

Research Paper

An optimised gravitational wave follow-up strategy with the Australian Square Kilometre Array Pathfinder

D. Dobie^{1,2}, T. Murphy¹, D. L. Kaplan³, S. Ghosh³, K. W. Bannister² and R. W. Hunstead¹

¹Sydney Institute for Astronomy, School of Physics, University of Sydney, NSW 2006, Australia, ²ATNF, CSIRO Astronomy and Space Science, PO Box 76, Epping, NSW 1710, Australia and ³Department of Physics, University of Wisconsin-Milwaukee, Milwaukee, WI 53201, USA

Abstract

The detection of a neutron star merger by the Advanced Laser Interferometer Gravitational-Wave Observatory and Advanced Virgo gravitational wave detectors, and the subsequent detection of an electromagnetic counterpart have opened a new era of transient astronomy. With upgrades to the Advanced Laser Interferometer Gravitational-Wave Observatory 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 is expected to become fully operational in early 2019, and its 30 deg² 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 Australian Square Kilometre Array Pathfinder 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.

Keywords: gravitational waves – instrumentation: interferometers – radio continuum: general

(Received 13 September 2018; revised 18 January 2019; accepted 27 February 2019)

1. Introduction

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

Observations over the first 100 d suggested that the observed X-ray and radio emission originate from the same source: optically thin synchrotron emission (Mooley et al. 2018*a*). 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. 2018) and temporal index $\delta_1 = 0.84 \pm 0.05$, peaked 149 \pm 2 d post-merger (Dobie et al. 2018) and has declined according to a similar power-law with temporal index $\delta_2 = -2.2 \pm 0.2$ or

steeper (Mooley et al. 2018*c*; van Eerten et al. 2018) out to 300 d 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; Hotokez et al. 2016). Instead, this radio emission has been proposed to be produced by mildly relativistic quasi-spherical outflow (a ‘cocoon’; Gottlieb, Nakar, & Piran 2018; 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. 2018; Resmi et al. 2018; Mooley et al. 2018*b*). 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.

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. 2018*b*) and improve standard-siren constraints on H_0 (Abbott et al. 2017; Hotokezaka et al. 2018)

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

Author for correspondence: D. Dobie, Email: ddob1600@uni.sydney.edu.au

Cite this article: Dobie D, Murphy T, Kaplan DL, Ghosh S, Bannister KW and Hunstead RW. (2019) An optimised gravitational wave follow-up strategy with the Australian Square Kilometre Array Pathfinder. *Publications of the Astronomical Society of Australia* 36, e019, 1–10. <https://doi.org/10.1017/pasa.2019.9>

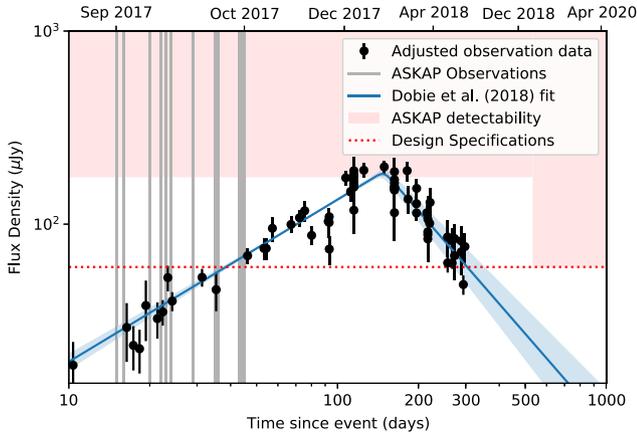


Figure 1. Detectability of GW170817 with ASKAP. Radio observations (from Hallinan et al. 2017; Mooley et al. 2018a; Dobie et al. 2018; Mooley et al. 2018c; Alexander et al. 2018; Margutti et al. 2018; van Eerten et al. 2018) scaled to 1.4 GHz based on a spectral index of $\alpha = -0.58$ are noted in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

(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.

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 arcsec–3 arcmin at frequencies from 700 to 1 800 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 our use of ASKAP observing time.

2. Searching for radio emission

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, 2017c). Imaging observations began 2 d after the event, covering the 90% confidence region using the strategy presented in this paper (Dobie et al. 2017a, 2017b); see Andreoni et al. (2017) for further details.

