delta_1 = 0.84$, peaks $t_{\rm peak}$ days post-merger, and fades with a temporal index of δ_2, we fit a smoothed broken power law[2] using the `Astropy` modeling package (Astropy Collaboration et al., 2018) that behaves as $S_\nu \propto t^{\delta_1}$ for $t \lesssim t_{\rm peak}$ and $S_\nu \propto t^{\delta_2}$ for $t \gtrsim t_{\rm peak}$ with a smooth transition around $t \approx t_{\rm peak}$. We do not expect to see any variability due to interstellar scintillation, due to the source size (Hallinan et al., 2017).

We have fit the lightcurve allowing the smoothing factor to freely vary and find a minor preference for small smoothing factors down to 0.001, corresponding to a transition of 0.3 days either side of the break. To approximate our observing cadence near the peak of the lightcurve we use a smoothing factor of 0.02 (corresponding to a <20 day transition) which produces no significant changes in fit parameters.

---

[2] `http://docs.astropy.org/en/stable/api/astropy.modeling.powerlaws.SmoothlyBrokenPowerLaw1D.html`

Figure 2.3 shows the two dimensional joint confidence region as a function of t_{peak} and δ_2, where we indicate the best-fit values, $\delta_2 = -1.6 \pm 0.2$ and $t_{\text{peak}} = 149 \pm 2$ days, and the 90% confidence region. The best fit has $\chi^2 = 41.6$ for 35 degrees-of-freedom. For a radio light curve that is continuing to rise, the temporal index would remain the same, $\delta_2 = \delta_1$, which we indicate with the dashed line in Figure 2.3. Comparing the $\chi^2(\delta_2 = \delta_1)$ to the minimum χ^2 for $\delta_2 = -1.6$, we find a change of 380 for one additional parameter and can exclude a light curve that continues to rise at greater than 5σ significance using an F-test. We further find a change of χ^2 of 35 from $\delta_2 = 0$ to the best-fit value $\delta_2 = -1.6 \pm 0.2$, leading to a declining light curve. Preliminary reduction of further observations confirms the observed trend.

2.3.3 Interpreting the Radio Light Curve

The observed light curve turns over and declines with no evidence for a steep rise coming with an energetically-dominant off-axis jet (Nakar & Piran, 2018), but a weaker jet may still be present. The relatively sharp peak in the radio light curve implies that the energy injection has reduced substantially (or stopped), or that the ejecta has collected mass comparable to its own. The former scenario would be relevant for a successful jet (e.g., Kasliwal et al., 2017; Lazzati et al., 2018; Mooley et al., 2018a; Margutti et al., 2018b; Troja et al., 2018; D'Avanzo et al., 2018) or a low energy choked-jet cocoon (e.g., Kasliwal et al., 2017; Gottlieb et al., 2018b; Piro & Kollmeier, 2018; Mooley et al., 2018a), while the latter would be relevant in the case of an isotropic fireball (i.e., dynamical ejecta; Nakar & Piran, 2011; Mooley et al., 2018a; D'Avanzo et al., 2018; Hotokezaka et al., 2018).

While no substantial degree of linear polarisation would be expected from isotropic dynamical ejecta, in the successful jet model the required asymmetry is built into the jet structure (the energy and speed of the various ejecta components are both functions of the angle from the jet axis; see e.g. Lazzati et al., 2018). Thus, the relevant emitting surface is never completely symmetric for misaligned observers, resulting in an appreciable degree of linear polarisation ($\sim 20\%$; Rossi et al., 2004). A detection of significant linear polarisation would thus point to a successful jet rather than isotropic dynamical ejecta (also see Gill & Granot, 2018).

The radio light curve can give the energy profile of the ejecta, but it is not sufficient for distinguishing between the contributions from radial and angular structures within the ejecta. Very Long Baseline Interferometry (VLBI) can, however, provide images at sub-milliarcsecond angular resolution, and thus constrain the geometry of the outflow. Distinguishing between the successful-jet, choked-jet cocoon and dynamical ejecta models is thus possible using VLBI observations.

The time of the radio peak is near the observed plateau on the X-ray light curve (Margutti et al., 2018b; Troja et al., 2018; Ruan et al., 2018b; D'Avanzo et al., 2018), and suggests that the X-rays peaked at the same time as the radio light curve. The turnover in the X-ray (and radio) light curve is therefore dynamical or geometric in origin, and the cooling break has (likely) not entered the X-ray band yet. This is consistent with the interpretation of D'Avanzo et al. (2018) and Margutti et al. (2018b) who find that the radio, optical and X-rays lie on the same power-law until day 150 post-merger.

The light curve of a relativistic jet afterglow will decay as t^{-p}, while in the non-relativistic regime the decline will be proportional to $t^{(15p-21)/10}$, with p the exponent on the distribution of electron energies, $N(E) \propto E^{-p}$ (Granot et al., 2002; Nakar & Piran, 2011). In the case of GW170817, $p = 2.17$ (e.g. Margutti et al., 2018b; Mooley et al., 2018a), so the expected decay slopes are $t^{-2.2}$ and $t^{-1.2}$. Our radio data are consistent with expectations for the mildly- or non-relativistic regimes. Based on the time and the flux density at the peak of the radio light curve, we can further calculate the isotropic equivalent energy (Nakar & Piran, 2018) as a few $\times$ 10^{50} erg for the cocoon scenario (also see Resmi et al. 2018) and a few $\times$ 10^{49} erg for the dynamical ejecta scenario. Both of those are lower than the isotropic-equivalent kinetic energies found for short GRBs (Fong et al., 2015).

If the peak of the light curve was dominated by an off-axis jet, then $\Gamma(\theta_{\rm obs} - \theta_{\rm jet}) \simeq 1$ (Nakar & Piran, 2018, where the bulk Lorentz factor of the jet is Γ, the off-axis angle of the observer is $\theta_{\rm obs}$, and the opening angle of the jet is $\theta_{\rm jet}$) implies that $(\theta_{\rm obs} - \theta_{\rm jet}) \simeq 20^\circ$, assuming that material with $\Gamma \simeq 3$ dominated the on-axis emission at peak. Therefore we can constrain $\theta_{\rm jet} \lesssim 8^\circ$ using the viewing angle constraint from LIGO/Virgo ($\theta_{\rm obs} < 28^\circ$; Abbott et al., 2017d).

Continued radio monitoring will be essential for constraining the decay index. A steep decline in the radio light curve would favor the scenario in which a successful jet broke out of the dynamical ejecta. Transition of the ejecta from the mildly-relativistic to the Newtonian regime would be characterized by deviation from a power-law decay and a change in spectral index, which could be detected with sensitive follow up observations. It is even possible for the ejecta to have angular structures that could cause the light curve to rise again: the early-time kilonova signal in the optical suggested the presence of ~ 0.05 $M_\odot$ material traveling at speeds of 0.1c to 0.3c which should give rise to a radio peak on timescales of a few years (Alexander et al., 2017a; Nakar & Piran, 2011, 2018). Finally, the full radio light curve of GW170817 will be crucial for calorimetry, since it will capture all of the energy in the ejecta. The total energy will further shed light into whether GW170817 is a standard short GRB viewed off-axis or it represents a distinct phenomenon.

2.4 Conclusions

We have presented new ATCA and VLA observations of GW170817 covering 125–200 days post-merger. Combined with previous radio observations these data show no evidence for spectral evolution, but they conclusively show that the radio counterpart peaked in brightness at 149 ± 2 days post-merger and is now declining. We use this to rule out emission being caused by highly energetic, quasi-isotropic outflow or highly energetic, highly-relativistic outflow but are not able to uniquely determine the geometry and structure of the outflow material. Continued radio monitoring will allow the temporal decay index to be accurately determined, although this may not be sufficient to establish the presence of a successful jet (Nakar & Piran, 2018) and degeneracies in the ejecta total energy and the density of the circum-merger environment may preclude confirmation of any particular model. Polarisation measurements and VLBI observations should be able to break this degeneracy and thus distinguish between the models (also see Gill & Granot, 2018).

Chapter 3

An Optimised Gravitational Wave Follow-Up Strategy with the Australian Square Kilometre Array Pathfinder

Dougal Dobie, Tara Murphy, David L. Kaplan, Shaon Ghosh,
Keith W. Bannister, Richard W. Hunstead

This chapter is reproduced from Dobie et al. (2019a), which was published in Publications of the Astronomical Society of Australia as "An optimised gravitational wave follow-up strategy with the Australian Square Kilometre Array Pathfinder". I wrote all text, produced all figures and wrote all code beyond an initial template script contributed by David Kaplan and Shaon Ghosh. Figures have been reformatted in a style consistent with the rest of this thesis.

The detection of a neutron star merger by the Advanced Laser Interferometer Gravitational-Wave Observatory (LIGO) and Advanced Virgo gravitational wave detectors and the subsequent detection of an electromagnetic counterpart has opened a new era of transient astronomy. With upgrades to the Advanced LIGO and Advanced Virgo detectors and new detectors coming online in Japan and India, neutron star mergers will be detected at a higher rate in the future, starting with the O3 observing run which will begin in early 2019. The detection of electromagnetic emission from these mergers provides vital information about merger parameters and allows independent measurement of the Hubble constant. The Australian Square Kilometre Array Pathfinder (ASKAP) is expected to become fully operational early 2019 and its 30 deg^2 field of view will enable us to rapidly survey large areas of sky. In this work we explore prospects for detecting both prompt and long-term radio emission from neutron star mergers with ASKAP and determine an observing strategy that optimises the use of telescope time. We investigate different strategies to tile the sky with telescope pointings in order to detect radio counterparts with limited observing time, using 475 simulated gravitational wave events. Our results show a significant improvement in observing efficiency when compared with a naïve strategy of covering the entire localisation above some confidence threshold, even when achieving the same total probability covered.

3.1 Introduction

The detection of gravitational waves from a neutron star merger (GW170817; Abbott et al., 2017d) by the Advanced Laser Interferometer Gravitational-Wave Observatory (LIGO) and the subsequent detection of counterparts across the electromagnetic spectrum (Abbott et al., 2017f) heralds a new era in astronomy. A short gamma-ray burst (sGRB; Abbott et al., 2017g; Goldstein et al., 2017) was detected 1.7 seconds post-merger with thermal kilonova emission detected in the following hours (Coulter et al., 2017; Abbott et al., 2017f, and references therein) and a non-thermal afterglow detected in X-rays (Evans et al., 2017; Haggard et al., 2017; Troja et al., 2017) and radio (Hallinan et al., 2017) in the following days.

Observations over the first hundred days suggested that the observed X-ray and radio emission originate from the same source: optically-thin synchrotron emission (Mooley et al., 2018a). The radio and X-ray light curves initially rose according to a power-law of the form $S_\nu(t, \nu) \propto t^{\delta_1} \nu^{\alpha}$ with spectral index $\alpha = -0.585 \pm 0.005$ (Margutti et al., 2018b) and temporal index $\delta_1 = 0.84 \pm 0.05$, peaked 149 ± 2 days post-merger (Dobie et al., 2018a) and has declined according to a similar power-law with temporal index $\delta_2 = -2.2 \pm 0.2$ or steeper (Mooley et al., 2018d; Troja et al., 2019a) out to 300 days post-merger.

These observations differ from the expected radio emission predicted prior to the event, with emission from a relativistic jet peaking within days of the merger and only detectable in a dense circum-merger medium, or late-time emission from sub-relativistic ejecta peaking on timescales of thousands of days (e.g. Metzger & Berger, 2012; Hotokezaka et al., 2016). Instead, this radio emission has been proposed to be produced by mildly relativistic quasi-spherical outflow (a "cocoon"; Gottlieb et al., 2018a; Nakar & Piran, 2018) with some contribution from an embedded off-axis structured jet (Lazzati et al., 2018; Alexander et al., 2018; D'Avanzo et al., 2018; Margutti et al., 2018b; Resmi et al., 2018; Mooley et al., 2018b). We emphasise that we do not know whether the afterglow of GW170817 is typical of the population, although this will be answered as more events are detected.

Monitoring of the radio lightcurve allows us to estimate factors such as the circum-merger density, and energetics of the merger outflow, and may yet be able to determine the fate of the jet (Nakar & Piran, 2018). Combining Very Long Baseline Interferometry (VLBI) observations with monitoring of the radio lightcurve can also determine the inclination angle of the merger (Mooley et al., 2018b) and improve standard-siren constraints on H_0 (Abbott et al., 2017e; Hotokezaka et al., 2019)

With upgrades to the Advanced LIGO and Advanced Virgo detectors, and the Kamioka Gravitational Wave Detector (KAGRA; Somiya, 2012; Aso et al., 2013) and LIGO-India (Unnikrishnan, 2013) coming online, the detection rate of gravitational wave events will increase substantially to tens and possibly hundreds of events per year, localised to tens of $\deg^2$ (Abbott et al., 2018). Observation of a large number of events will enable better understanding of the neutron star merger population, and determine whether GW170817 was typical or exceptional. In particular, radio observations will constrain parameters including the isotropic equivalent energy and the circum-merger density and also the radio emission mechanism.

Table 3.1: ASKAP design specifications and at the time of GW170817. Image RMS (σ) is calculated using the radiometer equation. Survey speed (SS) assumes a 100μJy image RMS and ignores telescope overheads. The telescope is expected to reach design specifications in early 2019.

	for GW170817	Design
Antennas	16	36
T_{sys} (K)	70	50
Bandwidth (MHz)	240	300
$\sigma_{\mathrm{10min.}}$ (μJy/beam)	300	82
$\sigma_{\mathrm{12hr.}}$ (μJy/beam)	35	12
SS (deg^{2}hr^{-1})	21	270
Frequency (GHz)	0.7–1.4	0.7–1.8

In this paper we compare possible observing strategies for detecting radio emission from gravitational wave events with the Australian Square Kilometre Array Pathfinder (ASKAP; Johnston et al., 2008). ASKAP is currently being commissioned and the final telescope will consist of 36 antennas, 12 m in diameter, separated by baselines ranging from 37 m to 6 km resulting in angular scales of $10''$–$3'$ at frequencies from 700–1800 MHz. ASKAP is designed for all-sky surveys and uses MkII phased-array feeds (PAFs, Hampson et al., 2012) consisting of a 6×6 array of beams, which produce a 30 deg^2 field of view (FoV). Combined with its angular resolution and sensitivity, this makes ASKAP capable of following up poorly localised gravitational wave events with no other electromagnetic counterpart that would take hundreds of pointings to cover with other gigahertz-frequency radio telescopes. We analyse observing strategies previously designed for use with other telescopes, and determine the best way to implement them to optimise ASKAP follow-up.

3.2 Searching for radio emission

3.2.1 ASKAP Follow-up of GW170817

We began follow-up observations of GW170817 with ASKAP 15 hours after the event, searching for coherent radio emission in fly's-eye mode (Bannister et al., 2017b,c). Imaging observations began two days after the event, covering the 90% confidence region using the strategy presented in this paper (Dobie et al., 2017a,b); see Andreoni et al. (2017) for further details.

Figure 3.1 shows the lightcurve of GW170817 from Dobie et al. (2018a) adjusted to 1.4 GHz based on a spectral index of $\alpha = -0.58$, along with the 5σ detectability limit for the current telescope configuration and the design specifications (see Table 3.1). The light curve peaked at the limit of ASKAP detectability and rapidly declined below this limit. However, if the telescope had been at full sensitivity for this event, it would have been detectable from 40 days post-merger and peaked at a flux density equivalent to a 15σ detection.

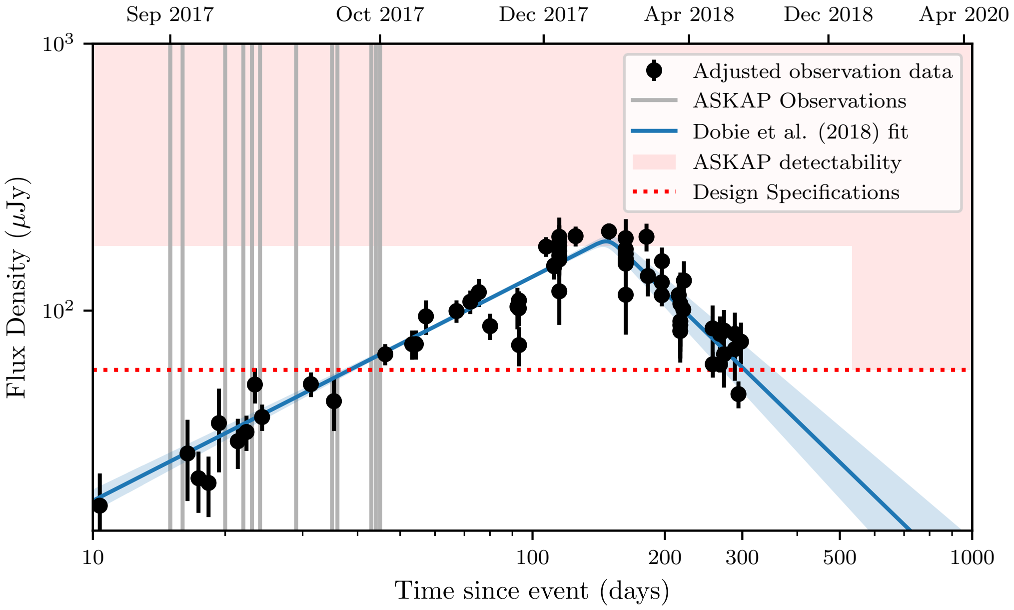

Figure 3.1: Detectability of GW170817 with ASKAP. Radio observations from Hallinan et al. (2017); Mooley et al. (2018a); Dobie et al. (2018a); Mooley et al. (2018d) scaled to 1.4 GHz based on a spectral index of $\alpha = -0.58$ are noted in black, while the smoothed broken power law fit from Dobie et al. (2018a) is shown in blue, with 1σ uncertainties shaded. The ASKAP 5σ detectability limits (see Table 3.1) for current (shaded) and design (dotted line) specifications are shown. Vertical lines denote ASKAP observations of GW170817.

3.2.2 Long-Term Synchrotron Emission

The focus of this paper is determining an observing strategy to detect the long-term radio emission (tens–thousands of days) from gravitational wave events, that optimises the use of telescope time. This emission could be produced, for example, by a relativistic jet launched after the merger (Granot et al., 2002) on timescales of days–weeks or the dynamical ejecta that caused the kilonova on timescales of months–years (Nakar & Piran, 2011; Hotokezaka & Piran, 2015).

The prompt gamma-ray emission from GW170817 is most likely driven by a jet, which produces a hot cocoon as it propagates into the merger ejecta (Gottlieb et al., 2018a). As the cocoon expands quasi-spherically into the circum-merger medium, the high velocity components produce the early-time radio emission we have detected. A jet that successfully propagated through and breaks out of the cocoon (a structured jet) may drive the late-time radio emission. Continuum radio monitoring of GW170817 has been unable to distinguish between emission arising solely from a cocoon or structured jet (Nakar & Piran, 2018). VLBI observations suggest that the late-time radio emission is produced by a relativistic jet (Mooley et al., 2018b), while polarisation measurements place constraints on the geometry and strength of the magnetic field within the jet (Corsi et al., 2018b).

It may also be possible to detect long-term synchrotron emission where searches for kilonovae are limited by intrinsic factors like inclination angle, or extrinsic factors such as dust-obscuration, proximity to the Galactic Plane or solar angle. For example, optical and X-ray monitoring of GW170817 was limited by space telescope solar angle constraints, and the kilonova emission vital for localising the event would not have been detectable by ground based optical telescopes if the event had occured 3 months later in the daytime.

3.2.3 Prompt Radio Emission

It has been suggested that neutron star mergers may produce prompt, coherent radio emission. This may occur through magnetic braking as the magnetic fields of the two neutron stars synchronise (Totani, 2013), an induced electromotive force in the final stage of the merger accelerating electrons to ultra-relativistic speeds (Wang et al., 2016). The resultant supramassive neutron star may also collapse into a black hole (Zhang, 2014; Falcke & Rezzolla, 2014) and produce a relativistic plasma outflow (Pshirkov & Postnov, 2010). This emission would possibly be manifest as a signal similar to a Fast Radio Burst (FRB, Lorimer et al., 2007; Thornton et al., 2013), although there may be other forms of short-timescale radio emission present at the time of the merger such as pulsar-like behaviour (Lipunov & Panchenko, 1996).

It may be possible to search for prompt radio emission from neutron star mergers with ASKAP, although this is not yet feasible due to observational constraints. During O2 the latency in receiving gravitational wave alerts was a few minutes and even after receiving the alert we are limited by observing overheads – ASKAP can slew at a rate of 3 $\text{deg}\,\text{s}^{-1}$ in azimuth and 1 $\text{deg}\,\text{s}^{-1}$ in elevation. In contrast, the dispersion-induced time delay is expected to be $< 10\,\text{s}$ at 700 MHz (Chu et al., 2016), the lowest frequency ASKAP band. However, the gravitational wave signal of the inspiral prior to GW170817 was detectable by Advanced LIGO $\sim 100\,\text{s}$ pre-merger (Abbott et al., 2017d). In the future it may be possible to detect events prior to the merger (Chan et al., 2018), thereby allowing for telescopes to be on-source as the merger occurs.

ASKAP can observe in fly's-eye mode, where the telescope array is operated as a set of individual dishes pointed independently (and not necessarily contiguously). Each antenna has a FoV of 30 deg^2, split into 36 beams in a 6×6 square, yielding a potential observing area of hundreds of deg^2, allowing entire gravitational wave localisation regions to be covered instantaneously. Using the full 36 antenna array in fly's-eye mode allows for simultaneous coverage of 3% of the sky. The probability of serendipitous detection of prompt radio emission can be increased by pointing the telescope at the region of sky where Advanced LIGO has maximum sensitivity. We note that due to the expected timescales for synchrotron radio emission discussed in Section 3.2.2, it is unlikely that low-latency observations in continuum mode will make any detection.

ASKAP has proven effective in detecting FRBs in fly's-eye mode, with 23 detected so far (Bannister et al., 2017a; Macquart et al., 2019; Shannon et al., 2018), making up over a quarter of all FRBs published to-date[1] It is capable of arcminute-localisation of coherent radio emission in fly's-eye mode, insufficient to identify the host galaxy, but small enough for follow-up with other telescopes (Mahony et al., 2018).

The detection of any form of coherent radio emission coincident with a gravitational wave event would aid in localising and constraining the event in a similar way to the concurrent detection of GW170817 and GRB170817A (Abbott et al., 2017g), and may shed light on the origin of FRBs and the emission mechanisms of neutron star mergers. This localisation could be improved by using distance constraints from the 3D gravitational wave skymap and estimating the distance of the coherent emission from its dispersion measure (DM; Ioka, 2003; Inoue, 2004), although current DM-distance estimates are quite inaccurate due to large uncertainties in the DM contribution from the Milky Way Halo (Dolag et al., 2015) and the host galaxy. Therefore while we can narrow the search for the host galaxy we are unlikely to be able to uniquely identify it solely from prompt emission detected in fly's-eye mode. There is also the prospect for an archival search for gravitational wave events coincident with ASKAP detections of FRBs (Abbott et al., 2016a).

3.3 An Optimised Pointing Strategy

3.3.1 Observing Gravitational Wave Source Localisations

The first four observed gravitational wave events were detected by the LIGO Hanford and LIGO Livingston detectors and localised to areas ranging between 520 and 1600 deg^2 (Abbott et al., 2017b,h). The first three-detector detection (using Advanced LIGO and Advanced Virgo), GW170814, was localised to 60 deg^2 (Abbott et al., 2017c), while GW170817 was localised to 28 deg^2 (Abbott et al., 2017d). During O3 the expected median 90% localisation area is 120–180 deg^2, with up to 20% of events having localisation areas smaller than 20 deg^2 (Abbott et al., 2018). Future observing runs, with improved sensitivity in existing detectors and the addition of the KAGRA and LIGO-India detectors, may be able to localise the majority of events to areas $< 12\ \mathrm{deg}^2$ (Abbott et al., 2018).

It is not always feasible to observe entire localisation regions due to the large amount of observing time required. Therefore an optimised observing strategy can significantly increase the probability of successfully detecting an electromagnetic counterpart and decrease the observing time required. Observing a smaller fraction of the localisation area can significantly reduce the search area. For networks of three or four gravitational wave detectors, observing the 50% confidence region reduces the search area by up to a factor of ten compared to observations of the 90% confidence region (Klimenko et al., 2011; Bartos et al., 2014). Using simulated skymaps from Singer et al. (2014) we find that the 50% confidence region for two and three detector networks is on average five times, and up to twenty times, smaller than the 90% confidence region.

---

[1] `http://www.frbcat.org/` (Petroff et al., 2016).

Salafia et al. (2017) used search strategies based on models of afterglow lightcurves and other parameters to determine spatial and temporal observing strategies that improve observing efficiencies. Other optimisation strategies use estimates of afterglow luminosity to determine how time should be allocated between candidate hosts (e.g. Coughlin & Stubbs, 2016; Arcavi et al., 2017b). However, current models of radio emission from neutron star mergers (see Section 3.2.2) are much less well constrained than kilonova models, and are also highly dependent on the circum-merger environment, making this approach less compelling for radio follow-up.

A strategy that has been implemented in follow-up using telescopes with a small (∼ arcminute) FoV is targeted observations of specific galaxies in the localisation volume (e.g. Evans et al., 2016; Arcavi et al., 2017b). We discuss applying a galaxy targeted technique for ASKAP in Section 3.3.4.

3.3.2 Tiling Strategies for Wide-Field telescopes

Ghosh et al. (2016) discuss possible pointing strategies for wide-field optical telescopes, and consider each pointing as a discrete, pre-determined tile placed on the sky. Gravitational wave sky-maps contain the three-dimensional (Right Ascension, Declination, distance) probability distribution function (PDF) of the event using the Hierarchical Equal Area isoLatitude Pixelization (HEALPix[2]) projection, which divides the sky into pixels of equal surface area. The coverage of a single tile is then the sum of the PDF across all pixels contained in it.

The sensitivity of a radio telescope is not uniform across the FoV. Traditional radio interferometers with single pixel feeds have primary beams with highest sensitivity at the beam centre, and decreasing radially within the central lobe. An observing strategy using rectangular tiles will therefore result in non-uniform sensitivity across the localisation region. In contrast, the 6×6 square arrangement of beams in the ASKAP PAF produces a FoV that is accurately represented by a square tile. The sensitivity across a single pointing is not uniform, but uniform sensitivity across the interior beams can be achieved by interleaving multiple pointings. We represent the ASKAP FoV as a 5.5×5.5 deg and allow a 0.5 deg overlap between tiles to ensure uniform sensitivity around the edges. In this section we outline four possible tiling strategies that could be used for ASKAP observations.

Contour Covering tiles

The simplest strategy to cover a gravitational wave localisation region is to cover entirely a given probability contour, for example, 90%. We consider the minimum area containing 90% of the event PDF (the 90% contour) and select all tiles from a pre-computed grid that contains any part of the contour, to achieve at least 90% coverage. This strategy is sub-optimal because it covers significant extraneous area if the shorter axis of the localisation ellipse is the width of a few tiles or smaller.

---

[2] `http://healpix.sourceforge.net/`

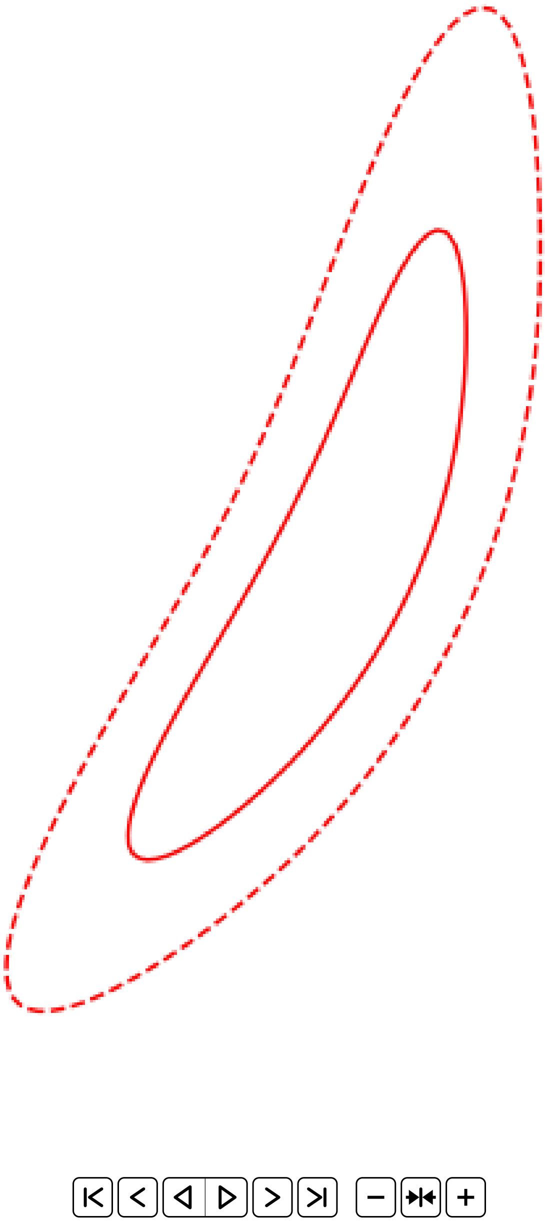

Figure 3.2: Illustration of the greedy ranked tiling strategy. Red lines correspond to 50% (solid) and 90% (dashed) localisation contours for an illustrative skymap. The gravitational wave skymap is covered with a pre-defined, overlapping grid of tiles (grey), ranked by their total enclosed probability. The highest ranked tile (blue) is selected and the probability in the region enclosed by it is set to zero. The tiles are then re-ranked and the next tile chosen, until the desired probability target is reached. Please use Adobe Acrobat to view the animated graphic.

Ranked tiles

A more effective way to minimise the extraneous area observed is the method of ranked tiles. Each tile is ranked by its sky coverage and tiles are chosen such that the desired confidence level is achieved with the minimum number of tiles. For this paper we implement both of the above strategies adapted for the ASKAP FoV as defined in Ghosh et al. (2016).

Greedy ranked tiles

While the approach above worked well for a single set of non-overlapping tiles, newer facilities have the ability to have multiple sets of fixed tiles, or even to not use fixed sets of tiles at all. Therefore Ghosh et al. (2017) introduced the method of greedy ranked tiles, which builds upon the ranked tiles strategy and allows overlapping tiles to be used without double-counting the probability contained in the overlapping region. The probability in the highest ranked tile is added to the cumulative probability coverage and then the area covered by the tile is zeroed out on the skymap. The process is then repeated until the desired overall coverage is achieved. An animation of this process is shown in Figure 3.2. We use the same strategy defined by Ghosh et al. (2017), adapted for the ASKAP FoV.

Shifted ranked tiles

Ghosh et al. (2016) proposed that the ranked tiles strategy could be further optimised by using a grid of tiles free to be shifted in strips of constant declination. In this paper our strategy differs slightly as we only shift groups of adjacent tiles together, rather than entire strips of declination. Our implementation of the shifted ranked tiles strategy is defined as follows:

1. Calculate the ranked tiles given by the primary sky grid and select the tiles that cover a given confidence region (for this example we use 90%).

2. Find the tiles that are adjacent to those selected above and include those in our sample. This gives the ranked tiles with a one-tile buffer, which we call the tile test-set. All other tiles are discarded for computational efficiency.

3. Split the test-set into groups of constant declination. For each declination strip, group adjacent tiles together. This is required as some gravitational wave skymaps are multimodal (e.g. bimodal, multiple lobes, crosses) and shifting entire strips of declination for these events is suboptimal.

4. For each group of adjacent tiles, shift the tiles in increments of Right Ascension. Re-calculate the ranked tiles for the test-set and record the number of tiles required for at least 90% coverage and the actual coverage percentage at each shift. The optimal shift is then the shift which produces the required coverage in the least number of tiles. If multiple shifts produce the required coverage with the same number of tiles, choose the shift that maximises the coverage percentage. Repeat for each strip of declination in the test-set.

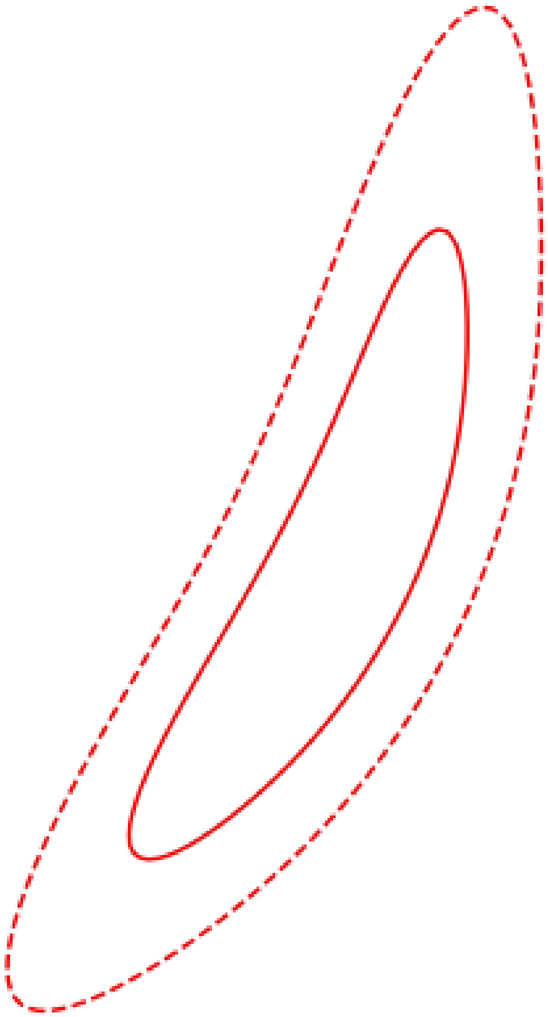

Figure 3.3: Illustration of the shifted ranked tiling strategy. Red lines correspond to 50% (solid) and 90% (dashed) localisation contours for an illustrative skymap. The gravitational wave skymap is covered with a single pre-defined grid of non-overlapping tiles (light grey). The ranked tiles required to cover the desired probability level are selected (dark grey) and grouped into strips of constant Declination. Each group is iteratively shifted in Right Ascension and the shift resulting in the desired probability enclosed within the minimum number of tiles. The set of optimally shifted tiles are then ranked and any extraneous tiles removed, leaving the optimal set of tiles for the entire localisation (blue). Please use Adobe Acrobat to view the animated graphic.

5. Apply the optimal shift in each declination strip to the test-set and re-calculate the ranked tiles.

Because the ranked tile grid is periodic, optimisation only needs to be performed for one tile-width. An animation of this process is shown in Figure 3.3.

3.3.3 An Observing Strategy for ASKAP

Both the greedy and shifted ranked tile strategies are more complex and require more computation time compared to the other two strategies. The added burden of these strategies may outweigh any benefits they provide for some telescopes, but not for others, and here we consider an observing strategy designed specifically for ASKAP. Three main factors dictate the way we search for gravitational wave afterglows with ASKAP: our wide FoV, our transient detection strategy and our large observational latency. In optimising our observing strategy, we must consider all three factors.

Field of View

ASKAP has a FoV of 30 deg^2 — significantly larger than all other GHz-frequency radio telescopes and larger than many optical telescopes involved in gravitational wave follow-up (Abbott et al., 2017f). The sky localisation capability of a telescope does not scale linearly with the FoV, as telescopes with a FoV comparable to the dimensions of the localisation region cover significantly more area that is extraneous to the localisation contour. The wide FoV of ASKAP allows localisation regions to be covered with fewer pointings than other telescopes, meaning that improving our observing strategy by a small number of pointings can result in significantly larger fractional improvements.

Detection of transients and variables

Transients surveys with optical telescopes typically use image differencing, where a deep reference image is subtracted from the observed image and removes all objects that are constant with time. The difficulty in image subtraction lies in matching parameters, including the instrument calibration and point-spread function, between the two science and observed images. This is easier if both images cover the same region of sky, so surveys for optical transients gain significant advantage from using a pre-defined grid of pointings.

However, in radio astronomy image differencing is rarely used due to the incomplete $u - v$ coverage of interferometers which causes the point-spread function (the dirty beam) to be more complex than a simple Airy disc. This introduces image artefacts that can easily be mistaken for transient sources in a difference image. Transient searches with ASKAP use the Variables and Slow Transients pipeline (Murphy et al., 2013), which detects transients via a process of source-finding, cross-matching and lightcurve analysis. Therefore there is no advantage in using a pre-defined grid of pointings. Improved transient detection methods such as image subtraction of LST-aligned observations are being tested with ASKAP (Stewart et al. in prep.).

Latency

Calculating an optimal pointing strategy can be time consuming. Ghosh et al. (2017) find that calculating the shifted ranked tiles strategy for some skymaps can take up to an hour, which is significant when searching for kilonova emission from a counterpart, which may peak within hours of the merger. For any prompt follow-up we require a strategy that can be computed on timescales of seconds.

The time taken to compute the pointing strategy is not a concern when searching for long-term radio emission, which may not peak until tens–hundreds of days post-merger. Additionally, the current process of detecting gravitational wave events has multiple stages of localisation and parameter estimation, with each stage being more accurate but slower to compute. Currently an event can be detected in seconds, rapidly localised in minutes and full parameters calculated between hours and days post-merger (Singer et al., 2014). For GW170817 the event was initially localised to a 90% credible region of 31 deg^2 four hours post-merger (LIGO Scientific Collaboration & Virgo Collaboration, 2017a), which was refined to 28 deg^2 two months post-merger (LIGO Scientific Collaboration & Virgo Collaboration, 2017b) and finally 16 deg^2 nine months post-merger (Abbott et al., 2019a).

When searching for long-term emission, the lack of latency constraints allows us to wait for the improved parameter estimates, which will typically result in a smaller localisation region and therefore less required observing time. The localisation region can also change significantly once the full parameter estimate has been computed, as it also includes further noise removal and better detector calibration. For example, the full parameter localisation of GW170814 was 40% smaller than the rapid localisation with only a small overlap between the two (Fig. 3 of Abbott et al., 2017c). This advantage is not applicable to searches for prompt radio emission.

3.3.4 Galaxy Targeting

The observing time required for follow-up of an event can be reduced by targeting galaxies within the localisation volume. This technique was employed in the follow-up of GW170817 where the localisation region was initially narrowed to 54 candidate host galaxies (e.g. Cook et al., 2017) before the optical counterpart was detected in the third ranked galaxy, NGC 4993.

We use the Galaxy List for the Advanced Detector Era (GLADE, Dálya et al., 2018), an all-sky galaxy catalogue aimed at supporting electromagnetic follow-up of gravitational wave events. The catalogue contains 3,262,883 objects, although measurements of blue luminosity are only available for half of them. GLADE is complete to a distance of 37 Mpc and is 54% complete to the minimum Advanced LIGO O3 detector horizon of 120 Mpc (Dálya et al., 2018).

A single ASKAP pointing can contain hundreds of galaxies. Averaging across the sky, excluding the Galactic Plane ($|b| < 10\,\mathrm{deg}$), we find a typical ASKAP pointing contains 169 GLADE galaxies within a distance of 200 Mpc and with a blue luminosity greater than $10^9\,L_{B,\odot}$. GLADE is $\sim 40\%$ complete at 200 Mpc, so we expect a typical ASKAP pointing to contain 420 galaxies subject to these constraints.

Targeting individual galaxies is therefore inefficient and instead we focus observations on regions containing a high density of candidate host galaxies by convolving the gravitational wave skymap with a galaxy catalogue using the method described by Evans et al. (2016). For each pixel, $P_{\rm GW}(D)$ is the probability that the event occurred in that pixel at a distance D. This is calculated assuming a Gaussian distribution and the mean and standard deviation given by the 3D skymap.

We determine the completeness of the galaxy catalogue as a function of distance, $C(D)$, as in White et al. (2011), from the B-luminosity for each galaxy. We then compare the cumulative blue luminosity of galaxies in the catalogue with the expected value, using the blue luminosity density derived from the SDSS survey by Kopparapu et al. (2008), $(1.98 \pm 0.16) \times 10^{-2} L_{10}$ Mpc^{-3} (where $L_{10} = 10^{10} L_{B,\odot}$) as a reference. Within ~ 35 Mpc the calculated completeness exceeds 100% (i.e. the expected value is an underestimate) and in this case we set $C(D) = 1$.

We now consider each galaxy contained in the pixel and calculate $P_{\rm g}(D)$, the total probability that there is a galaxy at a given distance, D, weighted by its blue luminosity. The GLADE catalogue does not include distance uncertainty for most galaxies, and where this is the case we assume a 10% uncertainty. The combined probability for each pixel is then given by

$$P_p = P_{\rm GW}(D) P_{\rm g}(D) C(D) + P_{\rm GW}(D)(1 - C(D)), \tag{3.1}$$

where the first factor corresponds to the event occurring in a catalogued galaxy, and the second corresponds to the event occurring in a non-catalogued galaxy. The convolved skymap is then normalised such that $\sum_{\rm allsky} P_p = 1$, so that the tiling efficiency can be easily compared between simulated events.

The catalogue of simulated events that we employ for testing our algorithms uses randomly generated event positions distributed homogeneously across the sky. To test the galaxy-targeting strategy we use the same method as Evans et al. (2016) and first consider whether or not the event occurred in a galaxy in our catalogue. We generate a random number, $0 < x < 1$, and compare it with the completeness of the galaxy catalogue at the distance of the event. If $x \geq C(D)$ then the event is considered to have occurred in an uncatalogued galaxy and we simply convolve the skymap and catalogue using equation 3.1. If $x < C(D)$, the event is treated as having occurred in a catalogued galaxy. We randomly choose a galaxy from the catalogue, weighted by $P_{\rm g}(D)$, the probability that the galaxy is located at the distance of the merger. We then rotate the positions of every galaxy in the catalogue so that the position of the selected host galaxy is consistent with the location of the simulated event.

