LOFAR observations of gravitational wave merger events: O3 results and O4 strategy

K. Gourdji [1,2]★ A. Rowlinson [3,4] R. A. M. J. Wijers [4] J. W. Broderick [5] and A. Shulevski [3]

[1]*Centre for Astrophysics and Supercomputing, Swinburne University of Technology, Hawthorn, VIC 3122, Australia*
[2]*OzGrav: ARC Centre of Excellence for Gravitational Wave Discovery, Hawthorn, VIC 3122, Australia*
[3]*ASTRON, the Netherlands Institute for Radio Astronomy, Postbus 2, NL-7990 AA Dwingeloo, the Netherlands*
[4]*Anton Pannekoek Institute for Astronomy, University of Amsterdam, Science Park 904, NL-1098 XH Amsterdam, the Netherlands*
[5]*International Centre for Radio Astronomy Research, Curtin University, GPO Box U1987, Bentley, WA 6845, Australia*

ABSTRACT

The electromagnetic counterparts to gravitational wave (GW) merger events hold immense scientific value, but are difficult to detect due to the typically large localization errors associated with GW events. The Low-Frequency Array (LOFAR) is an attractive GW follow-up instrument owing to its high sensitivity, large instantaneous field of view, and ability to automatically trigger on events to probe potential prompt emission within minutes. Here, we report on 144-MHz LOFAR radio observations of three GW merger events containing at least one neutron star that were detected during the third GW observing run. Specifically, we probe 9 and 16 per cent of the location probability density maps of S190426c and S200213t, respectively, and place limits at the location of an interesting optical transient (PS19hgw/AT2019wxt) found within the localization map of S191213g. While these GW events are not particularly significant, we use multi-epoch LOFAR data to devise a sensitive wide-field GW follow-up strategy to be used in future GW observing runs. In particular, we improve on our previously published strategy by implementing direction-dependent calibration and mosaicing, resulting in nearly an order of magnitude increase in sensitivity and more uniform coverage. We achieve a uniform 5σ sensitivity of 870 μJy beam^{-1} across a single instantaneous LOFAR pointing's 21 deg^2 core, and a median sensitivity of 1.1 mJy beam^{-1} when including the full 89 deg^2 hexagonal beam pattern. We also place the deepest transient surface density limits yet on time-scales of the order of month for surveys between 60 and 340 MHz (0.017 deg^{-2} above 2.0 mJy beam^{-1} and 0.073 deg^{-2} above 1.5 mJy beam^{-1}).

Key words: gravitational waves – techniques: interferometric – radio continuum: transients – black hole–neutron star mergers – neutron star mergers.

1 INTRODUCTION

Gravitational wave (GW) instruments have been detecting GW events from compact object mergers since 2015 (Abbott et al. 2016). The first GW merger event to consist of at least one neutron star (NS), GW170817, was detected in the second observing run, in 2017 (Abbott et al. 2017a). The event was temporally consistent with a short gamma-ray burst (GRB) detected by the *Fermi* Gamma-ray Burst Monitor 1.7 s later (Goldstein et al. 2017). The binary neutron star (BNS) merger's exceptionally small 30 deg^2 localization area facilitated swift identification of an optical counterpart (Coulter et al. 2017). This initiated a large-scale multiwavelength observing campaign of the electromagnetic (EM) counterpart that was eventually detected across the full EM spectrum (Abbott et al. 2017b). This rich trove of data enabled several significant scientific advancements across a multitude of disciplines in astrophysics. Radio data in particular provided an unprecedented detailed view of a relativistic short GRB jet and confirmed that GW170817 launched a successful jet that was viewed off-axis (Mooley et al. 2018a).

In general, radio emission is predicted to arise at various stages of merging compact object systems that are composed of at least one NS. Coherent emission has been predicted on time-scales ranging from seconds before the merger to years after should the merger produce a stable NS remnant (see Rowlinson & Anderson 2019; Gourdji et al. 2020, for an overview of merger-related models). Immediately following an NS merger, relativistic ejecta may be launched, which will shock the surrounding medium and power a synchrotron radio afterglow visible on time-scales ranging from days to several months (Paczyński & Rhoads 1993). Such radio emission was observed for the first time from a GW event with binary NS merger GW170817 (Hallinan et al. 2017). Broad-band radio monitoring of this afterglow was key to understanding the geometry, energy, kinematics, and total energy of the relativistic outflow (e.g. Hallinan et al. 2017; Mooley et al. 2018a, b). Notably, the geometry constraints broke degeneracies in GW distance measurements, leading to a significant improvement in the derived Hubble constant constraint (Hotokezaka et al. 2019).

