

Strategies for Finding Prompt Radio Counterparts to Gravitational Wave Transients with the Murchison Widefield Array

D. L. Kaplan^{1,9}, T. Murphy^{2,3}, A. Rowlinson^{4,5}, S. D. Croft^{6,7}, R. B. Wayth^{3,8} and C. M. Trott^{3,8}

¹Department of Physics, University of Wisconsin–Milwaukee, Milwaukee, WI 53201, USA

²Sydney Institute for Astronomy, School of Physics, The University of Sydney, NSW 2006, Australia

³ARC Centre of Excellence for All-sky Astrophysics (CAASTRO)

⁴Anton Pannekoek Institute for Astronomy, University of Amsterdam, Science Park 904, 1098 XH Amsterdam, The Netherlands

⁵ASTRON, The Netherlands Institute for Radio Astronomy, Postbus 2, 7990 AA, Dwingeloo, The Netherlands

⁶University of California, Berkeley, Astronomy Dept., 501 Campbell Hall #3411, Berkeley, CA 94720, USA

⁷Eureka Scientific, Inc., 2452 Delmer Street Suite 100, Oakland, CA 94602, USA

⁸International Centre for Radio Astronomy Research, Curtin University, Bentley, WA 6102, Australia

⁹Email: kaplan@uwm.edu

(RECEIVED June 29, 2016; ACCEPTED September 2, 2016)

Abstract

We present and evaluate several strategies to search for prompt, low-frequency radio emission associated with gravitational wave transients using the Murchison Widefield Array. As we are able to repoint the Murchison Widefield Array on timescales of tens of seconds, we can search for the dispersed radio signal that has been predicted to originate along with or shortly after a neutron star–neutron star merger. We find that given the large, 600 deg^2 instantaneous field of view of the Murchison Widefield Array, we can cover a significant fraction of the predicted gravitational wave error region, although due to the complicated geometry of the latter, we only cover $>50\%$ of the error region for approximately 5% of events, and roughly 15% of events will be located $<10^\circ$ from the Murchison Widefield Array pointing centre such that they will be covered in the radio images. For optimal conditions, our limiting flux density for a 10-s long transient would be 0.1 Jy , increasing to about 1 Jy for a wider range of events. This corresponds to luminosity limits of $10^{38-39} \text{ erg s}^{-1}$ based on expectations for the distances of the gravitational wave transients, which should be sufficient to detect or significantly constrain a range of models for prompt emission.

Keywords: gamma-ray burst: general, gravitational waves, methods: observational, radio continuum: general

1 INTRODUCTION

In 2015 September, the LIGO/Virgo Consortium (LVC) began its O1 science run that resulted in the first detection of gravitational waves (GWs) (Abbott et al. 2016c, also see Abbott et al. 2016a for a second event). Together with the GW analysis, the LVC sent private alerts to the electromagnetic (EM) follow-up community (Abbott et al. 2016b, 2016f) to identify coincident EM transients (e.g., Lipunov et al. 2016; Evans et al. 2016b; Troja et al. 2016; Ackermann et al. 2016; Morokuma et al. 2016; Savchenko et al. 2016; Connaughton et al. 2016). This identification is complicated by the very large uncertainty regions for the GW events: As discussed in Singer et al. (2014) and Kasliwal & Nissanke (2014), the error regions for these events can cover hundreds of square degrees (especially for the initial sensitivities of the detectors). Moreover, they need not be compact or simply connected. Whilst identifying an EM counterpart would greatly enhance

the utility of the GW signal (e.g., Phinney 2009; Singer et al. 2014; Chu et al. 2016; Branchina & De Domenico 2016) and would enable a range of new physical and astrophysical tests, it is not a simple task (e.g., Kasliwal & Nissanke 2014; Cowperthwaite & Berger 2015).

The EM counterparts span a range of models at a range of wavelengths, see Metzger & Berger (2012), Singer et al. (2014), Kasliwal & Nissanke (2014), and Chu et al. (2016), amongst other recent publications. At low radio frequencies, telescopes such as the Murchison Widefield Array (MWA; Tingay et al. 2013), the Low Frequency Array (LOFAR; van Haarlem et al. 2013), and the Long Wavelength Array (LWA; Ellingson et al. 2009) have a number of advantages over optical/infrared searches: They have fields of view of hundreds to thousands of square degrees; unlike the optical/near-infrared sky which has a large number of transients present in every field (e.g., Cowperthwaite & Berger 2015), the radio sky is relatively quiet at these frequencies (Karastergiou et al. 2015;

Tingay et al. 2015; Stewart et al. 2016; Rowlinson et al. 2016; Polisensky et al. 2016) with very few unrelated transients (e.g., Hotokezaka et al. 2016) to distract from those associated with the GW event; and many of the low-frequency facilities have no moving elements and so in principle can respond within seconds to an external trigger.