3.4 Comparison of Pointing Strategies

3.4.1 Coverage of Simulated Skymaps

We tested our pointing strategy using 475 skymaps from the 2016 O2 scenario described in Singer et al. (2014). This consists of the two Advanced LIGO detectors and Advanced Virgo operating on an independent and random (i.e. downtime is not correlated between detectors) 80% duty cycle, with a neutron star merger range of 108 Mpc for the two Advanced LIGO detectors and 36 Mpc for Advanced Virgo. This is greater than the achieved sensitivity during O2, but below the expected O3 sensitivity of 120–170 Mpc for Advanced LIGO and 65–85 Mpc for Advanced Virgo (Abbott et al., 2018).

We consider all 475 skymaps and compare the outcome when covering them with contour-covering, simple ranked tiles, greedy ranked tiles or the shifted ranked tile approach. We have considered probability coverage targets ranging from 20% to 95%, as approaching 100% coverage drastically increases computation time while producing diminishing returns. The contour-covering and greedy ranked tiles methods typically take seconds to compute for a single skymap, while the shifted ranked method takes minutes. Computing the reusable set of tiles for both ranked tiles strategies has an additional overhead of 10–20 hours depending on the desired resolution. We evaluated each method based on the average probability density per tile it achieves, and the fraction of events it detects. We do not consider whether the radio emission from the event would have been detectable, as the models for possible emission vary by orders of magnitude in temporal behaviour and peak flux density, and instead consider an event to be detected if it occurred within a tile.

We use the ranked and contour-covering tiles strategies as a baseline for comparing the shifted and greedy ranked tiles methods, both for the original gravitational wave skymaps and the convolved galaxy-targeted skymaps. The galaxy targeting strategy was then tested using the greedy optimisation strategy.

3.4.2 Comparing Advanced Ranked Tiling Strategies

Figure 3.4 shows the average probability coverage per tile as a function of target probability for all four tiling strategies using the gravitational wave skymaps. As expected, the contour-covering strategy is the least efficient, while the greedy and shifted ranked tile strategies outperform the simple ranked tile strategy by 2.4 and 3.2 percentage points per tile on average. At 90% probability coverage this corresponds to the shifted ranked tile strategy producing ∼10% more efficient coverage per tile compared to the greedy ranked tile strategy.

We also consider our ability to detect events. Figure 3.5 shows the fraction of all 475 events that were detected as a function of probability coverage for both the shifted and greedy tiles. Due to the wide FoV of ASKAP, even an optimal tiling strategy results in extraneous coverage, so we have plotted the detection fraction as a function of both the target and achieved probability coverage averaged across all events. Both tiling strategies produce similar results across the entire probability

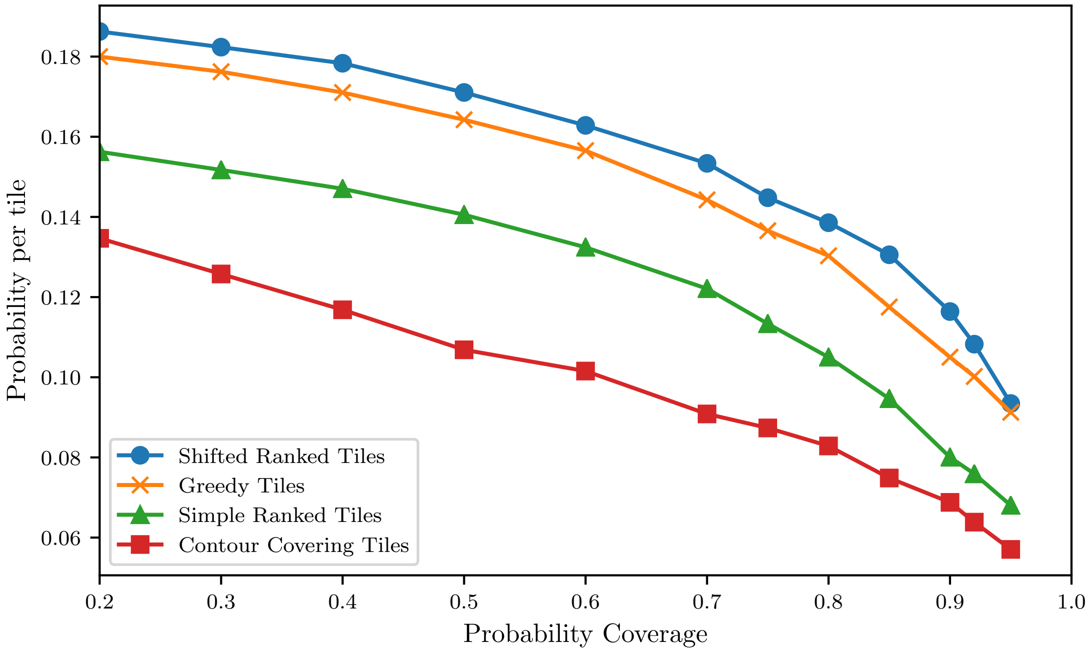

Figure 3.4: Covered probability per tile, for all four tiling strategies averaged across their application to 475 simulated skymaps. The two optimised strategies significantly outperform the contour covering and simple ranked tiled strategies.

range. As expected, the detection fraction is approximately equal to the covered probability, confirming that Advanced LIGO-Virgo skymaps do accurately recreate the probability distribution function of the event.

As expected, the shifted ranked tiles strategy slightly outperforms the greedy ranked tiles strategy. For most skymaps they will achieve identical results, but in some cases the greedy strategy produces a sub-optimal strategy as it only considers the next-most optimal tile at each step rather than the final overall tiling. This can result in gaps between tiles, or cases where more tiles are used than necessary. For example, if there is a strip of probability 2 tiles wide the greedy strategy will place the first tile at the centre of the strip, resulting in 3 tiles being required to cover the localisation (e.g. compare the tiling in Figures 3.2 and 3.3). The shifted ranked tiles strategy does not have these issues, as it considers the placement of independent groups of adjacent tiles. Each iteration computes the total probability coverage of the final overally tiling, and by considering adjacent tiles there are no gaps.

The actual gain in telescope time is dependent on the localisation area, the desired probability coverage, the total number of observing epochs and the time allocated per epoch. However, if we consider a probability coverage of 90% and an allocation of 1 hour per pointing we can estimate the gain in telescope time for specific events. Simulated event 493048 is detected with the LIGO-Hanford and Advanced Virgo detectors and has a 90% localisation of 135 deg^2, comparable the median detector horizon for O3. The simple ranked tiles strategy requires 15 tiles to cover the event, while the shifted and greedy ranked tiles strategies require 10 and 11 tiles respec-

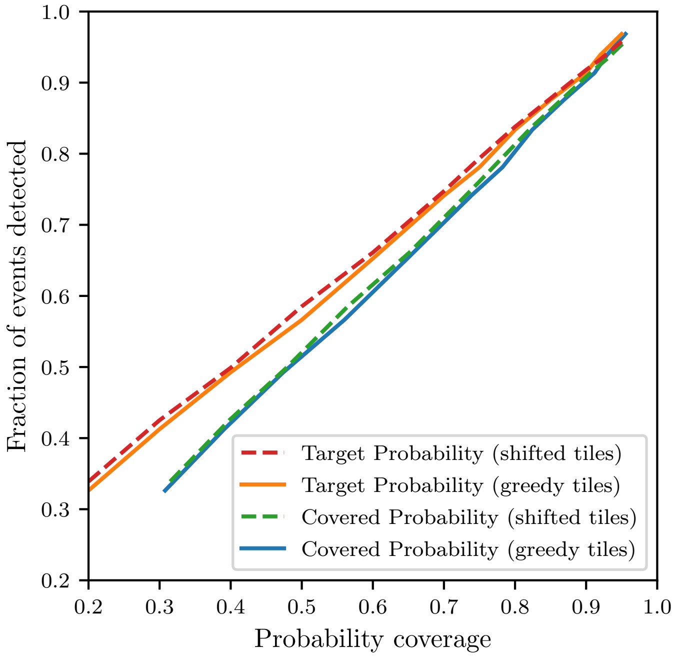

Figure 3.5: Fraction of events detected as a function of target probability coverage (orange/red) and actual probability coverage (blue/green) for both the greedy tiles (solid) and shifted tiles (dashed) methods. The actual probability coverage is calculated by taking the mean coverage for all simulations.

tively, corresponding to savings of 5 hours per epoch. The total savings increase with the time per pointing and the total number of epochs.

As a result, the shifted ranked tiles strategy should be used when the required observing latency is longer than the typical computation time (on the order of minutes). For low-latency observations the greedy ranked tiles strategy is preferable due to its shorter computation time.

3.4.3 Galaxy Targeting Strategy

We applied the greedy and shifted ranked tile strategies to all 250 simulated skymaps that have 3D localisations. Figure 3.6 shows the fraction of events covered by both strategies when applied to the galaxy targeted skymap, and the greedy strategy applied to the original skymap, as a function of probability coverage of the original skymap. There is no appreciable difference between the two strategies, although for mid-range probability coverages, original skymap marginally outperforms the galaxy targeted skymap, and for higher probability coverages (> 0.9) the galaxy targeted skymap performs marginally better. We propose that this negligible difference is the result of the incompleteness of the galaxy catalogue, as well as the wide ASKAP FoV that is more suitable for surveying large areas of the sky than targeting individual sources. For most events the optimal tiling strategy produced was nearly identical for the two different skymaps. We expect that the effectiveness of using a galaxy targeting strategy for most telescopes will significantly improve in the near future,

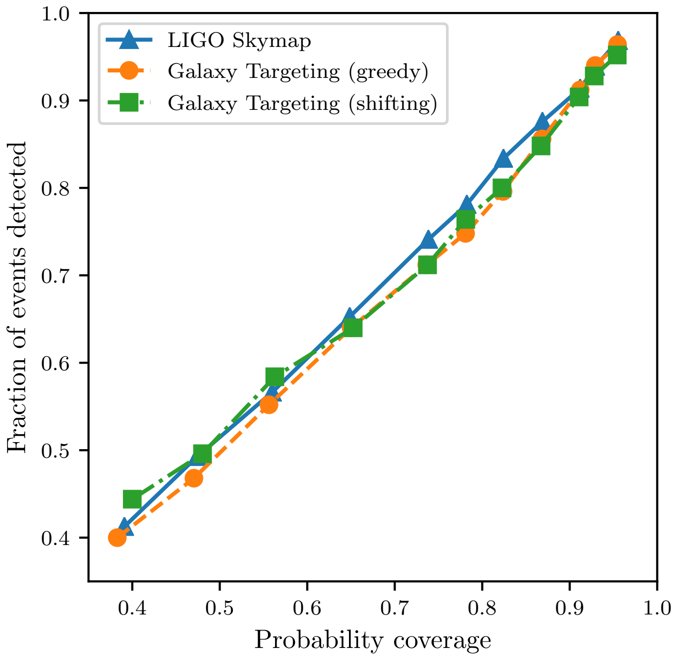

Figure 3.6: Comparison of sky coverage using the galaxy targeted probability map and the original gravitational wave skymap.

as wider and deeper galaxy surveys produce more complete catalogues. However, it is unclear whether this improvement will affect ASKAP.

We note that improvements to gravitational wave detectors will result in larger detector horizons. Advanced LIGO is expected to detect neutron star mergers at a distance of 190 Mpc (Abbott et al., 2018), while next generation detectors like Cosmic Explorer may detect neutron star mergers at $z = 6$ (Abbott et al., 2017a). Distant events will have larger localisation volumes along with less-complete catalogues of potential host galaxies, making a galaxy targeting approach less feasible.

3.5 Detectability of Events with ASKAP

3.5.1 Sensitivity and Angular Resolution

Table 3.1 shows the expected timeline and specifications as ASKAP moves towards design sensitivity. ASKAP should reach design specifications by the start of O3 in early 2019 and will be capable of observing entire localisation regions to sub-mJy levels in hours, or achieving a 12 μJy image RMS in a 12 hour pointing.

While GW170817 would have been detectable with ASKAP 40–300 days post-merger (Figure 3.1), we emphasise that this event may not be typical of the population of binary neutron star systems expected in the future. The observed luminosity of the afterglow increases with the isotropic equivalent energy, and the circum-merger density, n, the velocity of the outflow and the fraction of the internal energy deposited into the outflowing electrons and the accompanying magnetic field. These parameters also determine the temporal evolution of the afterglow, although the exact scaling factors are model-dependant and not necessarily trivial to compute (e.g. Sari et al., 1998; Nakar & Piran, 2011; Lazzati et al., 2017; Gottlieb et al., 2018a; Lamb et al., 2018). The expected radio emission is also anisotropic, and highly dependent on the merger inclination angle, with on-axis events typically more luminous. The horizon of gravitational wave detectors is also larger for on-axis events, known as the gravitational wave Malmquist effect (Schutz, 2011).

Typical values for most of the source parameters are unclear, although it is reasonable to assume that inclination angle is uniformly distributed. The circum-merger density of GW170817 is estimated as $n \sim 10^{-4}$ cm^{-3}, significantly lower than the median circum-burst density of short GRBs ($3 - 15 \times 10^{-3}$ cm^{-3}), with densities $n \sim 1$ cm^{-3} not uncommon (Fong et al., 2015). However, we expect approximately half of all neutron star mergers to have circum-merger densities $n > 0.1$ cm^{-3} (Hotokezaka et al., 2016). Even at the limit of the Advanced LIGO detection range, these events may still be radio loud, peaking at a flux density of tens of mJy (see Chapter 6), which is readily detected in a short ASKAP observation.

The 10 arcsec angular resolution of ASKAP at 1.4 GHz is large compared to the $1 - 2$ arcsec achievable with the Australia Telescope Compact Array (ATCA) or the Jansky Very Large Array (VLA), and we will not always be able to resolve the GW counterpart from its host galaxy. Short GRBs have physical offsets between 0.5–75 kpc from the host galaxy centre (median 5 kpc; Berger, 2014), which Hotokezaka et al. (2016) found corresponds to 70% of radio counterparts at 200 Mpc contaminated by the galaxy core at ASKAP's angular resolution. In comparison, GW170817 was offset from the centre of NGC 4993 by 10.6 arcsec (or 2 kpc; Coulter et al., 2017), which is at the limit of resolution with ASKAP at design specifications. Events close to the host nucleus will still be detectable with ASKAP if the radio counterpart has a peak flux density greater than that of the nucleus, which may be the the case for a majority of events (Hotokezaka et al., 2016).

Once an event has been detected with ASKAP, we expect the bulk of the monitoring campaign will be performed with the ATCA, the VLA and MeerKAT, due to their superior sensitivity and angular resolution. This will enable the lightcurve to be much more tightly constrained than is possible with ASKAP alone.

3.5.2 Detection of False-Positives

A major concern in the search for optical counterparts to gravitational wave events is the detection of false positives. Nissanke et al. (2013b) found that there could be tens to thousands of optical false positive detections per gravitational wave event. The rate of radio transients, on the other hand, is comparatively low, and the chance of a false-positive detection in the form of an unrelated radio transient is minimal. Bannister et al. (2011) found one transient at 843 MHz in a study spanning timescales of 1 day to 20 years, corresponding to an areal transient rate of 7.5×10^{-4} deg^{-2} for flux densities $>$14 mJy. Other works found the 1.4 GHz radio transient areal density to be $<$0.1 deg^{-2} ($>$1 mJy on timescales of 0.5 to 7 years; Hancock et al., 2016), $<$0.37 deg^{-2} ($>$0.2 mJy on timescales of 1 week to 3 months; Mooley et al., 2013) and $<$0.01 deg^{-2} ($>$1.5 mJy on daily timescales; Bhandari et al., 2018).

False positives may also be detected as a result of an insufficiently sensitive reference image, although ultimately a reference image will be available from the Evolutionary Map of the Universe (EMU, Norris et al., 2011), an ASKAP survey of the sky south of $+30$ deg to an image RMS of $\sim 10\,\mu$Jy at 1.3 GHz, along with our own observations soon after the event. Until EMU is completed we will rely on the NRAO VLA Sky Survey (NVSS, Condon et al., 1998) which covers the sky north of -40 deg with an image RMS of ~ 0.45 mJy at 1.4 GHz, the VLA Sky Survey (VLASS[3]) which will cover the sky north of -40 deg to an image RMS of 67 μJy at 2-4 GHz, or the Sydney University Molonglo Sky Survey (SUMSS, Bock et al., 1999; Mauch et al., 2003) which covers the sky south of -30 deg to an image RMS of $\sim$ 1–2.5 mJy at 843 MHz. The age of these surveys may also produce false positives, with long-term AGN variability caused by refractive interstellar scintillation being detected as transient behaviour on short timescales.

3.5.3 Expectations for LIGO-Virgo's Third Observing Run

The O3 observing run is expected to begin in April 2019, with a median detector horizon of 120-170 Mpc (Advanced Virgo: 65-85 Mpc) for neutron star mergers (Abbott et al., 2018). Given the current best estimate of the neutron star merger rate, $1.54^{+3.2}_{-1.22}\,\mathrm{Mpc}^{-3}\mathrm{Myr}^{-1}$ (Abbott et al., 2017d), we expect there to be 1–50 significant neutron star merger triggers during O3. ASKAP is capable of observing the sky south of $+30$ deg, corresponding to 75% of the whole sky. We expect to be able to observe a similar fraction of gravitational-wave triggers, although the real fraction may be lower given the geographically dependent sensitivity of gravitational wave detectors, any possible anisotropy in the distribution of neutron star mergers, and other detector constraints.

The localisation capability of Advanced LIGO for O3 and beyond is discussed in Section 3.3.1, with the median 90% probability containment approximately 150 deg^2 for O3, which is observable with a handful of ASKAP pointings.

The distance to GW170817 (40 Mpc) is significantly less than the expected O3 detector horizon, and was also well localised (90% containment within 16 deg^2) which allowed electromagnetic emission from the event to be detected by telescopes targeting candidate host galaxies. For events with a larger localisation volume, the number of candidate host galaxies may be large enough that they cannot all be targeted in a reasonable time. ASKAP will follow-up all observable neutron star merger events, with a focus on those poorly localised events with no detected electromagnetic counterpart.

---

[3] `https://science.nrao.edu/science/surveys/vlass`

3.6 Conclusions

The first detection of a neutron star merger, GW170817, was initially localised to 16 deg^2 through its gravitational wave signal. This enabled an optical counterpart to be identified, and the event localised to a galaxy, within hours. However, future events are likely to have much larger localisation volumes (hundreds of deg^2) and detecting the electromagnetic counterparts to these events will require wide-field telescopes. Once ASKAP reaches design specifications we expect it to play an important role in the follow-up of poorly localised events (particularly where no other electromagnetic counterpart has been detected) due to its wide FoV and high survey speed. We will be able to use ASKAP to localise events that may not be observable at other wavelengths due to factors including solar angle and dust obscuration.

We have discussed prospects for ASKAP detecting two forms of radio emission from neutron star mergers (prompt coherent emission that may be similar to an FRB, and long-term emission similar to a standard sGRB afterglow or GW170817). We have compared four different methods of tiling localisation regions with telescope pointings, building on previous work using widefield optical telescopes. We find that our implementation of the shifted ranked tiles method (see Section 3.3.2) outperforms previously investigated tiling methods for long-term ASKAP follow-up, while the greedy ranked tiles method (see Section 3.3.2 & Ghosh et al., 2017) is preferable for prompt follow-up due to its faster computation time. We also find that there is no significant advantage to using a galaxy-targeting approach.

Applying these pointing strategies to follow-up of gravitational wave events will optimise the use of telescope time and maximise the chance of a detection, localising the event and allowing us to better understand its properties.

Chapter 4

ASKAP Follow-up of GW190814

Dougal Dobie, Adam Stewart, Tara Murphy, Emil Lenc,
Ziteng Wang, David L. Kaplan, Igor Andreoni, Julie Banfield,
Ian Brown, Alessandra Corsi, Kishalay De, Daniel A. Goldstein,
Gregg Hallinan, Aidan Hotan, Kenta Hotokezaka,
Amruta D. Jaodand, Viraj Karambelkar, Mansi M. Kasliwal,
David McConnell, Kunal Mooley, Vanessa A. Moss,
Jeffrey A. Newman, Daniel A. Perley, Abhishek Prakash,
Joshua Pritchard, Elaine M. Sadler, Yashvi Sharma,
Charlotte Ward, Matthew Whiting, Rongpu Zhou

This chapter is reproduced from Dobie et al. (2019b), which was published in the Astrophysical Journal as "An ASKAP search for a radio counterpart to the first high-significance neutron star-black hole merger LIGO/Virgo S190814bv". After being confirmed as a real detection by LIGO (Abbott et al., 2020b) the event was renamed GW190814. The text has been updated to reflect the name change and improved parameter estimates.
I wrote the majority of the text, carried out all ATCA observations and performed the transients analysis jointly with Adam Stewart. Igor Andreoni, Kishalay De, Daniel Goldstein and Daniel Perley carried out analysis of the optical observations and wrote Section 4.4.2. Amruta Jaodand analysed the X-ray observations and wrote Section 4.4.3 with contributions from myself and David Kaplan. Kenta Hotokezaka provided afterglow modelling, performed the parameter analysis and wrote Section 4.5.3 with me

We present results from a search for a radio transient associated with the LIGO/Virgo source GW190814, a likely neutron star-black hole (NSBH) merger, with the Australian Square Kilometre Array Pathfinder. We imaged a 30 deg^2 field at ΔT=2, 9 and 33 days post-merger at a frequency of 944 MHz, comparing them to reference images from the Rapid ASKAP Continuum Survey observed 110 days prior to the event. Each epoch of our observations covers 89% of the LIGO/Virgo localisation region. We conducted an untargeted search for radio transients in this field, resulting in 21 candidates. For one of these, AT2019osy, we performed multi-wavelength follow-up and ultimately ruled out the association with GW190814. All other candidates are likely unrelated variables, but we cannot conclusively rule them out. We discuss our results in the context of model predictions for radio emission from neutron star-black hole mergers and place constrains on the circum-merger density and inclination angle of the merger. This survey is simultaneously the first large-scale radio follow-up of an NSBH merger, and the most sensitive widefield radio transients search to-date.

4.1 Introduction

On 14 August 2019 the LIGO and Virgo collaborations detected the compact binary merger GW190814 with the LIGO Livingston (L1), LIGO Hanford (H1) and Virgo (V1) gravitational wave detectors (Abbott et al., 2020b). The primary component is a black hole with mass $23.2^{1.1}_{1.0}\,M_{\odot}$, but the nature of the secondary component is unclear as its mass ($2.59^{+0.08}_{-0.09}\,M_{\odot}$) lies in the gap between the most massive neutron stars and the lightest black holes (Farr et al., 2011; Özel et al., 2012; Zhang et al., 2019; Cromartie et al., 2019). Therefore the expected nature of any electromagnetic radiation from the merger (if any) is unclear.

The preferred skymap at the time of observation, `LALInference.v1.fits.gz`, has a 90% localisation region of 23 deg^2 and a sky-averaged distance estimate of 267 ± 52 Mpc. High-energy observations (Molkov et al., 2019; Kocevski et al., 2019; Pilia et al., 2019; Sugizaki et al., 2019; Palmer et al., 2019) found no evidence for a coincident short gamma-ray burst (GRB). Optical observations found numerous candidate counterparts that have since been ruled out with further photometric and spectroscopic observations (Andreoni et al., 2020).

While the low probability of remnant matter (LIGO Scientific Collaboration and Virgo Collaboration et al., 2019a) may suggest that the merger produced no electromagnetic counterpart, the lack of optical counterparts may also be explained by intrinsic factors such as inclination angle, mass ratio, remnant lifetime or a lack of polar ejecta (Kasen et al., 2017), or extrinsic factors like dust-obscuration. In this case, radio emission may be the only way to localise this event.

We performed follow-up of GW190814 with the Australian Square Kilometre Array Pathfinder (ASKAP; Johnston et al., 2008). In Section 4.3 we discuss our untargeted radio transients search. In Section 4.4 we summarise multi-wavelength follow-up of candidate counterpart AT2019osy that was initially detected in this search.

4.2 Observations & Data Reduction

We observed a target field centred on (J2000) coordinates $\alpha = 00^{\mathrm{h}}50^{\mathrm{m}}37\overset{\mathrm{s}}{.}5$, $\delta = -25^{\circ}16'57\overset{\mathrm{s}}{.}37$ at $\Delta T = 2$, 9 and 33 days post-merger with ASKAP. This target field, shown in Figure 4.1 at $\Delta T = 2$ days, covers 89% of the skymap probability.

Epoch	SBID	Start (UTC)	Int. time (h:m:s)	ΔT (d)	% Flagged	Sensitivity (μJy)	Beam Size
0	8582	2019-04-27 04:59:14	00:15:00	−110	26	270	$10.2'' \times 14.9''$
1	9602	2019-08-16 14:10:27	10:39:25	2	25	35	$10.0'' \times 12.3''$
2	9649	2019-08-23 13:42:59	10:39:01	9	26	39	$11.8'' \times 12.4''$
3	9910	2019-09-16 12:08:34	10:38:42	33	32	39	$9.8'' \times 12.1''$

Table 4.1: Details of our ASKAP observations for each scheduling block ID (SBID). All observations were carried out with 288 MHz of bandwidth centered on a frequency of 944 MHz and 33 of 36 antennas. Typically 26% of the data was flagged due to RFI or correlator drop-outs. The ASKAP images from our follow-up observations are available from the CSIRO ASKAP Science Data Archive[1] under project code AS111.

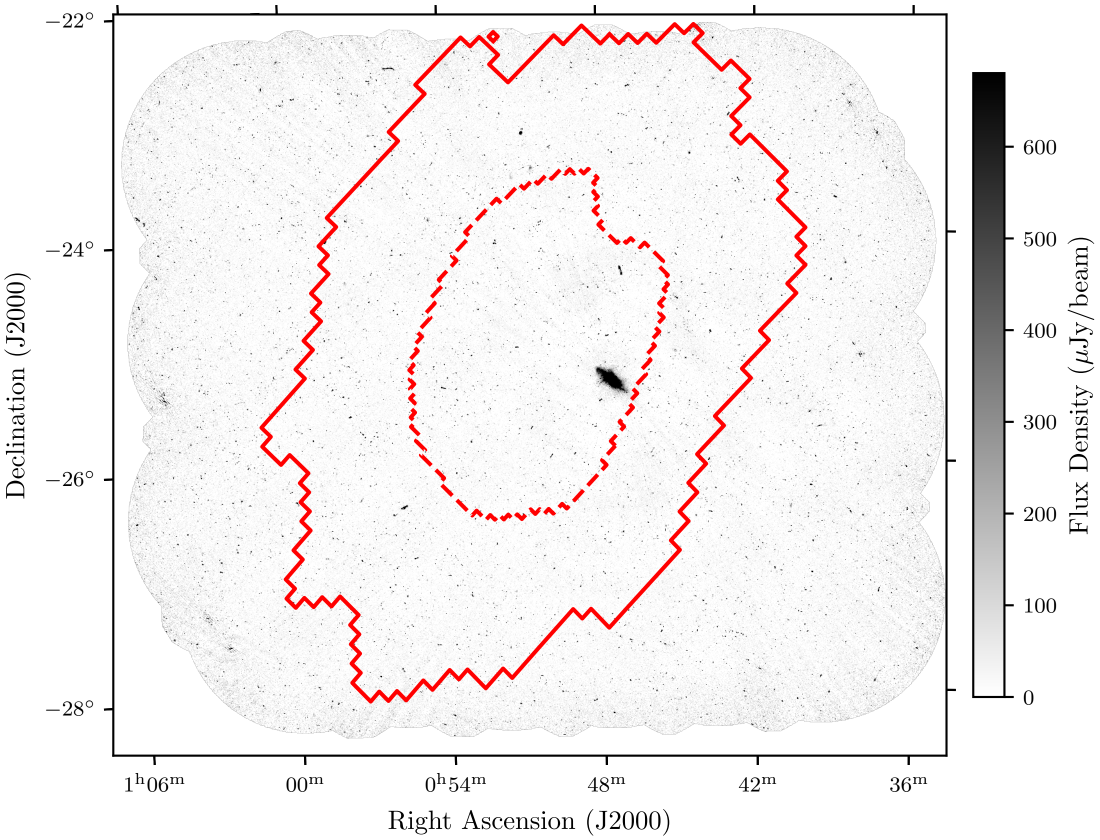

Figure 4.1: ASKAP image of the localisation region of GW190814 centered on 00:50:37.5, −25:16:57.371 observed 2 days post-merger. The 30deg^2 field of view covers ∼ 89% of the localisation region, with 50% (90%) contours shown in red dashed (solid) lines. The large object near the centre of the image is the radio-emitting starburst galaxy NGC 253. Note: there is a secondary lobe of the localisation towards the south-east that is outside the ASKAP footprint.

Table 4.1 gives a summary of our ASKAP observations. Data were observed using 36 beams arranged in a closepack36 footprint with beam spacing of 0.9 degrees. The field was tracked for a nominal time of 10.5 hrs and 288 MHz of bandwidth was recorded with a center frequency of 944 MHz. Typical sensitivity was $\sim 39\,\mu$Jy with a beam size of $\sim 12''$.

We imaged the data with the ASKAPsoft pipeline version 0.24.4 (Whiting et al., 2017), using a set of parameters optimised for deep continuum fields. Each beam was imaged independently and then combined using a linear mosaic. Multi-frequency synthesis with two Taylor terms was used, along with Multi-scale CLEAN using scales up to 27 pixels in size. Visibilities were weighted using Wiener preconditioning with a robustness parameter of zero. Two major cycles of self–calibration were used to refine the antenna gain solutions derived from observations of PKS B1934–638 in each beam (see McConnell et al., 2016, for a description of the ASKAP beamforming and calibration process). We also used pre-release data from the 888 MHz Rapid ASKAP Continuum Survey (RACS) as a reference epoch.

The astrometric accuracy and flux scaling of each epoch is consistent with every other epoch. The median flux ratio of compact sources for any two of the ASKAP observations is consistent with 1 within uncertainties. The median RA offset is 0.09–0.36$''$and the median declination offset is 0.02–0.2$''$(smaller than the pixel size) with a typical standard deviation of 0.7$''$and 0.6$''$respectively.

4.3 Untargeted Search for Radio Transients

To search for a radio counterpart to GW190814, we performed an untargeted search for transients and highly variable sources using the LOFAR Transients Pipeline (TraP; Swinbank et al., 2015). We ran TraP with source detection and analysis thresholds of 5σ and 3σ respectively, using the force beam option to make the Gaussian shape fit parameters for all sources the same as the restoring beam.

We selected candidates by identifying sources that were significant outliers in both variability metrics calculated by TraP: η, which is the weighted reduced χ^2, and the variability index V (equivalent to the fractional variability). This was done by fitting a Gaussian function to the distributions of both metrics in logarithmic space, with σ thresholds chosen to be $\eta > 1.5\sigma_\eta$ and $V > 1.0\sigma_V$, equating to values of $\eta > 2.73$ and $V > 0.18$. The thresholds were adapted from Rowlinson et al. (2019a), which gives approximate recall and precision rates of 90% and 50% respectively.

This resulted in 285 transient or variable candidates, which was reduced to 89 sources after manual inspection to remove imaging artefacts and components of complex extended sources.

Table 4.2: Candidate counterparts from an untargeted search of the S190814bv localisation region. Non-detections are denoted by 3σ upper-limits based on the local noise measured by `BANE` (Hancock et al., 2018). The angular separation and redshift of the corresponding optical source are shown.

Name	RA (deg)	Dec (deg)	S_0 (mJy)	S_1 (mJy)	S_2 (mJy)	S_3 (mJy)	V_{int}	η_{int}	offset ($''$)	z
ASKAP J004033.2−233530	10.13813	−23.5917	4.700 ± 0.454	4.517 ± 0.062	4.732 ± 0.069	6.648 ± 0.068	0.22	306	···	···
ASKAP J004054.8−273246	10.22816	−27.5463	< 1.1	0.498 ± 0.069	0.525 ± 0.076	0.272 ± 0.078	0.32	3.29	13.4	0.19 ± 0.05
ASKAP J004150.3−270632	10.45977	−27.1090	< 1.0	0.656 ± 0.058	0.536 ± 0.063	0.436 ± 0.064	0.20	3.32	···	···
ASKAP J004424.5−265522	11.10216	−26.9230	< 1.2	0.281 ± 0.055	0.437 ± 0.060	0.475 ± 0.060	0.26	3.26	···	···
ASKAP J004825.7−264137	12.10704	−26.6937	< 0.75	0.384 ± 0.053	0.615 ± 0.057	0.614 ± 0.057	0.25	5.94	···	···
ASKAP J004916.8−270745	12.32005	−27.1292	< 0.88	0.586 ± 0.049	0.725 ± 0.053	0.954 ± 0.055	0.25	12.6	16.8	0.38 ± 0.13[b]
ASKAP J005234.9−264144	13.14558	−26.6956	< 0.73	0.379 ± 0.050	0.380 ± 0.055	0.226 ± 0.054	0.27	2.75	···	···
ASKAP J005304.8−255451	13.27001	−25.9144	< 1.1	0.230 ± 0.050	0.375 ± 0.054	0.214 ± 0.053	0.33	2.75	···	···
ASKAP J005426.1−253833	13.60866	−25.6425	< 0.72	0.274 ± 0.053	0.487 ± 0.059	0.273 ± 0.059	0.36	4.51	17.9	0.33 ± 0.11
ASKAP J005434.6−280235[a]	13.64412	−28.0431	< 0.70	3.399 ± 0.097	1.337 ± 0.103	1.264 ± 0.104	0.61	149	11.5	0.21 ± 0.11
ASKAP J005523.7−250403	13.84868	−25.0675	< 0.86	0.972 ± 0.053	0.753 ± 0.060	0.669 ± 0.060	0.20	7.85	···	···
ASKAP J005547.4−270433	13.94764	−27.0759	< 0.80	0.399 ± 0.055	0.598 ± 0.059	0.557 ± 0.059	0.20	3.45	0.1	0.0733[c]
ASKAP J005606.9−255300	14.02875	−25.8835	< 0.80	0.623 ± 0.052	0.899 ± 0.059	1.011 ± 0.059	0.24	13.3	9.2	0.26 ± 0.14
ASKAP J005618.1−273012	14.07556	−27.5035	2.006 ± 0.559	1.770 ± 0.066	2.613 ± 0.070	2.050 ± 0.069	0.20	39.4	11.1	0.18 ± 0.09
ASKAP J005709.0−243659	14.28753	−24.6165	< 0.78	0.890 ± 0.054	0.611 ± 0.060	0.489 ± 0.059	0.31	13.5	14.2	0.22 ± 0.10
ASKAP J005709.7−250751	14.29030	−25.1310	< 0.81	0.654 ± 0.054	0.814 ± 0.062	0.447 ± 0.062	0.29	8.85	···	···
ASKAP J005729.6−231608	14.37350	−23.2690	< 0.98	0.620 ± 0.060	0.803 ± 0.065	0.495 ± 0.064	0.24	5.76	···	···
ASKAP J005809.0−273407	14.53757	−27.5688	< 0.79	0.849 ± 0.068	0.602 ± 0.072	0.552 ± 0.073	0.24	5.25	···	···
ASKAP J010004.6−231155	15.01934	−23.1988	< 0.79	1.002 ± 0.067	0.767 ± 0.073	0.642 ± 0.070	0.23	7.15	···	···
ASKAP J010258.6−265119[a]	15.74436	−26.8555	< 0.87	< 0.099	0.261 ± 0.091	0.232 ± 0.098	0.45	3.75	···	···
ASKAP J010534.6−231604	16.39415	−23.2680	< 0.85	< 0.087	0.485 ± 0.140	0.718 ± 0.146	0.58	3.36	···	···

[a] Ruled out as a counterpart

[b] There are 3 optical sources within 20$''$of this candidate. The two closest have a photometric redshift that is inconsistent with the distance to S190814bv.

[c] Spectroscopic redshift.

4.3.1 Analysis of Candidates for Possible Association with GW190814

The 89 variable sources were filtered to remove those that were not consistent with the predicted emission of GW190814[2], which should not exhibit more than a single rise and decline on these timescales (Hotokezaka et al., 2016), according to the following criteria:

1. Sources that showed a decline between epochs 1 and 2, followed by a rise between epochs 2 and 3. 41 sources were excluded.
2. Sources detected in RACS epoch 0 where epochs 1 and 2 had lower integrated flux values than epoch 0. 3 sources were excluded.

We then searched the GLADE catalogue (GLADE; Dálya et al., 2018) for galaxies in the localisation volume within $20''$(or ~ 20 kpc at the estimated distance of GW190814 LIGO Scientific Collaboration and Virgo Collaboration et al., 2019a) of a variable source. We found one candidate (ASKAP J005547.4−270433) that is near 2dFGRS TGS211Z177, a catalogued galaxy with $z = 0.0738$ (Colless et al., 2001). This source was the only strong candidate after epoch 2 and prior to the acquisition of epoch 3 we performed multi-wavelength follow-up which we discuss in Section 4.4. We excluded two candidates that matched with a GLADE galaxy $> 3\sigma$ beyond the estimated distance to GW190814 (267 ± 52 Mpc LIGO Scientific Collaboration and Virgo Collaboration et al., 2019a).

We crossmatched the 42 remaining candidates with the Photometric Redshifts for the Legacy Surveys (PRLS) catalogue (Dey et al., 2019; Zhou et al., 2020). We excluded 22 variable sources that had all optical matches at distances differing by $> 3\sigma$ from the estimated distance to GW190814. This left 7 sources with at least one crossmatch within the localisation volume and 13 sources with no reliable distance estimate (see Table 4.2).

4.4 Follow-up of ASKAP J005547.4−270433

4.4.1 Radio Observations

We carried out follow-up observations of ASKAP J005547.4–270433 (hereafter AT2019osy) with the ATCA (C3278, PI: Dobie) using two 2 GHz bands centered on 5.5 and 9 GHz at 14, 22 and 34 days post-merger. We reduced the data using the same method as Dobie et al. (2018a) using PKS B1934–638 and B0118–272 as flux and phase calibrators respectively.

We also carried out VLA observations (VLA 18B-320, PI: Frail) on 2019 Aug 28 and Sep 09. Standard 8 bit WIDAR correlator setups were used for L and S bands, and 3 bit setups for C and X bands to obtain a contiguous frequency coverage between $1 - 12$ GHz. 3C48 and J0118–2141 were used as the flux and phase calibrators

---

[2] Continued monitoring of this field to November 2020 has shown that all 89 sources are unrelated to GW190814.

Table 4.3: Radio observations of AT2019osy. Observations with the ATCA and VLA were carried out with maximum baselines of 6 km and 40 km respectively.

Telescope	ΔT (days)	Frequency (GHz)	Flux Density (μJy)
ASKAP	2	0.943	376 ± 33
ASKAP	9	0.943	550 ± 34
VLA	13	1.5	409 ± 34
		3.0	301 ± 21
		6.0	213 ± 11
		10.0	187 ± 11
ATCA	14	5.0	369 ± 23
		6.0	335 ± 19
		8.5	307 ± 15
		9.5	278 ± 14
ATCA	22	5.0	380 ± 21
		6.0	353 ± 17
		8.5	299 ± 14
		9.5	234 ± 14
VLA	25	1.5	303 ± 48
		3.0	317 ± 21
		6.0	220 ± 10
		10.0	150 ± 10
ASKAP	33	0.943	513 ± 34
ATCA	34	5.0	348 ± 17
		6.0	349 ± 14
		8.5	320 ± 15
		9.5	275 ± 14

respectively. The data were processed using the NRAO CASA pipeline and imaged using the `clean` task in CASA.

A summary of our observations is given in Table 4.3. We find a constant flux density offset[3] of $\sim 40\%$ between the initial ATCA and VLA observations and no evidence for radio variability beyond the initial rise observed with ASKAP.

4.4.2 Optical Observations

We conducted optical imaging of AT2019osy with the Dark Energy Camera (DECam, Flaugher et al., 2015) on the 4m Blanco telescope under NOAO program ID 2019B-0372 (PI: Soares-Santos). Images including the location of AT2019osy were taken in i and z bands nightly from 2019-08-15 to 2019-08-18 and on 2019-08-21 (UT) and reduced in real-time (Goldstein et al., 2019). A detailed offline analysis of the subtraction images zooming in on the location around AT2019osy, reveals no robust point source at this location to a depth of $i > 21.2$mag and $z > 20.0$mag on UT 2019-08-15 (the night of the merger) increasing linearly in limiting magnitude to $i > 23.5$mag and $z > 23.5$mag on UT 2019-08-21 (consistent with independent analysis by Herner et al. 2019). We also analyzed the DECam images using `The Tractor` image modeling software (Lang et al., 2016) and found that a model with an exponential galaxy profile with a point source at the galaxy nucleus is required to fit the data, both before and after GW190814. This suggests that there is no optical transient temporally coincident with GW190814 but possibly some underlying nuclear variability.

On 2019-08-22 UT, we observed AT2019osy in the near infrared using the Wide-field Infrared Camera (WIRC, Wilson et al., 2003) with the 200-inch Hale telescope at Palomar Observatory for a total of 10 minutes exposure time (De et al., 2019). The WIRC data were reduced and stacked using a custom pipeline (De et al., 2020). No counterpart to AT2019osy was detected down to an AB limiting magnitude of $J > 21.5$ (5σ).

We also obtained a spectrum of the host galaxy of AT2019osy using the Double Beam Spectrograph (Oke & Gunn, 1982) on the Palomar 200-inch Hale Telescope (P200), which we reduced using `pyraf-dbsp` (Bellm & Sesar, 2016). The spectrum is dominated by red continuum that is likely primarily associated with the host galaxy; no obvious broad features are evident. We identify several narrow emission lines (Hα; [NII]$\lambda\lambda$6548,6583, [SII]$\lambda\lambda$6716,6731, and marginal [OII]λ3727) at a common redshift of 0.0733, consistent within 2-sigma of the LVC distance constraint. Hβ and [OIII]λ5007 are not detected in the spectrum. We measure a flux ratio of log[NIIλ6583/Hα]=0.2, indicating at least partial contribution by an AGN (Kauffmann et al., 2003).