In addition to relativistic outflows, the merger will eject neutron-rich matter, which will also shock the surrounding medium to produce a long-lasting but fainter synchrotron afterglow, in addition to an optical kilonova (Nakar & Piran 2011). The time-scales of this afterglow will depend largely on the velocity distribution of

★ E-mail: kgourdji@swin.edu.au

the ejected matter. This ejecta is expected to have a distribution of mostly sub-relativistic velocities but with a relativistic fast tail that could contribute to the aforementioned early-time synchrotron emission (Kyutoku, Ioka & Shibata 2014). The majority slower moving ejecta will cause a radio afterglow that will peak at lower radio frequencies and at later times (year time-scales; Hotokezaka & Piran 2015; Hotokezaka et al. 2016). This type of afterglow has yet to be observed and searches for an associated re-brightening in GW170817 are ongoing (e.g. Balasubramanian et al. 2022).

The Low-Frequency Array (LOFAR) has been active in GW follow-up since the first GW observing run in 2015 (see Gourdji et al. 2022, for a summary). As we demonstrate in this paper, LOFAR's large instantaneous field of view and high sensitivity (5σ of 870 μJy across ~ 21 deg^2) make it an attractive tool for GW follow-up, where events containing an NS are localized to areas ranging from tens to, more typically, thousands of square degrees (LIGO Scientific Collaboration 2021a, b). These large localization areas challenge EM counterpart identification and indeed, thus far, GW170817 is the only event with a secure EM counterpart. Furthermore, the relatively low frequencies at which LOFAR observes for GW follow-up (between 115 and 189 MHz, for single-beam pointings) can help to constrain the location of the synchrotron self-absorption frequency break in the highly time-dependent and degenerate afterglow spectrum (Sari, Piran & Narayan 1998). Additionally, LOFAR's low-latency rapid response mode enables us to trigger on GW events and collect data within minutes, which probes potential prompt coherent emission (e.g. Rowlinson et al. 2019, 2021).

In this paper, we report on late-time LOFAR radio follow-up of three GW merger events containing at least one NS detected during the third GW observing. We use and build on the strategies developed in Gourdji et al. (2022) to search for radio transients in wide-field LOFAR data corresponding to the radio afterglow of GW merger events. Here, we present a major improvement in reduction strategy relative to that presented in Gourdji et al. (2022): we perform direction-dependent calibration (DDC) and multibeam mosaicing, resulting in a factor of 7 increase in median sensitivity across the beam pattern and more uniform coverage. While we do not find any significant radio transients, we present our strategy for the fourth observing run (O4) as well as the deepest wide-field slow transient survey at low frequencies.

In Section 2, we provide details about the merger events we observed, the LOFAR observations, and data reduction. In Section 3, we present results from our transient search. We conclude and discuss our strategy for the upcoming GW observing run (O4) in Section 4.

2 GW EVENTS, LOFAR OBSERVATIONS, AND DATA REDUCTION

The third GW observing run (O3) took place between 2019 April 1 and 2020 March 27, with a hiatus during the month of October 2019. The Kamioka Gravitational Wave Detector (KAGRA) joined the GW detector network on 2020 February 25. Following offline GW analysis of O3 data by the LIGO Scientific, Virgo, and KAGRA (LVK) collaboration, there are 79 significant GW merger events (LIGO Scientific Collaboration 2021a, b). These correspond to GW events with a false alarm rate (FAR) less than 2 per year and probability of being astrophysical (p_{astro}) greater than 0.5. Of these significant events, seven are consistent with containing at least one neutron star [one BNS and six NS–black hole (NSBH) candidates].[1]

[1]Two of the NSBH candidates could also be BBH systems with a low-mass BH component.