Whilst most expectations for transients associated with GW sources at radio wavelengths have concentrated on late-time radio afterglows and remnants (Metzger, Williams, & Berger 2015; Hotokezaka & Piran 2015; Morsony, Workman, & Ryan 2016; Hotokezaka et al. 2016; Palliyaguru et al. 2016), which only peak after hundreds or thousands of days at 150 MHz and can be quite faint (depending on the parameters of the explosion and the circumburst medium), there are models that predict a prompt, coherent radio transient from the GW event (e.g., Lipunov & Panchenko 1996; Usov & Katz 2000; Pshirkov & Postnov 2010; Totani 2013; Zhang 2014; Wang et al. 2016; Metzger & Zivancev 2016) which may be related to the phenomenon of fast radio bursts (FRBs; Lorimer et al. 2007; Thornton et al. 2013) at least some of which may be cosmological in origin (Keane et al. 2016, although see Williams & Berger 2016; Vedantham et al. 2016). Searches for direct connections between GW events and FRBs are proceeding largely¹ through searches for GW events associated with individual FRBs (e.g., Abbott et al. 2016e) since the GW detectors have quasi-all sky sensitivity. But this strategy can be reversed: Given their wide fields of view and very fast response times (Kaplan et al. 2015), low-frequency facilities might be ideal for finding such prompt emission (Chu et al. 2016; Howell et al. 2015) triggered instead by the GW signal. We then must optimise the follow-up procedure to maximise the prospects of a discovery without time for human-aided decision making.

Strategies to aid follow-up have been studied in the optical/near-infrared regime (Rana et al. 2016; Chan et al. 2015) where the signals are likely to be faint, relatively short in duration, and may be quite red (Metzger et al. 2010; Metzger & Berger 2012; Barnes & Kasen 2013; Metzger & Fernández 2014; Kasen, Fernández, & Metzger 2015). In the optical/near-IR, the search can be aided by using prior information on host galaxies and likely distances to help reduce the search volume (e.g., Gehrels et al. 2016; Singer et al. 2016). Strategies have also been studied in the X-ray regime (e.g., Evans et al. 2016a), looking to directly probe the association between GW events and short gamma-ray bursts.

In this paper, we present and compare concrete strategies for low-frequency radio follow-up to search for prompt radio emission from a GW transient, where we use the MWA as an example to determine the likely sensitivity and success rate. Unlike in the optical/near-infrared, where a limited time window nonetheless allows limited sky coverage (e.g., Chan et al. 2015), if we are searching for a prompt

(duration \lesssim ms) radio signal, we are limited to only a single pointing, and so we must optimise where that is with limited information. We discuss this in the context of simulated GW error regions from the first couple of years of GW observations, using two and three detectors (based on the simulated events from Singer et al. 2014). The MWA occupies a middle ground in the current generation of low-frequency arrays: It has a considerably wider field of view than the more sensitive LOFAR, but it is pointed, unlike the LWA which observes the whole (visible) sky. The MWA can respond on timescales of seconds to external triggers, which is currently not possible with LOFAR (A. Rowlinson 2016, personal communication) and which is not needed with the LWA's all-sky coverage. We discuss a search that focusses on standard imaging techniques (Tingay et al. 2015; Kaplan et al. 2015; Rowlinson et al. 2016) and not more rapid 'beam-formed' data (e.g., Coenen et al. 2014; Karastergiou et al. 2015; Tremblay et al. 2015), which, although it can be more sensitive to fast signals, is much more computationally intensive to process.

2 SEARCH METRICS

To assess the success or failure of our pointing strategies, we looked at a number of different metrics and computed these for simulated GW events based on realistic expectations for the first two years of GW detector operations (Singer et al. 2014)². These simulated events included the large uncertainty region that will be communicated rapidly to EM observers, as well as the actual event locations and distances. Therefore, for a given pointing strategy, we computed the distribution of flux density and luminosity sensitivities for each simulated event. These were compared with model predictions (such as Pshirkov & Postnov 2010; Totani 2013; Zhang 2014). We also computed the separation between our pointing centre and the event location or the fraction of the total probability map covered by the observations (Abbott et al. 2016b, 2016f), but these are of limited use for a wide-field aperture array. This is because unlike optical observations with a finite but relatively uniformly covered field of view (limited slightly by vignetting), the sensitivity for an aperture array over the sky is controlled primarily by the primary beam and varies considerably over the area imaged, with some sensitivity even down to the horizon (e.g., Neben et al. 2015).