---

[3] The flux densities of nearby sources and the calibrator source J0118-2141 between the ATCA and the VLA are consistent with the flux offset of $\sim 40\%$ seen in AT2019osy. This offset can partially be explained by resolution effects, and detailed investigation of it is ongoing.

4.4.3 X-ray Observations

We observed the field of AT2019osy, starting at 2019-09-23 10:30:48 UT for 20 ks with the `Chandra` ACIS-S instrument (S3 chip) and very faint data mode. The data were analyzed with `CIAO` (v 4.11; Fruscione et al., 2006) and calibration was carried out with `CALDBv4.8.4.1`. We reprocessed the primary and secondary data using the `repro` script, created X-ray images for the 0.3–8 keV range. No sources were visible near AT2019osy (verified with both `wavdetect` and `celldetect`), with a maximum count rate of $2.85\times10^{-4}\,\mathrm{s}^{-1}$. Assuming a neutral hydrogen column density $N_{\mathrm{H}} = 1.8 \times 10^{20}\,\mathrm{cm}^{-2}$ and a power-law model with index $n = 1.66$ (corresponding to the observed radio spectral index of -0.4), this count rate yields a 0.3–8 keV unabsorbed flux upper limit of $3.2\times10^{-15}\,\mathrm{erg\,cm^{-2}\,s^{-1}}$ (as reported in Jaodand et al., 2019) or an unabsorbed luminosity of $4.2\times10^{40}\,\mathrm{erg\,s^{-1}}$.

4.4.4 Source Classification

AT2019osy exhibits no significant radio variability beyond the initial rise and there is no evidence for a coincident optical transient. The coincident galaxy is edge-on, likely with significant dust obscuration towards the nucleus, and therefore the optical spectrum is consistent with an AGN within a star-forming galaxy. The inferred radio and X-ray luminosity of AT2019osy along with the small offset from the optical centroid of 2dFGRS TGS211Z177 suggests that the source is a variable low-luminosity AGN (Ballo et al., 2012) and unrelated to GW190814.

4.5 Discussion

4.5.1 Candidate Classification

We find 21 candidate counterparts to GW190814 above our detection threshold of $170\,\mu$Jy, corresponding to $\lesssim 1\%$ of observed sources. This is consistent with the expected rate of AGN variability from Radcliffe et al. (2019), who find $\sim 2\%$ of μJy-level sources exhibit significant variability likely attributable to the presence of an AGN. Additionally, the expected level of compact source variability caused by refractive interstellar scintillation along this line of sight is $\sim 35\%$ (Cordes & Lazio, 2002), comparable to V_{int} for all but three sources which we discuss below.

We classify ASKAP J005434.6–280235 as a variable AGN based on follow-up observations (De et al., 2019; Dobie et al., 2019d). ASKAP J010258.6–265119 is coincident centrally between two large radio lobes and hence likely associated with core emission from a radio galaxy. ASKAP J010534.6–231604 is coincident ($< 1''$) with WISE J010534.64–231605.5 (Cutri & et al., 2012), which is likely a variable AGN at a distance of $z \sim 1$ (Glowacki et al., 2017).

While we cannot conclusively rule the sources in Table 4.2 out as counterparts to GW190814, they are likely AGN exhibiting a combination of intrinsic and extrinsic variability. Of course, at most one candidate can be the actual counterpart, and there is nothing yet to distinguish any of these from the others[4].

---

[4] Continued monitoring of this field to November 2020 has shown that all are sources are unrelated to GW190814.

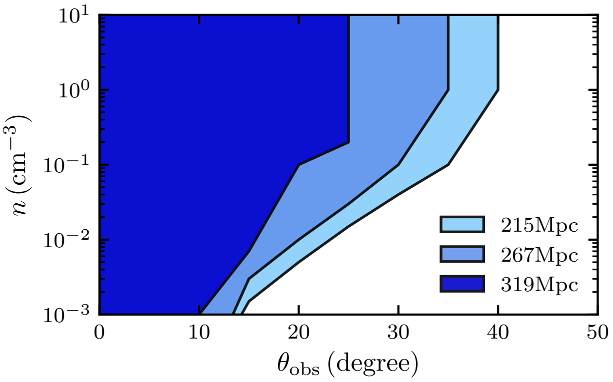

Figure 4.2: Radio constraints on viewing angle and circum-merger density for a merger with isotropic equivalent energy 10^{51} erg, an initial jet opening angle of 10° and microphysics parameters $\epsilon_e = 0.1$, $\epsilon_B = 0.01$ and $p = 2.2$, assuming that the merger occured in the 89% of the localisation region we observed. Shaded regions correspond to parts of the parameter space that are ruled out by our radio constraints for a range of distances corresponding to 1σ either side of the median. The inclination angle of the merger estimated from the gravitational wave signal is $\theta_{obs} = 46^{+17}_{-11}$ deg (Abbott et al., 2020b).

4.5.2 Radio Transient Rates

Our follow-up of GW190814 is the most sensitive widefield radio transients search to-date, approximately an order of magnitude more sensitive compared to previous searches with comparable areal coverage (Hobbs et al., 2016) and approximately an order of magnitude more areal coverage than previous searches at comparable sensitivities (Mooley et al., 2013).

We have found 4 transient candidates (i.e. sources with a prior constraining non-detection) in total; the three sources discussed in Section 4.5.1 and ASKAP J005104.2–230852, which was ruled out as a candidate to GW190814 based on the redshift of nearby optical sources. This source is coincident ($< 0.6''$) with WISE J005104.13-230851.8, which is likely a variable AGN. We therefore place an upper-limit on the 943 MHz radio transients surface density of 0.05 deg^{-2} for sources above 170 μJy at 95% confidence.

4.5.3 Non-Detection of a Radio Afterglow

Predicted radio lightcurves from NSBH mergers span a large range of flux densities and timescales (e.g. Piran et al., 2013; Lamb & Kobayashi, 2016; Bhattacharya et al., 2019). If the radio emission is dominated by the outflowing dynamical ejecta the lightcurve will peak on timescales of years, whereas if the emission is jet-dominated the lightcurve will peak at comparably lower flux densities on timescales of days–months (Hotokezaka et al., 2016). In each scenario the lightcurve is also dependent on the merger energetics, circum-merger density and inclination angle, each of which can change both the peak time and flux density by an order of magnitude. The merger energetics are determined by the mass ratio, the spin of the black hole (Abbott et al., 2020b) and the unknown neutron star equation of state (Kyutoku et al., 2011; Foucart, 2012).

We place a 5σ upper limit on the 943 MHz radio emission from GW190814 of 170 μJy at $\Delta T = 2$, 9 and 33 days post-merger. We compute off-axis afterglow lightcurves based on a top-hat jet model (see Hotokezaka & Piran (2015) for details) assuming an isotropic equivalent energy $E_{\rm iso} = 10^{51}$ erg (typical of short GRB afterglows; Fong et al., 2015), an initial jet opening angle of $\theta_{\rm j} = 10^\circ$ and microphysics parameters $\epsilon_e = 0.1$, $\epsilon_B = 0.01$ and $p = 2.2$. By comparing these lightcurves to the observed upper limits we can constrain the merger inclination angle, $\theta_{\rm obs}$, and circum-merger density, n. Figure 4.2 shows these constraints, assuming that the merger occured within the 89% of the localisation region covered by our observations. We can rule out the part of the parameter space typically occupied by short GRBs, assuming that their inclination angle is smaller than the opening of the angle of the jet (Fong et al., 2015). Under a more conservative assumption of the isotropic equivalent energy ($E_{\rm iso} = 10^{50}$ erg) we can only rule out a small part of the parameter space around $\theta_{\rm obs} = 10^\circ$ and $n = 1\,{\rm cm}^{-3}$.

In comparison, if we scale the non-thermal lightcurve of GW170817 to 943 MHz based on a spectral index of $\alpha = -0.575$ (Mooley et al., 2018d; Hajela et al., 2019) and place it at a distance comparable to GW190814, we find a peak flux density of $\sim 5\,\mu$Jy, well below our detection threshold. The non-thermal emission from GW170817 did not peak until ~ 150 d post-merger (Dobie et al., 2018a), and further observations on timescales of months–years post-merger will enable us to place tighter constraints on the circum-merger density and inclination angle. This may be useful in improving the gravitational wave localisation (Corley et al., 2019).

4.6 Conclusions

We have performed widefield radio follow-up of the NS-BH merger GW190814 with the Australian Square Kilometre Array Pathfinder. We cover 89% of the sky localisation with a single 30 deg^2 pointing centered on the localisation maxima. We found 21 candidate counterparts and performed comprehensive multi-wavelength follow-up of one, AT2019osy. The number of candidates is consistent with the expected rate of AGN variability. Most exhibit variability that is consistent with that expected from interstellar scintillation and are therefore unlikely to be related to GW190814

The non-detection of a radio counterpart allows us to place constraints on the circum-merger density, n, and inclination angle of the merger, $\theta_{\rm obs}$, if it occurred within the area covered by our observations. Under the assumption of $E_{\rm iso} = 10^{51}$ erg, $\theta_{\rm j} = 10^\circ$, $\epsilon_e = 0.1$, $\epsilon_B = 0.01$ and $p = 2.2$, we constrain $\theta_{\rm obs} > 10^\circ$ for all n at the extreme of the probability distribution of distance to the event. We will be able to place tighter constraints on these merger parameters once inclination angle estimates from gravitational wave strain data are released publicly.

As well as probing different parameters to optical searches, radio observations of future events may detect a gravitational wave counterpart where optical follow-up is inhibited by observing constraints, or intrinsic properties of the merger. We have demonstrated that it is possible to perform comprehensive follow-up of gravitational wave events with ASKAP, due to its large field of view that enables a survey speed significantly faster than comparable radio facilities.

Chapter 5

Constraining Properties of Neutron Star Merger Outflows with Radio Observations

Dougal Dobie, David L. Kaplan,
Kenta Hotokezaka, Tara Murphy, Adam Deller,
Gregg Hallinan, Samaya Nissanke

This chapter is reproduced from Dobie et al. (2020a), which was published in Monthly Notices of the Royal Astronomical Society as "Constraining Properties of Neutron Star Merger Outflows with Radio Observations". I wrote the majority of the text, produced all figures, and wrote all code. Kenta Hotokezaka produced the afterglow simulations and wrote Section 5.2. Adam Deller made significant contributions to the VLBI detectability criteria.

The jet opening angle and inclination of GW170817 – the first detected binary neutron star merger – were vital to understand its energetics, relation to short gamma-ray bursts, and refinement of the standard siren-based determination of the Hubble constant, H_0. These basic quantities were determined through a combination of the radio lightcurve and Very Long Baseline Interferometry (VLBI) measurements of proper motion. In this paper we discuss and quantify the prospects for the use of radio VLBI observations and observations of scintillation-induced variability to measure the source size and proper motion of merger afterglows, and thereby infer properties of the merger including inclination angle, opening angle and energetics. We show that these techniques are complementary as they probe different parts of the circum-merger density/inclination angle parameter space and different periods of the temporal evolution of the afterglow. We also find that while VLBI observations will be limited to the very closest events it will be possible to detect scintillation for a large fraction of events beyond the range of current gravitational wave detectors. Scintillation will also be detectable with next generation telescopes such as the Square Kilometre Array, 2000 antenna Deep Synoptic Array and the next generation Very Large Array, for a large fraction of events detected with third generation gravitational wave detectors. Finally, we discuss prospects for the measurement of the H_0 with VLBI observations of neutron star mergers and compare this technique to other standard siren methods.

5.1 Introduction

The first detection of gravitational waves and electromagnetic radiation from a neutron star merger (GW170817; Abbott et al., 2017d,f,g) has given insight into high energy astrophysics, nuclear physics and cosmology. Observations of the radio lightcurve of GW170817 were able to place constraints on merger parameters including the isotropic equivalent energy of the merger, the density of the surrounding environment, and the jet opening angle (Hallinan et al., 2017; Kim et al., 2017; Mooley et al., 2018a; Dobie et al., 2018a; Alexander et al., 2018; Margutti et al., 2018b; Mooley et al., 2018d; Resmi et al., 2018; Troja et al., 2018, 2019a; Wu & MacFadyen, 2018; Hajela et al., 2019; Ziaeepour, 2019). However, observations of the radio lightcurve alone were unable to distinguish between two competing models for the geometry of the outflow (Nakar & Piran, 2018), although the steep decline of the lightcurve did slightly favour the presence of a jet (Troja et al., 2018; Lamb et al., 2018). This tension was not resolved until Very Long Baseline Interferometry (VLBI) observations of the afterglow detected superluminal motion, suggesting that the late-time radio emission in GW170817 was jet-dominated (Mooley et al., 2018b; Ghirlanda et al., 2019).

The observation of superluminal motion has also placed tighter constraints on the inclination angle of the merger. In turn, this helped break the distance-inclination degeneracy (Finn & Chernoff, 1993; Nissanke et al., 2010) in the gravitational wave observations, which contributed to most of the error budget in the initial standard siren measurement of the Hubble constant, H_0, using GW170817 (Abbott et al., 2017e). This allowed for a measurement of the H_0 with a precision of 7% (Hotokezaka et al., 2019) compared to 17% using the gravitational wave data alone (Abbott et al., 2017e). VLBI observations of ~ 15 similarly favourably oriented events with comparable signal-to-noise as GW170817 (combined with improvements in jet modelling and calibration of gravitational wave detectors) will allow for a measurement of H_0 with sufficient precision ($<2\%$) and accuracy (Mukherjee et al., 2019) to potentially resolve the discrepancy (Verde et al., 2019) between current estimates from cosmic microwave background power spectrum measurements (Planck Collaboration et al., 2020) and distance ladder observations (Riess et al., 2018, 2019; Reid et al., 2019). In comparison, it will take tens to hundreds of events to achieve a similar level of precision with gravitational wave observations alone (Schutz, 1986; Del Pozzo, 2012; Messenger & Read, 2012; Taylor et al., 2012b; Farr et al., 2019) and gravitational wave observations with independent redshift measurements (Dalal et al., 2006; Nissanke et al., 2010, 2013a; Chen et al., 2018; Vitale & Chen, 2018; Mortlock et al., 2019; Feeney et al., 2019; Soares-Santos et al., 2019). We discuss prospects for measurements of H_0 using gravitational waves in more detail in Section 5.5.3.

Table 5.1: Estimates of the observing angle, θ_{obs}, jet opening angle, θ_j, and circum-merger density, $n_0 = n/(1\,\mathrm{cm}^{-3})$, microphysics parameters, ϵ_e, ϵ_B and isotropic equivalent energy, $E_{\mathrm{iso},0} = E_{\mathrm{iso}}/(\mathrm{erg})$, of GW170817 using Gaussian jet (GJ), boosted fireball (BF), power-law jet (PLJ) and other structured jet (SJ) models. We also include the time post-merger of the latest observation covered by the fit. We have calculated the estimate of n from Hotokezaka et al. (2019) assuming an isotropic equivalent energy of $E_{\mathrm{iso}} = 10^{52}$ erg.

Reference	Model	θ_{obs} (deg)	θ_j (deg)	$\log_{10} n_0$	$\log_{10} \epsilon_e$	$\log_{10} \epsilon_B$	$\log_{10} E_{\mathrm{iso},0}$	Last obs. (days)
Ghirlanda et al. (2019)*	GJ	$15.0^{+1.5}_{-1.0}$	$3.4^{+1.0}_{-1.0}$	$-3.6^{+0.7}_{-0.7}$	-	$-3.9^{+1.7}_{-1.5}$	$52.4^{+0.6}_{-0.7}$	289
Hajela et al. (2019)†	BF	$30.4^{+4.0}_{-3.4}$	$5.9^{+1.0}_{-0.7}$	$-2.61^{+0.42}_{-0.63}$	$-0.75^{+0.43}_{-0.62}$	$-2.63^{+0.89}_{-1.2}$	$52.33^{+0.6}_{-0.55}$	743
Hotokezaka et al. (2019)*†‡	PLJ	$16.62^{+1.1}_{-0.57}$	$3.44^{+0.57}_{-0.57}$	$-4.03^{+0.17}_{-0.19}$	-	-	-	294
	GJ	$17.19^{+1.1}_{-0.57}$	$2.75^{+0.17}_{-0.17}$	$-4.06^{+0.19}_{-0.2}$	-	-	-	294
Lamb et al. (2019a)	SJ	$20.6^{+1.7}_{-1.7}$	$4.01^{+0.57}_{-0.57}$	$-3.3^{+0.6}_{-1.0}$	$-1.3^{+0.6}_{-0.7}$	$-2.4^{+1.4}_{-0.9}$	$52.0^{+0.6}_{-0.9}$	358
	GJ	$19.5^{+1.1}_{-1.1}$	$5.16^{+0.57}_{-0.57}$	$-4.1^{+0.5}_{-0.5}$	$-1.4^{+0.5}_{-0.6}$	$-2.1^{+0.8}_{-1.0}$	$52.4^{+0.4}_{-0.5}$	358
Lazzati et al. (2018)	SJ	$33.0^{+4.0}_{-2.5}$	~ 5	$-2.38^{+0.48}_{-0.21}$	$-1.222^{+0.067}_{-0.079}$	$-2.48^{+0.21}_{-0.4}$	-	198
Lin et al. (2019)	GJ	$25.2^{+8.0}_{-5.7}$	$4.6^{+1.7}_{-1.1}$	$-2.5^{+1.1}_{-1.1}$	$-1.28^{+0.81}_{-1.2}$	$-4.1^{+1.4}_{-1.2}$	$52.38^{+0.93}_{-0.9}$	360
Resmi et al. (2018)	GJ	$26.9^{+8.6}_{-4.6}$	$6.9^{+2.3}_{-1.7}$	$-2.68^{+0.88}_{-1.0}$	$-4.37^{+1.1}_{-0.48}$	$-0.66^{+0.13}_{-0.45}$	$51.76^{+0.52}_{-0.39}$	152
Ryan et al. (2020)	GJ	$22.9^{+6.3}_{-6.3}$	$4.0^{+1.1}_{-1.1}$	$-2.70^{+0.95}_{-1.0}$	$-1.4^{+0.7}_{-1.1}$	$-3.96^{+1.1}_{-0.74}$	-	391
	PLJ	$25.2^{+6.9}_{-7.4}$	$2.86^{+0.57}_{-0.57}$	$-2.6^{+1.1}_{-1.1}$	$-1.24^{+0.73}_{-1.2}$	$-3.76^{+1.1}_{-0.87}$	-	391
Troja et al. (2019a)	GJ	$29.^{+11}_{-12}$	$4.6^{+1.7}_{-2.3}$	$-2.37^{+0.84}_{-1.3}$	$-1.13^{+0.53}_{-0.88}$	$-4.18^{+0.85}_{-0.58}$	-	391
Wu & MacFadyen (2019)	BF	$30.3^{+7.0}_{-4.0}$	~ 5	$-2.0^{+0.7}_{-1.0}$	$-1.0^{+0.6}_{-0.9}$	$-3.6^{+1.3}_{-1.4}$	-	260

* Incorporates centroid motion measurements

† Does not incorporate optical data

‡ Does not incorporate X-ray data

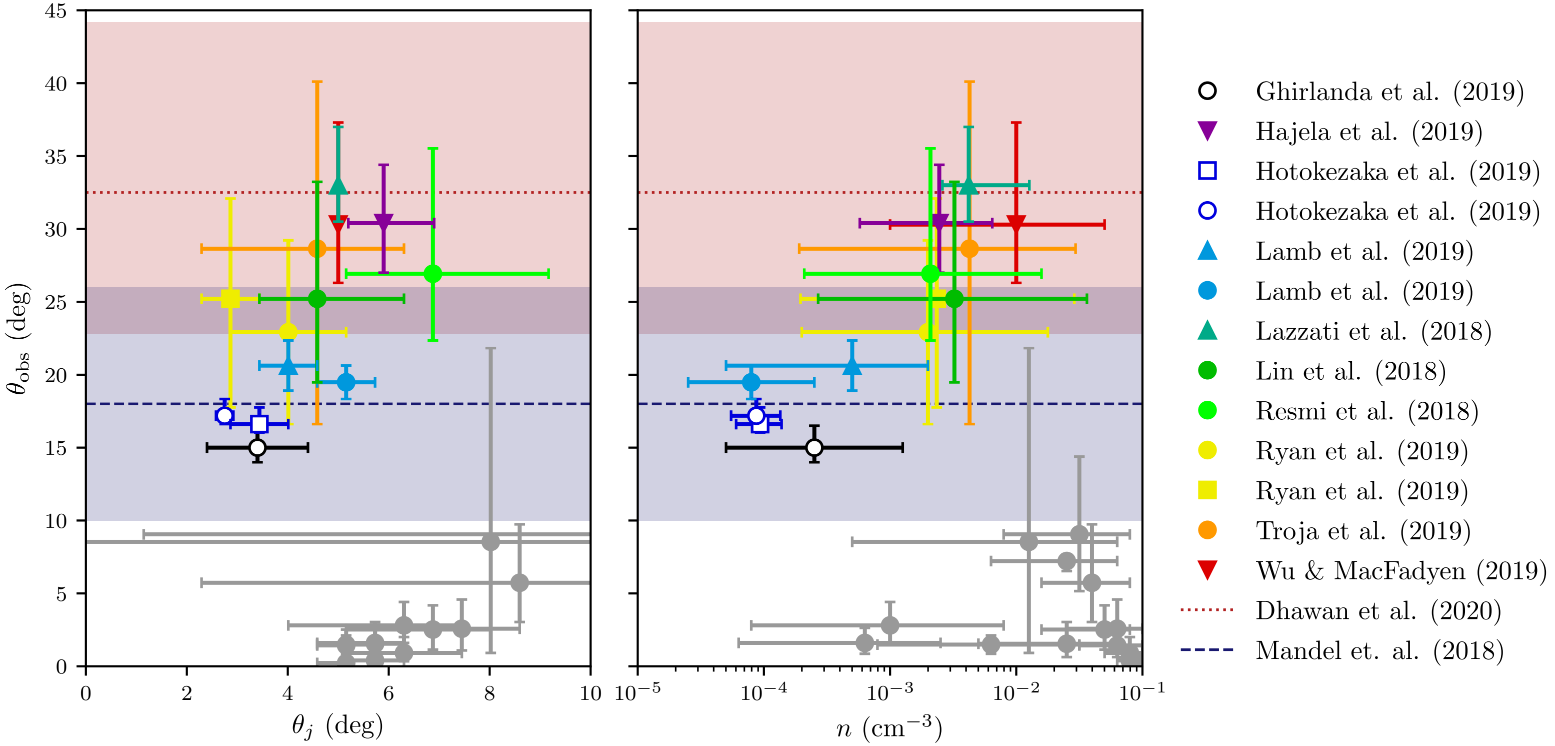

Figure 5.1: Estimates of the observing angle, $\theta_{\rm obs}$, of GW170817 as a function of jet opening angle, θ_j, and circum-merger density, n, for a range of Gaussian jet (circles), power-law jet (squares), boosted fireball (down triangles) and structured jet (up triangles) afterglow models. Open markers denote measurements that incorporate measurements of centroid motion, closed markers are lightcurve modelling, and horizontal lines correspond to estimates of $\theta_{\rm obs}$ (1σ uncertainty shaded) that are independent of the non-thermal afterglow. Grey markers represent parameter estimates for previously observed short gamma ray bursts (Wu & MacFadyen, 2019).

The broadband radio lightcurve of GW170817 was sampled at a cadence that was sufficient to constrain the spectral and temporal behaviour of the source. Depending on the neutron star merger detection rates in future observing runs with the Laser Interferometer Gravitational-Wave Observatory (LIGO) and Virgo detectors (Abbott et al., 2018), it may not be possible to perform radio monitoring at a similar high cadence for future events, which may result in an undersampled (and possibly unconstraining) lightcurve. Additionally, some events may not be localised until hundreds of days post-merger if detected via radio emission alone (Dobie et al., 2019a), in which case the early-time behaviour of the source will be unknown.

Even for events with a well-sampled lightcurve the information obtained about the properties of the jet and the surrounding environment is somewhat degenerate (Nakar & Piran, 2018), although it may be possible to infer the qualitative merger geometry (e.g. Troja et al., 2018; Lamb et al., 2018). However, techniques like VLBI and polarisation measurements will be important in understanding and tightly constraining merger energetics and outflow geometry quantitatively. The angular size, and therefore physical size, of the source may also be measured through observations of interstellar scintillation (Goodman, 1997) which has previously been used to constrain the size of gamma-ray burst (GRB) outflows (e.g. Frail et al., 1997; Chandra et al., 2008). Understanding the size of afterglows will be an important factor in understanding their physical behaviour (e.g. Granot et al., 2018; Lazzati et al., 2018; Nakar & Piran, 2018), and may place constraints on merger inclination in scenarios where VLBI cannot.

In this paper we discuss the detectability of expansion and motion of outflow from neutron star mergers through VLBI and observations of interstellar scintillation and implications for understanding the neutron star merger population and constraining the Hubble Constant.

5.2 The Geometry of Radio Afterglows

The observed size and motion of the afterglow of a relativistic blast wave have been studied in the literature (e.g. Sari 1998; Granot et al. 1999; Gill & Granot 2018; Mooley et al. 2018b; Xie et al. 2018; Hotokezaka et al. 2019). For instance, the observed size of a relativistic blast wave seen by an on-axis observer is $\sim \alpha\Gamma c T$, where Γ is the Lorentz factor of the blast wave, T is the observer time, and α is a numerical factor determined by the outflow's dynamics. For an on-axis blast wave decelerating in the ISM the size is analytically given by Sari (1998);

$$\theta_{\rm S} \approx 20\,\mu{\rm as}\left(\frac{n}{1\,{\rm cm}^{-3}}\right)^{-1/8}\left(\frac{E_{\rm iso}}{10^{52}\,{\rm erg}}\right)^{1/8}\left(\frac{T}{1\,{\rm day}}\right)^{5/8}\left(\frac{d}{100\,{\rm Mpc}}\right)^{-1}, \tag{5.1}$$

where n is the ISM density, $E_{\rm iso}$ is the isotropic-equivalent kinetic energy, and d is the distance to the source. This estimate is valid for $\Gamma \gtrsim 1/\theta_j$, $1/\theta_v$, where θ_j is the jet half-opening angle and θ_v is the viewing angle from the jet axis.

Unlike gamma-ray bursts which are preferentially seen on-axis because of Doppler beaming, afterglows of gravitational-wave mergers are most likely to be seen from a direction far from the axis of a collimated jet. By fitting both spatial information from VLBI observations and the afterglow lightcurve, Mooley et al. (2018b) find that the jet in GW170817 was observed with a viewing angle of $\theta_v \approx 15°$–$25°$ and estimate the jet half-opening angle as $\theta_j \lesssim 5°$. These estimates are comparable to values inferred by independent measurements using a variety of models of the non-thermal afterglow (see Table 5.1), although we note that these fits generally infer somewhat larger viewing angles. Additionally, Dhawan et al. (2020) find $\theta_v = 32.5^{+11.7\circ}_{-9.7}$ using thermal afterglow modelling, while Mandel (2018) estimates $\theta_v = 18 \pm 8°$ independent of the merger afterglow. Figure 5.1 shows the parameter space occupied by these estimates in comparison to previously observed short GRBs.

When observed from these angles the flux center of the radio afterglow exhibits motion perpendicular to the line of sight and the apparent velocity may exceed the speed of light. The afterglow light curve arising from a decelerating jet typically peaks when $\theta_v \approx 1/\Gamma$. The apparent velocity at the peak is estimated as

$$\beta_{\rm app,max} \approx \Gamma\beta \tag{5.2}$$

corresponding to

$$\approx 1.7 \cot\theta_{\rm v}\ \mu{\rm as/day} \left(\frac{d}{100\,{\rm Mpc}}\right)^{-1}, \tag{5.3}$$

where we assume that the opening angle of the emitting region is much less than the viewing angle.

The observed size and motion may be different from the above estimates when an outflow is seen from off-axis and the outflow has some structure , i.e., structured jets. A structured jet arises from the interaction of the jet with the merger ejetcta and the central engine activity (Gottlieb et al., 2018b; Xie et al., 2018; Lazzati et al., 2018). Structured jets are often modeled by using a simple analytic function, e.g., a power law or Gaussian function. In this work, motivated by the afterglow observations of GW170817, we calculate the synchrotron radio flux, centroid motion, and source size assuming a structured jet model described by a power-law function for the angular distribution of kinetic energy

$$E(\theta) = \frac{E_{\rm iso}}{1 + (\theta/\theta_{j,c})^{3.5}}, \tag{5.4}$$

where θ is the polar angle from the jet axis, $\theta_{j,c}$ and $E_{\rm iso}$ are the half opening angle and isotropic-equivalent energy of the core of a jet. The initial Lorentz factor of a jet is also assumed to have a power-law distribution

$$\Gamma(\theta) = 1 + \frac{\Gamma_c}{1 + (\theta/\theta_{j,c})^5}, \tag{5.5}$$

where Γ_c is the Lorentz factor of the jet's core. With these distributions, we solve the radial expansion of a jet in a uniform ISM density. In the following, we use $\theta_{j,c} = 0.05\,{\rm rad}$, $E_{\rm iso} = 10^{52}\,{\rm erg}$, and $\Gamma_c = 600$, for which the afterglow light curve and superluminal motion are consistent with the observed data of GW170817 (Hotokezaka et al., 2019). With these jet dynamics, we calculate the radio flux arising from the shock produced by a jet expanding in the ISM with the standard synchrotron afterglow model (Sari et al., 1998). We note that our results depend only weakly on the choice of the analytic function of the structure as long as the properties of the jet core is fixed. We set the fraction of the shock energy that goes into the electrons and the magnetic field to be $\epsilon_e = 0.1$ and $\epsilon_B = 0.01$ respectively, and set the electron energy distribution to a power law with index $p = 2.16$, comparable to GW170817 (Mooley et al., 2018d; Hajela et al., 2019).

We assume that radio observations occur at frequencies above the characteristic frequency of the slowest electrons, and below the synchrotron cooling break, such that the radio spectrum is reasonably modeled with a single (negatively sloped) power-law component. We set the spectral index to $\alpha = -0.585$ as observed in GW170817 (Alexander et al., 2018; Mooley et al., 2018d; Troja et al., 2019a).

In our afterglow models we use a 2D Cartesian coordinate system with origin at the position of the merger and define the y coordinate to be along the merger axis. To determine the size and position of the afterglow emission region we define the flux weighted quantity

$$\langle Q \rangle = \frac{\int Q \nu dS}{\int \nu dS} \tag{5.6}$$

where S is the flux at the position and ν is the frequency. The centroid of the emission region is then given by $\langle x \rangle$, and the size of the emission region in the x and y directions is $\sqrt{\langle x \rangle^2 - \langle x^2 \rangle}$ and $\sqrt{\langle y^2 \rangle}$ respectively.

To ensure that this work can be generalised to all future events, we also provide scaling relations for all relevant quantities in terms of the energetics and microphysics parameters. The flux density scales according to

$$S \propto E_{\rm iso} (\epsilon_B n)^{(p+1)/4} \epsilon_e^{p-1}, \tag{5.7}$$

while the size and centroid of the emission and the time since merger region scale as $(E_{\rm iso}/n)^{1/3}$. We consider values of $E_{\rm iso} = 4 \times 10^{49}$, 1.8×10^{51}, 4.5×10^{52} erg, corresponding to the range and median isotropic equivalent gamma ray energy of short GRBs (Fong et al., 2015). We consider $\epsilon_e = 0.01$, 0.3, 0.5 and $\epsilon_B = 10^{-4}$, 2×10^{-3}, 2×10^{-2} based on the estimates and 1-σ uncertainties of these parameters for GW170817 (Hajela et al., 2019).

5.3 Observations of Interstellar Scintillation

Interstellar scintillation is the observed extrinsic variability of radio sources induced by inhomogeneities in the electron density along their line of sight (Armstrong et al., 1995). The induced variability of a radio source can be quantified with the scattering strength, given by

$$\xi = 2.6 \times 10^3 \mathrm{SM}^{0.6} D^{0.5} \nu^{-1.7} \tag{5.8}$$

where SM is the scattering measure (which describes the cumulative contribution of the inhomogeneities along the line of sight), D is the distance to the equivalent phase screen in kpc, and ν is the observing frequency in GHz (Narayan, 1992; Taylor & Cordes, 1993; Walker, 1998).

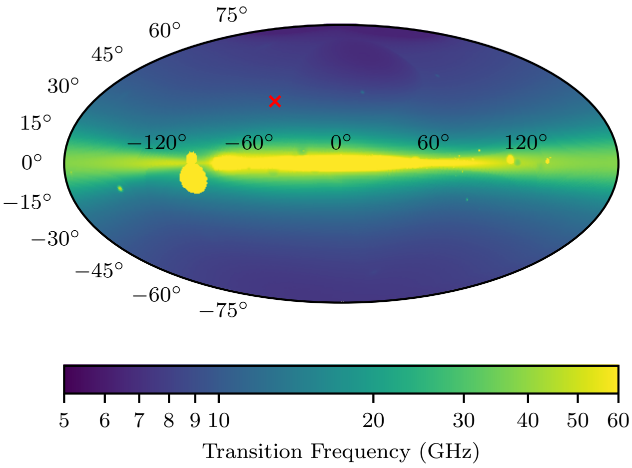

Figure 5.2: Transition frequency as a function of Galactic coordinates for extragalactic sources. The red cross denotes the location of GW170817.

Walker (1998, 2001) describes scaling relations for various scintillation parameters at an observing frequency, ν, as a function of the transition frequency, ν_0, the frequency at which the scattering strength is unity, and the size of the Fresnel zone at the transition frequency given by $\theta_{\mathrm{F}0} = 8/\sqrt{D\nu_0}$ and the observing frequency. The scattering strength itself depends on the distribution of electrons along the line of sight, which we assume is dominated by electrons within the Milky Way for simplicity. The variability timescale is dependent on the transverse speed of the phase screen relative to the observer and source across the line of sight, and the spatial distribution of electrons along the line of sight. We adopt the same approximation as Walker (1998) and assume that the variability timescale at the transition frequency is 2 hours in all directions based on a phase screen transverse velocity of $50\,\mathrm{kms}^{-1}$ (Rickett et al., 1995).

Scintillation occurs in two main regimes; weak scattering ($\nu > \nu_0$, where the phase changes in the observed radio signal introduced by the interstellar medium are negligible) and strong scattering ($\nu < \nu_0$, where the phase changes are the dominant cause of the observed variability). Additionally, there are two forms of strong scattering; refractive scintillation which is characterised by broadband variability on timescales of days, and diffractive scintillation which is characterised by narrow-band variability on much shorter timescales. Scintillation is typically observable in sources that are smaller than the size of the relevant scattering disk (θ_{F} for weak scattering, and some multiple of θ_{F} for strong scattering), although Narayan (1992) provides scaling relations for sources that are larger than the scattering disk.

We use NE2001 (Cordes & Lazio, 2002, 2003), a model for the Galactic electron distribution based on independent measurements of the dispersion measure (DM) and distance of 269 pulsars, to retrieve values of ν_0 and $\theta_{\mathrm{F}0}$ along various lines of sight. NE2001 has been widely used for scattering calculations in the past, but other electron density models including YMW16 (Yao et al., 2017), which uses similar techniques to NE2001, and RISS19 (Hancock et al., 2019a), which uses H_α measurements to trace free electrons, may also be suitable and will produce qualitatively similar results. We note that both NE2001 and YMW16 only consider the contribution of Galactic electrons, and therefore an extragalactic measure like RISS19 should provide more accurate estimates. However, RISS19 is less tested against observations than other models and therefore we choose to use NE2001 for easy comparison to the literature

The analysis in the following sections focuses on the synchrotron emission produced by the forward shock of the merger outflow, however at very early times ($t \lesssim 1$ day), when the outflow is extremely compact, the radio emission may be dominated by the reverse shock. See Lamb & Kobayashi (2019) for a discussion of the effects of scintillation on emission from the reverse shock.

5.3.1 Prospects for Detection of Scintillation

We perform a qualitative analysis of the detectability of all forms of scintillation at a range of frequencies. We divide the radio spectrum into four parts: low frequencies (∼300 MHz), mid frequencies (0.8–2 GHz), high frequencies (2–10 GHz) and millimetre wavelengths (>20 GHz) based on the capabilities of existing radio facilities. In general, scintillation is easier to detect in sources that have small angular sizes, i.e. more distant, off-axis, events. However, in the case of relativistic outflows both of these properties correspond to lower observed flux densities, and therefore make the emission and any variability more difficult to detect. Events occuring in dense environments are more luminous and remain compact for longer.

Weak scattering

At observing frequencies $\nu > \nu_0$ (or equivalently $\xi < 1$) we will observe weak scattering where the dominant cause of phase changes is path-length variations.

The size of the scattering disk (which determines whether an object can be treated as a point source) is simply the Fresnel zone, θ_{F}, which at the transition frequency is $< 5\,\mu$as for all positions on the sky ($\theta_{\mathrm{F0}} < 0.5\,\mu$as near the Galactic plane), and scales as $\nu^{-1/2}$. This is typically smaller than the angular size of the source, θ_{s}, by the time it becomes luminous enough to be detectable (e.g. see Figure 5.3). The modulation index of weak scattering is strongly dependent on observing frequency, and is given by

$$m_{\mathrm{W}} = \left(\frac{\nu_0}{\nu}\right)^{17/12} \min\left[1, \frac{\theta_{\mathrm{F}}}{\theta_{\mathrm{s}}}\right]^{7/6} \tag{5.9}$$

Hence, as we move above the transition frequency and further into the weak scattering regime both the modulation index and size of the Fresnel zone decrease.

Additionally, early-time observations at frequencies above the transition frequency are less feasible compared to those at lower frequencies due to the negative spectral index of the afterglow. Combining these three factors we find that detecting variability due to weak scattering from radio afterglows will be quite difficult with current and planned radio telescopes and do not consider it any further.

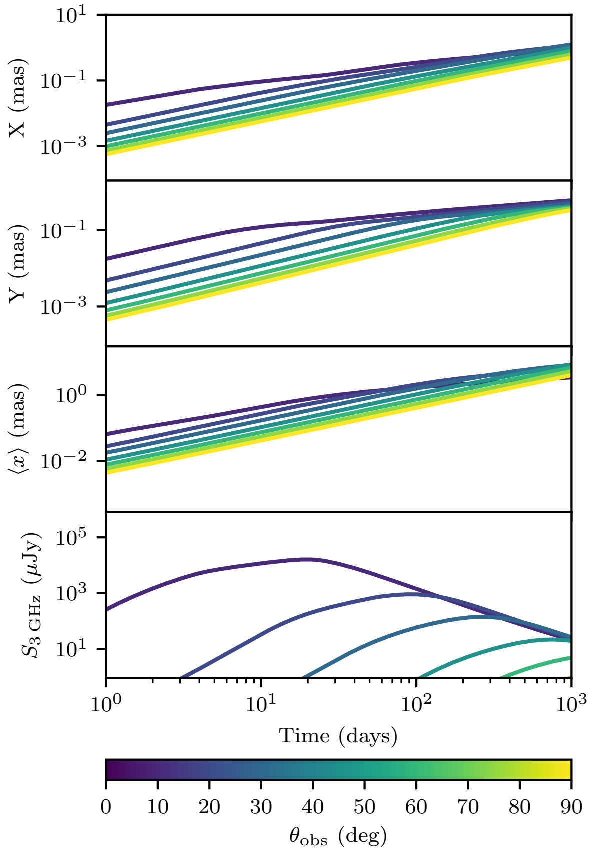

Figure 5.3: Angular size (top 2 rows), centroid offset (row 3) and 3 GHz flux density (bottom) for a neutron star merger jet model at a distance of 40 Mpc with a circum-merger density of $n = 10^{-3}\,\mathrm{cm}^{-3}$ for a range of inclination angles.

Table 5.2: Specifications of radio facilities we consider in the detection of refractive and diffractive scintillation. ν_{obs} is the center observing frequency, BW is bandwidth, t_{obs} is the proposed observing time and σ_{min} is the corresponding image sensitivity. For existing facilities we use noise estimates based on the achieved sensitivity in observations of GW170817 (e.g. Mooley et al., 2018d), taking into account increased noise do to host galaxy emission, while we use thermal noise estimates for the SKA (Braun et al., 2019).

Telescope	ν_{obs} (GHz)	BW (GHz)	t_{obs} (hours)	σ_{min} (μJy)
VLA-S	3.0	1.5	3	2.4
VLA-C	6.0	4.0	3	1.5
ATCA-CX	7.2	4.0	12	8
SKA-1 Mid	1.4	0.77	3	0.45

Diffractive Interstellar Scintillation (DISS)

For observing frequencies $\nu < \nu_0$ we are in the strong scattering regime, where two forms of scintillation are present. Diffractive scintillation is fast (minutes–hours), narrow-band variability typically observed in compact Galactic objects like pulsars (e.g. Rickett, 1977; Cordes, 1986; Phillips & Clegg, 1992; Goodman, 1997; Bell et al., 2016). We must consider the detectability of such variations on the small frequency- and time-scales involved, which require consideration of the telescope sensitivities on those scales. For optimal detectability we require our time resolution to be less than the variability timescale and the channel width to to be comparable to the frequency scale, which is given by

$$\Delta\nu = \nu \left(\frac{\nu}{\nu_0}\right)^{17/5}. \tag{5.10}$$

The size of the scattering disk is given by

$$\theta_d = \theta_{F0} \left(\frac{\nu}{\nu_0}\right)^{6/5}, \tag{5.11}$$

Low frequencies ($\leq 300\,\mathrm{MHz}$) are in the strong scattering regime for the entirety of the sky. However, the diffractive scintillation bandwidth for compact extragalactic sources at low frequency is typically $\sim 3\,\mathrm{kHz}$, smaller than the continuum channel width of current generation telescopes at this frequency.