Figure 1 shows the lightcurve of GW170817 from Dobie et al. (2018) adjusted to 1.4 GHz based on a spectral index of $\alpha = -0.58$,

Table 1. ASKAP design specifications and at the time of GW170817. Image RMS (σ) is calculated using the radiometer equation. Survey speed (SS) assumes a $100 \mu\text{Jy}$ 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
$\sigma_{10 \text{ min}}$ ($\mu\text{Jy}/\text{beam}$)	300	82
$\sigma_{12 \text{ h}}$ ($\mu\text{Jy}/\text{beam}$)	35	12
SS ($\text{deg}^2 \text{ h}^{-1}$)	21	270
Frequency (GHz)	0.7–1.4	0.7–1.8

along with the 5σ detectability limit for the current telescope configuration and the design specifications (see Table 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 d post-merger and peaked at a flux density equivalent to a 15σ detection.

2.2. Long-term synchrotron emission

The focus of this paper is determining an observing strategy to detect the long-term radio emission (tens to 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 to weeks or the dynamical ejecta that caused the kilonova on timescales of months to 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. 2018). 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. 2018).

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 occurred 3 months later in the daytime.

2.3. Prompt radio emission

It has been suggested that neutron star mergers may produce prompt, coherent radio emission. This could 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 LIGO 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 deg s^{-1} in azimuth and 1 deg 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 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 LIGO has maximum sensitivity. We note that due to the expected timescales for synchrotron radio emission discussed in Section 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. 2018; Shannon et al. 2018), making up over a quarter of all FRBs published to-date^a. It is capable of arcminute-localisation of coherent radio emission in fly’s-eye mode, which is insufficient to positively 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 plausibly 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. 2016).

3. An optimised pointing strategy

3.1. Observing LIGO localisation regions

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, 2017h). 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 $120\text{--}180 \text{ deg}^2$, 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 $<12 \text{ deg}^2$ (Abbott et al. 2018).

It is not always feasible to observe entire localisation regions due to the large amount of observing time required, and 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 with 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.

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 (e.g. Coughlin & Stubbs 2016; Arcavi et al. 2017) use estimates of afterglow luminosity to determine how time should be allocated between candidate hosts. However, current models of radio emission from neutron star mergers (see Section 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 (\sim arcminute) FoV is targeted observations of specific galaxies in the localisation volume (e.g. Evans et al. 2016; Arcavi et al. 2017). We discuss applying a galaxy-targeted technique for ASKAP in Section 3.4.

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, predetermined tile placed on the sky. LIGO 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^b) projection, which divides the sky into pixels of equal surface area.

^a<http://www.frbcatalog.org/> (Petroff et al. 2016).

^b<http://healpix.sourceforge.net/>.

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 an 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.

3.2.1. 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 suboptimal because it will cover significant extraneous area if the shorter axis of the localisation ellipse is the width of a few tiles or smaller.

3.2.2. 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).

3.2.3. 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 2. We use the same strategy defined by Ghosh *et al.* (2017), adapted for the ASKAP FoV.

3.2.4. 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:

Figure 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 predefined, 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.

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 an additional one-tile buffer, which we call the test-set of tiles. 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, and 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.
5. Apply the optimal shift in each declination strip to the test-set and recalculate 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.

Figure 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 predefined 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).

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.

3.3.1. Field of view

ASKAP has a FoV of 30 deg^2 —significantly larger than all other gigahertz-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.

3.3.2. 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 predefined 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 surveys with ASKAP use the Variables and Slow Transients pipeline (VAST; Murphy et al. 2013), which detects transients via a process of source-finding, cross-matching, and lightcurve analysis, meaning there is currently no advantage in using a predefined grid of pointings. Improved transient detection methods such as using image subtraction for LST-aligned observations are being tested with ASKAP (Stewart et al., in preparation).

3.3.3. 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 optical/ultraviolet/infrared 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 to 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 (The LIGO Scientific Collaboration & the Virgo Collaboration 2017a), which was refined to 28 deg^2 two months post-merger (The LIGO Scientific Collaboration & the Virgo Collaboration 2017b) and finally 16 deg^2 nine months post-merger (The LIGO Scientific Collaboration et al. 2018).