The flux density sensitivity of the MWA observations at the positions of the GW events was computed using the tile area at 150 MHz from Tingay et al. (2013), along with a receiver temperature of 50 K. To that, we added a predicted sky temperature, computed by integrating the Global Sky Model (de Oliveira-Costa et al. 2008, interpolated to 150 MHz) over the MWA tile response. We do not account for the additional contribution of the Sun. We assume a 10σ detection threshold with 30 MHz bandwidth over a 10-s observation. The search duration is determined from likely dispersion measures: For an event at a redshift $z \approx 0.05$ which is a typical horizon for

¹ Even without a direct connection, current and future population studies (Keane & Petroff 2015; Law et al. 2015; Rane et al. 2016; Li et al. 2016) may be able to argue statistically for or against a connection (e.g., Thornton et al. 2013; Zhang 2016; Callister, Kanner, & Weinstein 2016).

² Given at: <http://www.ligo.org/scientists/first2years/>.

the current detectors, Ioka (2003) and Inoue (2004) predict an extragalactic dispersion measure $DM \approx 50 \text{ pc cm}^{-3}$. This can be added to dispersion measures of $50\text{--}100 \text{ pc cm}^{-3}$ from the Milky Way and the host galaxy (e.g., Cordes & Lazio 2002; Keane et al. 2016; Kaplan et al. 2015) for a total of $150\text{--}250 \text{ pc cm}^{-3}$. With the 30 MHz bandwidth of the MWA centred at 150 MHz, this gives an event duration of 10–20 s (Lorimer & Kramer 2004; Kaplan et al. 2015; Rowlinson et al. 2016). We assume no additional loss of sensitivity due to the complex nature of the Galactic synchrotron emission, but this is likely reasonable given the short duration of the expected signals. The resulting flux density is converted to a luminosity using the simulated event distance.

3 POINTING STRATEGIES

To point the MWA, we change the delays for individual dipole antennas on each tile. The whole array (128 tiles) can be pointed together, or it can be split into subarrays, but only a single pointing is current possible for each tile. The pointing is generally done at a series of discrete steps about 7° apart. The nominal field of view is about 600 deg^2 at 150 MHz (Tingay et al. 2013). Our first goal was to determine a pointing strategy: When a HEALPIX (Górski et al. 2005) sky probability map is received, where do we point the MWA and do we point as a single array or use subarrays?

We consider several simple strategies:

- (1) Zenith pointing.
- (2) Point towards the maximum of the probability accessible (i.e., above the horizon) in the map.
- (3) Point towards the maximum of the probability weighted by $\cos^2 Z$ accessible (i.e., over the horizon) in the map, where Z is the zenith angle.
- (4) Maximise the overlap between the MWA primary beam pattern and the GW probability map.

The first of these serves as a benchmark. In addition, the sensitivity of the MWA is maximum at zenith and the primary beam at that pointing has been characterised considerably better than for other pointings. Finally, this strategy has the benefit that no decisions are necessary, so the MWA can re-point as soon as a GW alert is received. Moreover, meridian drift-scans are amongst the most common observational programme (e.g., for the MWA Transients Survey and the Galactic and Extragalactic All-Sky MWA Survey; Bell et al. 2016; Wayth et al. 2015), so if we did not interrupt an ongoing observation this would be the most likely result.

The second strategy is also simple. We simply identify the point in the GW probability map (which is sent along with the alert announcements) that has the highest value and which is also above the horizon. This is also relatively fast to compute, although it does require parsing of the GW probability map and not just knowledge of an alert.

The third strategy is very similar to the second, except that we account for the overall envelope of a Hertzian dipole

which is the basic component of an MWA tile (Tingay et al. 2013; Sutunjo et al. 2015; Neben et al. 2015). This down-weights observations close to the horizon.

Finally, the fourth strategy examines the overlap between the LVC GW sky probability and the MWA's pointing pattern. Specifically, it tries to maximise:

$$I_{\text{MWA}} = \int d\Omega P_{\text{LVC}}(\alpha, \delta) B_{\text{tile}}(\alpha, \delta), \quad (1)$$

where P_{LVC} is the normalised sky probability returned in the LVC HEALPIX map as a function of sky position (α, δ) , and $B_{\text{tile}}(\alpha, \delta)$ is the individual tile response for the MWA, normalised to 1 at the zenith. Constructed in this way, we maximise I_{MWA} by choosing the best discrete pointing B_{tile} .

The implementation of the four strategies proceeds as follows. Strategy 1 is fixed to the zenith, so no computation is necessary. For the other strategies, we first compute the MWA horizon. If the integral of P_{LVC} (weighted by $\cos^2 Z$ for strategy 3, or by B_{tile} for strategy 4) above the horizon is less than some threshold (currently 2%), we do not consider the target worthwhile, and do not return a pointing position. Otherwise, strategies 2 and 3 return the discretised pointing closest to the maximum of P_{LVC} (strategy 2) or $P_{\text{LVC}} \cos^2 Z$ (strategy 3). For strategy 4, we iterate through the range of discrete tile pointings. For each one, we compute the normalised MWA tile beam pattern sampled on the HEALPIX grid³. We then identify the pointing position that maximises the integral of $P_{\text{LVC}} \times B$. If that integral is less than a threshold (again currently 2%), we again do not consider the target worthwhile, and do not return a pointing position. Otherwise, we return the optimal target position α, δ , along with (if desired) the beamformer delays and the integrated probability.