These are GW190425 (BNS), GW190814, GW190426_152155, GW191219_163120, GW200105_162426, GW200115_042309, and GW200210_092254 (NSBH). All events but GW191219_163120 and GW200210_092254 were detected during real-time analysis and were issued low-latency public alerts via Gamma-ray Coordinate Network (GCN) Notices. No EM counterparts were robustly identified for any of the O3 candidates. We conducted late-time radio follow-up on a total of three merger events. One of these three GW events is considered a significant GW merger candidate following offline analysis by LVK: NSBH candidate GW190426_152155 (referred to as S190426c for the remainder of this paper). The events observed with LOFAR are S190426c (NSBH), S191213g (BNS), and S200213t (BNS). Note that the latter two events are not included in the final LVK catalogue, which was published approximately a year and half after their initial detections.

All LOFAR observations were conducted in the HBA Dual Inner configuration (24 core stations and 14 remote stations). Each multibeam observation was centred at an observing frequency of 144.24 MHz and has a total bandwidth of 13.48 MHz. The two single-beam observations for S191213g were centred at 152.15 MHz with a 72.22 MHz bandwidth. Each subband contains 64 channels of width 3.051 kHz. An integration time of 2 s was used for S190426c and S200213t, and 1 s for S191213g. Each observation was flanked by a 10 min calibrator scan of 3C48 and 3C196, in that order, except for the S200213t field, where 3C295 was used instead of 3C48.

2.1 GW events and LOFAR observations

2.1.1 S190426c

S190426c (also known as GW190426_152155) is an NSBH star merger candidate that was detected by all three LIGO/Virgo detectors on 2019 April 26 at 15:21:55 UT (LIGO Scientific Collaboration & Virgo Collaboration 2019). A preliminary trigger alert was issued about 25 min later, stating that the event had an FAR of once per 1.63 yr and was likely to be a merger system containing at least one NS and to have produced a merger remnant. This information, along with the event's location and the availability of an appropriate calibrator source, met the triggering criteria of our LOFAR rapid response strategy.[2] Unfortunately, a contemporaneous LOFAR observatory software update prevented triggered rapid-response observations of this event, which would have enabled us to be on-source within 5 min (see e.g. Rowlinson et al. 2019). For late-time follow-up, we centred a single pointing comprising seven overlapping LOFAR beams with centres separated by 2.8°, each with a radius of about 4.5°, on the peak of the location probability density map to produce a mosaic encapsulating 9 per cent of the location probability density map (see Fig. 1a). Given that this field had not yet been observed by the LOFAR Two-metre Sky Survey (LoTSS; Shimwell et al. 2017, 2019), we obtained 4 h of interferometric LOFAR data 1 week post-merger on 2019 May 3, to create a reference image for transient searches in subsequent observations. We conducted a follow-up 4-h observation approximately 1 month post-merger on 2019 May 24. The LVK later updated the classification of the event upon offline analysis and demoted the event from a 14 per cent chance of being terrestrial to a 58 per cent chance. For this reason as well as the high FAR, we cancelled further LOFAR follow-up of S190426c.

[2]We decreased the maximum allowed FAR to 10^{-14} Hz in our triggering criteria partway through O3, to reduce the number of events that would trigger LOFAR. Hence, the other two GW events did not trigger LOFAR rapid response observations.