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 (Figure 3 of Abbott et al. 2017c). This advantage is not applicable to searches for prompt radio emission.

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; however, 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 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 ($|b| < 10$ 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 LIGO skymap with a galaxy catalogue using the method described by Evans *et al.* (2016). For each pixel, $P_{\text{GW}}(D)$ 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 gravitational wave skymap.

We determine the completeness of the galaxy catalogue as a function of distance, $C(D)$, as in White, Daw, & Dhillon *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 Sloan Digital Sky Survey (SDSS) by Kopparapu *et al.* (2008), $(1.98 \pm 0.16) \times 10^{-2} L_{10} \text{ Mpc}^{-3}$ (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 $P_g(D)$, 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

$$P_p = P_{\text{GW}}(D)P_g(D)C(D) + P_{\text{GW}}(D)(1 - C(D)), \quad (1)$$

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 $\sum_{\text{allsky}} P_p = 1$, 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 (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 $P_g(D)$, 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.

4. Comparison of pointing strategies

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 h 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, for both the original LIGO skymaps and the convolved galaxy-targeted skymaps. The galaxy-targeting strategy was then tested using the greedy optimisation strategy.

4.2. Comparing advanced ranked tiling strategies

Figure 4 shows the average probability coverage per tile as a function of target probability for all four tiling strategies using the LIGO 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 $\sim 10\%$ more efficient coverage per tile compared to the greedy ranked tile strategy.

We also consider our ability to detect events. Figure 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

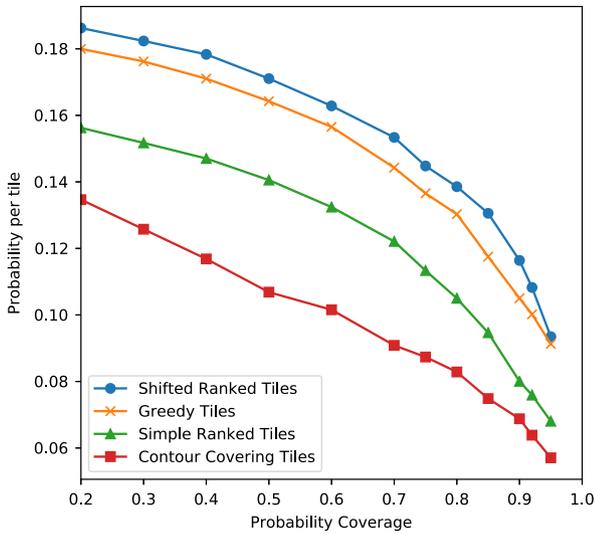


Figure 4. Covered probability per tile, for all four tiling strategies averaged across their application to 475 simulated skymaps. As expected, the two optimised strategies significantly outperform the contour covering and simple ranked tiled strategies.

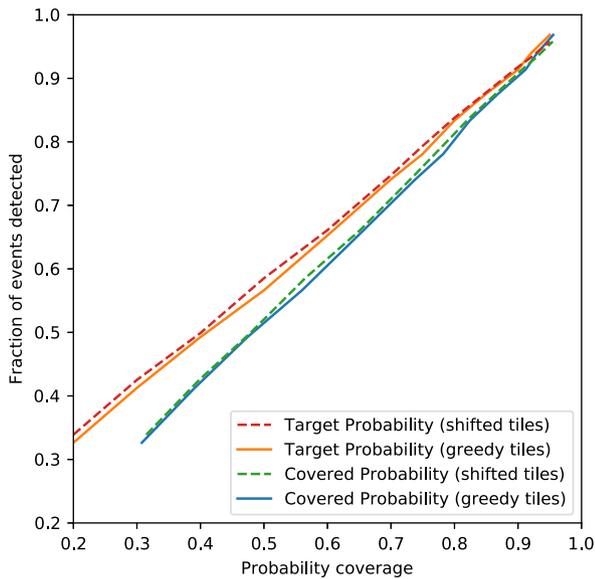


Figure 5. Fraction of events detected as a function of target probability coverage (orange/red) and actual probability coverage (blue/green) for both the greedy tiles (solid) and shifted tiles (dashed) methods. The actual probability coverage is calculated by taking the mean coverage for all simulations.