4 EVENT EXPECTATIONS

We sought to predict our event coverage and flux density/luminosity sensitivities to actual LIGO transients using the simulated events from 2016 (including potentially both LIGO sites as well as Virgo) given at <http://www.ligo.org/scientists/first2years/> (Singer et al. 2014). There are 475 simulated signals, covering a realistic range of signal-to-noise ratio and sky position. Examples of this are shown in Figure 1. It is clear that even with a reasonable probability coverage it is possible to miss the actual transient. We see two qualitatively different results. In the first, especially where only two GW detectors see the transient, the large, elongated uncertainty region means that there is a substantial chance of missing the actual transient regardless of strategy, even with the MWA's large field of view, but on average strategies 2–4 will see some fraction of the transients discussed below. The

³ For speed, we can resample the LVC HEALPIX grid from the initially fine resolution down to a coarser resolution suitable for the MWA. For example, often the LVC maps are returned with $\text{NSIDE} = 2048$, corresponding to a pixel size of 0.029° . This is often smaller than a single MWA pixel, and with 50×10^6 pixels, the calculation can be slow. Instead we resample (conserving probability) down to $\text{NSIDE} = 64$, corresponding to a pixel size of 0.92° (and 49 152 pixels).




Figure 1. Sky probability map of simulated LVC transients from Singer et al. (2014). The colour is proportional to the \log_{10} of the probability. The black lines are the MWA horizon. The MWA half-power beams are shown by the blue lines (strategy 1: zenith), red dashed lines (strategy 2: maximum probability), green dotted lines (strategy 3: maximum probability weighted by $\cos^2 Z$), and thick magenta lines (strategy 4: maximum I_{MWA}). The white stars are the actual event locations. For the event on the left, the GW signal was only recovered by two detectors with a net signal-to-noise ratio of 14.7, leading to a large error region. In contrast, the event on the right the GW signal was recovered by three detectors with a net signal-to-noise ratio of 21.8, which greatly improves the localisations and leads to very similar pointings for strategies 2–4. The images are Mollweide projections of the celestial sphere, labelled in Right Ascension and Declination, and centred on the local sidereal time at the MWA. For the event on the right, we also show a zoom around the position of the event.

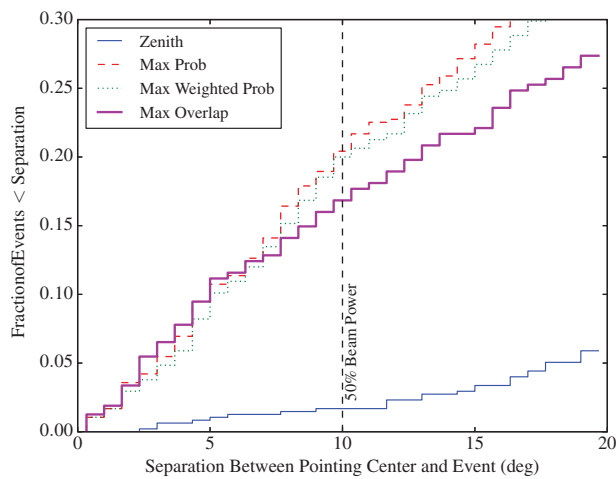


Figure 2. Cumulative histogram of θ for the simulated 2016 events, assuming observations at 150 MHz: blue lines (strategy 1: zenith), red dashed lines (strategy 2: maximum probability), green dotted lines (strategy 3: maximum probability weighted by $\cos^2 Z$), and thick magenta lines (strategy 4: maximum I_{MWA}). The vertical line is the half-power point for 150 MHz.

second type of result has a small enough uncertainty region that all pointing strategies give substantially similar results, and we are just limited by the horizon and the sensitivity of the MWA.