The scintillation timescale of diffractive scintillation is given by

$$t_d = 2\,\mathrm{hour} \left(\frac{\nu}{\nu_0}\right)^{6/5} \max\left[1, \frac{\theta_\mathrm{S}}{\theta_d}\right], \tag{5.12}$$

and in this situation is typically 1 minute. Therefore detecting DISS requires sub-minute observations split into single channels. The estimated sensitivity for this type of observation with current generation low-frequency telescopes is on the order of hundreds of mJy. In comparison, the flux density of GW170817 would have been $< 1\,\mathrm{mJy}$ at its peak.

The size of the scattering disk at low frequencies is significantly smaller than the Fresnel zone, meaning that any variability will have a low amplitude since the modulation index of diffractive scintillation is given by

$$m_d = \min\left[1, \frac{\theta_d}{\theta_\mathrm{S}}\right]. \tag{5.13}$$

Therefore, with the exception of bright events, DISS from gravitational wave afterglows will not be detectable at low frequencies with current radio facilities.

Frequencies around 1 GHz accessible with telescopes such as the Australian Square Kilometre Array Pathfinder (ASKAP; Johnston et al., 2008), the Karl G. Jansky Very Large Array (VLA) and MeerKAT[1] are in the strong scattering regime for the

---

[1] http://www.ska.ac.za/gallery/meerkat/

entirety of the sky. The typical diffractive scintillation timescale at these frequencies is a few minutes. Zic et al. (2019) demonstrate the capability of ASKAP to detect short-timescale variability with high spectral resolution by performing dynamic spectroscopy of UV Ceti (a bright, well-known flare star), achieving a sensitivity equivalent to $\sim 12\,\mathrm{mJy}$ in a 10 second integration is per 1 MHz channel. However, current telescopes do not have the instantaneous sensitivity required to detect scintillation from sources at distances comparable to the LIGO horizon, which will likely peak below 1 mJy in this frequency regime (Dobie et al., 2019a).

Additionally, the characteristic frequency scale of diffractive scintillation at these frequencies is $< 1\,\mathrm{MHz}$ across the sky. Current generation gigahertz-frequency telescopes have channel widths of 1 MHz and therefore do not have sufficient spectral resolution to detect extragalactic diffractive scintillation.

Frequencies up to 10 GHz are in the strong scattering regime away from the Galactic poles ($|b| \lesssim 40°$). At 3 GHz typical values for the scintillation characteristic bandwidth and timescale are tens of MHz and tens of minutes respectively. The sensitivity of a 5 minute observation with the VLA at 3 GHz using 5 MHz of bandwidth is $300\,\mu\mathrm{Jy}$. At frequencies closer to 10 GHz the characteristic bandwidth and timescale of diffractive scintillation both increase to values of $\sim 1\,\mathrm{GHz}$ and ~ 1 hour. Both of these represent reasonable prospects of detecting scintillation, and we perform a more quantitative analysis in Section 5.3.3.

Millimetre wavelengths are only in the strong scattering regime close to the Galactic plane, with characteristic bandwidths of a few GHz and timescales comparable to those at lower frequencies. We expect that the spectral index of the radio afterglow will be negative (Sari et al., 1998; Berger, 2014), although this may not be true in all cases. Therefore in general, sources will be more difficult to detect at these frequencies compared to observations at lower frequencies. However, emission from mergers occuring in environments that are more dense than the typical short GRB circum-burst density ($n \sim 10^{-2}\,\mathrm{cm}^{-3}$; Fong et al., 2015), may be detectable with relatively short integrations. Therefore we do not perform any further analysis, but do not rule out the possibility of detecting DISS with millimetre observations.

Refractive Interstellar Scintillation (RISS)

Refractive scintillation manifests as slower broadband changes and is observed in pulsars as well as compact extragalactic sources like quasars. For refractive interstellar scintillation the size of the scattering disk, θ_r, is given by

$$\theta_r = \theta_{F0} \left(\frac{\nu_0}{\nu}\right)^{11/5}, \tag{5.14}$$

the modulation index, m_r, is given by

$$m_r = \left(\frac{\nu}{\nu_0}\right)^{17/30} \min\left[1, \left(\frac{\theta_r}{\theta_\mathrm{S}}\right)^{7/6}\right], \tag{5.15}$$

and the variability timescale, t_r, is given by

$$t_r = 2\,\mathrm{hour} \left(\frac{\nu_0}{\nu}\right)^{11/5} \max\left[1, \frac{\theta_\mathrm{S}}{\theta_r}\right]. \tag{5.16}$$

The modulation index for RISS is typically $<10\%$ for a compact source at low frequencies, and even lower for sources larger than the size of the scattering disk. We therefore do not expect to detect any form of interstellar scintillation from GW radio afterglows with current low frequency telescopes.

Away from the Galactic plane the modulation index for refractive scintillation from compact sources ranges from 0.3–0.45 for mid-frequencies, with a typical timescale of a few days. The size of the scattering disk is $\theta_r \approx 0.5\,\mathrm{mas}$ for $\theta_{F0} \sim 3\,\mathrm{mas}$ and $\nu_0 \sim 10\,\mathrm{GHz}$ (typical values away from the Galactic plane), meaning that the radio afterglow behaves as a compact source until late times. At high frequencies the modulation index for compact sources is larger but the scattering disk is smaller and therefore the point-source approximation does not apply for as long. We investigate prospects for detecting RISS at mid and high frequencies in Section 5.3.3.

At mm-wavelengths, the timescale of refractive scintillation is comparable to t_d. For similar reasons as in Section 5.3.1 we do not perform further analysis, but do not rule out the possibility of detecting RISS with mm observations for events near the Galactic plane.

5.3.2 A Generalised Metric for Detecting Scintillation

Our ability to detect variability is strongly dependent on the exact observing strategy, most notably the total number of observations, observation sensitivity and observing cadence. We therefore define a generalised detectability metric that allows us to place broad estimates on the range at which scintillation may be detected for a combination of merger parameters. We emphasise that this metric should not be used in preparing follow-up observations of individual events and instead more detailed, event-specific calculations should be performed.

For scintillation to be detectable we require $m > m_{\mathrm{detect}}$ and $S > S_{\mathrm{detect}}$ where m is the modulation index of the source due to scintillation, S is the flux density, while m_{detect} and S_{detect} are the minimum detectable values for each of those quantities respectively. We define both detectability thresholds in terms of the image RMS, σ; $S_{\mathrm{detect}} = 5\sigma$ and $m_{\mathrm{detect}} = 5\sigma/S$. For some events scintillation may be detected on timescales of days–weeks in the form of inter-observation variability, but other events may exhibit intra-observation variability. As such we define

$$\sigma = \begin{cases} \sigma_{\mathrm{min}} & \text{for } t_{\mathrm{obs}} \geq t_s \\ \sigma_{\mathrm{min}}\sqrt{t_{\mathrm{obs}}/t_s} & \text{for } t_{\mathrm{obs}} < t_s \end{cases}, \tag{5.17}$$

where σ_{min} is the minimum reasonable image RMS achievable with a telescope, requiring an observation time of t_{obs} and t_s is the scintillation timescale.

In the case of diffractive scintillation we also have constraints based on the scintillation bandwidth, the telescope channel width, frequency and bandwidth. For scintillation to be detectable we require 10 samples across the scintillation bandwidth and define the effective bandwidth as $\Delta\nu_{\mathrm{eff}} = \Delta\nu/10$. We require that the telescope channel width (typically 1 MHz) is less than $\Delta\nu_{\mathrm{eff}}$. We then correct the image RMS defined in (5.17) for the fractional bandwidth, scaling it as $\Delta\nu_{\mathrm{eff}}^{-1/2}$.

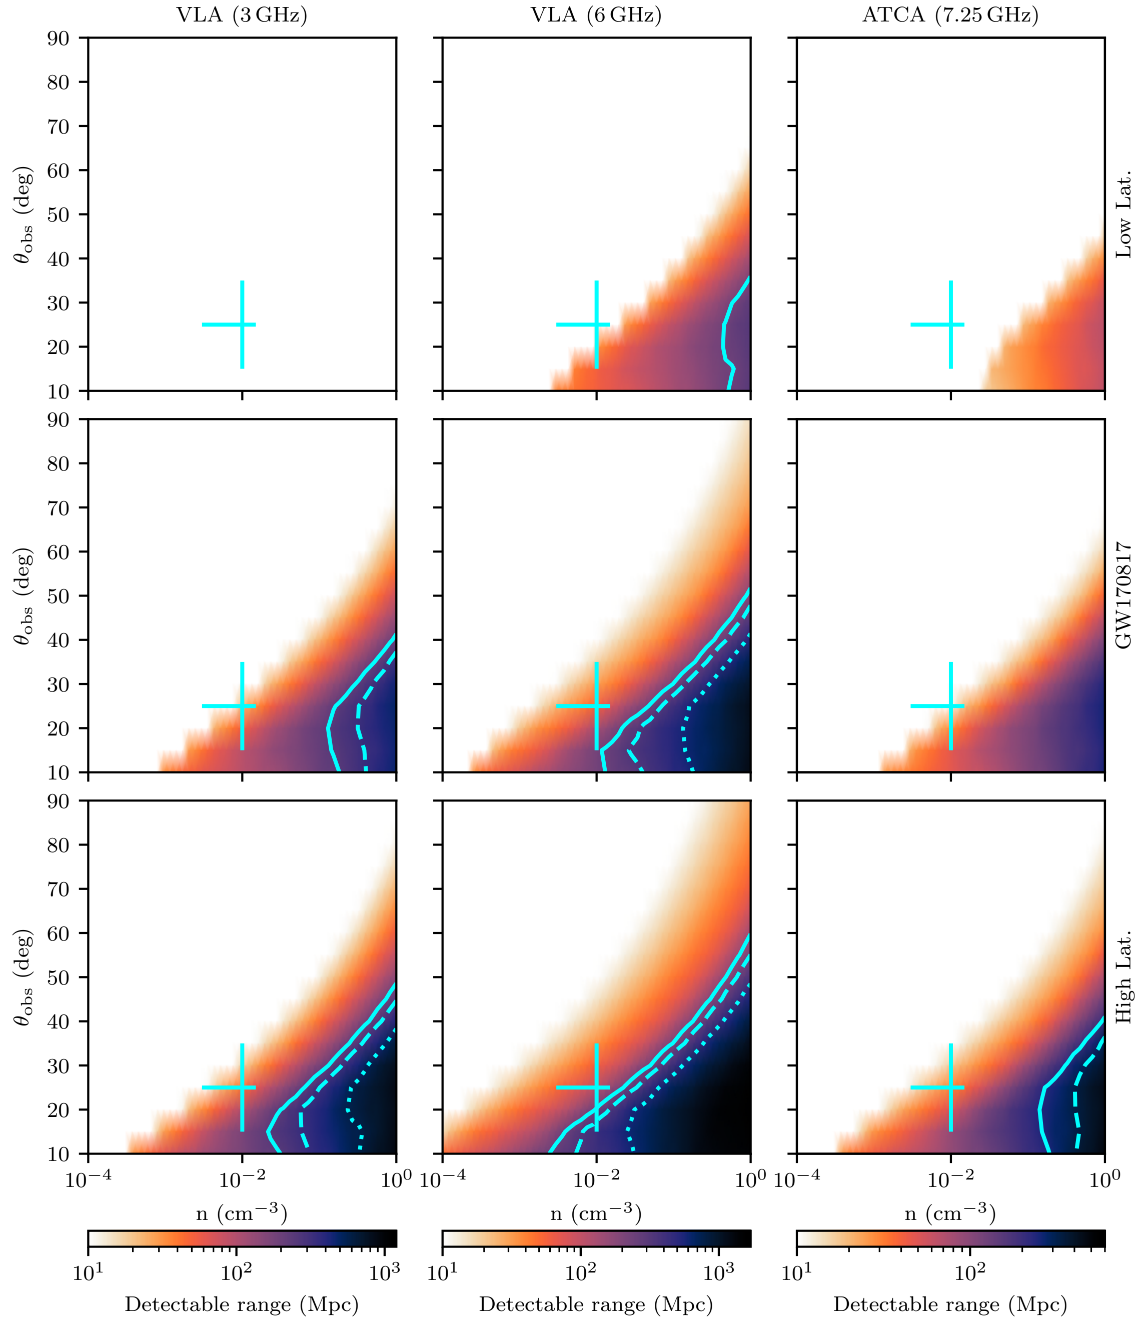

Figure 5.4: Maximum distance at which diffractive scintillation is detectable for a range of electron density parameters. Top: typical scintillation parameters at low Galactic latitudes ($10° < |b| < 20°$), $\nu_0 = 18.1\,$GHz and $\theta_{F0} = 1.5\,\mu$as. Middle: $\nu_0 = 10.3\,$GHz and $\theta_{F0} = 2.9\,\mu$as corresponding to the line of sight to GW170817. Bottom: typical scintillation parameters at high Galactic latitudes ($60° < |b| < 70°$), $\nu_0 = 7.82\,$GHz and $\theta_{F0} = 3.9\,\mu$as. This is shown for observations with the VLA at 3 GHz (left, not detectable for low $|b|$), 6 GHz (middle) and observations with the ATCA at 7.25 GHZ (right). The jagged edge is an artefact of using simulating models with steps in inclination angle of 5 deg. Contours corresponding to the inclination angle dependent LIGO horizon for O3 (135 Mpc, solid), design specifications (190 Mpc, dashed), and A+ (330 Mpc, dotted) are shown in blue. The blue cross corresponds to the typical short GRB circum-merger density and estimates for the inclination angle of GW170817.

Finally, we require that the above conditions are satisfied for at least thirty days (since observers have minimal a priori knowledge of when scintillation will be detectable) and ten scintillation timescales (to allow for a sufficient number of observations to characterise the variability as being produced by scintillation).

Variability caused by scintillation is more easily detectable for sources with small angular sizes and large flux densities. More distant events have smaller angular sizes (scaling as D^{-1}) and lower flux densities (scaling as D^{-2}) and therefore scintillation is detectable for a range of distances, and not simply out to a horizon distance. However, the minimum of that range is typically $< 10\,\mathrm{Mpc}$ so for brevity we simply quote the maximum detectable distance as our detectability metric due to the expected low rate of events occuring at such small distances (Abbott et al., 2019b).

5.3.3 Scintillation Detectability

Detectability of Diffractive Scintillation

Here we consider follow-up with 3 telescope configurations:

1. VLA follow-up in S band (3 GHz) with $t_{\mathrm{obs}} = 3\,\mathrm{hr}$ and $\sigma_{\mathrm{min}} = 2.4\,\mu\mathrm{Jy}$;
2. VLA follow-up in C band (6 GHz) with $t_{\mathrm{obs}} = 3\,\mathrm{hr}$, $\sigma_{\mathrm{min}} = 1.5\,\mu\mathrm{Jy}$;
3. ATCA follow-up in the CX band (7.25 GHz) with $t_{\mathrm{obs}} = 12\,\mathrm{hr}$ and $\sigma_{\mathrm{min}} = 8\,\mu\mathrm{Jy}$

The stated sensitivity reflects typical values achieved during the follow-up of GW170817 (Mooley et al., 2018d), which was limited by radio emission from the host galaxy. It may be possible to achieve better sensitivity for events with a less luminous host galaxy.

Figure 5.4 shows the maximum detectable distance at which scintillation is detectable with each of these telescope configurations as a function of circum-merger density and observing angle for a range of scintillation parameters corresponding to a range of Galactic latitudes. As expected, on-axis events in denser environments are detectable to a larger distance as they have higher flux densities but also remain compact for longer. Almost half of the parameter space (low density, off-axis events) is inaccessible with current radio facilities. We find that away from the Galactic plane ($|b| > 30^\circ$) the dependence of scintillation detectability on Galactic latitude is minimal, but not negligible.

We now compare the maximum detectable distance to the LIGO detector horizons which are averaged across the sky and inclination angles. The best detector horizon achieved to-date in O3 is 135 Mpc, while the expected range for O3 and design sensitivity is 150 and 190 Mpc respectively (Abbott et al., 2018). Additionally, the planned A+ upgrade that will be online for the O5 run scheduled in 2025 will increase the detector horizon to 330 Mpc (Abbott et al., 2018).

To find the dependence of the horizon on inclination angle we average equation (3.31) from Finn & Chernoff (1993) across the antenna patterns. We find

$$\mathcal{R}(\theta_{\rm obs}) \approx 0.658\overline{\mathcal{R}}\sqrt{1 + 6\cos^2\theta_{\rm obs} + \cos^4\theta_{\rm obs}} \tag{5.18}$$

where $\mathcal{R}(\theta_{\rm obs})$ is the inclination angle dependent range and $\overline{\mathcal{R}}$ is the gravitational wave detector horizon.

The detectable range of diffractive scintillation with the VLA extends beyond the LIGO horizon for dense, on-axis events away from the Galactic plane, while the range of the ATCA is typically 3–5 times lower. Events occurring in similar environments to GW170817 ($n \sim 10^{-3}\,\mathrm{cm}^{-3}$; see Figure 5.1) will not exhibit diffractive scintillation detectable with any current radio telescopes. However, events with circum-merger densities comparable to typical short GRB circum-burst densities ($n \sim 10^{-2}\,\mathrm{cm}^{-3}$; Fong et al., 2015) may exhibit scintillation detectable with the VLA C band receiver at distances of 100 Mpc for $\theta_{\rm obs} < 30^\circ$.

We also run the simulation for the range of energetics and microphysics parameters stated in Section 5.2, varying each parameter individually and keeping the remaining parameters at the fiducial value. The detectable range is most influenced by the isotropic equivalent energy, with the typical range varying by a factor of 0.002–3, although the best and worst cases are factors of 9 and 0.001. Varying ϵ_e and ϵ_B changes the detectable range by factors of 0.1-5 and 0.03–1.7 respectively. These factors are only very weakly dependent on the scintillation and telescope parameters, and any variance is negligible compared to the uncertainty in our models.

Detectability of Refractive Scintillation

We consider the three observing scenarios defined above, applied to refractive scintillation. Figure 5.5 shows the results of this analysis. We find that the overall trend of scintillation being detectable to larger distances for events that are on-axis and occur in dense environments holds true for both forms of strong scattering. However the fraction of the $\theta_{\rm obs}$–n parameter space accessible is much larger for refractive scintillation than diffractive scintillation, and the detectability range tends to be larger. We find that refractive scintillation from GW170817 may have been detectable assuming a robust, high-cadence, follow-up plan had been in place, which would have likely required knowing precise merger parameters a priori.

As in Section 5.3.3 we run the simulation for a range of energetics and microphysics parameters. Again, the detectable range is most influenced by the isotropic equivalent energy, with the typical range varying by a factor of 0.01–2.5, although the best and worst cases are factors of 4 and 0.005. Varying ϵ_e and ϵ_B changes the detectable range by factors of 0.2–3.2 and 0.1–1.5 respectively.

Future Prospects: Next generation radio telescopes

A large fraction of the parameter space exhibits refractive scintillation detectable with the VLA beyond the LIGO detector horizons. However, scintillation-induced variability is not detectable from low-density off-axis events. We therefore turn our focus to the Square Kilometre Array (SKA), which will have a sensitivity orders

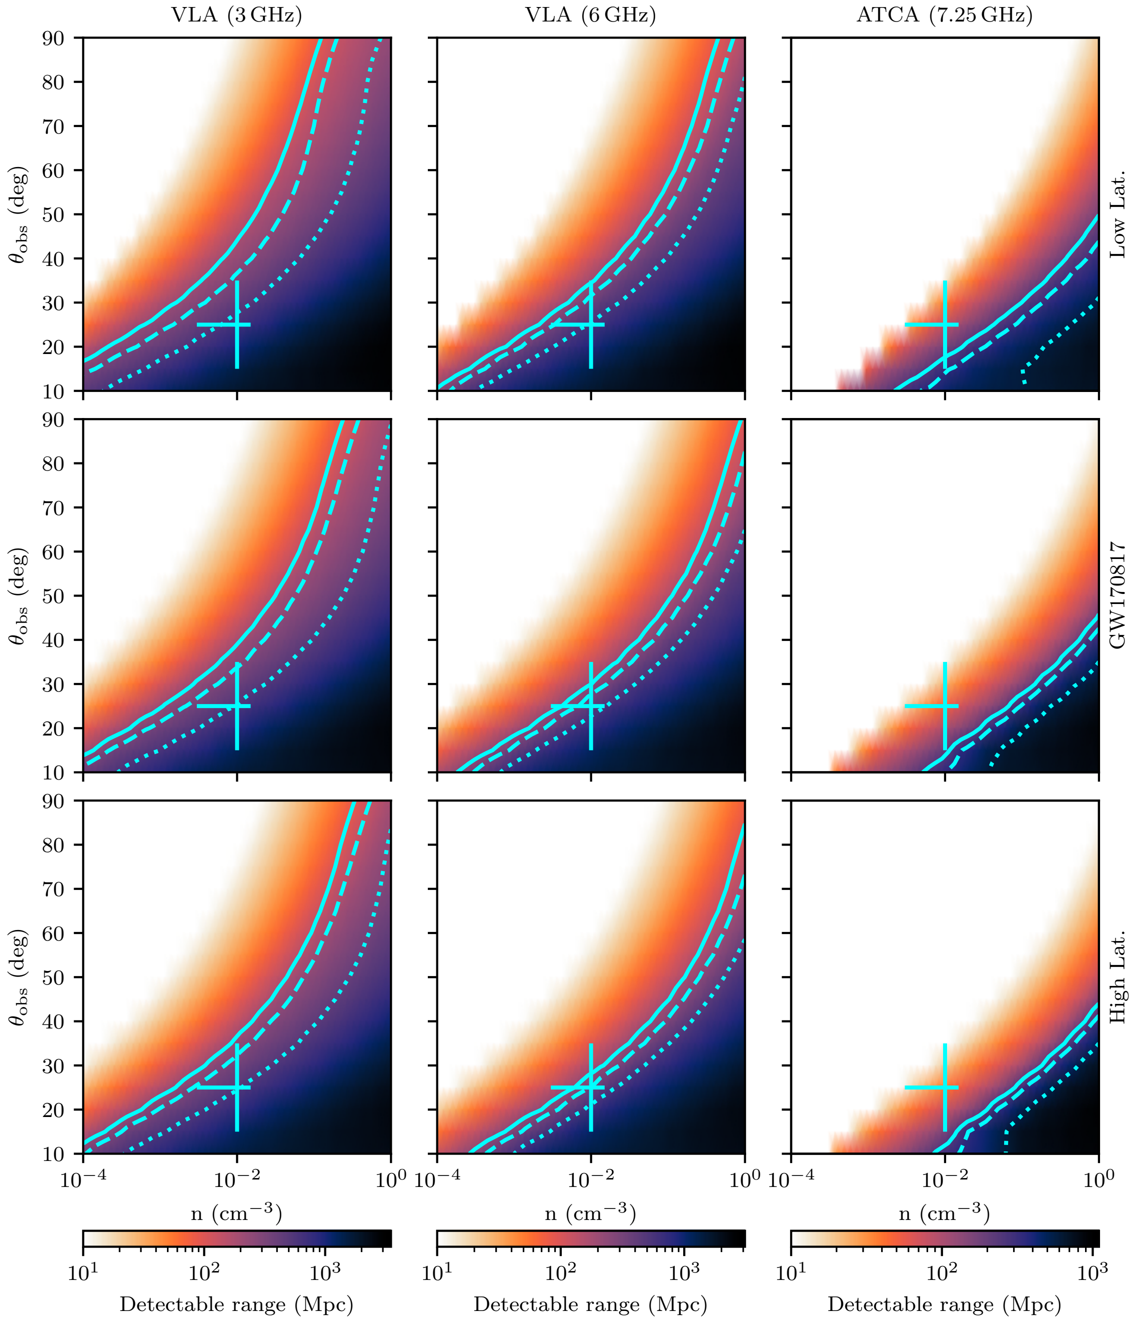

Figure 5.5: Maximum distance at which refractive scintillation is detectable for a range of electron density parameters. Top: typical scintillation parameters at low Galactic latitudes ($10^\circ < |b| < 20^\circ$), $\nu_0 = 18.1\,$GHz and $\theta_{F0} = 1.5\,\mu$as. Middle: $\nu_0 = 10.3\,$GHz and $\theta_{F0} = 2.9\,\mu$as corresponding to the line of sight to GW170817. Bottom: typical scintillation parameters at high Galactic latitudes ($60^\circ < |b| < 70^\circ$), $\nu_0 = 7.82\,$GHz and $\theta_{F0} = 3.9\,\mu$as. This is shown for observations with the VLA at 3 GHz (left), 6 GHz (middle) and observations with the ATCA at 7.25 GHZ (right). The jagged edge is an artefact of using simulating models with steps in inclination angle of 5°. Contours corresponding to the inclination angle dependent LIGO horizon for O3 (135 Mpc, solid), design specifications (190 Mpc, dashed), and A+ (330 Mpc, dotted) are shown in blue. The blue cross corresponds to the typical short GRB circum-merger density and estimates for the inclination angle of GW170817.

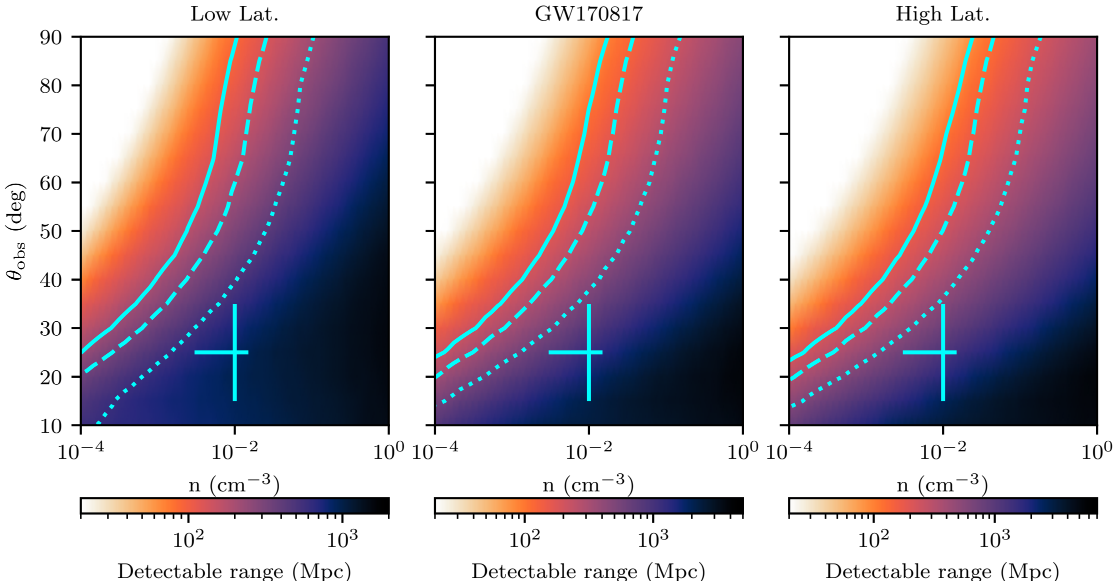

Figure 5.6: Maximum distance at which refractive scintillation is detectable with the Square Kilometre Array for a range of electron density parameters. Left: typical scintillation parameters at low Galactic latitudes ($10° < |b| < 20°$), $\nu_0 = 18.1\,\mathrm{GHz}$ and $\theta_{F0} = 1.5\,\mu$as. Middle: $\nu_0 = 10.3\,\mathrm{GHz}$ and $\theta_{F0} = 2.9\,\mu$as corresponding to the line of sight to GW170817. Right: typical scintillation parameters at high Galactic latitudes ($60° < |b| < 70°$), $\nu_0 = 7.82\,\mathrm{GHz}$ and $\theta_{F0} = 3.9\,\mu$as. Contours corresponding to the inclination angle dependent LIGO horizon for O3 (135 Mpc, solid), design specifications (190 Mpc, dashed), and A+ (330 Mpc, dotted) are shown in blue. The blue cross corresponds to the typical short GRB circum-merger density and estimates for the inclination angle of GW170817.

of magnitude better than existing radio telescopes. We consider observations with the SKA-1 (mid) array at $\nu = 1.4\,\mathrm{GHz}$, assuming a bandwidth of 770 MHz and a sensitivity of $1.2\,\mu$Jy in a 3 hour integration (Braun et al., 2019).

Figure 5.6 shows that it will be possible to detect refractive scintillation from all but a small minority of low density, off-axis events detected by LIGO/Virgo with the SKA. Most events will exhibit detectable scintillation out to Gpc distances, although generally not beyond the horizon of third generation gravitational wave detectors which will come online in the 2030s-2040s (Reitze et al., 2019) and have detection horizons of tens–hundreds of Gpc (Sathyaprakash et al., 2013). Like the current situation, the most dense and on-axis events will be detectable beyond the gravitational wave detector horizon.

We have also applied the criteria outlined in Section 5.3.3 to the SKA. We find that no events will exhibit diffractive scintillation detectable with SKA continuum observations due to the scintillation bandwidth being smaller than the continuum channel width. However, we find that using spectral line observing, and assuming a channel width given by $\Delta\nu/\nu = 10^{-4}$ the SKA has a detectability horizon ~ 5 times larger than observations in VLA C band. As well as having a larger horizon, these observations will allow tighter constraints to be placed on source sizes at early times, as the scattering disk is almost six times smaller.

Table 5.3: Estimated parameters for 3 VLBI observing scenarios. The HSA consists of the VLBA, the phased VLA, the GBT and Arecibo (Ar). $S_{\rm noise}$ is the estimated thermal noise in the observation, $\theta_{\rm B}$ is the approximate beam size, $\Delta\alpha\cos\delta$ and $\Delta\delta$ are the systematic uncertainties in R.A. and Dec. respectively, and $\theta_{\rm sys}$ is the systematic astrometric uncertainty we use for this analysis, estimated by taking the geometric mean of $\Delta\alpha\cos\delta$ and $\Delta\delta$.

	LBA	HSA	HSA (no Ar)
ν (GHz)	4.8	4.5	4.5
Obs. time (h)	12	2	8
Dec (°)	-30	20	50
$S_{\rm noise}$ (μJy)	20	3.2	3.1
$\theta_{\rm B}$ (mas)	15	3	3
$\Delta\alpha\cos\delta$ (μas)	80	60	80
$\Delta\delta$ (μas)	100	80	100
$\theta_{\rm sys}$ (μas)	90	70	90

The dependence on energetics and microphysics parameters for these ranges is comparable to the values in Section 5.3.3 and 5.3.3.

Two other major radio facilities are expected to come online on similar timescales. The 2000 antenna Deep Synoptic Array (DSA; Hallinan et al., 2019) will have a 1 hr continuum sensitivity of 1 μJy and while it will be more suited to discovering radio emission from compact object mergers due to it's large field of view and high survey speed, it will also be capable of observing scintillation from events within $\sim 80\%$ of the estimated range of the SKA. The next-generation VLA (ngVLA; McKinnon et al., 2019; Corsi et al., 2019a,b) will improve the sensitivity and resolution of the VLA by a factor of 10, corresponding to a detector horizon that is ~ 3 times larger than the existing VLA.

5.4 VLBI Observations

VLBI observations of GW170817 were important in determining the geometry of the merger, constraining both the emission model (jet-dominated) and the inclination angle ($\sim 20°$) via observations of a positional shift in the source centroid (Mooley et al., 2018b; Ghirlanda et al., 2019). In this section we discuss prospects for directly imaging outflow structure using VLBI imaging and expand on the work of Duque et al. (2019) by determining a detection metric driven by observing constraints.

The parameters of any VLBI observation are strongly dependent on the declination of the target source. Northern Hemisphere sources are accessible with the High Sensitivity Array (HSA), consisting of the Green Bank Telescope, the phased VLA, Arecibo and the Very Long Baseline Array (VLBA). The longest baseline spans 8611 km (Mauna Kea–Saint Croix), corresponding to a best achievable angular resolution of 0.8 mas at 9 GHz, which is the highest frequency available to Arecibo.

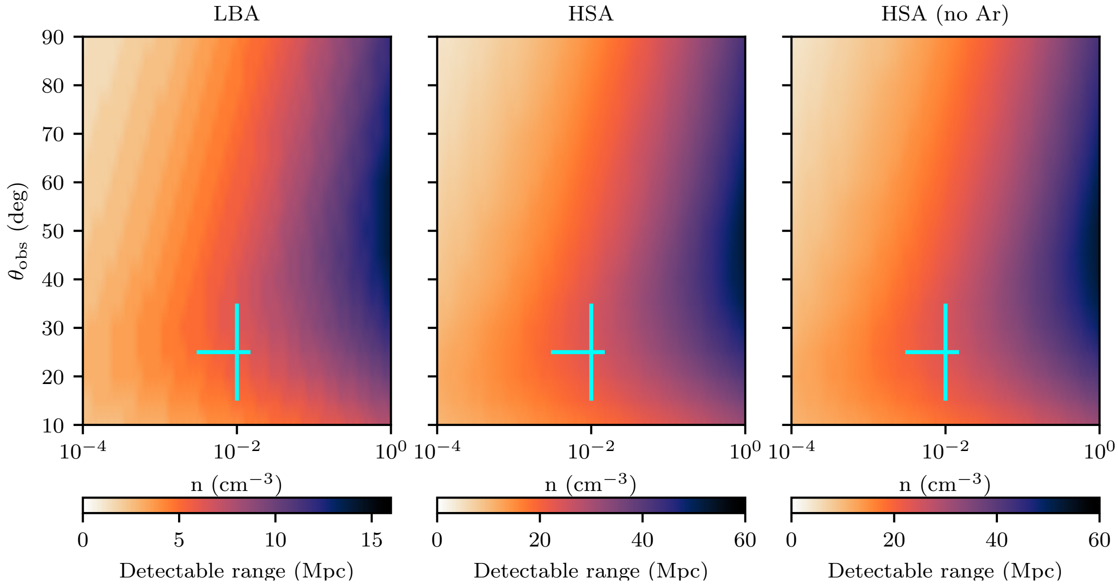

Figure 5.7: Maximum distance at which outflow structure can be resolved for a range of circum-merger densities and merger inclination angles. Left: observations at −30° declination with the LBA. Middle: observations at +20° declination with the full HSA. Right: observations at +50° declination using the HSA without Arecibo. The blue cross corresponds to the typical short GRB circum-merger density and estimates for the inclination angle of GW170817.

However, observations with Arecibo longer than 2 hours are only possible between declinations of +5 and +30 °[2]. Using the European VLBI Network (EVN) calculator[3] the flux density sensitivity of the array in a 2 hour observation ranges from 2.9–3.6 μJy for observing bands from 1.4–9 GHz. Observations without Arecibo allow for longer integration times, so we additionally consider an 8 hour observation, achieving a sensitivity of 3.1 μJy.

Southern Hemisphere sources are accessible with the Long Baseline Array (LBA), consisting of the phased ATCA and the Parkes, Mopra, Hobart, Ceduna and Tidbinbilla telescopes. The maximum baseline of the LBA is $\sim$ 1700 km, or five times smaller than that of the HSA. The sensitivity of a 12 hour observation is typically 20 μJy. The Hartebeestok telescope in South Africa can also be included to achieve an angular resolution comparable to the HSA but we do not consider it in this work due to the lack of intermediate baseline lengths. In determining the effective range of VLBI observations we consider three scenarios, outlined in Table 5.3;

1. LBA observations of a source at −30° declination;
2. HSA observations of a source at +20° declination and;
3. Observations of a source at +50° declination using the HSA without Arecibo

---

[2] `http://www.naic.edu/~astro/aovlbi/`

[3] `http://www.evlbi.org/cgi-bin/EVNcalc`

5.4.1 Resolving Outflow Structure

For an afterglow to be resolvable with VLBI, it must be both sufficiently bright and sufficiently large. For sources near the detection threshold, only sizes comparable to or larger than the VLBI synthesized beam can be measured with any degree of confidence, while brighter sources can be resolved even when smaller than the VLBI synthesized beam. The size of the source in the VLBI image is given by adding the source size, $\theta_{\rm S}$, and the VLBI beam size, $\theta_{\rm B}$, in quadrature;

$$\Theta = \sqrt{\theta_{\rm B}^2 + \theta_{\rm S}^2}. \tag{5.19}$$

The uncertainty in the fit of Θ given by

$$\sigma_\Theta = \frac{\sqrt{2}\Theta}{\rho} \tag{5.20}$$

where ρ is the signal-to-noise ratio of the source given by a simplified form of Equation 41 from Condon (1997);

$$\rho^2 = \frac{\Theta^2}{4\theta_{\rm B}^2}\left[1 + \left(\frac{\theta_{\rm B}}{\Theta}\right)^2\right]^3 \frac{S^2}{\sigma^2} \tag{5.21}$$

where we have assumed the beam is a circular Gaussian. Equation (5.21) assumes that the signal-to-noise ratio is dominated by the thermal noise. However, the sensitivity of observations of bright sources are dominated by phase errors which limit the dynamic range of the image. We therefore impose an additional cutoff of $S/\sigma < 100$, corresponding to a phase error of $\sim 5^\circ$ (Perley, 1999). The uncertainty in the source size can then be found using the equations of Murphy et al. (2017);

$$\sigma_{\theta_{\rm S}} = \sigma_\Theta \left[1 - \left(\frac{\theta_{\rm B}}{\Theta}\right)^2\right]^{-1/2}. \tag{5.22}$$

For the source to be resolved and its angular size measured we require that $\theta_{\rm S} > 2\sigma_{\theta_{\rm S}}$ and $\rho > 5$, and that both of these criteria are true for a minimum of 30 days.

In general the ability to resolve objects with large angular diameters may be limited by the structure and surface brightness of the object. However, we find that most events that are resolvable will only be slightly larger than the size of the beam, and therefore more detailed surface brightness considerations are not required.

Figure 5.7 shows the maximum distance at which events are resolvable as a function of circum-merger density and inclination angle. The detectability range is strongly dependent on merger parameters, and typical values are 20 Mpc and 5 Mpc using the HSA and LBA respectively (maximum detectable range 52 Mpc and 14 Mpc). Therefore only the very closest events will have resolvable outflows. We also note that while events with low circum-merger densities will expand to have larger physical sizes than those with high circum-merger densities, and the corresponding source luminosity is much lower. We find that events occuring in denser environments will be resolved easier, as the dominant factor is the signal-to-noise of the source.

As expected, the merger energetics are the most dominant parameter, with the median range varying by a factor of 0.2–1.5 compared to the range for the fiducial parameters. Dense, on-axis events are the least affected by varying merger energetics, varying by a factor of 0.3–1.2, compared to low density off-axis events that vary by a factor of 0.08–1.9. Microphysics parameters are generally less dominant, varying by factors of 0.5–1.6 and 0.4–1.2 for ϵ_e and ϵ_B respectively.

5.4.2 Astrometric Accuracy

Our ability to measure centroid motion is strongly affected by the precision with which we can measure the position of the source in each epoch. There is no analytic solution to the astrometric accuracy of VLBI observations, and we are therefore limited to estimates using numerical simulations. Pradel et al. (2006) provide values of systematic astrometric accuracy of VLBA and EVN observations for a range of declinations from −25° to +85° taking into account uncertainty in the Earth's orientation, calibrator position[4] and antenna positions, as well as uncertainties induced by the troposphere. For GW170817, Mooley et al. (2018b) estimate the systematic contribution to astrometric uncertainty, $\theta_{\rm sys}$, with the HSA to be 0.15 mas and 0.5 mas in RA and Dec. respectively after taking into consideration ionospheric effects. However, the precision of these observations were dominated by the low elevation of the source from the VLBA. In this work we use more optimistic values, outlined for each of the three cases in Table 5.3.

The astrometric accuracy also has a statistical component given by

$$\theta_{\rm stat} = \frac{\theta_{\rm B}}{\sqrt{8\ln 2}\rho} \tag{5.23}$$

and the total astrometric uncertainty is then

$$\theta_{\rm total} = \sqrt{\theta_{\rm sys}^2 + \theta_{\rm stat}^2}. \tag{5.24}$$

5.4.3 Detectability of Centroid Motion

The intrinsic factors influencing the detectability of centroid motion are the merger luminosity and the magnitude of the observed source offset. A denser circum-merger medium will produce a more luminous afterglow that moves slower, while it will be harder to detect centroid motion in more distant events which have lower flux densities and smaller angular offsets. On-axis events have larger centroid offsets at early times when the outflow is relativistic as the apparent velocity is dominated by superluminal motion. However, as the outflow decelerates the apparent velocity becomes dominated by the transverse component of the physical velocity of the outflow, which is higher for off-axis events.

Figure 5.3 shows the offset of the afterglow centroid from the merger location and its flux density as a function of time for an event at a distance of 40 Mpc with a circum-merger density of $n = 10^{-3}\,{\rm cm}^{-3}$, comparable to the parameters of GW170817 (see

---

[4] Systematic errors in calibrator position will affect both epochs equally and can be ignored for our purposes.