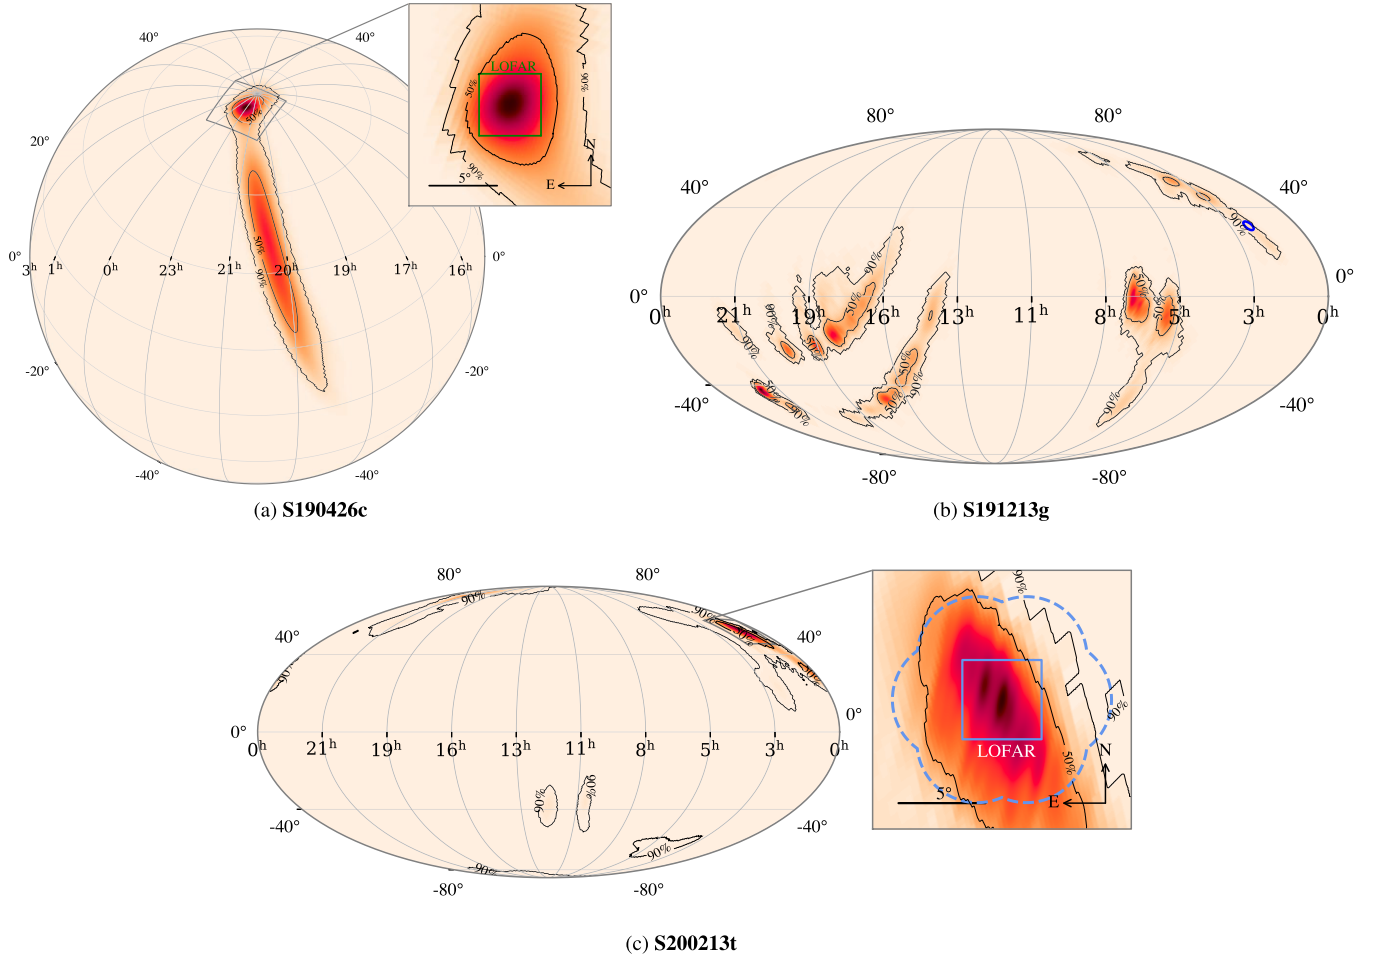**Figure 1.** Sky location probability density maps for the three GW events studied. (a) The corresponding LOFAR field ($4.6° \times 4.6°$) is represented by a green box in the inset figure. The 50 and 90 per cent credible regions cover 235 and 1119 deg^2, respectively. There is a 9 per cent chance that the GW event is located within the LOFAR field. The localization data come from GWTC-2.1 and correspond to products derived using the GSTLAL search pipeline (LIGO Scientific Collaboration & Virgo Collaboration 2021). (b) The blue circle marks the location of the single LOFAR beam centred on the location of optical transient AT2019wxt. The localization map is that derived by LALInference. (c) The full 89 deg^2 LOFAR mosaic is represented by the dashed blue line in the inset. The blue box in the inset corresponds to the $4.6° \times 4.6°$ inner core of the LOFAR mosaic, where the sensitivity is uniform to within about 5 per cent. The 50 and 90 per cent credible regions cover 129 and 2325 deg^2, respectively. The core of the LOFAR mosaic captures 16 per cent of the probability density map. The localization map is that derived by LALInference. These maps were created using LIGO.SKYMAP.