We show the results in Figures 2–4. In Figure 2, we plot the separation between the MWA pointing centre for each event and each strategy with the actual event location. For strategies 2 and 3, roughly 20% of the events have separations $< 10^\circ$, which is the half-power point at 150 MHz. This decreases to roughly 15% of the events for strategy 4, and $< 3\%$ for strategy 1 (the control). In many cases, the pointings will be similar for strategies 2–4 (as in the right panel of Figure 1),

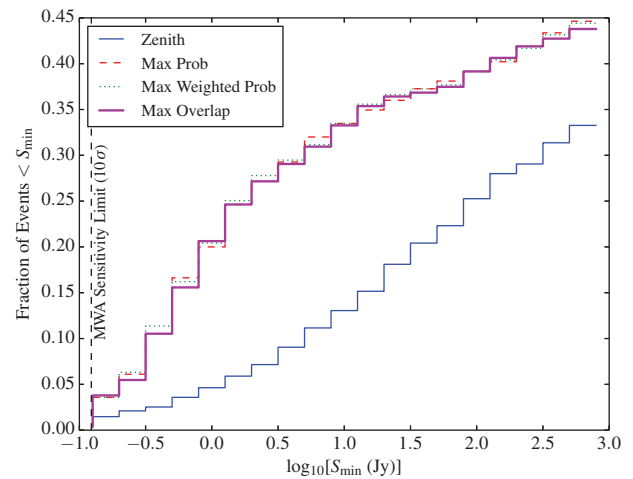


Figure 3. Cumulative histogram of limiting flux density (in Jy) for the simulated 2016 events, assuming observations at 150 MHz: blue lines (strategy 1: zenith), red dashed lines (strategy 2: maximum probability), green dotted lines (strategy 3: maximum probability weighted by $\cos^2 Z$), and thick magenta lines (strategy 4: maximum I_{MWA}). This assumes a 10σ detection over 30 MHz of bandwidth in a 10-s integration. The sky temperature has been computed by integrating the Global Sky Model (de Oliveira-Costa et al. 2008, interpolated to 150 MHz) over the MWA tile response and assumes an additional 50 K for the receiver temperature. The vertical line shows the nominal 10σ sensitivity limit from Tingay et al. (2013).

which accounts for the very similar distribution of events at separations $< 5^\circ$.

Using these results, we can also compute the expected flux density sensitivity of the MWA observations at the positions of the GW events, as described in Section 2. For the coldest parts of the sky away from the Galactic plane, we get a limiting flux density of about 0.1 Jy. However, given the influence of Galactic synchrotron emission and the limited collecting




Figure 4. Cumulative histogram of limiting luminosity νL_ν (in erg s^{-1}) for the simulated 2016 events, assuming observations at 150 MHz: blue lines (strategy 1: zenith), red dashed lines (strategy 2: maximum probability), green dotted lines (strategy 3: maximum probability weighted by $\cos^2 Z$), and thick magenta lines (strategy 4: maximum I_{MWA}). This assumes a 10σ detection over 30 MHz of bandwidth in a 10 s integration. The sky temperature has been computed by integrating the Global Sky Model (de Oliveira-Costa et al. 2008, interpolated to 150 MHz) over the MWA tile response and assumes an additional 50 K for the receiver temperature. The dashed vertical line shows the nominal 10σ sensitivity limit from Tingay et al. (2013) at a distance of 100 Mpc, whilst the dotted vertical line shows the predicted luminosity from Pshirkov & Postnov (2010).

area away from zenith, only 5% of the simulated events are close to that limit. If we consider the 15% of events that are with the half-power point, a more typical limiting flux density is 1 Jy (Figure 3). This can be compared with predictions from e.g., Pshirkov & Postnov (2010), who claim something like $S \approx 6$ Jy at a distance of 100 Mpc and a frequency of 150 MHz, so we should be able to see events like those in about 15% of the cases. Once again we see little difference between strategies 2–4.

We can also compute the limiting luminosity using the simulated GW event distances, finding limits of $10^{38-39} \text{ erg s}^{-1}$ (Figure 4). As a representative comparison, we use the signal predicted by Pshirkov & Postnov (2010): 6 Jy at 100 Mpc, at a frequency of 150 MHz (this assumes an intrinsic spin-down luminosity of $\dot{E} = 10^{50} \text{ erg s}^{-1}$, efficiency scaling exponent $\gamma = 0$, and that the burst is scattered to a duration of ~ 10 s). For roughly 30% of the events would we be able to see the such a signal. This agrees with the estimates presented in Kaplan et al. (2015), where the prompt emission predicted by various models (such as Pshirkov & Postnov 2010; Totani 2013; Zhang 2014) is compared against the sensitivity of the MWA for prompt searches. Overall, they find that the sensitivity of the MWA should be sufficient for events at a redshift of $z = 0.05$, given the available predictions. In Figure 5, we show the event-by-event comparison for the different strategies. Not unexpectedly, strategy 1 performs poorly. Comparing strategies 2 and 3 to 4, the spread is a lot smaller for the better events (those with luminosity limits $\lesssim 10^{39} \text{ erg s}^{-1}$),

since those tend to have smaller uncertainty regions that are well covered by all three strategies. For the remaining events, the results change significantly whether strategy 2/3 or 4 is used, but there is not a global preference for one or the other.