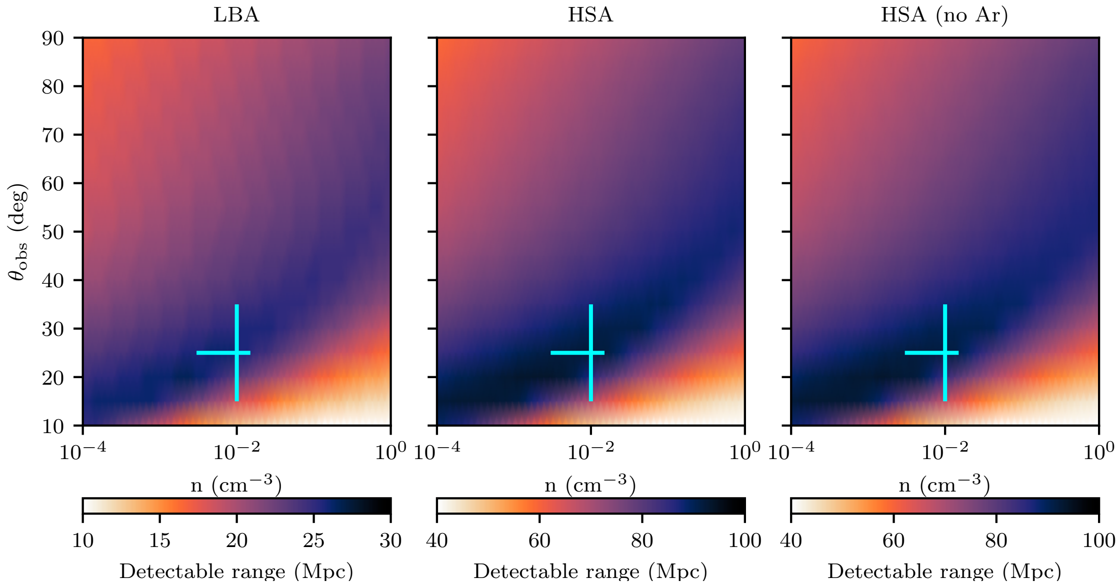

Figure 5.8: Maximum distance at which centroid motion can be detected for a range of circum-merger densities and merger inclination angles. Left: observations at -30° declination with the LBA. Middle: observations at +20° declination with the full HSA. Right: observations at +50° declination using the HSA without Arecibo. The blue cross corresponds to the typical short GRB circum-merger density and estimates for the inclination angle of GW170817.

Figure 5.1), for a range of inclination angles. The maximum detectable offset is the offset between the centroid position of the afterglow at the first and last times the flux density of the afterglow is above the detection threshold of the telescope.

We calculate the detectable distance, D, by scaling the flux density of the afterglow by D^{-2} and the angular offset by D^{-1}. We then calculate the centroid offset between each combination of times, t_1 and t_2, given by

$$\langle X \rangle = |\langle x \rangle(t_2) - \langle x \rangle(t_1)| \tag{5.25}$$

and define that offset as detectable if

$$\langle X \rangle > 5\left[\theta_{\rm total}(t_1) + \theta_{\rm total}(t_2)\right] \tag{5.26}$$

Finally, for an event to be considered detectable we additionally impose the constraint that there must be at least a 30 day buffer around both t_1 and t_2, as observers will have minimal a priori knowledge of the optimal times to observe.

Figure 5.8 shows the detectable range as a function of circum-merger density and inclination angle. We find that the detectability of centroid motion is less dependent on merger parameters than the resolvability of the outflow. Typical ranges are 20 Mpc and 80 Mpc for the LBA and VLBA respectively, only 20% lower than the maximum range. Unlike our ability to resolve outflow structure, the most dense and on-axis events have the lowest detectability ranges. While these events have the highest peak luminosities, the decline of the lightcurve is only weakly dependent on either parameter (e.g. see Figure 5.3). Denser events have lower initial velocity, and also decelerate faster meaning that the apparent velocity of events becomes dominated by the physical velocity of the outflow (which is higher for off-axis events) at earlier times.

Again, merger energetics dominate the microphysics parameters, with the median range varying by a factor of 0.1–1.3. Regions with larger ranges saw a greater decrease for lower values of $E_{\rm iso}$, and a smaller increase for higher values of $E_{\rm iso}$. The median range varied by a factor of 0.5–1.3 for ϵ_e and 0.4–1.1 for ϵ_B.

Superluminal motion will not be detectable for a large fraction of events within the LIGO O3 detector range of 135 Mpc (Abbott et al., 2018). As current gravitational wave detectors reach design sensitivity, a decreasing fraction of the more distant events will exhibit centroid motion that is detectable by current VLBI facilities.

We note that the limiting factor of VLBI observations in general is the maximum baseline, which is constrained by the size of the Earth. While larger baselines can be achieved using space-VLBI satellites like *RadioAstron* these facilities do not have sufficient sensitivity for gravitational wave follow-up (Kardashev et al., 2013). Adding new telescopes to existing VLBI arrays and/or increasing observing bandwidth would improve sensitivity, increasing the detection range of events with low circum-merger densities and large inclination angles at current observing frequencies. The higher sensitivity could also be used to facilitate observations at higher frequency, making it possible to discern the smaller motion of more distant events (which would have otherwise been too faint at higher frequency). However, for many events, VLBI information will not provide significant additional constraints. We discuss the implications for the prospects for determining the value of H0 using VLBI techniques in Section 5.5.3.

5.5 Implications

5.5.1 Determining Source Size and Geometry

We have discussed two possible methods for determining the physical size of outflow from neutron star mergers. Observations of scintillation from afterglows will allow us to place constraints on a larger fraction of events than direct imaging with VLBI due to the detectable range being a factor of 2–3 larger. However we caution that both techniques have a role to play in understanding outflow geometry.

Scintillation observations enable us to place constraints on source size at early times and are more effective for dense, on-axis events, as the source remains compact for longer. In contrast, VLBI observations are more useful at later times when the outflow has expanded and the emission centroid has shifted away from the merger location. This distinction will be important for follow-up of events in the future. The late-time VLBI observations of GW170817 constrained the geometry of the merger outflow once it became jet-dominated (tens of days post-merger; Mooley et al., 2018b; Ghirlanda et al., 2019) but the early-time behaviour (and therefore the properties of the cocoon ejecta) is still not well understood.

As well as giving insight into different stages of source evolution, both techniques probe a different part of the parameter space. As the circum-merger density increases scintillation becomes more detectable, while either VLBI phenomena becomes less detectable. Comparing Figure 5.5 and 5.7 shows that while the detectability range for direct imaging of the outflow is only a few tens of megaparsecs, these observations can constrain source sizes for a significant fraction of the parameter space that is inaccessible for scintillation observations with current radio facilities.

Any constraints on source size will be useful in determining the structure of the merger outflow - e.g. Hallinan et al. (2017) used the absence of scintillation-induced variability to rule out the observed radio emission from GW170817 being produced by subrelativisitic ejecta. By obtaining a larger sample of neutron star mergers with well-understood outflow structure we may be able to shed light on the central engine that drives the outflow and produces short GRBs. Many models propose ways in which this central engine may be formed (e.g. see Fryer et al., 2019, and references therein), but the exact mechanism is still unknown.

5.5.2 Observations of Superluminal Motion

The measurement of centroid motion with VLBI provides strong evidence for the presence of a successful jet, and is a robust and direct method to determine the velocity of that jet (see also Linial & Sari 2019 for an alternative method). Combining the observations of superluminal motion and afterglow light curve data, one can determine the jet half-opening angle and viewing angle (Mooley et al., 2018b; Hotokezaka et al., 2019). If the lightcurve data only samples a single segment of the synchrotron spectrum (Sari et al., 1998), the energetics of a jet still degenerates with the circum-merger density. However, if the synchrotron cooling break and/or self-absorption frequency are measured, this degeneracy can be broken. These parameters may be measured via multi-wavelength observations spanning radio to X-rays, although this was not possible in the follow-up of GW170817 because the synchrotron cooling break remained above the observable X-ray band (Alexander et al., 2018). Therefore, VLBI observations of neutron star mergers in conjunction with multi-wavelength lightcurve monitoring can potentially allow us to robustly determine the jet's energetics and structure, viewing angle, as well as the density of the surrounding medium.

In addition to determining merger parameters, centroid motion observations may potentially be used to infer the merger progenitors. For example, the mass ejection (and therefore electromagnetic emission) from neutron star-black hole mergers may by highly anisotropic in comparison to NS mergers which are expected to be axisymmetric (Kyutoku et al., 2013; Foucart et al., 2014, 2016).

While the link between neutron star mergers and short GRBs has been clearly demonstrated by the detection of GW170817 (Abbott et al., 2017g; Goldstein et al., 2017), the details of the relationship remains unclear. The inferred on-axis lightcurve of GW170817 is consistent with the known population of short GRBs occuring at cosmological distances (Wu & MacFadyen, 2019; Salafia et al., 2019), suggesting that all short GRBs may be produced by the same mechanism and the observed diversity in the short GRB population is caused by extrinsic properties including the viewing angle (Lipunov et al., 2001; Rossi et al., 2002; Fong et al., 2015; Salafia et al., 2020). Direct measurement of viewing angle and outflow energetics of a larger sample of events using VLBI observations will allow this claim to be tested and place tight constraints on the short GRB luminosity function.

VLBI observations of centroid motion can also be combined with radio lightcurve monitoring to infer the opening angle of the jet produced by the merger, as was done for GW170817 (Mooley et al., 2018b). A sample of events with measured jet opening angles will constrain the inverse beaming fraction of the GRBs, and thereby establish whether neutron star mergers are responsible for the entire short GRB population. Understanding the typical jet opening angle will also improve estimates of the rate of joint GRB-GW detections (Howell et al., 2019; Beniamini et al., 2019), and inform future multi-messenger observing strategies.

5.5.3 Hubble Constant

We will be able to detect centroid motion in most events accessible with the HSA within 80 Mpc, while the effective range for the rest of the sky is 20 Mpc. Using the inferred neutron star merger rate (assuming a Gaussian mass distribution) from the first two LIGO observing runs ($1090^{+1720}_{-800}\,\mathrm{Gpc}^{-3}\mathrm{yr}^{-1}$; Abbott et al., 2020a), we find that in future observing runs 1–2 events per year will be useful for combined radio-GW measurements of H_0. A second event would reduce the uncertainty in the measurement from Hotokezaka et al. (2019) to $\sim 5\%$, but resolving the tension between distance ladder and CMB measurements requires a precision of $< 2\%$ which will not be achievable for decades.

In comparison, achieving this precision with only gravitational wave data from a localised event requires $\sim$ 100 more events based on the precision achieved for GW170817 (Abbott et al., 2017e). Including lightcurve modelling in this analysis can yield a improved precision (e.g. Guidorzi et al., 2017; Doctor, 2020; Dhawan et al., 2020), although the total number of events required to resolve the H_0 tension is still ~ 100.

Using the same inferred merger rate we expect to detect ~ 25 neutron star mergers per year during the fourth LIGO observing run based on a detector horizon of 190 Mpc, and $\sim$ 125 neutron star mergers per year in subsequent runs based on a detector horizon of 330 Mpc (Abbott et al., 2018). If every merger is localised to a host galaxy then a measurement with sufficient precision to potentially resolve the H_0 tension will be achieved within the next decade.

As we have found during LIGO/Virgo O3a[5] it is unlikely that every detected merger will be localised as easily as GW170817. We also note that the signal-to-noise of a merger detection is strongly dependent on merger inclination (5.18). Higher significance detections will have smaller localisation volumes and therefore face-on mergers occuring within the nominal 80 Mpc VLBI range will be more likely to be localised to a host galaxy. However, as more gravitational wave detectors come online

---

[5] O3a is the first part of the third LIGO/Virgo observing run, from April-October 2019

(Abbott et al., 2018), inclination angle measurements from gravitational wave data alone will improve and VLBI observations may become less useful. We caution that measurement of H_0 using a population of events with VLBI-constrained inclination angles requires careful consideration of selection biases (Mortlock et al., 2019), which are not yet well understood, but will be as our sample of EM-bright gravitational wave events grows.

The uncertainty in the peculiar motion of the host galaxy of GW170817 is one of the largest errors in the current combined radio-GW measurement of H_0 (Hotokezaka et al., 2019; Mukherjee et al., 2019; Howlett & Davis, 2020). However, this uncertainty can be significantly reduced if similar measurements are done for merger events at farther distances. As shown in Figure 5.8, the centroid motion can be measured by VLBI up to $\sim 100\,\mathrm{Mpc}$ in the case of the favorable density and viewing angle. Thus, detecting the centroid motion of the jet in such GW events will be particularly important for combined radio-GW measurements of H_0.

Another way to measure H_0, first proposed by Schutz (1986) is the 'dark siren' method where BBH localisation volumes are convolved with galaxy catalogues to get a probabilistic measurement of H_0 (Del Pozzo, 2012). This method was applied to GW170817 ignoring the knowledge of the host galaxy (Fishbach et al., 2019), and achieves a similar result to Abbott et al. (2017e). Soares-Santos et al. (2019) also apply this method to GW170814, and find $H_0 = 75^{+40}_{-32}\,\mathrm{km}\,s^{-1}\,\mathrm{Mpc}^{-1}$. Chen et al. (2018) find that this method will only achieve a precision of $\sim 10\%$ within the next decade, while Nair et al. (2018) find that third generation GW detectors may achieve a precision of 7% with only 25 BBH mergers. However, it will be possible to measure the redshift of BBH merger with third generation GW detectors to as good as 8% (Messenger & Read, 2012), which will allow direct measurement of H_0 from gravitational wave events alone. The degeneracy between redshift and merger chirp mass may also be overcome with a large population of mergers and observational constraints on mass distributions (Taylor et al., 2012b; Farr et al., 2019).

While the non-VLBI methods discussed above have larger uncertainties on a per-event basis, the larger number of available events reduces the uncertainty contribution from host galaxy peculiar velocities, which should be randomly oriented and therefore partially cancel each other out (Howlett & Davis, 2020). It will also be possible to combine different standard siren measurements together, and while VLBI measurements may only be possible for $\sim 10\%$ of localised mergers, they will contribute more than that to the sensitivity of the overall measurement.

In general, we caution that these estimates of detection rates have large uncertainties due to small number statistics, the large uncertainty in the neutron star merger rate, and the even larger uncertainty in the distribution of circum-merger densities. They also rely on the assumptions of our detectability outlined in Section 5.4.3 and the afterglow models from Section 5.2.

5.6 Conclusions

In this paper we have discussed prospects for constraining the properties of neutron star mergers through observations of scintillation-induced variability and by using high resolution VLBI measurements to both detect motion of the emission centroid and directly image outflow structure. We find that while VLBI observations provide more direct measurements of source properties they are only feasible for the very closest events, while the scintillation technique can be applied to most events detected with current GW detectors it only provides indirect constraints on source size. Both techniques probe different parts of the merger parameter space, with VLBI measurements suited to events occuring in less dense environments. Additionally, both techniques probe different timescales and therefore where possible should be used in conjunction with one another to completely understand the source structure as the afterglow evolves. We also discuss prospects for measuring H_0 and resolving the tension between current competing measurements and find that while gravitational waves provide a completely independent technique that does not rely on distance ladders or complex statistical inferences, it will likely take at least a decade to achieve a precision that is comparable with current techniques. This improvement in precision relies not only on observing a larger population of mergers with VLBI-constrained inclination angles, but improvements in both hydrodynamic jet models and gravitational wave detector calibration.

Chapter 6

Prospects for Gravitational Wave Follow-up in the Era of Next-Generation Radio Telescopes

Dougal Dobie, Tara Murphy, David L. Kaplan,
Kenta Hotokezaka, Juan Pablo Bonilla Ataides,
Elizabeth K. Mahony, Elaine Sadler

I produced all figures and tables, wrote all code and wrote all text with significant feedback from all co-authors. Kenta Hotokezaka provided the afterglow simulations.

The detection of gravitational waves from a neutron star merger, GW170817, marked the dawn of a new era in time domain astronomy. Monitoring of the radio emission produced by the merger, including high-resolution radio imaging, enabled measurements of merger properties including the energetics and inclination angle. In this work we summarise the capabilities of current and future radio and gravitational wave facilities. We consider 3 observing strategies to identify future mergers – targeting galaxies within the merger localisation, widefield follow-up and deep monitoring of known counterparts. We find that while planned radio facilities like the Square Kilometre Array will be capable of detecting mergers at gigaparsec distances, no facilities are sufficiently sensitive to detect mergers at the horizon of proposed third-generation gravitational wave detectors that would operate starting in the 2030s. We also outline other contributions radio observations can make including host galaxy studies, placing constraints on the geometry of the outflow and combined measurements of the Hubble Constant.

6.1 Introduction

The 2015 discovery of gravitational waves produced by a binary black hole merger by the Laser Interferometer Gravitational Wave Observatory (LIGO) kickstarted a new era of astronomy (Abbott et al., 2016b). Despite comprehensive follow-up observations carried out using telescopes across the electromagnetic spectrum, no electromagnetic counterpart was detected (Abbott et al., 2016e,d), nor were any coincident neutrinos (Adrián-Martínez et al., 2016). Almost two years later LIGO/Virgo detected gravitational waves from a neutron star merger (Abbott et al., 2017d), which was accompanied by the contemporaneous detection of a short gamma-ray burst (Abbott et al., 2017g; Goldstein et al., 2017). Approximately 10 hours later a candidate optical counterpart was discovered (Coulter et al., 2017; Soares-Santos et al., 2017; Valenti et al., 2017; Arcavi et al., 2017a; Tanvir et al., 2017; Lipunov et al., 2017) in the galaxy NGC 4993. Over the following hours multi-wavelength observations detected ultraviolet and near-infrared emission from the source (Abbott et al., 2017f), while X-ray and radio emission was detected 9 and 16 days post-merger respectively (Troja et al., 2017; Hallinan et al., 2017). The delay in detecting the X-ray and radio emission stems from different physical origins – the optical emission is produced by a kilonova, while the X-ray and radio emission originates from the relativistic outflow launched by the merger (Metzger & Berger, 2012).

Continued radio monitoring of GW170817 revealed a lightcurve that gradually rose over the following months (Alexander et al., 2017a; Mooley et al., 2018a; Margutti et al., 2018b), before peaking approximately 150 days post-merger (Dobie et al., 2018a) and rapidly fading (Alexander et al., 2018; Troja et al., 2018; Mooley et al., 2018d). This comprehensive monitoring campaign with radio, optical and X-ray telescopes has allowed tight constraints to be placed on the spectral and temporal evolution of the non-thermal afterglow (Hajela et al., 2019; Makhathini et al., 2020; Troja et al., 2020; Hajela et al., 2020). In turn, these constraints have enabled physical properties of the merger, including the energetics, circum-merger density and inclination angle, to be inferred (e.g. Nakar et al., 2018). Further constraints on the outflow geometry were obtained through Very Long Baseline Interferometry (VLBI) observations (Mooley et al., 2018b; Ghirlanda et al., 2019), while the non-detection of linearly polarised radio emission constrains properties of the outflow's magnetic field (Corsi et al., 2018b).

During the third LIGO/Virgo Observing run 56 candidate compact object mergers were detected, including 11 candidates that likely have at least one neutron star component[1]. However, no convincing electromagnetic counterparts were detected over the duration of the run, although Graham et al. (2020) claim to have detected an optical counterpart to the candidate binary black hole merger S190521g[2]. The reason for the lack of detections remains unclear, but may be related to insufficiently sensitive optical follow-up (Coughlin et al., 2020), or overestimates of the true luminosity function of kilonovae (Kasliwal et al., 2020).

---

[1] `https://gracedb.ligo.org/superevents/public/O3/`

[2] `https://gracedb.ligo.org/superevents/S190521g/view/`

There has been significant research into the detectability of the radio afterglow of compact object mergers (e.g. Feng et al., 2014; Hotokezaka et al., 2016; Lazzati et al., 2017; Duque et al., 2019; Gottlieb et al., 2019; Lin & Totani, 2020). The joint detection of GW170817 and GRB170817A confirmed the relationship between neutron star mergers and short GRBs (Beniamini et al., 2019; Wu & MacFadyen, 2019) that has been predicted for decades (Eichler et al., 1989; Narayan et al., 1992), and therefore studies into the detectability of short GRB radio afterglows are also of relevance (e.g. Ghirlanda et al., 2013, 2014; Metzger et al., 2015b; Zhang et al., 2015). In this work we build upon previous studies by quantifying prospects for detecting afterglows with all major current and future GHz-frequency radio facilities using observing strategies tailored to the specifications of each facility. We have expanded our study beyond current gravitational wave detectors, and compare the range to the horizon of detectors that will be built over the coming decades. We also outline the various synergies between radio telescopes and gravitational wave detectors beyond simple detection and monitoring of the afterglow.

In Section 6.2 we summarise potential observing strategies for radio follow-up and the ability of existing and planned radio telescopes to carry out follow-up. In Section 6.3 we outline the detection and localisation capabilities of existing and planned gravitational wave detectors. Section 6.4 includes a discussion of the benefits and limitations of radio follow-up, and quantifies prospects for the detection of radio counterparts using a simple afterglow model. In Section 6.5 we discuss the other contributions that radio observations can make in the broader context of multi-messenger astronomy including constraints on the properties of mergers and their environments, measurements of H_0 and shedding light on the nature of the central engine that drives the outflow.

6.2 Radio Telescope Capabilities

In this section we outline the capabilities of radio interferometers operating at GHz frequencies and the emission that is expected to be detectable with them. We do not consider using single-dish radio telescopes to search for counterparts due to their comparably low angular resolution. We do not include a detailed discussion of low-frequency ($\leq 300\,\mathrm{MHz}$) telescopes including the LOw-Frequency ARray (LOFAR; van Haarlem et al., 2013), the Murchison Widefield Array (MWA; Tingay et al., 2013), the Long Wavelength Array (LWA; Taylor et al., 2012a) and the Owens Valley Long Wavelength Away (OVRO-LWA) as the model we use does not account for synchrortron self-aborption and therefore does not accurately represent the expected radio emission at these frequencies. We also exclude the Canadian Hydrogen Intensity Mapping Experiment (CHIME; CHIME/FRB Collaboration et al., 2018) and the upgraded Molonglo Observatory Synthesis Telescope (UTMOST; Bailes et al., 2017) from this discussion, as their primary purpose is not traditional imaging. Table 6.1 lists the specifications of each telescope in our analysis.

Table 6.1: Capabilities of existing and planned radio facilities including observing frequency (ν), bandwidth ($\Delta\nu$), field of view (Ω), angular resolution (θ).

Facility	Band	ν (GHz)	$\Delta\nu$ (GHz)	Ω (deg^2)	θ (arcsec)	Dec. limit (deg)
ATCA[a]	C/X	8.0	8.0	0.01	2	$< +30$
GMRT	B3	0.4	0.2	1.4	8	> -50
	B4	0.7	0.3	0.4	4	> -50
VLA[b]	L	1.5	1.0	0.12	6	> -30
	S	3.0	1.5	0.06	2.7	> -30
	C	6.0	4.0	0.01	1.3	> -30
Apertif	L	1.4	0.3	6	15	> -20
ASKAP	Band 1	0.9	0.3	30	15	$< +30$
MeerKAT	L	1.4	0.7	0.8	7	$< +30$
DSA-2000	–	1.35	1.3	10.6	3.5	> -30
SKA-1	Band 2	1.43	0.4	0.8	0.6	$< +30$
ngVLA	Band 1	2.4	2.3	0.13	0.002	> -30
SKA-2	Band 2	1.43	0.4	0.8	0.6	$< +30$

[a] Assumes a 6 km array configuration

[b] Assumes B configuration

6.2.1 Afterglow modelling

To assess the detectability of radio afterglows we use the same power-law jet model as Dobie et al. (2020a), and we summarise the key characteristics of the model below. The kinetic energy is distributed according to

$$E(\theta) = \frac{E_{\rm iso}}{1 + (\theta/\theta_{j,c})^{3.5}}, \tag{6.1}$$

where θ is the polar angle from the jet axis, $\theta_{j,c}$ and $E_{\rm iso}$ are the half opening angle and isotropic-equivalent energy of the core of the jet respectively. Here we assume $\theta_{j,c} = 0.05\,{\rm rad}$, $E_{\rm iso} = 10^{52}\,{\rm erg}$ based on the afterglow of GW170817 (Makhathini et al., 2020; Troja et al., 2020). We use the standard synchrotron afterglow model of Sari et al. (1998) to calculate the radio flux produced by the radial expansion of a jet into a uniform medium. The initial Lorentz factor of the jet is given by

$$\Gamma(\theta) = 1 + \frac{\Gamma_c}{1 + (\theta/\theta_{j,c})^{5}}, \tag{6.2}$$

where Γ_c is the Lorentz factor of the jet's core. We assume the fraction of shock energy distributed to the electrons and magnetic field of the jet to be $\epsilon_e = 0.1$ and $\epsilon_B = 0.01$ respectively, and that electron energy distribution follows a power law with index $p = 2.16$. The observed flux density at a frequency ν also scales according to

$$S_\nu \propto E_{\rm iso}(\epsilon_B n)^{(p+1)/4}\epsilon_e^{p-1}, \tag{6.3}$$

while the temporal evolution scales proportional to $n^{-1/3}$. These relations can be used to scale the results of this work to generalised energetics and microphysics parameters. However, since this model does not consider synchrotron self-absorption, the predicted flux density is overestimated for $n > 1\,\mathrm{cm}^{-3}$ and $\theta_{\mathrm{obs}} < 10$ deg. We therefore do not consider parameters outside of these limits in this work.

We emphasise that this model is by no means comprehensive – it is a single, generalised, model for one of many components of the radio emission that may be produced by compact binary mergers (e.g., see Hotokezaka & Piran, 2015). Below we summarise some of the other forms of emission that may be expected, but are not considered further here.

At early times ($t \lesssim 10$ days post-merger) the dominant source of radio emission may be the 'reverse shock' (Sari & Piran, 1995; Resmi & Zhang, 2016; Lamb & Kobayashi, 2019), which propagates from the outflow towards the site of the merger. This emission is dependent on similar parameters to the forward shock, including the inclination angle. Reverse-shock emission has been detected in long GRBs (e.g. Kulkarni et al., 1999; Kobayashi, 2000; Kobayashi & Sari, 2000; Laskar et al., 2013; Perley et al., 2014; Alexander et al., 2017b; Laskar et al., 2019; Rhodes et al., 2020) but may be difficult to detect in short GRBs due to the required follow-up latency (Lloyd-Ronning, 2018). However, observations of GRB 160821B suggest that a reverse shock component is responsible for part of the early time radio emission (Troja et al., 2019b; Lamb et al., 2019b). We stress that even if reverse shock emission may be difficult to detect as a distinct component, its contribution to the overall emission may be significant enough that forward-shock jet models (including the one we use in this paper) alone may not accurately describe the early-time lightcurve evolution.

While most mergers likely launch a jet, the fate of the jet is dependent on factors including the energetics of the merger and the density of the surrounding environment. As the jet propagates into the surrounding ejecta, it may form a cocoon, which will produce a distinct signal as it expands and breaks out of the ejecta (Nagakura et al., 2014; Murguia-Berthier et al., 2014; Lazzati et al., 2017; Nakar & Piran, 2017; Gottlieb et al., 2018a). Beyond the initial rise and decline, this geometry may also produce a double-peak in the lightcurve (Barkov et al., 2018).

Kilonovae (e.g., Metzger et al., 2010; Metzger & Berger, 2012; Kasen et al., 2015; Metzger et al., 2015a; Barnes et al., 2016), the source of the early-time optical and infrared emission from GW170817 (Pian et al., 2017; Kasen et al., 2017; Smartt et al., 2017), may also give rise to radio emission. This emission is produced by the sub-relativistic expansion of the ejecta associated with the kilonova and is expected to be fainter, and peak at later times, than the non-thermal synchrotron emission. The timescale of this peak ranges from months to decades and the peak luminosity is similarly uncertain (e.g., Nakar & Piran, 2011; Piran et al., 2013; Hotokezaka et al., 2018; Kathirgamaraju et al., 2019b; Margalit & Piran, 2020). Kilonova ejecta are expected to expand quasi-isotropically, and therefore this form of radio emission may be detectable even when the emission associated with the relativistic ejecta is not, due to the inclination and jet opening angles of the merger.

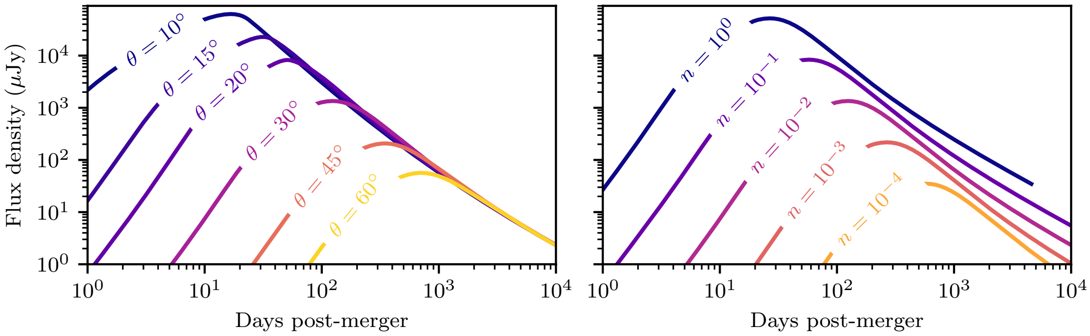

Figure 6.1: Comparison of lightcurve properties for a range of inclination angle, θ, and circum-merger density, n, assuming a merger distance of 40 Mpc and an observing frequency of 1.4 GHz for the merger parameters outlined in Section 6.2.1. Top: Lightcurves for a range of inclination angles, with $n = 10^{-2}\,\mathrm{cm}^{-3}$, corresponding to typical short GRB circum-burst density. Bottom: Lightcurves for a range of circum-merger densities, with $\theta = 30$ deg.

6.2.2 Follow-up Strategies

Targeting known galaxies

One survey strategy adopted in the follow-up of GW170817 was to target known galaxies in the localisation volume (e.g., Andreoni et al., 2017; Arcavi et al., 2017a; Coulter et al., 2017; Díaz et al., 2017; Evans et al., 2017; Valenti et al., 2017). This was also the predominant strategy used in radio follow-up, as at the time no facilities had sufficient survey speed to perform a widefield search.

Currently the best catalogues for this purpose include the Galaxy List for the Advanced Detector Era (GLADE; Dálya et al., 2018), Census of the Local Universe (CLU; Cook et al., 2019), Photometric Redshifts for the Legacy Surveys (PRLS; Zhou et al., 2020), the WISE × SuperCOSMOS Photometric Redshift Catalog (Bilicki et al., 2016), and the 2-MASS photometric redshift galaxy catalogue (2MPZ; Bilicki et al., 2014). GLADE and CLU are compilations of existing surveys (and encompass observations from 2MPZ), while PRLS uses the Dark Energy Camera Legacy Survey (Dey et al., 2019) and data from the *Wide-field Infrared Survey Explorer* (WISE; Wright et al., 2010). Additionally, the Mass AssociatioN for GRavitational waves ObserVations Efficiency (MANGROVE; Ducoin et al., 2020) combines GLADE with estimates of stellar mass from WISE, and Artale et al. (2020b) compute host probabilities from estimates of stellar mass and star formation rate.

Many works explore the prospects for this search method, e.g. Gehrels et al. (2016) demonstrated that simple galaxy targeting can reduce the required number of pointings by a factor of 10–100 compared to an untargeted search, while Coughlin et al. (2018) find that search efficiency can further be improved by a factor of 2. Some strategies expand on basic galaxy targeting by preferencing closer galaxies (Arcavi et al., 2017b), or by using assumptions about the afterglow emission (Coughlin & Stubbs, 2016; Salafia et al., 2017). Widefield telescopes may be able to observe multiple galaxies in a single pointing, and this follow-up may be optimised by convolving the gravitational wave skymap with a galaxy catalogue (Evans et al., 2016).

The limiting factor of this technique is the completeness of the survey catalogue, which for current surveys is lacking at distances comparable to the planned LIGO detector horizon of 330 Mpc (Abbott et al., 2018). While PRLS is deeper than the other surveys, it only covers $\sim 16\%$ of the sky and therefore cannot be utilised for all events. We discuss the issue of galaxy catalogue completeness and its impact on follow-up in 6.4.1. For more distant events it is therefore natural to consider searches that only use the localisation obtained from the gravitational wave signal.

Unbiased searches

Telescopes with larger survey speeds are capable of performing unbiased searches for radio counterparts by observing entire localisation regions. This technique has the obvious advantage of not being limited by the completeness of existing galaxy catalogues, while on the other hand, can result in coverage of a large amount of extraneous area and an increase in false-positives (which we discuss in Section 6.4.1). However, this can be partly mitigated by restricting candidates to known galaxies within the localisation volume.

There are multiple ways to optimise an unbiased search (e.g. Ghosh et al., 2016, 2017; Dobie et al., 2019a; Gupte & Bartos, 2020) but for the purposes of this work we simply assume that a telescope can observe a given sky area without any assumption of exact observing strategies.

Monitoring known counterparts

The final type of follow-up involves searching for radio emission from known counterparts detected at other wavelengths, which allows for maximised sensitivity as all available telescope time can be allocated to a single pointing.

The detection of a known counterpart can also help inform the required cadence and sensitivity of radio observations. Localising the counterpart to a galaxy provides accurate distance measurements (Hjorth et al., 2017) which can also be used to improve constraints on the merger inclination angle independent of afterglow models (Mandel, 2018). The detection of a short GRB can constrain the merger energetics (Abbott et al., 2017g; Goldstein et al., 2017), and a combination of multi-wavelength observations can be used to broadly infer the geometry of the merger (Kasliwal et al., 2017). Single-dish spectral line observations can also be used to constrain the density of the environment surrounding the merger (Hallinan et al., 2017).

Even if a radio counterpart is not detected, radio observations are still useful in placing constraints on the spectral properties of the afterglow (in combination with optical and X-ray follow-up) and inferring properties of the host galaxy, which we discuss further in Section 6.5.3.

6.2.3 Existing Facilities

Australia Telescope Compact Array (ATCA)

The Australia Telescope Compact Array is an East-West array consisting of six 22-m dishes with a maximum baseline of 6 km (depending on the array configuration). Each dish is equipped with a set of receivers that can sample frequencies from 1.1 to 105 GHz, with the System Equivalent Flux Density (SEFD) increasing towards higher frequencies (Wilson et al., 2011). In this paper we consider observations in the C/X band (4–12 GHz) which is generally the most sensitive due to the high Radio Frequency Interference (RFI) occupancy at lower frequencies[3].

Currently correlation is performed using the Compact Array Broadband Backend (Wilson et al., 2011), which allows observations in 2×2048 MHz windows. However, the Broadband Integrated-GPU Correlator for the Australia Telescope (BIGCAT) upgrade, which will be completed by December 2021, will double the available bandwidth to 8 GHz. As this upgrade will occur on a similar timeline to the start of the next LVC observing run, in this paper we assume that the upgrade is successfully commissioned and therefore the sensitivity will improve by a factor of $\sqrt{2}$ compared to the current level.

It is not feasible to carry out untargeted follow-up with the ATCA due to its low survey speed, even using on-the-fly mosaicing. However, it is useful for targeted follow-up, either of galaxies within the localisation volume of the gravitational wave event, or of counterparts detected at other wavelengths. Taking into account observing overheads, it's possible to observe ~ 50 galaxies to a detection threshold of $< 70\,\mu$Jy in a 12-hour observation (12 minutes per source), while a detection threshold of $\sim 15\,\mu$Jy can be achieved for follow-up of a single source. We note that in many cases this can be improved by optimising pointings such that multiple galaxies are observed within the same field of view[4], but to simplify the estimates in this paper we assume that we can simply observe 50 galaxies with 1 per pointing.

Karl. G. Jansky Very Large Array (VLA)

The Karl. G. Jansky Very Large Array (VLA) consists of twenty seven 25-m antennas with maximum baselines ranging from 36.4 km (A configuration) to 1.03 km (D configuration). In this paper we consider the B configuration which has baselines spanning 0.21–11.1 km, although the A and C (3.4 km maximum baseline) configurations would also be suitable. The telescope can observe at frequencies from 58 MHz to 50 GHz, but in this case we consider observations at S-band (3 GHz; offering the best sensitivity for the expected negative spectral index) and L-band (1.5 GHz; offering lower sensitivity but a twice-larger field of view).

Rana & Mooley (2019) propose an optimised galaxy targeting strategy for the VLA, improving the probability of detecting a radio counterpart by a factor of two compared to a simple approach. This strategy will enable observers to cover approxi-

---

[3] `https://www.narrabri.atnf.csiro.au/observing/rfi/monitor/rfi_monitor.html#atca`

[4] `https://github.com/ddobie/atca-ligo`

mately 200 galaxies to a sensitivity of 15 μJy. The VLA can also carry out untargeted searches using on-the-fly mosaicing, where the antennas are driven at a constant rate along a strip of sky, eliminating slew overheads which are dominated by the settling time of $\sim$ 7 seconds per pointing. This technique has already been applied to follow-up of two gravitational wave events, GW151226 (Abbott et al., 2016c) and GW190814 (Abbott et al., 2020b) covering 100 deg^2/10% (Mooley et al., 2018c) and 5 deg^2/50% (Mooley et al., 2019) of the localisation regions respectively.

Giant Metrewave Radio Telescope (GMRT)

The Giant Metrewave Radio Telescope (GMRT) consists of thirty 45-m antennas with a maximum baseline of 25 km. Recent upgrades to the receivers give the GMRT unparalleled sensitivity at sub-GHz frequencies, and here we consider observations in bands 3 and 4, centered on 400 and 700 MHz respectively. Observations in these lower bands are vital in understanding the spectral evolution of the afterglow, in particular, the evolution of the synchrotron self-absorption frequency.

Here we consider the GMRT as a dedicated follow-up instrument for localised mergers, as its combination of sensitivity and field of view is not conducive to untargeted searches. We consider 3 hour observations in both Band 3 and 4, corresponding to an image sensitivity of 15 and 20 μJy respectively.

Arcminute Microkelvin Imager Large Array (AMI-LA)

The Arcminute Microkelvin Imager Large Array (AMI-LA) consists of eight 12.8m antennas with a maximum baseline of 110 m operating at an observing frequency of 15 GHz. The AMI-LA Rapid Response Mode (ALARRM; Staley et al., 2013) enables the telescope to respond to GRB alerts within 2 minutes of the burst, allowing for tight constraints to be placed on early-time emission (e.g. Lamb & Kobayashi, 2019). However, the AMI-LA field of view limits the practicality of early-time follow-up of neutron star mergers to those with a simultaneous detection of a GRB by the *Neil Gehrels Swift Observatory*. While AMI-LA can also be used to monitor known counterparts, we do not consider it in our detectability analysis in Section 6.4.2, and omit it from Table 6.1, due to the superior sensitivity of the VLA.

6.2.4 Next-generation Facilities

Australian Square Kilometre Array Pathfinder (ASKAP)

The Australian Square Kilometre Array Pathfinder (ASKAP; Johnston et al., 2008, Hotan et al., submitted) is an array of thirty-six 12-m dishes with baselines ranging from 37 m to 6 km. ASKAP is designed for all-sky surveys between 700-1800 MHz (in this case we consider observations at 900 MHz) and uses MkII phased-array feeds (Hampson et al., 2012) consisting of 36 beams resulting in a 30 deg^2 field of view. While the SEFD is higher than other comparable telescopes, the large field of view results in a high survey speed, making it possible to search large areas of sky for a gravitational wave counterpart. However, this does mean that ASKAP is not useful for monitoring events where an electromagnetic counterpart has already been discovered. We therefore only consider the utility of ASKAP for localising events, and not monitoring them.

Similarly, ASKAP is not particularly useful for a galaxy-targeted search strategy since tens–hundreds of candidate host galaxies can be covered with a single pointing. Some widefield telescopes select their pointing strategy by convolving the localisation skymap with galaxy catalogues (e.g., Evans et al., 2016) but in Dobie et al. (2019a) we demonstrated that this strategy does not produce any appreciable benefits for ASKAP follow-up.For this paper we consider two strategies, a single deep ~ 10 hour pointing or a widefield strategy consisting of four ~ 3 hour pointings for events localised to $< 100\,\mathrm{deg}^2$ (corresponding to $\sim 60\%$ of neutron star mergers during O4, and the majority of mergers detected with 3G detectors Abbott et al., 2018; Hall & Evans, 2019). So far ASKAP has performed follow-up of S190510g and GW190814 (Dobie et al., 2019b, Dobie et al. in prep.) and both times the latter pointing strategy was used. We note that due to the comparably low angular resolution of ASKAP, some afterglows may be contaminated by emission from the host galaxy. We discuss this problem in Section 6.4.1.

Apertif

The Westerbork Synthesis Radio Telescope (WSRT) is an East-West array consisting of fourteen 25-m antennas with a maximum baseline of 2.7 km. Twelve of the antennas have been fitted with L-band (operating from 1–1.75 GHz) PAFs as part of the *APERture Tile In Focus* (Apertif) project, improving the telescopes field of view to $9.5\,\mathrm{deg}^2$ (Oosterloo et al., 2010; Adams & van Leeuwen, 2019). Similar to ASKAP, Apertif's wide field of view means it is more suited to widefield searches rather than targeted follow-up, while the observing frequency and maximum baseline mean it is also subject to the same host galaxy contamination issues. The sensitivity of Apertif is comparable to ASKAP, and we consider follow-up with one deep 12 hour pointing achieving a sensitivity of $25\,\mu$Jy or four pointings achieving a sensitivity of $50\,\mu$Jy.

MeerKAT

MeerKAT (Jonas et al., 2016) consists of sixty four 13.5-m diameter antennas with a maximum baseline of 8 km. fitted with 1.4 GHz receivers and has a $0.8\,\mathrm{deg}^2$ field of view. The receivers have a system equivalent flux density of ~ 430 Jy (Mauch et al., 2020), making MeerKAT the premier facility for searching for emission from known counterparts. Here we consider two follow-up strategies – a single deep 10 hour pointing achieving a sensitivity of $\sim 2\,\mu$Jy and ten 1 hour pointings covering $10\,\mathrm{deg}^2$ to a sensitivity of $\sim 7\,\mu$Jy. The former strategy applies to very well localised events (comparable to expectations for third generation detectors, see 6.3.2) and monitoring of known counterparts, while the latter strategy applies to untargeted follow-up of events localised to $\lesssim 10\,\mathrm{deg}^2$ or targeted follow-up of less-localised events using a galaxy catalogue convolution strategy (Evans et al., 2016).