averaged across all events. Both tiling strategies produce similar results across the entire probability range. As expected, the detection fraction is approximately equal to the covered probability, confirming that Advanced LIGO-Virgo skymaps do accurately recreate the PDF 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 suboptimal 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

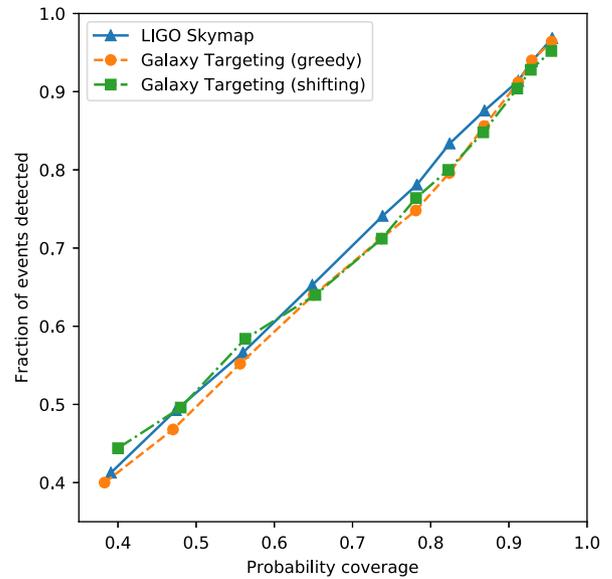


Figure 6. Comparison of sky coverage using the galaxy-targeted probability map and the original gravitational wave skymap.

(e.g. compare the tiling in Figures 2 and 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 overall 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 deg², comparable with 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, respectively, corresponding to savings of 5 h 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.

4.3. Galaxy-targeting strategy

We applied the greedy and shifted ranked tiles strategies to all 250 simulated skymaps that have 3D localisations. Figure 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 sky maps. We expect that the effectiveness of using a galaxy-targeting strategy for most telescopes will significantly improve in the near future, as wider and deeper galaxy surveys produce more complete catalogues. However, it is unclear whether this improvement will affect ASKAP.

Conversely, we note that improvements to gravitational wave detectors will result in larger detector horizons. 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.

5. Detectability of events with ASKAP

5.1. Sensitivity and angular resolution

Table 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 $12 \mu\text{Jy}$ image root mean square (RMS) in a 12-h pointing.

While GW170817 would have been detectable with ASKAP 40–300 d post-merger (Figure 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, the circum-merger density, n , the velocity of the outflow, 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.* 2018; Lamb, Mandel, & Resmi *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 $n \sim 10^{-4} \text{ cm}^{-3}$, significantly lower than the median circum-burst density of sGRBs ($3 - 15 \times 10^{-3} \text{ cm}^{-3}$), with densities $n \sim 1 \text{ cm}^{-3}$ not uncommon (Fong *et al.* 2015). However, we expect approximately half of all neutron star mergers to have circum-merger densities $n > 0.1 \text{ cm}^{-3}$ (Hotokez *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 (Dobie *et al.*, in preparation), 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 (JVLA), and we will not always be able to resolve the GW counterpart from its host galaxy. sGRBs have physical offsets between 0.5 and 75 kpc from the host galaxy centre (median 5 kpc; Berger 2014), which Hotokez *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 case for a majority of events (Hotokez *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 JVLA, and potentially 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.

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, Kasliwal, & Georgieva (2013) 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 d to 20 yr, corresponding to an areal transient rate of $7.5 \times 10^{-4} \text{ deg}^{-2}$ for flux densities $> 14 \text{ mJy}$. Other works found the 1.4 GHz radio transient areal density to be $< 0.1 \text{ deg}^{-2}$ ($> 1 \text{ mJy}$ on timescales of 0.5 to 7 yr; Hancock *et al.* 2016), $< 0.37 \text{ deg}^{-2}$ ($> 0.2 \text{ mJy}$ on timescales of 1 week to 3 months; Mooley *et al.* 2013), and $< 0.01 \text{ deg}^{-2}$ ($> 1.5 \text{ mJy}$ 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 $+30 \text{ deg}$ to an image RMS of $\sim 10 \mu\text{Jy}$ 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 1998) which covers the sky north of -40 deg with an image RMS of $\sim 0.45 \text{ mJy}$ at 1.4 GHz, the VLA Sky Survey (VLASS^c) which will cover the sky north of -40 deg to an image RMS of $67 \mu\text{Jy}$ at 2–4 GHz, or the Sydney University Molonglo Sky Survey (SUMSS, Bock, Large, & Sadler *et al.* 1999; Mauch *et al.* 2003) which covers the sky south of -30 deg to an image RMS of $\sim 1\text{--}2.5 \text{ mJy}$ 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.