We explored the frequency dependence of these limits by repeating the exercise above for observations at 120, 150, and 180 MHz, which are the range where the MWA’s sensitivity is optimised. At lower frequencies, our primary beam will be larger and we will cover more of the GW error region. Conversely, at higher frequencies, the sky noise is lower, so the same observation will reach a lower limiting flux density. The 20th percentile for the limiting flux density (corresponding to a nominal 1 Jy in Figure 3) is about 40% higher at 120 MHz compared to 150 MHz, and 20% higher at 150 MHz compared to 180 MHz. However, we must also correct for the intrinsic spectral index β (with $S_\nu \propto \nu^\beta$), which is predicted to vary between -1 and -2 (see e.g., Kaplan et al. 2015). If this spectral index is steeper than -1.5 , then the lower frequencies will dominate. This leads us to a preference for lower frequency observations, but the unknown spectral index makes this preference weak.

5 DISCUSSION & CONCLUSION

In the analysis above, we found that about 15–20% of the events would be within the MWA’s half-power point. Therefore, we would require follow-up of $\gtrsim 6$ events before we have one with a relatively sensitive observation down to luminosity limits of $\sim 10^{39} \text{ erg s}^{-1}$. Currently, the predicted rates of neutron star-neutron star inspirals are $0.4\text{--}400 \text{ yr}^{-1}$ with a mean of 40 yr^{-1} (Abadie et al. 2010; Dominik et al. 2015; Abbott et al. 2016d), so a single year of observing should be sufficient for one or more constraining observations if the rates are not too close to the most pessimistic case.

Comparing the four strategies outlined in Section 2, we can easily reject the control strategy 1 (zenith pointing), but the remaining strategies are largely comparable in performance. Individual events may be seen better with one or the other but with the limited information available from the prompt GW triggers we cannot know which will be best in advance.

Regardless of which strategy, the metrics in Section 2 rely on the MWA being sensitivity-limited rather than event-limited. In principle, we could have different tiles with different pointing locations, so as to cover the large LVC uncertainty region. However, the limited collecting area of the MWA drives us to point all of the tiles in a single subarray so as to achieve the most sensitive possible observation, rather than attempt to cover more of the GW error region at lower sensitivity (cf. Bloemen et al. 2015). This is because, unlike in the optical regime where prompt emission from gamma-ray bursts is a known (albeit rare) phenomenon (e.g., Vestrand et al. 2014), with a known luminosity function, prompt radio emission from a gamma-ray burst or a GW event has never been seen (Bannister et al. 2012; Kaplan et al. 2015) so we do not know if shallower observations will be adequate: In the best 30% of the cases where the MWA did cover the




Figure 5. Comparison of limiting luminosity νL_ν (in erg s^{-1}) for the simulated 2016 events for each strategy, assuming observations at 150 MHz: left (strategy 1: zenith), middle (strategy 2: maximum probability), right (strategy 3: maximum probability weighted by $\cos^2 Z$), all compared to strategy 4.

GW event with a reasonable sensitivity, our luminosity limits were only ~ 1 order of magnitude below model predictions. Splitting the MWA into subarrays would mean that all observations were less constraining. We can instead make up for the possibility of missing the GW event in a statistical sense by observing a larger number of events. At the same time, the increasing performance of the GW detectors will lead to a large number of targets with improving localisation. Therefore, we believe it best to stick with a single array, but this will be re-evaluated as the actual successes are evaluated. Similarly, we could experiment with other observational modes like splitting our 30-MHz bandpass into multiple sub-bands (as in Kaplan et al. 2015), which could be advantageous if a bright but frequency-dependent signal is expected. Given the degree of uncertainty about these models that is unlikely to be preferred at least to start, but as we gain experience, we may change our procedure.

We implemented strategy 4 for the MWA during the first LIGO observing run (O1; 2015 September to 2016 January) covering the first detection, GW 150914. However, this trigger was released after a considerable delay (several days) needed for human examination of the event. Therefore, we did not require any real-time decisions about strategy but instead could use multiple pointings to tile the GW error regions (Abbott et al. 2016b). We expect that as the LVC improves their internal vetting and pipelines their latency will improve to 90–120 s after the GW event (Singer et al. 2014; Cadonati, Astone, & van den Broeck 2014) or possibly better (Cannon et al. 2012; Chu et al. 2016) and this strategy will be employed.

It is worth noting that the first published GW signal is from a binary black hole system (Abbott et al. 2016c), which is not expected to have any EM signature (Abbott et al. 2016b, but see Connaughton et al. 2016). The rates of similar events will likely be quite high once LIGO reaches its full design sensitivity, approaching 1 d^{-1} . If this is the case, then we will

certainly have the opportunity to cover a sufficient number of error regions to search statistically for any associated EM emission, although the greater distances to the more massive systems will limit our sensitivity.