6.2.5 Future Facilities

Square Kilometre Array (SKA)

The Square Kilometre Array (SKA) will be the worlds largest radio telescope, with a mid-frequency array in South Africa and a low-frequency array in Western Australia. Here we focus on the mid-frequency array, which will be split into two stages.

SKA-1 mid will consist of the existing MeerKAT array and an additional one hundred and thirty three 15-m dishes, with the array expected to come online in the mid-2020s. Braun et al. (2019) outlines the anticipated array performance, and for the purposes of this paper we consider observations in Band 2, centered on 1.43 GHz as a compromise between maximising sensitivity and the expected negative spectral index of gravitational wave afterglows. We assume the same observing strategy as MeerKAT, with a 1 hour continuum sensitivity of 2 μJy.

We also consider observations with the SKA-2. While the design specifications for the SKA-2 are still uncertain, here we assume an order of magnitude sensitivity improvement over SKA-1 for the same observing strategies.

Next Generation Very Large Array (ngVLA)

The Next Generation Very Large Array (ngVLA; Murphy et al., 2018) is a planned replacement for the VLA operating between 1.2 and 116 GHz. The main array will consist of two hundred and fourteen 18-m dishes on baselines of up to 1000 km. We consider observations in the lowest frequency band, centered on 2.4 GHz due to the larger field of view and higher relative sensitivity (Selina et al., 2018). It is currently anticipated that Early Science observations will begin in 2028 with full operation from 2034 onwards[5].

The ngVLA will be capable of performing untargeted follow-up of well-localised events, observing 10 deg^2 to a sensitivity of 1 μJy in 10 hours (Corsi et al., 2019a). It has comparable instantaneous sensitivity to the SKA, making it the premier northern hemisphere facility for targeted follow-up of known counterparts. The ngVLA is also capable of observing at frequencies up to 93 GHz, compared to the SKA which will only observe up to 12.5 GHz (Braun et al., 2019). This will allow a better characterisation of the spectral properties of afterglows in conjunction with optical and X-ray telescopes, as was performed for GW170817 (Makhathini et al., 2020; Troja et al., 2020).

Deep Synoptic Array (DSA-2000)

The Deep Synoptic Array (DSA-2000) is a proposed telescope that will consist of two thousand 5-m dishes capable of simultaneously observing between 0.7-2 GHz and is scheduled to be fully operational by 2026 (Hallinan et al., 2019). The telescope is optimised for survey speed, with a 10.6 deg^2 instantaneous field of view and 2.5 Jy SEFD. Approximately 1 hour per day will be allocated to follow-up of gravitational wave events, with a focus on events that can be covered with a single pointing. In

---

[5] `https://ngvla.nrao.edu/page/faq#faq_16_content`

these cases, the entire localisation region can be covered to a 5σ detection threshold of $\sim 5\,\mu$Jy. Similar to ASKAP, we also consider a widefield strategy of twelve ~ 5 minute pointings for follow-up of events localised to ~ 100 deg^2.

While the DSA-2000 can cover a 10 deg^2 localisation ten times faster than the SKA or ngVLA, it is important to note that both telescopes can cover that area more efficiently. Gravitational wave localisations are generally irregular shapes and often multi-modal, while telescope fields of view are either circular or rectangular. Ghosh et al. (2016) find that it is significantly more efficient to cover localisation regions with a distributed group of multiple small-FoV telescopes than a single widefield telescope due to the lower extraneous coverage. Similarly, the SKA/ngVLA will use many small-FoV pointings and will therefore achieve more efficient coverage than the DSA-2000 strategy of fewer large-FoV pointings.

6.2.6 Serendipitous Observations

Advances in radio telescope technology will allow for numerous widefield surveys of the radio sky will be undertaken in the coming decades. Deep all-sky surveys will provide sensitive reference images for transient follow-up, and widefield transient searches will likely provide serendipitous coverage of gravitational wave events.

ASKAP

The Evolutionary Map of the Universe (EMU; Norris et al., 2011) will cover the sky South of +30 deg to a sensitivity of $10\,\mu$Jy at 1.3 GHz. This will provide the most sensitive map of the Southern radio sky to-date and be useful as a reference image for transient searches. The Rapid ASKAP Continuum Survey (RACS; McConnell et al., 2020) achieves a typical sensitivity of $250\,\mu$Jy at 900 MHz and has already been used as a reference image in follow-up of GW190814 (Dobie et al., 2019b).

The proposed Variables And Slow Transients (VAST; Murphy et al., 2013) survey is split into three main components, Wide, Deep and Galactic, and will span at least 5 years. VAST-Wide will observe an area of 10000 deg^2 to a detection threshold of 2.5 mJy on a daily cadence. VAST-Deep will achieve a detection threshold of $250\,\mu$Jy and will observe 10000 deg^2 7 times, and a single 30 deg^2 field daily. VAST-Galactic will observe 750 deg^2 of the galactic plane 64 times to a detection threshold of $500\,\mu$Jy, which will be useful for events where optical follow-up is hindered by extinction and a high rate of unrelated transients. Here we consider the VAST-Wide and low cadence VAST-Deep surveys, as the wide areal coverage makes them more conducive to this kind of search. We note that the ASKAP Survey Science observing strategy has not yet been finalised, and therefore the exact observing strategy used may vary.

MeerKAT

While MeerKAT has a higher instantaneous sensitivity than ASKAP, there are currently no plans to use it for a dedicated widefield untargeted transients survey (Fender et al., 2017). Instead, untargeted searches will be conducted using commensal data from other surveys, which we outline below

- The LADUMA survey (Holwerda et al., 2012) will observe a single 1 deg^2 field for 3424 hours, providing high sensitivity and a long time baseline to search for transients;
- The MeerKAT Fornax survey (Serra et al., 2016) will cover 12 deg^2 with two 4.5 hour epochs and a final sensitivity of 2 μJy;
- MHONGOOSE (de Blok et al., 2016) will cover a non-contiguous 30 deg^2 with 6–8 deep pointings and a final sensitivity of 0.15 μJy;
- MeerTIME (Bailes et al., 2018) will observe 1000 pulsars $\sim$ 20 times per year to a sensitivity of $\sim$ 75 μJy, while also observing 28 globular clusters less regularly to a sensitivity of 10 μJy.

The LADUMA, Fornax and MHONGOOSE surveys will observe at L-band ($\sim$ 1.4 GHz), while the MeerTIME observing strategy is not yet finalised. Overall, the proposed transients search with MeerKAT will search a smaller area of sky to a greater sensitivity compared to ASKAP. As this is less conducive to serendipitous coverage of multi-messenger events, we instead consider an idealised untargeted radio transient survey with 9 epochs separated by 4 months covering 5000 deg^2 to a sensitivity of 20 μJy. This corresponds to a total observing time of 3750 hours.

Deep Synoptic Array

The Cadenced All-Sky Survey (Hallinan et al., 2019) will observe 16 epochs of the sky North of -30 deg on a 4-month cadence to a detection threshold of 10 μJy, providing coverage of the majority of gravitational wave events. It will also ultimately provide a reference image with an rms noise of 500 nJy.

6.3 Gravitational Wave Detectors

6.3.1 Second Generation Detectors

Abbott et al. (2018) outlines current plans for future gravitational wave observing runs with current detectors which we summarise here for completeness.

Fourth Observing Run (O4; 2022-2023)

The Fourth Observing Run (O4) will run for one year with both LIGO detectors close to design sensitivity (190 Mpc). Advanced Virgo will have a binary neutron star range of 90–120 Mpc (comparable to the LIGO detector ranges during O3), while the sensitivity of KAGRA (KAGRA Collaboration et al., 2019, 2020) has a large uncertainty and estimates for its binary neutron star range is 25–130 Mpc. Assuming a KAGRA range of 80 Mpc the estimated number of detections is 10^{+52}_{-10}, with a median 90% localisation of 33 deg^2. We adopt a sky and inclination angle averaged detection range of 160 Mpc based on the minimum specifications for the LIGO detectors and reflecting the lower range of the other two detectors.

Table 6.2: Capabilities of gravitational wave detector networks made of the Hanford (H), Livingston (L), Virgo (V), Kagra (K), LIGO-India (I) detectors. Detectors improved by the A+ upgrade are denoted by a subscript + while LIGO-Voyager detectors are denoted by a subscript V.

Epoch	Facilities	Timeline	Range[a] (Mpc)	Localisation[b] (deg^2)
O4	H, L, V, K	2022–2023	190	35
O5	H_+, L_+, V_+, K	2025–2026	330	35
2G	H_+, L_+, V_+, K, I_+	2026	330	35
Voyager	H_V, L_V, V_V	2030	1100	70
3G	ET, CE, Voy	2040	10^5	10
	ET, 2CE		10^5	1

[a] Maximum range of any detector in the network
[b] Order of magnitude estimate for typical localisation

Fifth Observing Run (O5; 2024-2025)

The fifth observing run will begin after the A+ upgrade, which will increase the LIGO detector range to 330 Mpc. The Virgo detector will also undergo significant upgrades, and will operate with a binary neutron star range of 150–260 Mpc and KAGRA will operate with a range of at least 130 Mpc and possibly as high as 155 Mpc. For the purposes of this paper we assume a sky and inclination angle averaged detection range of 300 Mpc.

A Five Detector Network (2025+)

LIGO-India (Iyer et al., 2011) is expected to join operations in 2025 with a range of 330 Mpc. The geographical location of LIGO-India improves the localisation capability of the global network by a factor of 5–10, and in some cases may result in mergers being localised to areas as small as 1 deg^2.

6.3.2 Third Generation Detectors

LIGO Voyager (LIGO Scientific Collaboration, 2019) is a planned upgrade to the three existing LIGO facilities that will increase their range to 1.1 Gpc, and also improve the localisation of closer events. This upgrade is expected to occur by the end of the decade, and we assume an operational start date of 2030. Voyager is an intermediate step between the Advanced LIGO detectors and third generation detectors, however due to the significant improvements in the nominal detector range and the uncertainty in the design and timeline of the detector, we consider it a third generation detector for the purposes of this work.

The proposed Neutron Star Extreme Matter Observatory (NEMO; Ackley et al., 2020a) will bridge the gap between LIGO A+ and true third-generation detectors. While the addition of this detector will not significantly improve network sensitivity or localisation capabilities, it will make the detection of post-merger gravitational waves feasible by extending the observable frequency range. This will provide insight into the nature of merger remnants (Abbott et al., 2019d), and the properties of the jet produced by the merger (e.g., Beniamini et al., 2020). Constraints on both of these will inform both long-term and prompt radio follow-up efforts.

True third generation gravitational wave detectors will be an order of magnitude more sensitive than current detectors due to reduced quantum shot noise, improved mirror coatings and the placement of the detectors deep underground to reduce Newtonian noise. Both proposed detectors will be triple Michelson interferometers (Freise et al., 2009), as opposed to the L-shaped interferometers used in second generation detectors, which enables the measurement of the gravitational wave polarisation. The Einstein Telescope (ET; Punturo et al., 2010) will have 10 km arms, while Cosmic Explorer (CE) will have 40 km arms resulting in a higher sensitivity (Abbott et al., 2017a).

While a single third generation detector will be capable of detecting compact binary coalescences, precise localisation requires multiple detectors. Hall & Evans (2019) outline various observing scenarios based on combinations of Voyager, ET and CE detectors. For the purposes of this work we consider three simplified scenarios - the proposed network of three Voyager detectors, one Voyager detector with two 3G detectors (likely ET and CE) and three 3G detectors (likely ET and two CEs).

A network of one ET detector and two CE detectors will localise $> 10\%$ of neutron star coalescences at $z = 0.3$ (corresponding to $\gtrsim 50$ events per year) to $\sim 0.1\ \mathrm{deg}^2$ with a median localisation of $1\ \mathrm{deg}^2$ (Hall & Evans, 2019).

6.3.3 Space-based detectors

Space-based gravitational wave detectors operate on the same principle as their ground-based counterparts, but can operate at significantly lower (sub-Hz) frequencies, making them sensitive to massive black hole mergers and extreme mass ratio inspirals (e.g. the capture of a neutron star by a massive black hole). While electromagnetic follow-up of these newly accessible merger types is vital for understanding them, we currently lack quantitative predictions for their radio emission. We therefore exclude these from our analysis in Section 6.4, and instead qualitatively summarise other novel ways radio telescopes can be used in conjunction with space-based detectors.

There are currently two proposed space-based detectors, both consisting of 3 spacecraft arranged in an equilateral triangle carrying 2 W lasers. The Laser Interferometer Space Antenna (Amaro-Seoane et al., 2017) will use 2.5×10^9 m baselines, carrying 30 cm diameter telescopes. LISA is scheduled to launch in 2034 and begin science operations 2.5 years later. The Taiji Program in Space (Hu & Wu, 2017) will use longer baselines (3×10^9 m), and larger telescopes (40 cm) (Ruan et al., 2018a). Taiji is currently scheduled to launch in 2033. Ruan et al. (2020) propose operating both facilities as a network, thereby improving the localisation of mergers.

Kyutoku et al. (2019) have found that the number of pointings required for an all-sky search for binary pulsars with the SKA can be decreased by a factor of 10^6 by using neutron star binary localisations from LISA. Lau et al. (2020) and Andrews et al. (2020) estimate that LISA will detect 35–240 binary neutron stars, with the discrepancy being caused by differing assumptions for the merger rate of binary systems in the Milky Way. Joint radio-GW observations of short-period LISA binaries will enable the neutron star equation of state to be constrained, and determine the mass-radius relation to a precision of 0.2% (Thrane et al., 2020)

Space-based detectors will also allow for multi-band gravitational wave astronomy, providing early estimates of localisation and merger time for binary black holes that eventually merge in the LIGO band years later (Sesana, 2016). This early warning system will enable telescopes to be on-source as the merger happens, allowing theories linking Fast Radio Bursts (FRB; Lorimer et al., 2007; Thornton et al., 2013) to black hole mergers (e.g. Zhang, 2016; Liu et al., 2016) to be tested. We discuss prompt follow-up in Section 6.5.5.

6.4 Searching for Radio Afterglows

6.4.1 Benefits and Limitations

Advantages over other wavelengths

The ultraviolet/optical/infrared luminosity of kilonovae is dependent on parameters including the mass and velocity of the ejecta and the fraction of lanthanides produced (e.g. see Metzger, 2017b, and references therein). Additionally, in higher mass mergers (including neutron star-black hole mergers), the blue component of the kilonova that was vital in localising GW170817 may not be produced and even the red component may be suppressed (Hotokezaka et al., 2013b; Kasen et al., 2017). The lack of optical counterparts detected during O3 suggests that not all events will produce kilonovae comparable to GW170817 – Kasliwal et al. (2020) found the probability of zero detections in comprehensive optical follow-up of 13 mergers is only 4% if all kilonovae are as bright as GW170817, although Coughlin et al. (2020) suggest that these observations were not sensitive enough to constrain merger properties. Radio afterglows are mostly independent of these parameters, and therefore probe a different part of the merger parameter space to optical observations.

While the flux density of a radio afterglow is dependent on viewing angle, even significantly off-axis mergers will produce detectable emission. Conversely, GRBs are highly anisotropic and therefore most mergers will not produce a detectable GRB counterpart. In fact, the true rate of GRBs may be as much as 10^4 times higher than what we detect (Soderberg et al., 2006), and radio observations are one method of detecting off-axis bursts (Levinson et al., 2002). However, the inclination angle dependence of GW detector sensitivity means that the fraction of GW events with a GRB counterpart is higher than the fraction of detectable GRBs from the general population, and the rate of joint detections will be a few per year for O4 and tens per year for designed LIGO specifications (Howell et al., 2019; Saleem, 2020).

Extrinsic factors such as dust extinction and solar angle may also limit follow-up at other wavelengths. For example, comprehensive optical follow-up of GW170817 would not have been possible had the merger occured a month later, and even X-ray monitoring was hindered by ~ 90 days of sun avoidance. Radio telescopes are not limited by either of these factors.

False-positive rate

Radio searches are hindered by the discovery of false positives in the form of variable sources manifesting as transients, as well as unrelated transients. Previous untargeted searches for radio transients over long ($> 1\,$d) timescales (Mooley et al., 2013, 2016, O'Brien et al. in prep., Stewart et al. in prep.)have been dominated by variable Active Galactic Nuclei (AGN). Some variable AGN will be straightforward to immediately classify, but there exist many that cannot be classified without comprehensive broadband observations (e.g. Sadler et al., 2006; Lovell et al., 2008; Nyland et al., 2020). Searches will also be hindered by sources exhibiting extrinsic variability caused by interstellar scintillation, which can cause variabiliy of tens of percent at GHz frequencies (Cordes & Lazio, 2002), resulting in the same compact source being undetected in one epoch and detected in the next. This can be mitigated by comparing candidate counterparts to galaxy catalogues and ruling out any that are spatially consistent with the nucleus – as discussed in Section 6.5.1, neutron star mergers are likely to occur away from the galaxy nucleus. However, some telescopes may not provide sufficient astrometric accuracy to do this, necessitating follow-up observations with other facilities.

Searches consisting of multiple short-integration pointings may also discover short-duration radio transients such as flare stars (Helfand et al., 1999; Villadsen & Hallinan, 2019, Pritchard et al. in prep.) or even Fast Radio Bursts, although these will be easily ruled out as unrelated with follow-up observations.

As radio follow-up observations become more sensitive, we will also discover afterglows from other transients including tidal disruption events and a variety of GRBs. At a detection threshold of $10\,\mu$Jy the expected areal density of radio transients is a few per deg^2 at GHz frequencies (Metzger et al., 2015b). Since the emission from most radio transients originates from a synchrotron blast wave, even broadband radio observations may not be sufficient to rule out false positives. Instead long-term monitoring to determine the temporal evolution of the source will be required.

While widefield optical searches discover thousands of false positives, ruling them out is made easier by having a large sample of known optical transients with more distinct spectra and underlying physics, enabling the use of machine learning techniques for immediate and automatic classification (e.g. Mahabal et al., 2008; Bloom et al., 2012; Goldstein et al., 2015; Mahabal et al., 2019; Stachie et al., 2020). Comparably, few real radio transients have been detected in untargeted searches to-date (Stewart et al., 2016; Law et al., 2018) and current radio transient surveys rely on manual inspection of candidates and classification using follow-up observations and archival data. Planned widefield transients searches will allow the better characterisation of transient properties, which will ultimately enable the use of automated classification algorithms.

It is naturally preferable to confirm the association between any detected radio transient as soon as possible. It is possible (albeit, unlikely) that the radio afterglow may be detected early enough that optical emission from the kilonova is still detectable, observations of which are vital in constraining merger properties like the ejecta mass. More importantly, comprehensive broadband monitoring of the non-thermal afterglow that traces the rise, peak and decline of the lightcurve, is vital in constraining merger properties which we discuss in detail in Section 6.5.1.

Overall, we strongly emphasise the importance of designing follow-up strategies that are not only sufficient to detect afterglows, such as the logarithmic monitoring proposed by Hotokezaka et al. (2016), or the more detailed strategy in Carbone & Corsi (2018), but also use a cadence that enables false-positives to be ruled out in a timely manner.

Mergers with radio-loud hosts

Resolving the afterglow from any host galaxy emission may be a decisive factor in whether the afterglow is detectable, particularly in a widefield untargeted search. GW170817 occurred in a radio-loud host galaxy and was offset by 10.31 arcsec (Blanchard et al., 2017; Levan et al., 2017a) and therefore non-standard sourcefinding techniques could have been required to find it in an unbiased search depending on the angular resolution of the data. Host galaxy offsets for short GRBs range from 0.5-75 kpc with a median of 5 kpc (Fong et al., 2010; Fong & Berger, 2013; Berger, 2014). At 200(500) Mpc this corresponds to an angular offset of 10(4) arcsec, comparable to the angular resolution of some current radio telescopes. For events occuring at cosmological distances, detectable with third generation detectors (see 6.3.2), typical offsets will be comparable to the angular resolution of the SKA. Therefore the presence of nuclear emission from the host may complicate searches, but not make them impossible. However, radio emission from star formation regions can span a much larger volume (Linden et al., 2020) and may pose more of a problem.

By comparing the expected afterglow luminosity to the typical luminosity of both AGN and star forming galaxies, Hotokezaka et al. (2016) found that most mergers will not occur in galaxies that are sufficiently radio-bright to hinder radio follow-up efforts. This is true even for telescopes with angular resolution $\gtrsim 10$ arcsec. The detectability metrics we use in this work assume that future events either occur in radio-quiet galaxies, are significantly brighter than any host emission, or are sufficiently offset from the nucleus for this to not be a limiting factor.

Galaxy catalogue completeness

Most current radio telescopes lack the field of view and survey speed to carry out untargeted searches and are therefore restricted to the galaxy-targeted approach outlined in Section 6.2.2. The effectiveness of this strategy relies upon having a complete catalogue of galaxies within the localisation volume of each event.

There is currently no all-sky galaxy catalogue that approaches completeness at the current LIGO horizon. GLADE (Dálya et al., 2018) is complete to a distance of $\sim$

40 Mpc, $\sim 50\%$ complete at the nominal O4 horizon of 170 Mpc and $< 40\%$ complete at the design horizon of 330 Mpc based on cumulative B-luminosity. Additionally the galaxies in GLADE are not isotropically distributed – the median line of sight density is ~ 10 deg^{-2}, compared to < 1 deg^{-2} in the Galactic plane and $> 10^3$ deg^{-2} in fields covered by the HyperLEDA survey (Makarov et al., 2014).

Other existing surveys provide more complete samples along particular lines of sight, but not the entire sky. These smaller catalogues may still be useful for follow-up of specific events (e.g. the galaxy-targeted follow-up of GW190814 carried out by Gomez et al. (2019); Ackley et al. (2020b) could have used the PRLS survey rather than GLADE) but do not present a general solution for follow-up of all events.

Planned all-sky surveys (e.g. LSST Science Collaboration et al., 2009; Doré et al., 2016, among others) will drastically improve upon existing catalogues and make it feasible to target galaxies for events at Gpc distances. However, these surveys will not be finished within the near future, by which time facilities like ASKAP and MeerKAT will be fully operational and the DSA-2000 and SKA may be coming online. These telescopes are not suited to galaxy targeting due to their wide fields of view (compared to most existing facilities), and therefore it's unlikely that future galaxy catalogues will have a tangible impact on the pointing strategies used in radio follow-up, although they will be useful in determining whether discovered transients are associated with host galaxies within the localisation volume.

6.4.2 Follow-up of known mergers

To quantify prospects for detecting radio afterglows we use a simple detection metric – the flux density at the observing frequency must exceed five times the expected thermal noise, σ. This estimate does not include the noise due to source confusion, which will not be the limiting factor for the observations discussed here. To scale our model from the nominal rest frequency of our model, ν_0, to the observing frequency, ν, we assume the emission obeys a simple power law with spectral index $\alpha = (1-p)/2$. Since the model flux density, S_0, is calculated in the rest frame of the merger we also apply the standard K-correction (see e.g. Novak et al., 2017), resulting in a frequency-corrected flux density of

$$S = \frac{S_0}{(1+z)^\alpha}\left(\frac{\nu}{\nu_0}\right)^\alpha \tag{6.4}$$

where we convert between redshift, z, and luminosity distance, D_L, using ASTROPY.COSMOLOGY.WMAP9 (Hinshaw et al., 2013). We also compare our results to the horizon of gravitational wave detectors which is dependent on inclination angle, scaling as

$$\mathcal{R}(\theta_{\rm obs}) \approx 0.658\overline{\mathcal{R}}\sqrt{1+6\cos^2\theta_{\rm obs}+\cos^4\theta_{\rm obs}} \tag{6.5}$$

where $\mathcal{R}(\theta_{\rm obs})$ is the inclination angle dependent range and $\overline{\mathcal{R}}$ is the sky-averaged detector horizon found in Table 6.2 (Finn & Chernoff, 1993).

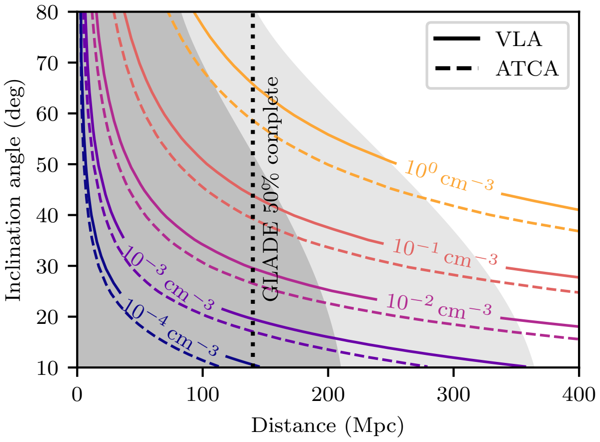

Figure 6.2: Maximum distance at which gravitational wave afterglows can be detected as a function of inclination angle for a range of circum-merger densities. Solid lines denote observations of 200 galaxies to a detection threshold of 75 μJy at 3 GHz with the VLA, while dashed lines denote observations with the ATCA targeting 50 galaxies to a detection threshold of 70 μJy at 8 GHz. The LIGO horizon for O4 and design specifications is shaded in dark and light grey respectively. The 50% completeness of the GLADE catalogue is also labelled.

Galaxy Targeting

Figure 6.2 shows the maximum distance at which an afterglow with can be detected for a range of circum-merger densities spanning 10^{-4}–1 cm^{-3}, using a galaxy-targeting approach with the ATCA and the VLA as outlined in Section 6.2.3.

We note that while this strategy will allow most events to be detected at distances comparable to the LIGO horizon, the incompleteness of existing galaxy catalogues at these distances makes this strategy only feasible for the closest mergers ($D_L \ll 100\,$Mpc). We do not consider applying this approach to any next generation facilities as their fields of view are large enough that they are more suited to untargeted searches.

Unbiased Searches

We therefore turn our focus to the unbiased searches described in Table 6.3, which we split into four broad categories.

Figure 6.3 shows the detectability of events in untargeted searches with current facilities. Most on-axis mergers, as well as most off-axis mergers occuring in dense environments, localised to $\leq$ 10 deg^2 detected with current gravitational wave facilities and the A+ upgrade will be detectable with MeerKAT and some will be detectable with the VLA and Apertif. However, we note that only a small fraction of events will be localised this well with these detectors. LIGO Voyager will have better localisation capabilities, and some events will produce afterglows that are detectable at the detector horizon.

Table 6.3: Capabilities of untargeted searches for radio afterglows for a range of telescopes and observing strategies, including observing frequency (ν), bandwidth ($\Delta\nu$), total areal coverage ($\Omega_{\rm total}$) and required observing time ($T_{\rm total}$)

Telescope	ν	Strategy (GHz)	$\Omega_{\rm total}$ (deg^2)	$S_{\rm detect}$ (μJy)	$T_{\rm total}$ (hr)
Apertif	1.4	deep	10	125	12
		wide	40	250	12
ASKAP	0.9	deep	30	175	10
		wide	300	550	10
DSA	1.35	deep	10	5	1
		wide	100	5	2.5
MeerKAT	1.4	wide	10	35	12
ngVLA	2.4	wide	10	5	10
		ultra-wide	100	25	10
SKA-1	1.43	wide	10	10	10
		ultra-wide	100	40	10
SKA-2	1.43	wide	10	1	10
		ultra-wide	100	4	10
VLA	1.5	wide	5	75	12

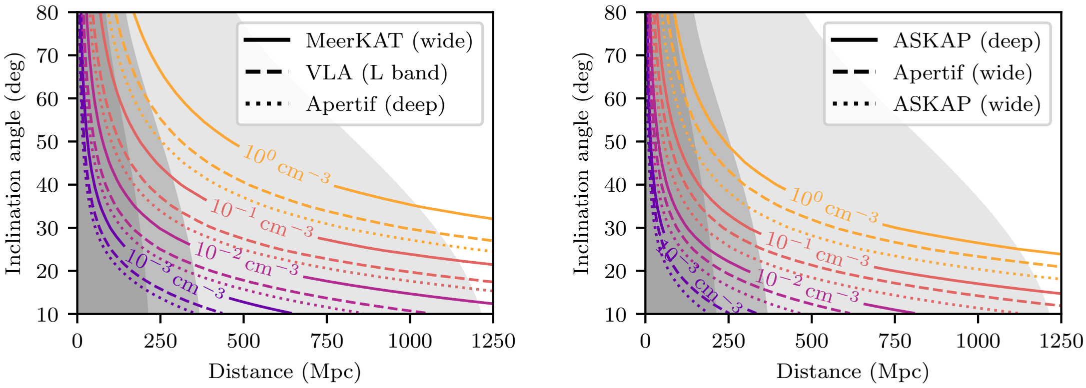

Figure 6.3: Similar to Figure 6.2. Left: unbiased observations of events localised to ≤ 10 deg^2 with MeerKAT, Apertif and the VLA. Right: unbiased observations with Apertif (covering 40 deg^2) and ASKAP (deep covering 30 deg^2 and wide covering 300 deg^2). Gravitational wave detector horizons for O4, A+ and Voyager are shown in increasingly light tones of grey.

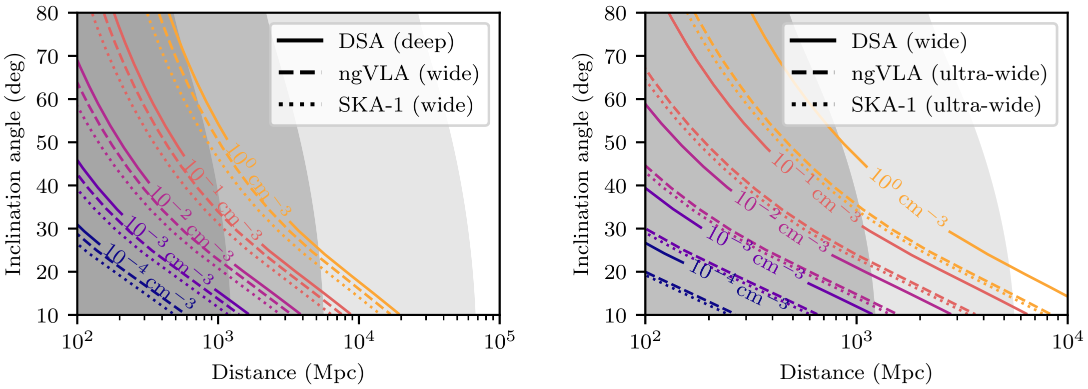

Figure 6.4: Similar to Figure 6.2. Unbiased observations of events localised to ≤ 10 deg^2 (left) and ≤ 100 deg^2 (right) with the SKA-1, DSA-2000 and ngVLA. The range of a nominal SKA-2 design is ~ 3 times further than the SKA-1. The detector horizons of Voyager, a preliminary 3G detector with a 5 Gpc horizon, and 3G detectors are shown in increasingly light tones of grey.

For events that are localised to tens of square degrees, comparable to the median localisation for 2G detectors (Abbott et al., 2018), we consider follow-up with ASKAP and Apertif. Figure 6.3 also shows the maximum distance at which afterglows will be detected, and we find that it is feasible to detect the afterglow produced by most on-axis mergers with current gravitational wave detectors.

Figure 6.4 shows the same metrics applied to the DSA-2000, ngVLA and SKA-1 compared to the horizons of Voyager and 3G detectors for events localised to $\leq$ 10 deg^2 and $\leq$ 100 deg^2. We find that while the majority of events detected with Voyager will be accompanied by detectable afterglows, it will not be possible to detect afterglows in widefield follow-up of the most distant events discovered by a complete 3G network. However, we note that the median localisation achievable with a complete 3G network is $\sim$ 1 deg^2. Therefore widefield searches will not be necessary for most events, and the targeted single-pointing strategy outlined in Section 6.4.2 may be a more useful metric.

Monitoring known counterparts

Figure 6.5 shows prospects for detecting radio emission from events that have been localised through an electromagnetic counterpart with current radio facilities. Assuming that neutron star mergers occur in comparably dense environments to short GRBs ($n \sim 10^{-2}\,\mathrm{cm}^{-3}$ Fong et al., 2015) we find that most neutron star mergers detected during O4, and a large fraction with the A+ configuration, should produce radio emission that is detectable with deep single pointing observations. However current facilities will not be sufficient for a comprehensive census of radio afterglows as we move towards the 3G era – only on-axis mergers occuring in dense environments ($n \gtrsim 10^{-1}\,\mathrm{cm}^{-3}$) will be detectable at the Voyager horizon.

The sensitivity of future radio telescopes will partially address this problem. Figure 6.6 shows the detectability horizon for the ngVLA and both phases of the SKA

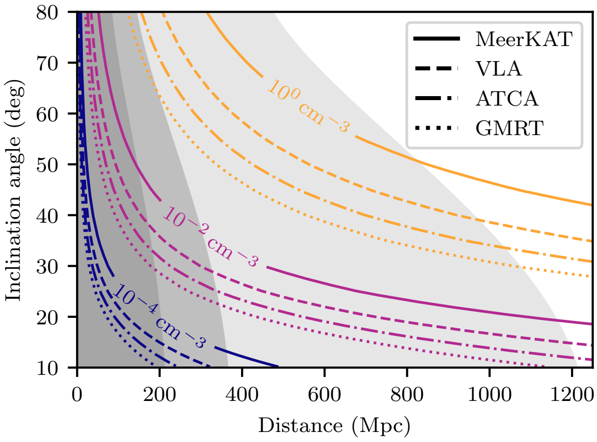

Figure 6.5: Similar to Figure 6.2 showing targeted single-pointing observations with existing radio telescopes. Gravitational wave detector horizons for O4, A+ and Voyager are shown in increasingly light tones of grey.

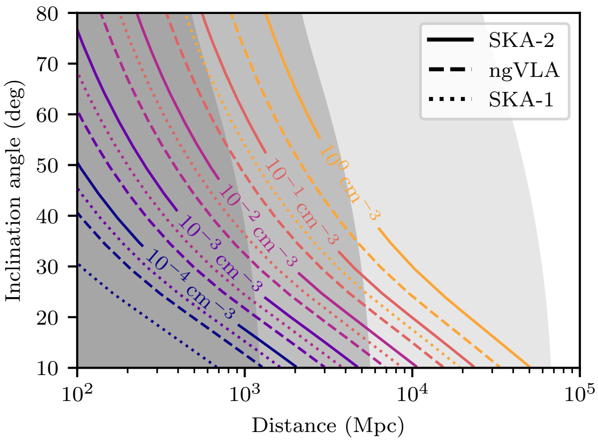

Figure 6.6: Similar to Figure 6.2 showing targeted single-pointing observations with planned radio telescopes. The range of the DSA-2000 is comparable to the range of the SKA-1. The detector horizons of Voyager, a preliminary 3G detector with a 5 Gpc horizon, and 3G detectors are shown in increasingly light tones of grey.

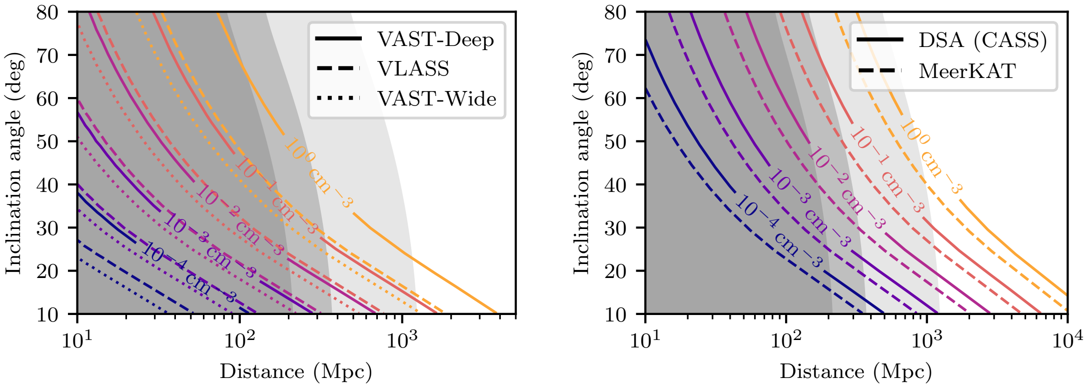

Figure 6.7: Detectability of radio afterglows for planned (left) and potential future (right) untargeted transients surveys. Detector horizons for O4, A+ and Voyager are shown in increasingly light shades of grey.

(with the DSA-2000 horizon comparable to that of the SKA-1). While most events detected with Voyager will be within range of future radio telescopes, the most distant events detected with 3G detectors will be well beyond the range of even the SKA-2. However, these detectors will detect thousands of events per year and therefore the limiting factor in obtaining a census of radio afterglows will be the amount of telescope time available for follow-up.

6.4.3 Serendipitous detections and orphan afterglows

We also consider the serendipitous detection of gravitational wave events, and orphan afterglows, in the transients surveys outlined in 6.2.6. Table 6.4 shows the properties of the surveys, and Figure 6.7 shows the application of the same detectability metrics as above. To determine the capability of surveys to obtain a complete sample of mergers occuring within their footprint, we also apply an additional constraint of the afterglow remaining detectable for a time corresponding to the survey cadence. This ensures that the afterglow will be detected in at least one epoch of the survey. The results of this change are shown in Figure 6.8, where we have excluded VLASS (as its slow cadence results in effectively zero range) and VAST-Wide (as its daily cadence results in no significant changes to the result above). Sensitivity remains the dominant limiting factor for off-axis events, while on-axis events occurring in denser environments are limited by survey cadence as their emission peaks at earlier times with a shorter turnover period.

While planned transients searches will not provide complete samples of afterglows due to their limited sky coverage, these results demonstrate that it is worthwhile carrying out targeted searches for afterglows within those datasets. These surveys may also make it feasible to search for counterparts to poorly localised events that do not individually warrant follow-up observations. We also note that once Voyager begins operations very few orphan afterglows will be detected, as only on-axis events occuring in the most dense environments will be detectable at distances beyond the Voyager horizon. However, this does not consider the dependence of detector sensitivity on sky position, nor the duty cycle of the detector network.

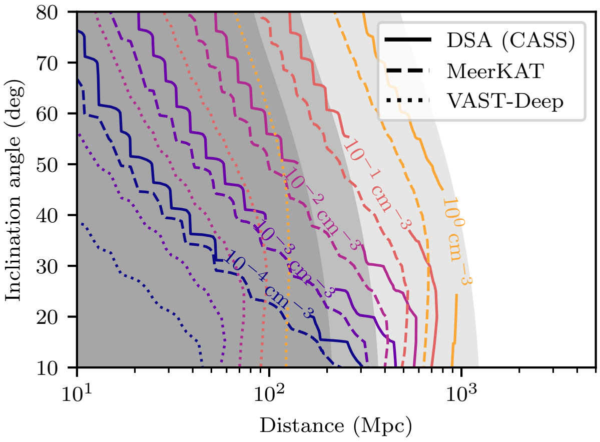

Figure 6.8: Detectability of a complete sample of radio afterglows in planned and potential future untargeted transients surveys. Detector horizons for O4, A+ and Voyager are shown in increasingly light shades of grey. As the inclination angle of the merger decreases the detectability becomes limited by the survey cadence rather than sensitivity. This effect is more pronounced for mergers occurring in denser environments.

6.5 Other contributions

6.5.1 Constraining merger properties

Observations of the non-thermal afterglow produced by mergers can be used to place constraints on the merger dynamics, the microphysics paramaters of the outflow, the inclination angle of the merger and the geometry of the outflow more generally. This can take two forms – interpreting the lightcurve in the context of standard afterglow models (Sari et al., 1998, 1999) to ascertain broad properties of the merger, or by fitting analytic models to the afterglow to constrain more detailed properties.

While the non-thermal afterglow extends from radio to X-rays, a dearth of sufficiently sensitive optical and X-ray facilities makes high cadence monitoring at these wavelengths impractical. The abundance of radio facilities allows for a higher observing cadence and therefore radio observations can be used to tightly constrain the temporal evolution of the lightcurve. However, multi-wavelength observations are still vital in determining whether any observed variability is driven by true temporal evolution, or spectral evolution of the source (e.g. the synchrotron cooling break moving into the radio band) manifesting as variability at a single frequency.

However, as noted by Nakar & Piran (2018), monitoring of the lightcurve is generally insufficient to distinguish between competing afterglow models and therefore other measurements are necessary to break the various degeneracies that are present. We outline the potential for other forms of radio observation to break these degeneracies in the section below.

Table 6.4: Ongoing, upcoming and idealised transients surveys. The fraction of the total sky covered by the survey is in column 3, while $S_{\rm det}$ corresponds to a 5σ detection threshold based on expected image noise. The MeerKAT survey is a theoretical idealised survey and there are no current plans to undertake it.

Survey	ν (GHz)	Sky coverage	Cadence (months)	$S_{\rm det}$ (μJy)
VLASS	3	0.82	32	600
VAST-Wide	0.9	0.23	daily	2500
VAST-Deep	0.9	0.23	8	250
MeerKAT	1.4	0.12	4	20
DSA (CASS)	1.35	0.75	4	10

6.5.2 Constraining outflow geometry

Radio observations can be used to gain unique insight into the structure of the merger outflow, beyond model-dependent inferences from the non-thermal lightcurve. High-cadence observations of afterglows can reveal variability induced by interstellar scintillation, which can be used to infer the physical size of the outflow and track its evolution (Frail et al., 1997; Goodman, 1997; Chandra et al., 2008). High resolution Very Long Baseline Interferometry (VLBI) observations can be used to measure the size of outflows and the proper motion of the emission centroid, which in turn can be used to infer the inclination angle of the merger (Mooley et al., 2018b; Ghirlanda et al., 2019; Nakar & Piran, 2020). Hotokezaka et al. (2019) used these improved inclination angle constraints to produce an improved measurement of H_0 compared to using gravitational wave data alone (Abbott et al., 2017e). We outline future prospects for these techniques in Dobie et al. (2020a).