5.3. Expectations for LIGO-Virgo O3

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, $1.54^{+3.2}_{-1.22} \text{ Mpc}^{-3} \text{ Myr}^{-1}$ (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 $+30 \text{ deg}$, corresponding to 75% of the whole sky. We expect to be able to observe a

^c<https://science.nrao.edu/science/surveys/vlass>.

similar fraction of LIGO 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 LIGO for O3 and beyond is discussed in Section 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.

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 SS. Importantly, 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 3.2.4) outperforms previously investigated tiling methods for long-term ASKAP follow-up, while the greedy ranked tiles method (see Section 3.2.3 and 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.

Author ORCIDs.  D. Dobie <https://orcid.org/0000-0003-0699-7019>.

Acknowledgements. D.D. acknowledges useful discussions with H. Qiu. We thank the referee for helpful comments and suggestions which improved this paper. D.D. is supported by an Australian Government Research Training Program Scholarship. TM acknowledges the support of the Australian Research Council through grant FT150100099. DLK was supported by NSF grant AST-1412421. SG was supported by NSF Award PHY-1607585. This research has made use of NASA's Astrophysics Data System Bibliographic Services. The Australian SKA Pathfinder is part of the Australia Telescope

National Facility which is funded by the Commonwealth of Australia for operation as a National Facility managed by CSIRO. The Murchison Radio-astronomy Observatory is managed by the CSIRO, who also provide operational support to ASKAP. The authors acknowledge the Wajarri Yamatji people as the traditional owners of the Observatory site.

Software: ASTROPY (Astropy Collaboration et al. 2018), MATPLOTLIB (Hunter 2007), and NUMPY (van der Walt et al. 2011).