As discussed in Kaplan et al. (2015) and Chu et al. (2016), the expected delay of the radio signal relative to the GW transient is tens of seconds up to several minutes, based on the simulated distances of the transients and the expected extragalactic plus galactic dispersion measures. The actual time of any prompt radio signal may also be shifted by up to tens of seconds (e.g., Zhang 2014), potentially in either direction. Given the fast, ≈ 16 s response time that the MWA can achieve (Kaplan et al. 2015), we can easily repoint to catch any prompt emission as long as the GW latency improves. Overall, we emphasise the need to transmit the trigger and react, as soon as possible, preferably well within 1 min.

We have demonstrated that the MWA can respond quickly to GW transients and cover a reasonable fraction of events with good sensitivity. The strategies outlined here are specifically applicable to the MWA, in that we have made use of the MWA’s location and primary beam pattern in assessing the follow-up prospects. They can be adapted for other facilities, but there other considerations may lead to different strategies. For instance, with a considerably smaller field of view but better instantaneous sensitivity splitting into subarrays may be more viable. This will also evolve as new data and new models for prompt emission become available. Overall, we believe that the MWA has a good combination of field of view, sensitivity, and operational flexibility that enables this science: The MWA has a much larger field of view compared to most pointed radio telescopes (e.g., Chu et al. 2016; Abbott et al. 2016b), but is more sensitive than some all-sky facilities (e.g., Ellingson et al. 2009). With roughly 1 yr of sensitive GW observations, we should be able to answer unambiguously which if any of the models for prompt emission are real.

ACKNOWLEDGEMENTS

We thank an anonymous referee and editor for helpful comments. DLK and SDC were supported by NSF grant AST-1412421. Parts of this research were conducted by the Australian Research Council Centre of Excellence for All-sky Astrophysics (CAASTRO), through project number CE110001020. CMT is supported by an ARC Discovery Early Career Researcher Project Grant, DE140100316.