Knowing the physical size of the outflow can be used to distinguish between emission models (e.g. quasi-spherical outflow vs. jet-dominated) that predict qualitatively similar lightcurves (Gill & Granot, 2018; Nakar & Piran, 2018). These techniques can also be used in conjunction with radio polarimetry to constrain the magnetic field structure of the outflow – for example, Corsi et al. (2018b) used a non-detection of linear polarisation to constrain models for the emission from GW170817. Future facilities like the SKA and the ngVLA will be capable of not only detecting linear polarisation from mergers, but tracking how it (and therefore the magnetic field) evolves (Corsi et al., 2018a).

6.5.3 Studying host galaxies

The formation channels and evolution of compact binaries is not well constrained, and even if an electromagnetic counterpart is detected, this provides limited information about progenitor formation (O'Shaughnessy et al., 2017). Instead, detailed studies of the host galaxy are required to reveal the complete history of the binary system and understand population properties. This kind of study has proved fruitful in investigations of other transient classes (e.g. Hamuy et al., 2000; Sullivan et al., 2006; Modjaz et al., 2008; Berger, 2009; Levesque et al., 2010; Bannister et al., 2019; Bhandari et al., 2020). While observational constraints are limited by a lack of events, a detailed study of the host galaxy of GW170817 was carried out (Blanchard et al., 2017; Im et al., 2017; Levan et al., 2017a; Pan et al., 2017; Ebrová et al., 2020; Contini, 2018) and simulations are already revealing potential relationships between merger and host properties (e.g. Toffano et al., 2019; Artale et al., 2019, 2020a). Specifically, radio observations are able to determine the presence of AGN and thereby determine whether stellar mass or star formation is the driving factor of merger formation (Zheng & Ramirez-Ruiz, 2007; Corsi et al., 2017). They can also measure star formation rate in galaxies with optical obscuration (Perley & Perley, 2013). Similar studies of GRB host galaxies have been carried out in the past, placing constraints on these properties and determining redshift scaling relations (e.g. Klose et al., 2019; Stanway et al., 2014; Zhang et al., 2018).

A large, unbiased, sample of mergers and associated host galaxy properties will also aid in follow-up of future events (e.g. Adhikari et al., 2020), by allowing galaxy catalogues to be weighted by more robust metrics than mass-analogues such as blue luminosity (see 6.2.2).

6.5.4 Dark siren measurements of H_0

Another independent method of measuring H_0 is the 'dark siren', where gravitational wave event localisation volumes are convolved with a catalogue of galaxies with known redshift (Schutz, 1986; Del Pozzo, 2012). This method has been applied to GW170817 (Fishbach et al., 2019), GW190814 (Abbott et al., 2020b) and the binary black hole merger GW170814 (Abbott et al., 2017c; Soares-Santos et al., 2019), although it will require a combined measurement of hundreds of events to reach a precision comparable to current techniques (Chen et al., 2018).

One limiting factor of this technique is the completeness of the galaxy sample, which can reduce the precision of such measurements by as much as a factor of ~ 2 (Gray et al., 2020). Fishbach et al. (2019) and Abbott et al. (2020b) use the GLADE galaxy catalogue (Dálya et al., 2018), while Soares-Santos et al. (2019) use the Dark Energy Survey (DES) galaxy catalogue, with both measurements using spectroscopic redshifts from archival surveys where possible, and photometric redshifts elsewhere. Two of the events events were ideal for this kind of measurement – GW170817 well-localised and occurred very nearby, where GLADE is complete enough to be representative of the distribution of galaxies in the volume, while GW170814 is one of the most well-localised and closest BBH mergers detected and the entirety of the localisation region falls within DES fields (which only cover $\sim 12\%$ of the sky).

A comprehensive dark siren measurement requires a deep spectroscopic catalogue of a large fraction of the localisation region. Spectroscopy is highly preferable, as the uncertainty of photometric redshift measurements is generally larger and is also often underestimated by as much as $\Delta z = 0.05$ (Hildebrandt et al., 2008). Additionally, photometric redshift estimates may be subject to additional biases, particularly at low redshift (Bilicki et al., 2014). This catalogue could come from targeted follow-up observations, or archival data as used in searches to-date. An alternative to optical spectroscopy is to use radio spectral line observations to map the distribution of neutral hydrogen and trace recent star formation history in the localisation volume.

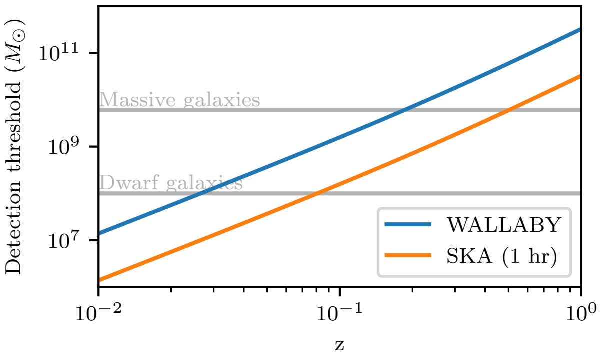

Figure 6.9: M_{HI} detection threshold for WALLABY (blue) and the SKA (orange) as a function of redshift. Typical values of M_{HI} for dwarf and massive galaxies are denoted by horizontal lines.

The Widefield ASKAP L-band Legacy All-sky Blind surveY (WALLABY Koribalski et al., 2020) will observe the sky South of declination +30 deg to a sensitivity of 1.6 mJy per 4 km/s channel. Figure 6.9 shows the HI mass sensitivity of ASKAP observations, calculated using the HI Fidelity calculator[6] (Meyer et al., 2017). Assuming a 5σ detection threshold, WALLABY observations will detect all galaxies with $M_{\rm HI} \gtrsim 10^9 M_\odot$ (comparable to most massive gas-rich galaxies) within the LIGO horizon. A dedicated follow-up campaign of a well-localised event like GW190814 could even provide a complete sample of gas-rich dwarf galaxies ($M_{\rm HI} \sim 10^8 M_\odot$) within the localisation volume. The SKA will have an instantaneous sensitivity that far exceeds that of ASKAP, albeit with a much smaller field of view. We also consider 1 hr integrations with the SKA that will achieve a sensitivity approximately an order of magnitude better than ASKAP.

It is beyond the scope of this work to establish a statistical framework that accounts for all the selection biases associated with this measurement, and quantify prospects for its use. However, we note that Artale et al. (2020a) show that the compact object merger rate in local galaxies ($z < 1$) positively correlates with both star formation and stellar mass, both of which are traced by neutral hydrogen. Therefore, while radio spectral line surveys will not provide complete galaxy samples, they will be subject to different selection effects than magnitude-limited optical surveys and may result in a more useful sample for dark siren measurements.

6.5.5 Prompt Radio Emission

Neutron star mergers may produce coherent radio emission in the final stages of the merger (e.g. Totani, 2013; Wang et al., 2016) or even up to hundreds of seconds post-merger (e.g. Pshirkov & Postnov, 2010; Zhang, 2014; Metzger & Zivancev, 2016). This emission will likely be similar to an FRB, although it may manifest in other forms (Lipunov & Panchenko, 1996).

---

[6] `https://hifi.icrar.org/`

While these merger models cannot explain the existence of repeating FRBs (Spitler et al., 2016; CHIME/FRB Collaboration et al., 2019b,a), others propose that repeating FRBs are produced by young magnetars formed by neutron star mergers (Yamasaki et al., 2018; Margalit et al., 2019). This is consistent with the recent detection of an FRB associated with the Galactic magnetar SGR 1935+2154 (Bochenek et al., 2020; CHIME/FRB Collaboration et al., 2020). However, this type of burst from ordinary magnetars likely makes up a small part of the overall FRB population which is dominated by more active magnetars formed via mechanisms including neutron star mergers (Margalit et al., 2020). Other FRBs have also shown behaviour consistent with being produced by magnetars (e.g. DeLaunay et al., 2016; Beloborodov, 2017; Metzger et al., 2017; Levin et al., 2020). The existence of multiple formation channels and sub-populations of FRB (Caleb et al., 2018) would also explain the discrepancy between the all-sky FRB and neutron star merger rates (Lawrence et al., 2017; Ravi, 2019).

Prompt follow-up of alerts from GRB and gravitational wave detectors has so-far been limited to low-frequency dipole arrays like the MWA (Kaplan et al., 2015, 2016), LOFAR (Rowlinson et al., 2019b) and the LWA (Anderson et al., 2018; Callister et al., 2019). These telescopes benefit from large fields of view that enable coverage of large fractions of the localisation area, no moving parts allowing for rapid re-pointing of the telescope, and large dispersion-induced delays on the arrival time of any associated signal. However, there have been no FRBs detected with these telescopes to-date (Sokolowski et al., 2018; Chawla et al., 2020), and therefore, despite the observational difficulties associated with it, follow-up with GHz-frequency telescopes may be necessary for the detection of prompt radio emission from mergers.

Low-mass mergers produce a signal that is detectable for tens of seconds pre-merger with current detectors (Chan et al., 2018), and it is therefore possible to distribute ‘negative-latency’ alerts (i.e. prior to the actual merger). The improved sensitivity of third generation detectors will allow the detection of some mergers with a negative latency of minutes (Kapadia et al., 2020). While this will provide lead-time for telescopes to get on-source, localisations with current detectors will still be tens–hundreds of square degrees, making it difficult to cover enough sky with traditional radio facilities. James et al. (2019) outline prospects for MWA negative-latency follow-up, while the prompt follow-up capability of ASKAP has been studied by Wang et al. (2020). Both studies find that while negative-latency follow-up is not practical for the general population of events, it will be possible to place tight constraints on prompt emission from the closest, most localised events. As detectors improve and we move towards third-generation detectors, prompt follow-up will become more practical as more sensitive detectors will detect mergers earlier and localise them to smaller regions.

Apertif, DSA-2000 and CHIME will all be ideal for prompt follow-up due to their wide fields of view, while it may be possible to also perform prompt follow-up with MeerKAT and the SKA by splitting them into multiple sub-arrays.

6.6 Conclusions

In this work we have provided a complete overview of synergies between radio telescopes and gravitational wave detectors, encompassing both existing facilities and planned or proposed facilities spanning the coming decades. We demonstrate that while targeting potential host galaxies proved useful in the follow-up of GW170817, this method will be less feasible in future follow-up due to limited catalogue completeness at distances comparable to the horizon of gravitational wave detectors. Additionally, the larger field of view of future telescopes is more conducive to unbiased widefield searches that target the localisation region of the merger. We find that these searches with current facilities will be capable of detecting mergers at hundreds of Mpc, while future facilities will be able to detect mergers at Gpc distances. Widefield transients surveys will provide serendipitous coverage of events and may also detect afterglows of events beyond the detector horizon, and those that occur during detector downtime.

Radio observations can be used to place constraints on properties of the merger outflow and the circum-merger environment, and for events with counterparts detected at other wavelengths, we find that current radio facilities are capable of detecting some afterglows at the Voyager horizon, while future facilities will detect afterglows at distances up to tens of Gpc for the most on-axis events. However, lightcurve monitoring alone is insufficient to completely constrain the geometry of the merger, and we also discuss possible ways of breaking model degeneracies. Finally, we also outline other contributions that radio observations can make to the broader multi-messenger commmunity including host galaxy studies, measurements of the Hubble constant (through both improved inclination angle measurements and probabilistic estimates based on H_I emission within the localisation volume) and understanding of the central engine via searches for prompt, coherent radio emission associated with mergers.

Chapter 7

Conclusion

The first detection of gravitational waves and light from a neutron star merger, GW170817, heralds the dawn of the era of multi-messenger astronomy (Abbott et al., 2017f). This discovery confirmed the association between short GRBs and neutron star mergers (Abbott et al., 2017g), the origin of the heavy elements (Chornock et al., 2017; Kasliwal et al., 2017; Pian et al., 2017) and that gravitational waves propagate at the speed of light (Abbott et al., 2017g). It also enabled independent measurements of the Hubble constant (Abbott et al., 2017e; Hotokezaka et al., 2019) and various tests of General Relativity (Abbott et al., 2019c).

These insights were only possible due to comprehensive follow-up observations with telescopes across the entire electromagnetic spectrum. The gamma-ray emission can be used to probe the central engine of the merger, and compare the arrival time of gravitational waves and light. The optical emission is dependent on factors including the opacity of the surrounding material and the mass of the ejecta and was vital in localising the merger. Radio observations detect synchrotron emission produced by the relativistic outflow from the merger, provide unique insight into the geometry and dynamics of the outflow and the surrounding environment. As part of this thesis I carried out observations and analysis of the radio afterglow, including the confirmation that the emission had peaked and begun to decline (Chapter 2).

While observations of GW170817 have answered many questions, there is a limit to the knowledge that can be gained with a sample size of one. Observations of a larger sample of mergers will enable studies of the neutron star merger population. This will lead to more breakthroughs, including a better understanding of the relationship between short GRBs and neutron star mergers, tighter constraints on the expected rate of electromagnetic transients produced by neutron star mergers, and more precise standard-siren measurements of the Hubble constant.

Progress towards a larger sample of mergers will be spurred by improvements in both gravitational wave detectors and follow-up capabilities.

Planned upgrades to existing gravitational wave detectors, along with the commissioning of new ones, will improve the sensitivity of the network (Abbott et al., 2018; Hall & Evans, 2019). Not only will this expand the merger detection horizon and increase the rate of detections, but it will also improve the localisation of events. Smaller localisation regions will facilitate more efficient follow-up observations, and enable the discovery of electromagnetic counterparts to become commonplace.

The commissioning of new electromagnetic facilities will also improve our localisation capabilities and allow for deeper study of the properties of afterglows. The Transient High-Energy Sky and Early Universe Surveyor (THESEUS; Amati et al., 2018) will improve the characterisation of GRBs and enable mergers to be localised to arcminute precision in real-time. Optical and infra-red facilities such as the Vera C. Rubin Observatory (Margutti et al., 2018a), the Southern Gravitational-Wave Optical Transient Observer (GOTO; Gompertz et al., 2020) and the Dynamic REd All-sky Monitoring Survey (DREAMS; Soon et al., 2020) will improve the efficiency of searches for kilonovae. Untargeted surveys for transients with these facilities will also likely detect emission from neutron star mergers independent of any gravitational wave signals (Andreoni et al., 2019).

However, searches for kilonovae will not be sufficient to localise all mergers. Follow-up of some events may be hindered by observational constraints, or extrinsic factors such as dust obscuration. Mergers of massive neutron star binaries, or neutron star-black holes, may not produce detectable optical or infra-red emission, although radio emission is still expected. Radio telescopes are also subject to far fewer observational constraints than optical facilities, and radio emission is generally not hindered by material along the line of sight. Therefore, radio follow-up is a promising method for localising mergers where facilities operating at other wavelengths cannot.

The commissioning of ASKAP and MeerKAT has enabled unprecedented widefield searches for radio counterparts to mergers, and these capabilities will only be improved as facilities like the DSA-2000 and ultimately, the Square Kilometre Array, come online. I have used ASKAP to carry out the first widefield radio follow-up of a gravitational wave event, and placed tight constraints on the inclination angle of the merger and the density of its surrounding environment (Chapter 4). I have also developed an optimised follow-up strategy that is tailored to widefield radio telescopes (Chapter 3), and quantified prospects for localising mergers with all suitable existing and planned radio facilities (Chapter 6).

ASKAP and MeerKAT will also carry out widefield, GHz-frequency, transients surveys that will begin to probe the expected rate of radio transients for the first time. These facilities will detect the radio afterglows of neutron star mergers without independent of the detection of GRBs or gravitational waves, which are subject to inclination angle selection effects. This will allow measurement of the GRB beaming fraction, and in turn enable estimates of the rate of short GRBs in the Universe, and the true energy scale of the emission they produce.

The contribution of radio observations to the multi-messenger era is not limited to localisation. Broadband monitoring of the radio afterglow produced by mergers is necessary to constrain properties including the inclination angle, circum-merger density, and energetics of the jet. A large sample of merger afterglows will shed light on the typical microphysics parameters associated with these outflows, and potentially allow study of the properties of the central engine that drives them. However, measurements beyond simple lightcurve monitoring are required to break model degeneracies such as the relationship between the inclination and jet opening angle angles.

This could be achieved via VLBI imaging to directly image outflow structure and detect relativisitc motion, or by using high-cadence monitoring to detect scintillation and infer source sizes. In this thesis I have outlined prospects for these techniques and quantified the population of mergers they will provide useful constraints for (Chapter 5). The inclination angle constraints provided by VLBI observations in particular will be useful in improving the precision of standard siren measurements of the Hubble Constant.

We are entering a new era of transient astronomy. The combination of gravitational wave detectors, the most sensitive widefield optical telescopes ever constructed, and the first widefield radio telescopes will provide an unprecedented view of the dynamic sky. Much like observations of the planets and the stars shaped our notions of the Universe many millennia ago, the transient events that will be detected in the coming years will lead to a deeper understanding of the Universe, and our place in it.

Appendix A

Software

A.1 GCNBOT

The third LIGO/Virgo Observing Run (O3) was scheduled to run for one year, with alerts expected at any time of day. It was vital to be able to alert our team of observers in a clear and concise way that allowed rapid response to alerts, without overwhelming them with a barrage of uninteresting ones. As part of this thesis I developed software (`https://github.com/ddobie/gcnbot_public`) to parse these alerts, determine their feasibility for follow-up and distribute the information in human-readable form to our team via Slack and SMS.

A.2 Triggering the Murchison Widefield Array

The Murchison Widefield Array (MWA; Tingay et al., 2013; Wayth et al., 2018) is a low-frequency radio telescope, consisting of 128 tiles of 16 dipoles spread across a maximum baseline of 3 km. The telescope has no moving parts, meaning that the telescope can be rapidly repointed. As part of this thesis I developed software within the existing MWA automatic trigger system (Hancock et al., 2019b) to enable prompt follow-up of gravitational wave alerts and neutrino detections (`https://github.com/MWATelescope/mwa_trigger`).

This software was run throughout O3 with the following trigger criteria:

- alert received within 10 minutes of the merger;
- $\geq 50\%$ probability of the merger containing at least one neutron star;
- $\geq 10\%$ of the localisation currently observable with the MWA

Only one event during O3 satisfied this criteria. S191213g[1] (LIGO Scientific Collaboration and Virgo Collaboration et al., 2019b) is a likely binary neutron star merger detected at a distance of 201 ± 81 Mpc. Observations with the MWA began 25 seconds after receiving the alert, corresponding to a total post-merger latency of less than 6 minutes. Analysis of this data is ongoing.

---

[1] `https://gracedb.ligo.org/superevents/S191213g/view/`

A.3 Optimised ATCA Observing

As discussed in Section 6.2.3, the Australia Telescope Compact Array (ATCA) can be used to search for radio counterparts to gravitational wave events by targeting candidate host galaxies within the localisation volume. While the ATCA has a relatively small field of view, it is still possible to observe multiple galaxies within a single pointing. As part of this thesis I developed software (`https://github.com/ddobie/atca-ligo`) that ensures each target can be observed with a given sensitivity while minimising the total observing time.

The software uses farthest-point hierarchical clustering (Voorhees, 1986) to group galaxies based on minimising the maximum separation on the sky between all targets in the group. A pointing centre for each group is calculated, based on the smallest circle that encloses all targets in the group, which is either the midpoint of the two farthest apart targets, or the circumcentre of the triangle made up of those two targets and the target furthest away from that midpoint.

Starting with the root node of the hierarchical tree (i.e. the group containing all targets), the integration time required to reach an arbitrary sensitivity for all targets within the group is calculated using the ATCA primary beam response (Wieringa & Kesteven, 1992) for the most distant target from the pointing centre. This is compared to the integration time required to achieve that sensitivity if all targets within the group were observed separately. If separate observations requires less integration time, the algorithm progresses to the next level of the tree, otherwise the targets are grouped and removed from the tree. This process continues until all targets are assigned a group. A suitable phase calibrator for each group is found using the CABB-SCHEDULE-API[2], and then the groups and their associated phase calibrators are sorted into an order that minimises the required slew time between them using ATMOS (Sault et al., 1995). The software automatically generates a schedule file ready for immediate use.

This method is sub-optimal and does not prioritise observations of sources that set earlier, nor does it consider the slew-time saved observing multiple targets within a single pointing. However, it generally reduces the required observing time by XXX% compared to the naïve approach of simply running ATMOS on the list of candidate host galaxies.

---

[2] `https://github.com/ste616/cabb-schedule-api`

A.4 TreasureMapPy

The Gravitational Wave Treasure Map (`treasuremap.space`; Wyatt et al., 2020) allows astronomers to report details of their follow-up observations of gravitational wave events. The goal of the project is to foster collaboration between observers and maximise the efficiency of follow-up observations with all facilities. As part of this thesis I wrote TreasureMapPy (`https://github.com/ddobie/TreasureMapPy`), a Python module that allows users to easily upload pointing details to the Treasure Map. This software can be incorporated into existing follow-up pipelines to automatically share observation details without requiring human interaction.

A.5 VAST-TOOLS

The Variables And Slow Transients (VAST; Murphy et al., 2013) pilot survey consists of repeated observations of 113 pointings with ASKAP. There are 13 epochs of observation in total, although only 4 epochs are considered complete – due to scheduling constraints, test observations and technical issues that required fields to be re-observed, most fields were observed more a nominal four times.

As part of this thesis I developed VAST-TOOLS jointly with Adam Stewart. VAST-TOOLS allows users to rapidly determine if a list of coordinates are within the survey footprint, crossmatch coordinates with survey source catalogues, produce images of regions, build lightcurves of sources and calculate variability metrics. It has already been used in follow-up of neutrino events (Dobie et al., 2019c) and a short GRB (Dobie et al., 2020c).

Appendix B

Published Notices

B.1 GCN Circulars

GCN21625

```
TITLE:   GCN CIRCULAR
NUMBER:  21625
SUBJECT: LIGO/Virgo G298048: ASKAP follow up
DATE:    17/08/21 00:58:33 GMT
FROM:    Tara Murphy at U of Sydney  <tara.murphy@sydney.edu.au>

D. Dobie (University of Sydney), A. Hotan (CSIRO), K. Bannister (CSIRO),
D. Kaplan (UWM), T. Murphy (University of Sydney), C. Lynch (University
of Sydney), on behalf of the ASKAP/VAST collaboration.

We are observing the LIGO localisation region (LVC GCN 21513) with the
Australian Square Kilometre Array Pathfinder (ASKAP) at a central
frequency of 1345 MHz with a bandwidth of 192 MHz. As the array is still
undergoing commissioning, we are using a nominal 12 (of 36) antennas,
although other site activities may cause that number to vary from one
observation to the next.

Our observing strategy consists of 3 pointings containing approximately
90% of the LIGO localisation region. Each pointing consists of a 5.5 x 5.5
degree square grid of 36 beams, though the exact number captured may
also vary. The first pointing is centered at

RA = 13:09:21.60
Dec = -25:00:00.00

and contains ~48% of the localisation region and the positions of 35
target galaxies (Cook et al. LVC GCN 21519) including NGC 4993, the
host of the possible optical counterpart SSS17a
(Coulter et al. LVC GCN 21529). Observations began at 2017-08-19
05:34:32 (UT) and are ongoing.

Processing and analysis of the first pointing is underway.

Thank you to CSIRO staff for supporting these observations.
```

GCN21639

TITLE: GCN CIRCULAR
NUMBER: 21639
SUBJECT: LIGO/Virgo G298048: ASKAP observations of SSS17a and NGC 4993 at 1.345 GHz
DATE: 17/08/22 07:23:04 GMT
FROM: Tara Murphy at U of Sydney <tara.murphy@sydney.edu.au>

D. Dobie (University of Sydney), A. Hotan (CSIRO), K. Bannister (CSIRO),
T. Murphy (University of Sydney), D. Kaplan (UWM), C. Lynch (University of Sydney),
on behalf of the ASKAP/VAST collaboration.

We have observed the LIGO/Virgo localisation region (LVC GCN 21527) with the Australian Square Kilometre Array Pathfinder (ASKAP) at a central frequency of 1.345 GHz with a bandwidth of 192 MHz. Observations started at 2017-08-19 05:34 UT and ended at 2017-08-19 07:58 UT.

We do not detect any emission at the position of the possible optical counterpart SSS17a (Coulter et al. LVC GCN 21529), or its candidate host galaxy NGC 4993. The upper limit for NGC 4993 is consistent with a spectral index of -0.7 found by fitting the measured results from the ATCA (Bannister et al. LVC GCN 21559) between 8.5 and 21.2 GHz.

Further analysis of our observations is ongoing.

Thank you to CSIRO staff for supporting these observations.

GCN23139

TITLE: GCN CIRCULAR
NUMBER: 23139
SUBJECT: LIGO/Virgo GW170817: A steep decline in the radio light curve and prediction for the X-rays
DATE: 18/08/13 18:38:03 GMT
FROM: Tara Murphy at U of Sydney <tara.murphy@sydney.edu.au>

D. Dobie (University of Sydney), K. Mooley (Caltech),
T. Murphy (University of Sydney), D. Kaplan (UWM), E. Lenc (CSIRO),
A. Corsi (TTU), D. Frail (NRAO), report on behalf of a larger collaboration

Our continued observations of GW170817 with the ATCA and the VLA up to 300 days post-merger (Mooley et al. in prep) confirm the t^(-2) decline in the radio light curve initially reported in Mooley et al. 2018 (arXiv:1806.09693). Such a slope rules out a cocoon-dominated outflow at late times, and is instead the classic signature of a relativistic jet, consistent with the VLBI result from Mooley et al. 2018. The t^(-2) decline is also expected in the X-ray light curve, and may be confirmed by the Chandra observation carried out on 2018 Aug 10.

GCN25445

TITLE: GCN CIRCULAR
NUMBER: 25445
SUBJECT: LIGO/Virgo S190814bv: No radio counterpart detected in ASKAP observations
DATE: 19/08/22 07:57:20 GMT
FROM: Dougal Dobie at VAST <ddobie94@gmail.com>

Dougal Dobie (University of Sydney/CSIRO), Adam Stewart (University of Sydney),
Ziteng Wang (University of Sydney), Tara Murphy (University of Sydney),
Emil Lenc (CSIRO), David Kaplan (UWM), Aidan Hotan (CSIRO),
Kunal Mooley (NRAO, Caltech), Gregg Hallinan (Caltech),
David McConnell (CSIRO), Julie Banfield (CSIRO), Wasim Raja (CSIRO),
Matthew Whiting (CSIRO), Vanessa Moss (CSIRO), Igor Andreoni (Caltech)
and the OzGrav, JAGWAR and GROWTH collaborations.

We report observations of the localisation region of S190814bv (LVC, GCN 25324) with the Australian Square Kilometre Array Pathfinder (ASKAP) on 2019-08-16 at a central frequency of 943 MHz with a bandwidth of 288 MHz.

We have observed a single 30 sq. deg. field centered on
    RA = 00:50:37.5
    Dec = -25:16:57.4
which covers approximately 85% of the sky localisation from the LALInference skymap (GCN 25333), with a median rms of 34 uJy.

We have searched for radio emission within 5 arcseconds of the location of the 124 optical transients reported on the Transient Name Server by the DECam-GROWTH and DESGW teams between 2019-08-16 and 2019-08-22 as at 2019-08-22 03:00 UTC. We report coincident compact radio emission at the location of 14 of them:

Name	RA	Dec	Int. Flux (uJy)	err. (uJy)	Notes
AT 2019nqa	00:52:39.1	-25:00:15	258	30	
AT 2019nqy	00:56:23.2	-24:41:11	393	29	
AT 2019nqz	00:46:46.5	-24:20:12	870	30	(a)
AT 2019nsr	00:57:27.6	-26:16:44	290	36	(c)
AT 2019nto	00:42:03.5	-24:48:19	342	28	(c)
AT 2019nuk	00:54:57.9	-26:08:03	233	28	(b)
AT 2019nul	00:55:16.4	-26:56:35	204	28	(b)
AT 2019nun	00:56:48.7	-24:54:31	377	29	(b,c)
AT 2019nuo	00:56:03.9	-23:18:15	388	36	(c)
AT 2019nup	00:55:04.3	-26:46:12	446	33	(c)
AT 2019nzj	00:52:05.3	-26:11:03	759	29	
AT 2019nzn	00:55:19.9	-24:09:29	233	32	(c)
AT 2019oay	00:45:25.2	-25:53:43	348	31	
AT 2019ocs	01:00:11.4	-25:53:22	352	29	

(a) reported in GCN25391
(b) reported in GCN25393
(c) source possibly extended

We have also performed a preliminary search for transients using TraP (Swinbank et al. 2015), comparing this observation to the Rapid ASKAP Continuum Survey (RACS, [1]) at a detection threshold of 0.95 mJy, corresponding to 5 times the lowest rms of the RACS image.

We find one candidate transient located at
RA = 00:54:34.6 +/- 0.02 arcsec
Dec = -28:02:35.3 +/- 0.01 arcsec
which we note is outside the 95% confidence region of S190814bv. We measure a flux density of 3.4 mJy in this observation and a local rms noise of 0.25 mJy in the 888 MHz RACS image observed on 2019-04-26. We measured an integrated flux density of 0.74 mJy in the RACS image using TraP. We also note that there is a radio source coincident with this location in the image from the Very Large Array Sky Survey (VLASS) observed on 2019-06-29 with a flux density of ~1.6 mJy at 3 GHz.

We conducted follow-up of this source with the Australia Telescope Compact Array (ATCA) on 2019-08-21 with two 2048 MHz bands centered on 5.5 and 9 GHz. We measure preliminary flux densities of 2.88 +/- 0.03 mJy and 2.93 +/- 0.02 mJy at 5.5 and 9 GHz, with respective in-band spectral indices of +0.17 and -0.37.

Combining the near-contemporaneous ATCA and ASKAP measurements we find a flat spectral index. Based on these observations this candidate is likely to be an unrelated AGN.

The ASKAP observation is publicly available on the CSIRO ASKAP Science Data Archive [2] under Scheduling Block ID 9602.

Further analysis of this ASKAP observation is ongoing and further epochs are planned.

Thank you to CSIRO staff for supporting these observations.

[1] https://www.atnf.csiro.au/content/racs
[2] https://casda.csiro.au/

GCN25472

TITLE: GCN CIRCULAR
NUMBER: 25472
SUBJECT: LIGO/Virgo S190814bv: No radio counterpart to DG19wxnjc/AT2019npv
in ASKAP observations
DATE: 19/08/25 04:26:24 GMT
FROM: Dougal Dobie at VAST <ddobie94@gmail.com>

Dougal Dobie (University of Sydney/CSIRO), Tara Murphy (University of Sydney), Emil Lenc (CSIRO), David Kaplan (UWM), Aidan Hotan (CSIRO), Adam Stewart (University of Sydney) and the OzGrav, JAGWAR and GROWTH collaborations.

We report a non-detection of the optical transient DG19wxnjc/AT2019npv (GCN 25393) in two recent observations with the Australian Square Kilometre Array Pathfinder (ASKAP) at a frequency of 943 MHz.

The 3-sigma upper limits at the position of this source are 75 uJy on 2019-08-16 and 96 uJy on 2019-08-23.

Both observations were a 30 sq. deg. field of the localisation region of S190814bv (LVC, GCN 25324) centered on
    RA = 00:50:37.5
    Dec = -25:16:57.4
with a central frequency of 943 MHz with a bandwidth of 288 MHz, as reported in GCN 25445.

Further analysis of these ASKAP observations are ongoing and future epochs are planned.

Thank you to CSIRO staff for supporting these observations.

GCN25621

TITLE: GCN CIRCULAR
NUMBER: 25621
SUBJECT: LIGO/Virgo S190814bv: ATCA observation of ASKAP J005547-270433/AT2019osy
DATE: 19/09/03 01:20:02 GMT
FROM: Dougal Dobie at VAST <ddobie94@gmail.com>

Dougal Dobie (University of Sydney/CSIRO), Emil Lenc (CSIRO),
Ian Brown (UWM), Tara Murphy (University of Sydney), Adam Stewart
(University of Sydney), David Kaplan (UWM), Kunal Mooley (Caltech),
Gregg Hallinan (Caltech) and the OzGrav, JAGWAR and GROWTH collaborations.

We observed ASKAP J005547.4-270433/AT2019osy (GCN 25487) with the Australia Telescope Compact Array (ATCA) on 2019 August 29 from 14:00-22:00 UT.

We report flux densities of:

369+/-23 uJy at 5.0 GHz
335+/-19 uJy at 6.0 GHz
307+/-15 uJy at 8.5 GHz
278+/-14 uJy at 9.5 GHz

Fitting a power-law to these values we find a spectral index of alpha=-0.39+/-0.11, consistent with the value of -0.42 found by Mooley et al. (GCN 25539).

Comparing to the VLA observations reported by Mooley et al. (GCN 25539), we find that the flux density has increased by ~40% in 1.5 days across 4.5-10 GHz.

Additional epochs are required to separate intrinsic variability from contamination from propagation effects and systematics due to differing spatial resolution between the ATCA, VLA and ASKAP.

Thank you to CSIRO staff for supporting these observations.

GCN25691

```
TITLE:   GCN CIRCULAR
NUMBER:  25691
SUBJECT: LIGO/Virgo S190814bv: ATCA monitoring of ASKAP J005547-270433/AT2019osy
DATE:    19/09/10 00:30:43 GMT
FROM:    Dougal Dobie at VAST  <ddobie94@gmail.com>
```

Dougal Dobie (University of Sydney/CSIRO), Emil Lenc (CSIRO),
Ian Brown (UWM),Tara Murphy (University of Sydney), Adam Stewart
(University of Sydney), David Kaplan (UWM), Kunal Mooley (Caltech),
Gregg Hallinan (Caltech) and the OzGrav, JAGWAR and GROWTH collaborations.

We observed ASKAP J005547.4-270433/AT2019osy (GCN 25487) with the Australia Telescope Compact Array (ATCA) on 2019 September 6 from 12:30-19:30 UT.

We report flux densities of:
  380+/-21 uJy at 5.0 GHz
  353+/-17 uJy at 6.0 GHz
  299+/-14 uJy at 8.5 GHz
  234+/-14 uJy at 9.5 GHz

Which are consistent with the measurements obtained with the ATCA on 2019 August 29, 8 days prior (GCN 25487, GCN 25621, GCN 25539). Therefore the previously observed variability is unlikely to be related to S190814bv.

Thank you to CSIRO staff for supporting these observations.

GCN27516

TITLE: GCN CIRCULAR
NUMBER: 27516
SUBJECT: GRB 200405B: ATCA follow-up and ASKAP limits on pre-burst radio emission
DATE: 20/04/10 01:56:01 GMT
FROM: Dougal Dobie at VAST <ddobie94@gmail.com>

Dougal Dobie (USYD/CSIRO), Tara Murphy (USYD), James Leung (USYD/CSIRO), Adam Stewart (USYD), Joshua Pritchard (USYD), David Kaplan (UWM)

GRB 200405B (GCN 27497) occurred in a field that has been observed 5 times as part of the ASKAP Variables And Slow Transients (VAST; Murphy et al. 2013) pilot survey between 2019-08-27 and 2020-01-25. We have searched for radio emission from the BAT ground-calculated position and the 4 uncatalogued X-ray sources detected by Swift (GCN 27500) and find no radio counterparts to a detection limit of ~1.5 mJy at 888 MHz.

We also performed follow-up observations of all 5 positions with the ATCA between 2020-04-09 05:00-10:30 UTC with 2x2 GHz bands centered on 5.5 and 9 GHz. No radio emission was detected coincident with any of the sources, we list 3-sigma upper limits below

Source	5.5 GHz (uJy)	9 GHz (uJy)
BAT pos	81	54
Source 1	90	120
Source 2	84	96
Source 3	87	90
Source 4	90	114

We do detect a radio source at coordinates of 04:10:26.8, -51:31:55 (offset 7.2 arcsec from Source 4), coincident with WISEA J041026.82-513155.2, with a flux density of ~4 mJy at 5.5 GHz and ~7 mJy at 9 GHz. This source is also detected in the VAST pilot survey with a flux density of ~10 mJy. We do not consider this a candidate afterglow for the GRB.

Further observations with ATCA and as part of the VAST program are planned.

Thank you to CSIRO staff for supporting these observations during these especially difficult times.

B.2 Astronomer's Telegrams

ATCA Observations of AT2018cow (ATel #11795)

Dougal Dobie (University of Sydney), Vikram Ravi (Caltech), Anna Ho (Caltech),
Mansi Kasliwal (Caltech), Tara Murphy (University of Sydney)
on 29 Jun 2018; 15:52 UT
Credential Certification: Dougal Dobie (ddob1600@uni.sydney.edu.au)

We have observed the position of AT2018cow (ATel #11727) with the Australia Telescope Compact Array (ATCA) on 2018-06-26 from 09:00-14:00 UT at central frequencies of 5.5 GHz, 9 GHz with a bandwidth of 2 GHz in each band, and 34 GHz with a bandwidth of 4 GHz. Due to maintenance, only four of six antennas were available, and only three antennas were used at 34 GHz.

We detect a point source coincident with the location of AT2018cow (ATel #11727) and report preliminary flux densities of:

<0.18 mJy at 5.5 GHz (3-sigma upper limit);
~0.46 mJy at 9 GHz;
~5.6 mJy at 34 GHz.

Analysis is ongoing and subsequent epochs are planned.

Thank you to CSIRO staff for supporting these observations.

AT2018cow: Further ATCA monitoring (ATel #11818)

Dougal Dobie (University of Sydney), Vikram Ravi (Caltech), Anna Ho (Caltech),
Mansi Kasliwal (Caltech), Tara Murphy (University of Sydney)
on 5 Jul 2018; 05:33 UT
Credential Certification: Dougal Dobie (ddob1600@uni.sydney.edu.au)

We have observed the position of AT2018cow (ATel #11727) with the Australia Telescope Compact Array (ATCA) on 2018-06-28 from 08:30-14:00 UT at central frequencies of 5.5 GHz, 9 GHz, 16.7 GHz and 21.2 GHz with a bandwidth of 2 GHz in each band, and 34 GHz with a bandwidth of 4 GHz.

A point source coincident with the location of AT2018cow (ATel #11727) is detected in all bands. We report preliminary flux densities of:

~0.22 mJy at 5.5 GHz;
~0.52 mJy at 9 GHz;
~1.5 mJy at 16.7 GHz;
~2.3 mJy at 21.2 GHz;
~7.6 mJy at 34 GHz.

Analysis is ongoing and subsequent epochs are planned.

A flux-density scale error in our previous ATCA observations (ATel #11795) resulted in an incorrect measurement being reported at 9 GHz. The corrected flux density of AT2018cow at 9 GHz on 2018-06-26 is ~0.3 mJy.

Thank you to CSIRO staff for supporting these observations.

AT2018cow: Continued ATCA monitoring (ATel #11862)

Dougal Dobie (University of Sydney), Vikram Ravi (Caltech), Anna Ho (Caltech),
Mansi Kasliwal (Caltech), Tara Murphy (University of Sydney)
on 17 Jul 2018; 04:29 UT
Credential Certification: Dougal Dobie (ddob1600@uni.sydney.edu.au)

We have observed the position of AT2018cow (ATel #11727) with the Australia Telescope Compact Array (ATCA) on 2018-07-03 between 09:00-13:45 UT at central frequencies of 5.5 GHz and 9 GHz with a bandwidth of 2 GHz in each band.

We report preliminary flux densities of ~0.4 mJy at 5.5 GHz and ~1.0 mJy at 9 GHz.

We have also observed AT2018cow on 2018-07-05 from 13:30-15:30 UT at a central frequency of 34 GHz with a bandwidth of 4 GHz.

We report a preliminary flux density of ~10 mJy.

Analysis is ongoing and subsequent epochs are planned.

Thank you to CSIRO staff for supporting these observations.

ASKAP observations of blazars possibly associated with neutrino events IC190730A and IC190704A (ATel #12981)

Dougal Dobie (University of Sydney/CSIRO), David L. Kaplan (University of Wisconsin, Milwaukee), Adam Stewart (University of Sydney), Tara Murphy (University of Sydney), Emil Lenc (CSIRO), David McConnell (CSIRO), Aidan Hotan (CSIRO), Julie Banfield (CSIRO), Wasim Raja (CSIRO), Matthew Whiting (CSIRO)
on 2 Aug 2019; 07:23 UT
Credential Certification: Dougal Dobie (ddob1600@uni.sydney.edu.au)

GCN #24981 reports the detection of a high-energy astrophysical neutrino candidate, IC190704A, spatially coincident with the known blazar 1WHSP J104516.2+275133.

ATel #12967 reports the detection of a high-energy astrophysical neutrino candidate, IC190730A, spatially coincident with the known blazar PKS 1502+106.

We have obtained observations of both blazars prior to the detection of the corresponding neutrino candidates with the Australian Square Kilometre Array Pathfinder (ASKAP) as part of the Rapid ASKAP Continuum Survey (RACS). Observations were carried out at 888 MHz with a bandwidth of 288 MHz.

The field containing PKS 1502+106 was observed on 2019-04-24 and we find a flux density of ~1.68 Jy, which is consistent with archival measurements. Historically this source appears to exhibit minimal radio variability.

Two fields containing 1WHSP J104516.2+275133 were observed on 2019-04-21 and we find a flux density of ~4.5 mJy in each. This is broadly consistent with observations with the VLA on 2019-07-09 (ATel #12926) and archival data from FIRST. Comparing the RACS measurement to the observation as part of the VLA Sky Survey on 2019-06-08 suggests the blazar had a much steeper spectrum prior to the neutrino event and is exhibiting significant variability.

This work was done as part of the ASKAP Variables and Slow Transients (VAST) collaboration (Murphy et al. 2013, PASA, 30, 6).

ATCA observations of SN2020oi (ATel #13403)

Dougal Dobie (University of Sydney/CSIRO), Assaf Horesh (HUJI),
Andrew O'Brien (University of Wisconsin, Milwaukee), Tara Murphy (University of Sydney)
on 14 Jan 2020; 08:44 UT
Credential Certification: Dougal Dobie (ddob1600@uni.sydney.edu.au)

We have performed preliminary analysis of observations of SN2020oi (TNS Report 58241) with the Australia Telescope Compact Array on 2020-01-11.