References

- Abbott, B. P., et al. 2016, *PhRvD*, **93**, 122008
 Abbott, B. P., et al. 2017a, *Classical Quantum Gravity*, **34**, 044001
 Abbott, B. P., et al. 2017b, *PhRvL*, **118**, 221101
 Abbott, B. P., et al. 2017c, *PhRvL*, **119**, 141101
 Abbott, B. P., et al. 2017d, *PhRvL*, **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., Abbott, R., Abbott, T. D., Acernese, F., & Ackley, K. 2017h, *ApJ*, **851**, L35
 Abbott, B. P., et al. 2018, *Living Rev. Relativ.*, **21**, 3
 Alexander, K. D., et al. 2018, *ApJ*, **863**, L18
 Andreoni, I., et al. 2017, *PASA*, **34**, e069
 Arcavi, I., et al. 2017, *ApJ*, **848**, L33
 Aso, Y., Michimura, Y., Somiya, K., Ando, M., Miyakawa, O., Sekiguchi, T., Tatsumi, D., & Yamamoto, H. 2013, *PhRvD*, **88**, 043007
 Astropy Collaboration, et al. 2018, *AJ*, **156**, 123
 Bannister, K. W., Murphy, T., Gaensler, B. M., Hunstead, R. W., & Chatterjee, S. 2011, *MNRAS*, **412**, 634
 Bannister, K. W., et al. 2017a, *ApJ*, **841**, L12
 Bannister, K., Shannon, R., Hotan, A., James, C., Macquart, J.-P., Osłowski, S., & Farah, W. 2017b, *GCN*, 21562
 Bannister, K., Shannon, R., Hotan, A., James, C., Macquart, J.-P., Osłowski, S., & Farah, W. 2017c, *GCN*, 21671
 Bartos, I., et al. 2014, *MNRAS*, **443**, 738
 Berger, E. 2014, *ARA&A*, **52**, 43
 Bhandari, S., et al. 2018, *MNRAS*, **478**, 1784
 Bock, D. C.-J., Large, M. I., & Sadler, E. M. 1999, *AJ*, **117**, 1578
 Chan, M. L., Messenger, C., Heng, I. S., & Hendry, M. 2018, *PhRvD*, **97**, 123014
 Chu, Q., Howell, E. J., Rowlinson, A., Gao, H., Zhang, B., Tingay, S. J., Boër, M., & Wen, L. 2016, *MNRAS*, **459**, 121
 Condon, J. J., Cotton, W. D., Greisen, E. W., Yin, Q. F., Perley, R. A., Taylor, G. B., & Broderick, J. J. 1998, *AJ*, **115**, 1693
 Cook, D. O., Van Sistine, A., Singer, L., & Kasliwal, M. M. 2017, *GCN*, 21519
 Corsi, A., et al. 2018, *ApJ*, **861**, L10
 Coughlin, M., & Stubbs, C. 2016, *ExA*, **42**, 165
 Coulter, D. A., et al. 2017, *Sci*, **358**, 1556
 D'Avanzo, P., et al. 2018, *A&A*, **613**, L1
 Dálya, G., et al. 2018, *MNRAS*, **479**, 2374
 Dobie, D., Hotan, A., Bannister, K., Kaplan, D., Murphy, T., & Lynch, C. 2017a, *GCN*, 21625
 Dobie, D., Hotan, A., Bannister, K., Kaplan, D., Murphy, T., & Lynch, C. 2017b, *GCN*, 21639
 Dobie, D., et al. 2018, *ApJ*, **858**, L15
 Dolag, K., Gaensler, B. M., Beck, A. M., & Beck, M. C. 2015, *MNRAS*, **451**, 4277
 Evans, P. A., et al. 2016, *MNRAS*, **462**, 1591
 Evans, P. A., et al. 2017, *Sci*, **358**, 1565
 Falcke, H., & Rezzolla L. 2014, *A&A*, **562**, A137
 Fong, W., Berger, E., Margutti, R., & Zauderer, B. A. 2015, *ApJ*, **815**, 102
 Ghosh, S., Bloemen, S., Nelemans, G., Groot, P. J., & Price, L. R. 2016, *A&A*, **592**, A82
 Ghosh, S., Chatterjee, D., Kaplan, D. L., Brady, P. R., & Van Sistine, A. 2017, *PASP*, **129**, 114503
 Goldstein, A., et al. 2017, *ApJ*, **848**, L14
 Gottlieb, O., Nakar, E., & Piran, T. 2018, *MNRAS*, **473**, 576
 Granot, J., et al. 2002, *ApJ*, **570**, L61