REFERENCES

- Abadie, J., et al. 2010, *CQGra*, 27, 173001
- Abbott, B. P., et al. 2016a, *PhRvL*, 116, 241103
- Abbott, B. P., et al. 2016b, *ApJ*, 826, L13
- Abbott, B. P., et al. 2016c, *PhRvL*, 116, 061102
- Abbott, B. P., et al. 2016d, *LRR*, 19, 1a
- Abbott, B. P., et al. 2016e, *PhRvD*, 93, 122008
- Abbott, B. P., et al. 2016f, *ApJS*, 225, 8
- Ackermann, M., et al. 2016, *ApJ*, 823, L2
- Bannister, K. W., Murphy, T., Gaensler, B. M., & Reynolds, J. E. 2012, *ApJ*, 757, 38
- Barnes, J., & Kasen, D. 2013, *ApJ*, 775, 18
- Bell, M. E., et al. 2016, *MNRAS*, 461, 908
- Bloemen, S., Groot, P., Nelemans, G., & Klein-Wolt, M. 2015, in *ASP Conf. Ser.*, Vol. 496, *Living Together: Planets, Host Stars and Binaries*, eds. S. M. Rucinski, G. Torres, & M. Zejda (San Francisco: ASP), 254
- Branchina, V., & De Domenico, M. 2016, arXiv:1604.08530
- Cadonati, L., Astone, P., & van den Broeck, C. 2014, *The LSC-Virgo White Paper on Gravitational Wave Searches and Astrophysics*, Tech. Rep. T1400054-v7, LIGO, <https://dcc.ligo.org/LIGO-T1400054/public>
- Callister, T., Kanner, J., & Weinstein, A. 2016, *ApJ*, 825, L12
- Cannon, K., et al. 2012, *ApJ*, 748, 136
- Chan, M. L., Hu, Y.-M., Messenger, C., Hendry, M., & Heng, I. S. 2015, e-prints, arXiv:1506.04035
- Chu, Q., et al. 2016, *MNRAS*, 459, 121
- Coenen, T., et al. 2014, *A&A*, 570, A60
- Connaughton, V., et al. 2016, *ApJ*, 826L, 6C
- Cordes, J. M., & Lazio, T. J. W. 2002, *ArXiv Astrophysics e-prints*, astro-ph/0207156
- Cowperthwaite, P. S., & Berger, E. 2015, *ApJ*, 814, 25
- de Oliveira-Costa, A., et al. 2008, *MNRAS*, 388, 247
- Dominik, M., et al. 2015, *ApJ*, 806, 263
- Ellingson, S. W., et al. 2009, *IEEEP*, 97, 1421
- Evans, P. A., et al. 2016b, *MNRAS*, 460, L40
- Evans, P. A., et al. 2016a, *MNRAS*, 455, 1522
- Gehrels, N., et al. 2016, *ApJ*, 820, 136
- Górski, K. M., et al. 2005, *ApJ*, 622, 759
- Hotokozaka, K., et al. 2016, *ApJ*, accepted, arXiv:1605.09395
- Hotokozaka, K., & Piran, T. 2015, *MNRAS*, 450, 1430
- Howell, E. J., et al. 2015, *PASA*, 32, e046
- Inoue, S. 2004, *MNRAS*, 348, 999
- Ioka, K. 2003, *ApJ*, 598, L79
- Kaplan, D. L., et al. 2015, *ApJ*, 814, L25
- Karastergiou, A., et al. 2015, *MNRAS*, 452, 1254
- Kasen, D., Fernández, R., & Metzger, B. D. 2015, *MNRAS*, 450, 1777
- Kasliwal, M. M., & Nissanke, S. 2014, *ApJ*, 789, L5
- Keane, E. F., et al. 2016, *Nature*, 530, 453
- Keane, E. F., & Petroff, E. 2015, *MNRAS*, 447, 2852
- Law, C. J., et al. 2015, *ApJ*, 807, 16
- Li, L., Huang, Y., Zhang, Z., Li, D., & Li, B. 2016, arXiv:1602.06099
- Lipunov, V. M., et al. 2016, arXiv:1605.01607
- Lipunov, V. M., & Panchenko, I. E. 1996, *A&A*, 312, 937
- Lorimer, D. R., Bailes, M., McLaughlin, M. A., Narkevic, D. J., & Crawford, F. 2007, *Science*, 318, 777
- Lorimer, D. R., & Kramer, M. 2004, *Handbook of Pulsar Astronomy* (Cambridge: Cambridge University Press)
- Metzger, B. D., & Berger, E. 2012, *ApJ*, 746, 48
- Metzger, B. D., & Fernández, R. 2014, *MNRAS*, 441, 3444
- Metzger, B. D., et al. 2010, *MNRAS*, 406, 2650
- Metzger, B. D., Williams, P. K. G., & Berger, E. 2015, *ApJ*, 806, 224
- Metzger, B. D., & Zivancev, C. 2016, *MNRAS*, 461, 4435
- Morokuma, T., et al. 2016, *PASJ*, 68, L9
- Morsony, B. J., Workman, J. C., & Ryan, D. M. 2016, *ApJ*, 825, L24
- Neben, A. R., et al. 2015, *RaSc*, 50, 614
- Palliyaguru, N. T., et al. 2016, *ApJ*, submitted, arXiv:1608.06518
- Phinney, E. S. 2009, in *Astronomy*, Vol. 2010, *astro2010: The Astronomy and Astrophysics Decadal Survey*, arXiv:0903.0098
- Polisensky, E., et al. 2016, *ApJ*, submitted, arXiv:1604.00667
- Pshirkov, M. S., & Postnov, K. A. 2010, *Ap&SS*, 330, 13
- Rana, J., Singhal, A., Gadre, B., Bhalerao, V., & Bose, S. 2016, *ApJ*, submitted, arXiv:1603.01689
- Rane, A., et al. 2016, *MNRAS*, 455, 2207
- Rowlinson, A., et al. 2016, *MNRAS*, 458, 3506
- Savchenko, V., et al. 2016, *ApJ*, 820, L36
- Singer, L. P., et al. 2016, *ApJ*, submitted, arXiv:1603.07333
- Singer, L. P., et al. 2014, *ApJ*, 795, 105
- Stewart, A. J., et al. 2016, *MNRAS*, 456, 2321
- Sutinjo, A., et al. 2015, *RaSc*, 50, 52
- Thornton, D., et al. 2013, *Science*, 341, 53
- Tingay, S. J., et al. 2013, *PASA*, 30, 7
- Tingay, S. J., et al. 2015, *AJ*, 150, 199
- Totani, T. 2013, *PASJ*, 65, L12
- Tremblay, S. E., et al. 2015, *PASA*, 32, e005
- Troja, E., Read, A. M., Tiengo, A., & Salvaterra, R. 2016, *ApJ*, 822, L8
- Usov, V. V., & Katz, J. I. 2000, *A&A*, 364, 655
- van Haarlem, M. P., et al. 2013, *A&A*, 556, A2
- Vedantham, H. K., et al. 2016, *ApJ*, 824, L9
- Vestrand, W. T., et al. 2014, *Science*, 343, 38
- Wang, J.-S., Yang, Y.-P., Wu, X.-F., Dai, Z.-G., & Wang, F.-Y. 2016, *ApJ*, 822, L7
- Wayth, R. B., et al. 2015, *PASA*, 32, e025
- Williams, P. K. G., & Berger, E. 2016, *ApJ*, 821, L22
- Zhang, B. 2014, *ApJ*, 780, L21
- Zhang, B. 2016, *ApJ*, 827, L31