We report the detection of a point source spatially coincident with the supernova with a flux density of ~1 mJy at 21.2 GHz.

Analysis is ongoing and further observations are planned.

Thank you to CSIRO staff for supporting these observations.

Bibliography

Abbott B. P., et al., 2016a, Phys. Rev. D, 93, 122008

Abbott B. P., et al., 2016b, Phys. Rev. Lett., 116, 061102

Abbott B. P., et al., 2016c, Phys. Rev. Lett., 116, 241103

Abbott B. P., et al., 2016d, ApJS, 225, 8

Abbott B. P., et al., 2016e, ApJ, 826, L13

Abbott B. P., et al., 2017a, Classical and Quantum Gravity, 34, 044001

Abbott B. P., et al., 2017b, Phys. Rev. Lett., 118, 221101

Abbott B. P., et al., 2017c, Physical Review Letters, 119, 141101

Abbott B. P., et al., 2017d, Phys. Rev. Lett., 119, 161101

Abbott B. P., et al., 2017e, Nature, 551, 85

Abbott B. P., et al., 2017f, ApJ, 848, L12

Abbott B. P., et al., 2017g, ApJ, 848, L13

Abbott B. P., et al., 2017h, ApJ, 851, L35

Abbott B. P., et al., 2018, Living Reviews in Relativity, 21, 3

Abbott B. P., et al., 2019a, Physical Review X, 9, 011001

Abbott B. P., et al., 2019b, Physical Review X, 9, 031040

Abbott B. P., et al., 2019c, Phys. Rev. Lett., 123, 011102

Abbott B. P., et al., 2019d, ApJ, 875, 160

Abbott B. P., et al., 2020a, ApJ, 892, L3

Abbott R., et al., 2020b, ApJ, 896, L44

Acernese F., et al., 2015, Classical and Quantum Gravity, 32, 024001

Ackley K., et al., 2020a, PASA, 37, e047

Ackley K., et al., 2020b, A&A, 643, A113

Adams W. S., 1925, The Observatory, 48, 337

Adams E. A. K., van Leeuwen J., 2019, Nature Astronomy, 3, 188

Adams S. M., et al., 2017, GCN, 21816

Adhikari S., et al., 2020, ApJ, 905, 21

Adrián-Martínez S., et al., 2016, Phys. Rev. D, 93, 122010

Alexander K. D., et al., 2017a, ApJ, 848, L21

Alexander K. D., et al., 2017b, ApJ, 848, 69

Alexander K. D., et al., 2018, ApJ, 863, L18

Amaro-Seoane P., et al., 2017, arXiv e-prints, p. arXiv:1702.00786

Amati L., et al., 2018, Advances in Space Research, 62, 191

Anderson M. M., et al., 2018, ApJ, 864, 22

Andreoni I., et al., 2017, PASA, 34, e069

Andreoni I., et al., 2019, PASP, 131, 068004

Andreoni I., et al., 2020, ApJ, 890, 131

Andrews J. J., et al., 2020, ApJ, 892, L9

Arcavi I., et al., 2017a, Nature, 551, 64

Arcavi I., et al., 2017b, ApJ, 848, L33

Armstrong J. W., Rickett B. J., Spangler S. R., 1995, ApJ, 443, 209

Artale M. C., et al., 2019, MNRAS, 487, 1675

Artale M. C., et al., 2020a, MNRAS, 491, 3419

Artale M. C., et al., 2020b, MNRAS, 495, 1841

Aso Y., et al., 2013, Phys. Rev. D, 88, 043007

Astropy Collaboration et al., 2018, AJ, 156, 123

Bailes M., et al., 2017, PASA, 34, e045

Bailes M., et al., 2018, arXiv e-prints, p. arXiv:1803.07424

Ballo L., et al., 2012, A&A, 545, A66

Bannister K. W., et al., 2011, MNRAS, 412, 634

Bannister K. W., et al., 2017a, ApJ, 841, L12

Bannister K., et al., 2017b, GCN, 21562

Bannister K., et al., 2017c, GCN, 21671

Bannister K. W., et al., 2019, Science, 365, 565

Barkov M. V., et al., 2018, arXiv e-prints, p. arXiv:1805.08338

Barnes J., et al., 2016, ApJ, 829, 110

Bartos I., et al., 2014, MNRAS, 443, 738

Bell M. E., et al., 2016, MNRAS, 461, 908

Bellm E. C., Sesar B., 2016, pyraf-dbsp: Reduction pipeline for the Palomar Double Beam Spectrograph (ascl:1602.002)

Beloborodov A. M., 2017, ApJ, 843, L26

Beniamini P., et al., 2019, MNRAS, 483, 840

Beniamini P., et al., 2020, ApJ, 895, L33

Berger E., 2009, ApJ, 690, 231

Berger E., 2014, ARA&A, 52, 43

Berger E., Kulkarni S. R., Frail D. A., 2004, ApJ, 612, 966

Berger E., et al., 2005, Nature, 438, 988

Bhandari S., et al., 2018, MNRAS, 478, 1784

Bhandari S., et al., 2020, ApJ, 895, L37

Bhattacharya M., Kumar P., Smoot G., 2019, MNRAS, 486, 5289

Bilicki M., et al., 2014, ApJS, 210, 9

Bilicki M., et al., 2016, ApJS, 225, 5

Bionta R. M., et al., 1987, Phys. Rev. Lett., 58, 1494

Blanchard P. K., et al., 2017, ApJ, 848, L22

Blandford R. D., McKee C. F., 1976, Physics of Fluids, 19, 1130

Bloom J. S., Kulkarni S. R., Djorgovski S. G., 2002, AJ, 123, 1111

Bloom J. S., et al., 2012, PASP, 124, 1175

Bochenek C. D., et al., 2020, Nature, 587, 59

Bock D. C.-J., Large M. I., Sadler E. M., 1999, AJ, 117, 1578

Boella G., et al., 1997, A&AS, 122, 299

Brahe T., 1573, De nova et nullius aevi memoria prius visa stella iam pridem anno a nato Christo 1572 mense novembri primium conspecta.

Braun R., et al., 2019, arXiv e-prints, p. arXiv:1912.12699

Bucciantini N., et al., 2012, MNRAS, 419, 1537

CHIME/FRB Collaboration et al., 2018, ApJ, 863, 48

CHIME/FRB Collaboration et al., 2019a, Nature, 566, 235

CHIME/FRB Collaboration et al., 2019b, ApJ, 885, L24

CHIME/FRB Collaboration et al., 2020, Nature, 587, 54

Caleb M., Spitler L. G., Stappers B. W., 2018, Nature Astronomy, 2, 839

Callister T. A., et al., 2019, ApJ, 877, L39

Carbone D., Corsi A., 2018, ApJ, 867, 135

Cervantes-Cota J., Galindo-Uribarri S., Smoot G., 2016, Universe, 2, 22

Chan M. L., et al., 2018, Phys. Rev. D, 97, 123014

Chandra P., et al., 2008, ApJ, 683, 924

Chawla P., et al., 2020, ApJ, 896, L41

Chen H.-Y., Fishbach M., Holz D. E., 2018, Nature, 562, 545

Chornock R., et al., 2017, ApJ, 848, L19

Chu Q., et al., 2016, MNRAS, 459, 121

Cohen E., Piran T., Sari R., 1998, ApJ, 509, 717

Colgate S. A., 1968, Canadian Journal of Physics, 46, S476

Colless M., et al., 2001, MNRAS, 328, 1039

Condon J. J., 1997, PASP, 109, 166

Condon J. J., et al., 1998, AJ, 115, 1693

Contini M., 2018, A&A, 620, A37

Cook D. O., et al., 2017, GCN, 21519

Cook D. O., et al., 2019, ApJ, 880, 7

Cordes J. M., 1986, ApJ, 311, 183

Cordes J. M., Lazio T. J. W., 2002, arXiv e-prints, pp astro–ph/0207156

Cordes J. M., Lazio T. J. W., 2003, arXiv e-prints, pp astro–ph/0301598

Corley K. R., et al., 2019, MNRAS, 488, 4459

Corsi A., et al., 2017, arXiv e-prints, p. arXiv:1711.06873

Corsi A., et al., 2018a, Compact Binary Mergers as Traced by Gravitational Waves. p. 689

Corsi A., et al., 2018b, ApJ, 861, L10

Corsi A., et al., 2019a, arXiv e-prints, p. arXiv:1903.10589

Corsi A., et al., 2019b, BAAS, 51, 209

Costa E., et al., 1997, Nature, 387, 783

Coughlin M., Stubbs C., 2016, Experimental Astronomy, 42, 165

Coughlin M. W., et al., 2018, MNRAS, 478, 692

Coughlin M. W., et al., 2020, MNRAS, 497, 1181

Coulter D. A., et al., 2017, Science, 358, 1556

Cowperthwaite P. S., et al., 2017, ApJ, 848, L17

Cromartie H. T., et al., 2019, Nature Astronomy, p. 439

Cutri R. M., et al. 2012, VizieR Online Data Catalog, p. II/311

D'Avanzo P., et al., 2018, A&A, 613, L1

Dalal N., et al., 2006, Phys. Rev. D, 74, 063006

Dálya G., et al., 2018, MNRAS, 479, 2374

De K., et al., 2019, GCN, 25449

De K., et al., 2020, PASP, 132, 025001

DeLaunay J. J., et al., 2016, ApJ, 832, L1

Del Pozzo W., 2012, Phys. Rev. D, 86, 043011

Dey A., et al., 2019, AJ, 157, 168

Dhawan S., et al., 2020, ApJ, 888, 67

Díaz M. C., et al., 2017, ApJ, 848, L29

Dietrich T., Ujevic M., 2017, Classical and Quantum Gravity, 34, 105014

Dobie D., et al., 2017a, GCN, 21625

Dobie D., et al., 2017b, GCN, 21639

Dobie D., et al., 2018a, ApJ, 858, L15

Dobie D., et al., 2018b, The Astronomer's Telegram, 11795

Dobie D., et al., 2018c, The Astronomer's Telegram, 11818

Dobie D., et al., 2018d, The Astronomer's Telegram, 11862

Dobie D., et al., 2018e, GCN, 23139

Dobie D., et al., 2019a, PASA, 36, e019

Dobie D., et al., 2019b, ApJ, 887, L13

Dobie D., et al., 2019c, The Astronomer's Telegram, 12981

Dobie D., et al., 2019d, GCN, 25445

Dobie D., et al., 2019e, GCN, 25472

Dobie D., et al., 2019f, GCN, 25621

Dobie D., et al., 2019g, GCN, 25691

Dobie D., et al., 2020a, MNRAS, 494, 2449

Dobie D., et al., 2020b, The Astronomer's Telegram, 13403

Dobie D., et al., 2020c, GCN, 27516

Doctor Z., 2020, ApJ, 892, L16

Dolag K., et al., 2015, MNRAS, 451, 4277

Doré O., et al., 2016, arXiv e-prints, p. arXiv:1606.07039

Drout M. R., et al., 2017, Science, 358, 1570

Ducoin J. G., et al., 2020, MNRAS, 492, 4768

Duffell P. C., MacFadyen A. I., 2013, ApJ, 776, L9

Duque R., Daigne F., Mochkovitch R., 2019, A&A, 631, A39

Dyson F. W., Eddington A. S., Davidson C., 1920, Philosophical Transactions of the Royal Society of London Series A, 220, 291

Ebrová I., et al., 2020, A&A, 634, A73

Eichler D., et al., 1989, Nature, 340, 126

Evans P. A., et al., 2016, MNRAS, 462, 1591

Evans P. A., et al., 2017, Science, 358, 1565

Event Horizon Telescope Collaboration 2019, ApJ, 875, L1

Ezquiaga J. M., Zumalacárregui M., 2017, Phys. Rev. Lett., 119, 251304

Falcke H., Rezzolla L., 2014, A&A, 562, A137

Farr W. M., et al., 2011, ApJ, 741, 103

Farr W. M., et al., 2019, ApJ, 883, L42

Feeney S. M., et al., 2019, Phys. Rev. Lett., 122, 061105

Fender R., et al., 2017, arXiv e-prints, p. arXiv:1711.04132

Feng L., Vaulin R., Hewitt J. N., 2014, arXiv e-prints, p. arXiv:1405.6219

Finn L. S., Chernoff D. F., 1993, Phys. Rev. D, 47, 2198

Fishbach M., et al., 2019, ApJ, 871, L13

Flaugher B., et al., 2015, AJ, 150, 150

Fong W., Berger E., 2013, ApJ, 776, 18

Fong W., Berger E., Fox D. B., 2010, ApJ, 708, 9

Fong W., et al., 2015, ApJ, 815, 102

Foucart F., 2012, Phys. Rev. D, 86, 124007

Foucart F., et al., 2014, Phys. Rev. D, 90, 024026

Foucart F., et al., 2016, Phys. Rev. D, 94, 123016

Fox D. B., et al., 2005, Nature, 437, 845

Frail D. A., et al., 1997, Nature, 389, 261

Frail D. A., Waxman E., Kulkarni S. R., 2000, ApJ, 537, 191

Freiburghaus C., et al., 1999, ApJ, 516, 381

Freise A., et al., 2009, Classical and Quantum Gravity, 26, 085012

Fruscione A., et al., 2006, in Proc. SPIE. p. 62701V, doi:10.1117/12.671760

Fryer C. L., et al., 2019, European Physical Journal A, 55, 132

Galama T. J., et al., 1998, Nature, 395, 670

Gehrels N., et al., 2005, Nature, 437, 851

Gehrels N., et al., 2016, ApJ, 820, 136

Ghirlanda G., et al., 2013, MNRAS, 435, 2543

Ghirlanda G., et al., 2014, PASA, 31, e022

Ghirlanda G., et al., 2019, Science, 363, 968

Ghosh S., et al., 2016, A&A, 592, A82

Ghosh S., et al., 2017, Publications of the Astronomical Society of the Pacific, 129, 114503

Giacomazzo B., et al., 2015, ApJ, 809, 39

Gill A. E., 1965, Physics of Fluids, 8, 1428

Gill R., Granot J., 2018, MNRAS, 478, 4128

Gill R., Granot J., 2020, MNRAS, 491, 5815

Glowacki M., et al., 2017, arXiv e-prints, p. arXiv:1709.08634

Goldstein D. A., et al., 2015, AJ, 150, 82

Goldstein A., et al., 2017, ApJ, 848, L14

Goldstein D. A., et al., 2019, ApJ, 881, L7

Gomez S., et al., 2019, ApJ, 884, L55

Gompertz B. P., et al., 2020, MNRAS, 497, 726

Goodman J., 1997, New Astron., 2, 449

Gottlieb O., Nakar E., Piran T., 2018a, MNRAS, 473, 576

Gottlieb O., et al., 2018b, MNRAS, 479, 588

Gottlieb O., Nakar E., Piran T., 2019, MNRAS, 488, 2405

Gottlieb O., Nakar E., Bromberg O., 2021, MNRAS, 500, 3511

Graham M. J., et al., 2020, Phys. Rev. Lett., 124, 251102

Granot J., Sari R., 2002, ApJ, 568, 820

Granot J., Piran T., Sari R., 1999, ApJ, 513, 679

Granot J., et al., 2002, ApJ, 570, L61

Granot J., Guetta D., Gill R., 2017, ApJ, 850, L24

Granot J., De Colle F., Ramirez-Ruiz E., 2018, MNRAS, 481, 2711

Gray R., et al., 2020, Phys. Rev. D, 101, 122001

Guidorzi C., et al., 2017, ApJ, 851, L36

Gupte N., Bartos I., 2020, Phys. Rev. D, 101, 123008

Haggard D., et al., 2017, ApJ, 848, L25

Haggard D., et al., 2018, The Astronomer's Telegram, 11242

Hajela A., et al., 2019, ApJ, 886, L17

Hajela A., et al., 2020, Research Notes of the American Astronomical Society, 4, 68

Hall E. D., Evans M., 2019, Classical and Quantum Gravity, 36, 225002

Hallinan G., et al., 2017, Science, 358, 1579

Hallinan G., et al., 2019, in BAAS. p. 255 (`arXiv:1907.07648`)

Hamacher D. W., 2014, Journal of Astronomical History and Heritage, 17, 161

Hamacher D. W., 2018, The Australian Journal of Anthropology, 29, 89

Hampson G., et al., 2012, in 2012 International Conference on Electromagnetics in Advanced Applications. pp 807–809, doi:10.1109/ICEAA.2012.6328742

Hamuy M., et al., 2000, AJ, 120, 1479

Hancock P. J., et al., 2016, MNRAS, 461, 3314

Hancock P. J., Trott C. M., Hurley-Walker N., 2018, PASA, 35, e011

Hancock P. J., et al., 2019a, arXiv e-prints, p. arXiv:1907.08395

Hancock P. J., et al., 2019b, PASA, 36, e046

Helfand D. J., et al., 1999, AJ, 117, 1568

Herner K., et al., 2019, GCN, 25495

Hildebrandt H., Wolf C., Benítez N., 2008, A&A, 480, 703

Hinshaw G., et al., 2013, ApJS, 208, 19

Hirata K., et al., 1987, Phys. Rev. Lett., 58, 1490

Hjorth J., Bloom J. S., 2012, The Gamma-Ray Burst - Supernova Connection. Cambridge University Press, pp 169–190

Hjorth J., et al., 2005, Nature, 437, 859

Hjorth J., et al., 2017, ApJ, 848, L31

Ho A. Y. Q., et al., 2019, ApJ, 871, 73

Hobbs G., et al., 2016, MNRAS, 456, 3948

Holwerda B. W., Blyth S. L., Baker A. J., 2012, in Tuffs R. J., Popescu C. C., eds, IAU Symposium Vol. 284, The Spectral Energy Distribution of Galaxies - SED 2011. pp 496–499 (`arXiv:1109.5605`), doi:10.1017/S1743921312009702

Horesh A., et al., 2020, ApJ, 903, 132

Hotokezaka K., Piran T., 2015, MNRAS, 450, 1430

Hotokezaka K., et al., 2013a, Phys. Rev. D, 88, 044026

Hotokezaka K., et al., 2013b, ApJ, 778, L16

Hotokezaka K., Piran T., Paul M., 2015, Nature Physics, 11, 1042

Hotokezaka K., et al., 2016, ApJ, 831, 190

Hotokezaka K., et al., 2018, ApJ, 867, 95

Hotokezaka K., et al., 2019, Nature Astronomy, 3, 940

Howell E. J., et al., 2019, MNRAS, 485, 1435

Howlett C., Davis T. M., 2020, MNRAS, 492, 3803

Hu W.-R., Wu Y.-L., 2017, National Science Review, 4, 685

Hulse R. A., Taylor J. H., 1975, ApJ, 195, L51

Im M., et al., 2017, ApJ, 849, L16

Inoue S., 2004, MNRAS, 348, 999

Ioka K., 2003, ApJ, 598, L79

Iyer B., et al., 2011, Tech. Rep. M1100296-v2, LIGO-India. IndIGO

James C. W., et al., 2019, MNRAS, 489, L75

Jaodand A., et al., 2019, GCN, 25822

Johnston S., et al., 2008, Experimental Astronomy, 22, 151

Jonas J., et al., 2016, in MeerKAT Science: On the Pathway to the SKA. p. 1

KAGRA Collaboration et al., 2019, Nature Astronomy, 3, 35

KAGRA Collaboration et al., 2020, arXiv e-prints, p. arXiv:2008.02921

Kapadia S. J., et al., 2020, ApJ, 898, L39

Kaplan D. L., et al., 2015, ApJ, 814, L25

Kaplan D. L., et al., 2016, Publications of the Astronomical Society of Australia, 33

Kaplan D. L., et al., 2019, ApJ, 884, 96

Kardashev N. S., et al., 2013, Astronomy Reports, 57, 153

Kasen D., Fernández R., Metzger B. D., 2015, MNRAS, 450, 1777

Kasen D., et al., 2017, Nature, 551, 80

Kasliwal M. M., et al., 2017, Science, 358, 1559

Kasliwal M. M., et al., 2020, ApJ, 905, 145

Kathirgamaraju A., Barniol Duran R., Giannios D., 2018, MNRAS, 473, L121

Kathirgamaraju A., et al., 2019a, MNRAS, 484, L98

Kathirgamaraju A., Giannios D., Beniamini P., 2019b, MNRAS, 487, 3914

Kauffmann G., et al., 2003, MNRAS, 346, 1055

Kim S., et al., 2017, ApJ, 850, L21

Klebesadel R. W., Strong I. B., Olson R. A., 1973, ApJ, 182, L85

Klimenko S., et al., 2011, Phys. Rev. D, 83, 102001

Klose S., et al., 2019, ApJ, 887, 206

Kobayashi S., 2000, ApJ, 545, 807

Kobayashi S., Sari R., 2000, ApJ, 542, 819

Kocevski D., et al., 2019, GCN, 25326

Kopparapu R. K., et al., 2008, ApJ, 675, 1459

Koribalski B. S., et al., 2020, Ap&SS, 365, 118

Kouveliotou C., et al., 1993, ApJ, 413, L101

Kulkarni S. R., et al., 1999, ApJ, 522, L97

Kyutoku K., et al., 2011, Phys. Rev. D, 84, 064018

Kyutoku K., Ioka K., Shibata M., 2013, Phys. Rev. D, 88, 041503

Kyutoku K., Nishino Y., Seto N., 2019, MNRAS, 483, 2615

LIGO Scientific Collaboration 2015, Classical and Quantum Gravity, 32, 074001

LIGO Scientific Collaboration 2019, Tech. Rep. T1900409-v5, Instrument Science White Paper 2019. LIGO

LIGO Scientific Collaboration Virgo Collaboration 2017a, GCN, 21513

LIGO Scientific Collaboration Virgo Collaboration 2017b, GCN, 21983

LIGO Scientific Collaboration and Virgo Collaboration et al., 2019a, GCN, 25333

LIGO Scientific Collaboration and Virgo Collaboration et al., 2019b, GCN, 26402

LSST Science Collaboration et al., 2009, arXiv e-prints, p. arXiv:0912.0201

Lamb G. P., Kobayashi S., 2016, The Astrophysical Journal, 829, 112

Lamb G. P., Kobayashi S., 2019, MNRAS, 489, 1820

Lamb G. P., Mandel I., Resmi L., 2018, MNRAS, 481, 2581

Lamb G. P., et al., 2019a, ApJ, 870, L15

Lamb G. P., et al., 2019b, ApJ, 883, 48

Lang D., Hogg D. W., Mykytyn D., 2016, The Tractor: Probabilistic astronomical source detection and measurement (ascl:1604.008)

Laskar T., et al., 2013, ApJ, 776, 119

Laskar T., et al., 2019, ApJ, 884, 121

Lattimer J. M., Schramm D. N., 1974, ApJ, 192, L145

Lau M. Y. M., et al., 2020, MNRAS, 492, 3061

Law C. J., et al., 2018, ApJ, 866, L22

Lawrence E., et al., 2017, AJ, 154, 117

Lazzati D., et al., 2017, MNRAS, 471, 1652

Lazzati D., et al., 2018, Phys. Rev. Lett., 120, 241103

Lazzati D., Ciolfi R., Perna R., 2020, ApJ, 898, 59

Lee W. H., Ramirez-Ruiz E., 2007, New Journal of Physics, 9, 17

Levan A. J., et al., 2017a, ApJ, 848, L28

Levan A. J., et al., 2017b, GCN, 22207

Levesque E. M., et al., 2010, AJ, 139, 694

Levin Y., Beloborodov A. M., Bransgrove A., 2020, ApJ, 895, L30

Levinson A., et al., 2002, ApJ, 576, 923

Lien A., et al., 2016, ApJ, 829, 7

Lin H., Totani T., 2020, MNRAS, 498, 2384

Lin H., Totani T., Kiuchi K., 2019, MNRAS, 485, 2155

Linden S. T., et al., 2020, ApJS, 248, 25

Linial I., Sari R., 2019, MNRAS, 483, 624

Lipunov V. M., Panchenko I. E., 1996, A&A, 312, 937

Lipunov V. M., Postnov K. A., Prokhorov M. E., 2001, Astronomy Reports, 45, 236

Lipunov V. M., et al., 2017, ApJ, 850, L1

Liu T., et al., 2016, ApJ, 826, 82

Lloyd-Ronning N., 2018, Galaxies, 6, 103

Lorimer D. R., et al., 2007, Science, 318, 777

Lovell J. E. J., et al., 2008, ApJ, 689, 108

Lyman J. D., et al., 2018, Nature Astronomy, 2, 751

Macquart J. P., et al., 2019, ApJ, 872, L19

Macquart J. P., et al., 2020, Nature, 581, 391

Mahabal A., et al., 2008, Astronomische Nachrichten, 329, 288

Mahabal A., et al., 2019, PASP, 131, 038002

Mahony E. K., et al., 2018, ApJ, 867, L10

Makarov D., et al., 2014, A&A, 570, A13

Makhathini S., et al., 2020, arXiv e-prints, p. arXiv:2006.02382

Mandel I., 2018, ApJ, 853, L12

Margalit B., Piran T., 2015, Monthly Notices of the Royal Astronomical Society, 452, 3419

Margalit B., Piran T., 2020, MNRAS, 495, 4981

Margalit B., Berger E., Metzger B. D., 2019, ApJ, 886, 110

Margalit B., et al., 2020, ApJ, 899, L27

Margutti R., et al., 2017, ApJ, 848, L20

Margutti R., et al., 2018a, arXiv e-prints, p. arXiv:1812.04051

Margutti R., et al., 2018b, ApJ, 856, L18

Mauch T., et al., 2003, MNRAS, 342, 1117

Mauch T., et al., 2020, ApJ, 888, 61

McConnell D., et al., 2016, PASA, 33, e042

McConnell D., et al., 2020, PASA, 37, e048

McKinnon M., et al., 2019, in BAAS. p. 81

McMullin J. P., et al., 2007, in Shaw R. A., Hill F., Bell D. J., eds, Astronomical Society of the Pacific Conference Series Vol. 376, Astronomical Data Analysis Software and Systems XVI. p. 127

Meegan C. A., et al., 1992, Nature, 355, 143

Messenger C., Read J., 2012, Phys. Rev. Lett., 108, 091101

Metzger B. D., 2017a, arXiv e-prints, p. arXiv:1710.05931

Metzger B. D., 2017b, Living Reviews in Relativity, 20, 3

Metzger B. D., Berger E., 2012, ApJ, 746, 48

Metzger B. D., Zivancev C., 2016, MNRAS, 461, 4435

Metzger B. D., Quataert E., Thompson T. A., 2008, MNRAS, 385, 1455

Metzger B. D., et al., 2010, MNRAS, 406, 2650

Metzger B. D., et al., 2015a, MNRAS, 446, 1115

Metzger B. D., Williams P. K. G., Berger E., 2015b, ApJ, 806, 224

Metzger B. D., Berger E., Margalit B., 2017, ApJ, 841, 14

Meyer M., et al., 2017, PASA, 34, 52

Michelson A. A., Morley E. W., 1887, American Journal of Science, 34, 333

Modjaz M., et al., 2008, AJ, 135, 1136

Modjaz M., et al., 2016, ApJ, 832, 108

Molkov S., et al., 2019, GCN, 25323

Mooley K. P., et al., 2013, ApJ, 768, 165

Mooley K. P., et al., 2016, ApJ, 818, 105

Mooley K. P., et al., 2017, GCN, 22211

Mooley K. P., et al., 2018a, Nature, 554, 207

Mooley K. P., et al., 2018b, Nature, 561, 355

Mooley K. P., et al., 2018c, ApJ, 857, 143

Mooley K. P., et al., 2018d, ApJ, 868, L11

Mooley K. P., et al., 2019, GCN, 25690

Mortlock D. J., et al., 2019, Phys. Rev. D, 100, 103523

Mukherjee S., et al., 2019, arXiv e-prints, p. arXiv:1909.08627

Murguia-Berthier A., et al., 2014, ApJ, 788, L8

Murguia-Berthier A., et al., 2017, ApJ, 848, L34

Murphy T., et al., 2013, PASA, 30, e006

Murphy E. J., et al., 2017, ApJ, 839, 35

Murphy E. J., et al., 2018, The ngVLA Science Case and Associated Science Requirements. p. 3

Nagakura H., et al., 2014, ApJ, 784, L28

Nair R., Bose S., Saini T. D., 2018, Phys. Rev. D, 98, 023502

Nakar E., Piran T., 2011, Nature, 478, 82

Nakar E., Piran T., 2017, ApJ, 834, 28

Nakar E., Piran T., 2018, MNRAS, 478, 407

Nakar E., Piran T., 2020, arXiv e-prints, p. arXiv:2005.01754

Nakar E., Piran T., Granot J., 2002, ApJ, 579, 699

Nakar E., et al., 2018, ApJ, 867, 18

Narayan R., 1992, Philosophical Transactions of the Royal Society of London Series A, 341, 151

Narayan R., Paczynski B., Piran T., 1992, ApJ, 395, L83

Nathanail A., et al., 2020, MNRAS, 495, 3780

Nissanke S., et al., 2010, ApJ, 725, 496

Nissanke S., et al., 2013a, arXiv e-prints, p. arXiv:1307.2638

Nissanke S., Kasliwal M., Georgieva A., 2013b, ApJ, 767, 124

Norris R. P., et al., 2011, PASA, 28, 215

Novak M., et al., 2017, A&A, 602, A5

Nyland K., et al., 2020, arXiv e-prints, p. arXiv:2005.04734

O'Shaughnessy R., et al., 2017, MNRAS, 464, 2831

Oke J. B., Gunn J. E., 1982, PASP, 94, 586

Oosterloo T., Verheijen M., van Cappellen W., 2010, in ISKAF2010 Science Meeting. p. 43 (`arXiv:1007.5141`)

Ott C. D., 2009, Classical and Quantum Gravity, 26, 063001

Özel F., et al., 2012, ApJ, 757, 55

Paczynski B., 1986, ApJ, 308, L43

Paczyński B., 1993, in Akerlof C. W., Srednicki M. A., eds, Vol. 688, Texas/PASCOS '92: Relativistic Astrophysics and Particle Cosmology. p. 321, doi:10.1111/j.1749-6632.1993.tb43907.x

Paczynski B., Rhoads J. E., 1993, ApJ, 418, L5

Palmer D. M., et al., 2019, GCN, 25341

Pan Y. C., et al., 2017, ApJ, 848, L30

Panaitescu A., Kumar P., 2004, MNRAS, 350, 213

Panaitescu A., Mészáros P., 1999, ApJ, 526, 707

Perley R. A., 1999, in Taylor G. B., Carilli C. L., Perley R. A., eds, Astronomical Society of the Pacific Conference Series Vol. 180, Synthesis Imaging in Radio Astronomy II. p. 275

Perley D. A., Perley R. A., 2013, ApJ, 778, 172

Perley D. A., et al., 2014, ApJ, 781, 37

Petroff E., et al., 2016, PASA, 33, e045

Phillips J. A., Clegg A. W., 1992, Nature, 360, 137

Pian E., et al., 2017, Nature, 551, 67

Pilia M., et al., 2019, GCN, 25327

Piran T., 1994, in Fishman G. J., ed., American Institute of Physics Conference Series Vol. 307, Gamma-Ray Bursts. p. 495, doi:10.1063/1.45856

Piran T., Nakar E., Rosswog S., 2013, MNRAS, 430, 2121

Piro A. L., Kollmeier J. A., 2018, ApJ, 855, 103

Piro L., et al., 2019, MNRAS, 483, 1912

Planck Collaboration et al., 2020, A&A, 641, A6

Popper D. M., 1954, ApJ, 120, 316

Pradel N., Charlot P., Lestrade J. F., 2006, A&A, 452, 1099

Prix R., 2009, Gravitational Waves from Spinning Neutron Stars. Springer Berlin Heidelberg, p. 651, doi:10.1007/978-3-540-76965-1_24

Pshirkov M. S., Postnov K. A., 2010, Ap&SS, 330, 13

Punturo M., et al., 2010, Classical and Quantum Gravity, 27, 194002

Radcliffe J. F., et al., 2019, MNRAS, 490, 4024

Radice D., et al., 2018, ApJ, 869, L35

Radice D., et al., 2019, ApJ, 876, L9

Rana J., Mooley K. P., 2019, arXiv e-prints, p. arXiv:1904.07335

Ravi V., 2019, Nature Astronomy, 3, 928

Reid M. J., Pesce D. W., Riess A. G., 2019, ApJ, 886, L27

Reitze D., et al., 2019, in BAAS. p. 35 (`arXiv:1907.04833`)

Resmi L., Zhang B., 2016, ApJ, 825, 48

Resmi L., et al., 2018, ApJ, 867, 57

Rhoads J. E., 1999, ApJ, 525, 737

Rhodes L., et al., 2020, MNRAS,

Rickett B. J., 1977, ARA&A, 15, 479

Rickett B. J., et al., 1995, A&A, 293, 479

Riess A. G., et al., 1998, AJ, 116, 1009

Riess A. G., et al., 2018, ApJ, 861, 126

Riess A. G., et al., 2019, ApJ, 876, 85

Rossi E., Lazzati D., Rees M. J., 2002, MNRAS, 332, 945

Rossi E. M., et al., 2004, MNRAS, 354, 86

Rowlinson A., et al., 2013, MNRAS, 430, 1061

Rowlinson A., et al., 2019a, Astronomy and Computing, 27, 111

Rowlinson A., et al., 2019b, MNRAS, 490, 3483

Ruan W.-H., et al., 2018a, arXiv e-prints, p. arXiv:1807.09495

Ruan J. J., et al., 2018b, ApJ, 853, L4

Ruan W.-H., et al., 2020, Nature Astronomy, 4, 108

Ruderman M., 1975, in Bergman P. G., Fenyves E. J., Motz L., eds, Vol. 262, Seventh Texas Symposium on Relativistic Astrophysics. pp 164–180, doi:10.1111/j.1749-6632.1975.tb31430.x

Ryan G., et al., 2020, ApJ, 896, 166

Sadler E. M., et al., 2006, MNRAS, 371, 898

Salafia O. S., et al., 2015, MNRAS, 450, 3549

Salafia O. S., et al., 2017, ApJ, 846, 62

Salafia O. S., Ghisellini G., Ghirlanda G., 2018, MNRAS, 474, L7

Salafia O. S., et al., 2019, A&A, 628, A18

Salafia O. S., et al., 2020, A&A, 636, A105

Saleem M., 2020, MNRAS, 493, 1633

Sanders R. H., 2018, International Journal of Modern Physics D, 27, 1847027

Sari R., 1997, ApJ, 489, L37

Sari R., 1998, ApJ, 494, L49

Sari R., 1999, ApJ, 524, L43

Sari R., Piran T., 1995, ApJ, 455, L143

Sari R., Piran T., Narayan R., 1998, ApJ, 497, L17

Sari R., Piran T., Halpern J. P., 1999, ApJ, 519, L17

Sathyaprakash B., et al., 2013, Classical and Quantum Gravity, 30, 079501

Sault R. J., Teuben P. J., Wright M. C. H., 1995, in Shaw R. A., Payne H. E., Hayes J. J. E., eds, Astronomical Society of the Pacific Conference Series Vol. 77, Astronomical Data Analysis Software and Systems IV. p. 433 (`arXiv:astro-ph/0612759`)

Savchenko V., et al., 2017, ApJ, 848, L15

Schutz B. F., 1986, Nature, 323, 310

Schutz B. F., 2011, Classical and Quantum Gravity, 28, 125023

Selina R. J., et al., 2018, The ngVLA Reference Design. p. 15

Serra P., et al., 2016, in MeerKAT Science: On the Pathway to the SKA. p. 8 (`arXiv:1709.01289`)

Sesana A., 2016, Phys. Rev. Lett., 116, 231102

Shannon R. M., et al., 2018, Nature, 562, 386

Shepherd M. C., 1997, in Hunt G., Payne H., eds, Astronomical Society of the Pacific Conference Series Vol. 125, Astronomical Data Analysis Software and Systems VI. p. 77

Singer L. P., 2015, PhD thesis, California Institute of Technology

Singer L. P., Price L. R., 2016, Phys. Rev. D, 93, 024013

Singer L. P., et al., 2014, ApJ, 795, 105

Smartt S. J., et al., 2017, Nature, 551, 75

Soares-Santos M., et al., 2017, ApJ, 848, L16

Soares-Santos M., et al., 2019, ApJ, 876, L7

Soderberg A. M., et al., 2006, ApJ, 638, 930

Sokolowski M., et al., 2018, ApJ, 867, L12

Somiya K., 2012, Classical and Quantum Gravity, 29, 124007

Soon J., et al., 2020, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series. p. 1120307, doi:10.1117/12.2539594

Spitler L. G., et al., 2016, Nature, 531, 202

Stachie C., et al., 2020, MNRAS,

Staley T. D., et al., 2013, MNRAS, 428, 3114

Stanway E. R., Levan A. J., Davies L. J. M., 2014, MNRAS, 444, 2133

Stewart A. J., et al., 2016, MNRAS, 456, 2321

Sugizaki M., et al., 2019, GCN, 25329

Sullivan M., et al., 2006, ApJ, 648, 868

Swinbank J. D., et al., 2015, Astronomy and Computing, 11, 25

Symbalisty E., Schramm D. N., 1982, Astrophys. Lett., 22, 143

Takahashi K., Ioka K., 2020, arXiv e-prints, p. arXiv:2007.13116

Tanvir N. R., et al., 2017, ApJ, 848, L27

Taylor J. H., Cordes J. M., 1993, ApJ, 411, 674

Taylor J. H., Weisberg J. M., 1989, ApJ, 345, 434

Taylor G. B., et al., 2012a, Journal of Astronomical Instrumentation, 1, 1250004

Taylor S. R., Gair J. R., Mandel I., 2012b, Phys. Rev. D, 85, 023535

Thornton D., et al., 2013, Science, 341, 53

Thrane E., Osłowski S., Lasky P. D., 2020, MNRAS, 493, 5408

Tingay S. J., et al., 2013, Publications of the Astronomical Society of Australia, 30, e007

Toffano M., et al., 2019, MNRAS, 489, 4622

Totani T., 2013, PASJ, 65, L12

Troja E., Piro L., 2018, The Astronomer's Telegram, 11245

Troja E., et al., 2017, Nature, 551, 71

Troja E., et al., 2018, MNRAS, 478, L18

Troja E., et al., 2019a, MNRAS, 489, 1919

Troja E., et al., 2019b, MNRAS, 489, 2104

Troja E., et al., 2020, MNRAS, 498, 5643

Unnikrishnan C. S., 2013, International Journal of Modern Physics D, 22, 1341010

Valenti S., et al., 2017, ApJ, 848, L24

Verde L., Treu T., Riess A. G., 2019, Nature Astronomy, 3, 891

Villadsen J., Hallinan G., 2019, ApJ, 871, 214

Vitale S., Chen H.-Y., 2018, Phys. Rev. Lett., 121, 021303

Voorhees E. M., 1986, Information Processing & Management, 22, 465

Walker M. A., 1998, MNRAS, 294, 307

Walker M. A., 2001, MNRAS, 321, 176

Walsh D., Carswell R. F., Weymann R. J., 1979, Nature, 279, 381

Wang J.-S., et al., 2016, ApJ, 822, L7

Wang Z., et al., 2020, PASA, 37, e051

Wayth R. B., et al., 2018, PASA, 35, 33

Weber J., 1960, Physical Review, 117, 306

Weber J., 1969, Phys. Rev. Lett., 22, 1320

Webster B. L., Murdin P., 1972, Nature, 235, 37

White D. J., Daw E. J., Dhillon V. S., 2011, Classical and Quantum Gravity, 28, 085016

Whiting M., et al., 2017, in Lorente N. P. F., Shortridge K., Wayth R., eds, Astronomical Society of the Pacific Conference Series Vol. 512, Astronomical Data Analysis Software and Systems XXV. p. 431

Wieringa M., Kesteven M., 1992, ATCA Technical Document Series, 39

Williams R. D., Seaman R., 2006, in Gabriel C., et al., eds, Astronomical Society of the Pacific Conference Series Vol. 351, Astronomical Data Analysis Software and Systems XV. p. 637

Wilson J. C., et al., 2003, in Iye M., Moorwood A. F. M., eds, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 4841, Proc. SPIE. pp 451–458, doi:10.1117/12.460336

Wilson W. E., et al., 2011, MNRAS, 416, 832

Wright E. L., et al., 2010, AJ, 140, 1868

Wu Y., MacFadyen A., 2018, The Astrophysical Journal, 869, 55

Wu Y., MacFadyen A., 2019, ApJ, 880, L23

Wyatt S. D., et al., 2020, ApJ, 894, 127

Xie X., Zrake J., MacFadyen A., 2018, ApJ, 863, 58

Yamasaki S., Totani T., Kiuchi K., 2018, PASJ, 70, 39

Yao J. M., Manchester R. N., Wang N., 2017, ApJ, 835, 29

Zhang B., 2014, ApJ, 780, L21

Zhang B., 2016, ApJ, 827, L31

Zhang Z.-B., et al., 2015, Research in Astronomy and Astrophysics, 15, 237

Zhang Z. B., et al., 2018, ApJ, 865, 82

Zhang J., et al., 2019, MNRAS, 488, 5020

Zhao F.-Y., Strom R. G., Jiang S.-Y., 2006, Chinese Journal of Astronomy and Astrophysics, 6, 635

Zheng Z., Ramirez-Ruiz E., 2007, ApJ, 665, 1220

Zhou R., et al., 2020, MNRAS,

Ziaeepour H., 2019, MNRAS, 490, 2822

Zic A., et al., 2019, MNRAS, 488, 559

Zrake J., MacFadyen A. I., 2013, ApJ, 769, L29

de Blok W. J. G., et al., 2016, in MeerKAT Science: On the Pathway to the SKA. p. 7 (`arXiv:1709.08458`)

van Haarlem M. P., et al., 2013, A&A, 556, A2

von Kienlin A., et al., 2019, ApJ, 876, 89