- Haggard, D., Nynka, M., Ruan, J. J., Kalogera, V., Cenko, S. B., Evans, P. & Kennea, J. A. 2017, *ApJ*, **848**, L25
- Hallinan, G., et al. 2017, *Sci*, **358**, 1579
- Hampson, G., et al. 2012, in 2012 International Conference on Electromagnetics in Advanced Applications (IEEE), 807. doi:10.1109/ICEAA.2012.6328742
- Hancock, P. J., Drury, J. A., Bell, M. E., Murphy, T., & Gaensler, B. M. 2016, *MNRAS*, **461**, 3314
- Hotokezaka, K., & Piran, T. 2015, *MNRAS*, **450**, 1430
- Hotokezaka, K., Nissanke, S., Hallinan, G., Lazio, T. J. W., Nakar, E., & Piran, T. 2016, *ApJ*, **831**, 190
- Hotokezaka, K., Nakar, E., Gottlieb, O., Nissanke, S., Masuda, K., Hallinan, G., Mooley, K. P., & Deller, A. T. 2018, preprint, [arXiv:1806.10596](https://arxiv.org/abs/1806.10596)
- Hunter, J. D. 2007, *Comput. Sci. Eng.*, **9**, 90
- Inoue, S. 2004, *MNRAS*, **348**, 999
- Ioka, K. 2003, *ApJ*, **598**, L79
- Johnston, S., et al. 2008, *ExA*, **22**, 151
- Klimenko, S., et al. 2011, *PhRvD*, **83**, 102001
- Kopparapu, R. K., Hanna, C., Kalogera, V., O'Shaughnessy, R., González, G., Brady, P. R., & Fairhurst, S. 2008, *ApJ*, **675**, 1459
- Lamb, G. P., Mandel, I., & Resmi, L. 2018, *MNRAS*, **481**, 2581
- Lazzati, D., Deich, A., Morsony, B. J., & Workman, J. C. 2017, *MNRAS*, **471**, 1652
- Lazzati, D., Perna, R., Morsony, B. J., Lopez-Camara, D., Cantiello, M., Ciolfi, R., Giacomazzo, B., & Workman, J. C. 2018, *PhRvL*, **120**, 241103
- 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, *Sci*, **318**, 777
- Macquart, J. P., Shannon, R. M., Bannister, K. W., James, C. W., Ekers, R. D., & Bunton, J. D. 2018, preprint, [arXiv:1810.04353](https://arxiv.org/abs/1810.04353)
- Mahony, E. K., et al. 2018, *ApJ*, **867**, L10
- Margutti, R., et al. 2018, *ApJ*, **856**, L18
- Mauch, T., Murphy, T., Buttery, H. J., Curran, J., Hunstead, R. W., Piestrzynski, B., Robertson, J. G., & Sadler, E. M. 2003, *MNRAS*, **342**, 1117
- Metzger, B. D., & Berger, E. 2012, *ApJ*, **746**, 48
- Mooley, K. P., Frail, D. A., Ofek, E. O., Miller, N. A., Kulkarni, S. R., & Horesh, A. 2013, *ApJ*, **768**, 165
- 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*, **868**, L11
- Murphy, T., et al. 2013, *PASA*, **30**, e006
- Nakar, E., & Piran T. 2011, *Nature*, **478**, 82
- Nakar, E., & Piran T. 2018, *MNRAS*, **478**, 407
- Nissanke, S., Kasliwal, M., & Georgieva, A. 2013, *ApJ*, **767**, 124
- Norris, R. P., et al. 2011, *PASA*, **28**, 215
- Petroff, E., et al. 2016, *PASA*, **33**, e045
- Pshirkov, M. S., & Postnov, K. A. 2010, *Ap&SS*, **330**, 13
- Resmi, L., et al. 2018, preprint, [arXiv:1803.02768](https://arxiv.org/abs/1803.02768)
- Salafia, O. S., Colpi, M., Branchesi, M., Chassande-Mottin, E., Ghirlanda, G., Ghisellini, G., & Vergani, S. D. 2017, *ApJ*, **846**, 62
- Sari, R., Piran, T., & Narayan, R. 1998, *ApJ*, **497**, L17
- Schutz, B. F. 2011, *Classical Quantum Gravity*, **28**, 125023
- Shannon, R., et al. 2018, *Nature*, **562**, 386
- Singer, L. P., et al. 2014, *ApJ*, **795**, 105
- Somiya, K. 2012, *Classical Quantum Gravity*, **29**, 124007
- The LIGO Scientific Collaboration, & the Virgo Collaboration 2017a, GCN, 21513
- The LIGO Scientific Collaboration, & the Virgo Collaboration 2017b, GCN, 21983
- The LIGO Scientific Collaboration, the Virgo Collaboration, Abbott, B. P., Abbott, R., Abbott, T. D., Acernese, F., & Ackley, K., et al. 2018, preprint, [arXiv:1805.11579](https://arxiv.org/abs/1805.11579)
- Thornton, D., et al. 2013, *Sci*, **341**, 53
- Totani, T. 2013, *PASJ*, **65**, L12
- Troja, E., et al. 2017, *Nature*, **551**, 71
- Unnikrishnan, C. S. 2013, *Int. J. Mod. Phys. D*, **22**, 1341010
- Wang, J.-S., Yang, Y.-P., Wu, X.-F., Dai, Z.-G., & Wang, F.-Y. 2016, *ApJ*, **822**, L7
- White, D. J., Daw, E. J., & Dhillon, V. S., 2011, *Classical Quantum Gravity*, **28**, 085016
- Zhang, B. 2014, *ApJ*, **780**, L21
- van Eerten, E. T. H., Ryan, G., Ricci, R., Burgess, J. M., Wieringa, M., Piro, L., Cenko, S. B., & Sakamoto, T. 2018, preprint, [arXiv:1808.06617](https://arxiv.org/abs/1808.06617)
- van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, *Comput. Sci. Eng.*, **13**, 22