The second release of the catalogue of merger events from the first half of O3 (LIGO Scientific Collaboration 2021a) provides the latest GW analysis of this GW event. There, S190426c is considered to be a marginally significant candidate, with an FAR of 0.91 per year and a probability of being astrophysical and an NSBH of 14 per cent according to the GstLAL pipeline.[3]

2.1.2 S191213g

S191213g was detected by the LIGO–Virgo detectors on 2019 December 13 at 04:34:08 UT with an FAR of 1.1 per year and was deemed not to be a significant GW event in offline analysis (LIGO Scientific Collaboration 2021b). S191213g was classified as a BNS merger candidate and the Pan-STARRS collaboration identified an interesting optical transient, PS19hgw/AT2019wxt, in its localization

region on 2019 December 16 at 07:19:12 UTC (McBrien et al. 2019). This led S191213g to be one of the O3 events with the most GCN Circulars, with follow-up across the EM spectrum and by neutrino facilities. We placed a single LOFAR beam at the location of AT2019wxt (RA = 01:55:41.941, Dec. = +31:25:04.55, see Fig. 1b) and obtained 8-h observations 111 and 142 d post-merger. The field had already been processed by LoTSS and hence could be used as a reference. The optical transient was later classified as a likely Type IIb supernova following its photometric evolution and spectroscopic studies, and a more recent multiwavelength analysis by Shivkumar et al. (2022) classified AT2019wxt as an ultra-stripped supernova candidate.

2.1.3 S200213t

S200213t was a low signal-to-noise event detected on 2020 February 13 at 04:10:40 UT by only the Hanford LIGO detector with an FAR of 0.56 per year. As with S191213g, S200213t was detected in real-time as a BNS but did not pass the LVK threshold to be considered

[3]The low astrophysical NSBH probability is due to the low inferred rate of detectable NSBHs.

Table 1. Image RMS values in units of µJy. Values are provided for each field's epoch image (numbered below the field names). For S190426c and S200213t, the median RMS values across the respective fields are reported. For S191213g, the RMS measured at the centre of the image, where AT2019wxt is located, is given.

	S190426c		S191213g		S200213t		
Epoch	1	2	1	2	1	2	3
Mosaic core[a]	231	226	–	–	174	153	159
Full field[b]	–	–	–	–	216	191	190
AT2019wxt	–	–	63	70	–	–	–

[a] Inner 21.15 deg^2 core of seven-beam mosaics.
[b] 89.11 deg^2 seven-beam mosaics.

a merger candidate in their offline analysis of GW data. We used a single LOFAR pointing comprising seven overlapping beams (as described for S190426c above) centred on the peak of the location probability density map to probe 16 per cent of the map with the most sensitive part of the beam pattern (see Fig. 1c). We obtained 8-h observations 6, 26, and 95 d post-merger. We ceased further follow-up given that the event did not meet the criteria of a GW merger candidate in offline analysis (LIGO Scientific Collaboration 2021b).

2.2 Data reduction

Raw data were preprocessed by the observatory's data averaging pipeline following each observation. Here, the visibilities were flagged for RFI using AOFLAGGER, averaged in frequency and time, and bright nearby sources were subtracted where necessary using the pipeline's 'demix' option. S190426c was averaged to four channels per subband and 4 s, and nearby Cas A was demixed. S191213g was averaged to four channels per subband and 2 s. S200213t was averaged to 16 channels per subband and 2 s, to provide the flexibility to either demix or subtract Cas A from the field during our offline processing.

The averaged data were then calibrated using PREFACTOR (v3.0; de Gasperin et al. 2019), which corrects for direction-independent effects, and was used to demix Cas A from our S200213t data. Calibration solutions from 3C196 were used for the first two fields, whereas 3C295 was used for S200213t. Skymodels from the TIFR GMRT Sky Survey were used for phase calibration of the target fields. The direction-independent calibrated visibilities were averaged to two channels per subband and subbands were concatenated into groups of 10 and averaged in duration to 8 s. The PREFACTOR calibrated products were then passed through the DDF-PIPELINE[4] developed and used by the LoTSS team (Shimwell et al. 2019; Tasse et al. 2021), to deliver DDC and final image mosaics for scientific analysis.[5] The restoring beam of the final images has a radius of 3 arcsec and a pixel resolution of 1.5 arcsec. The flux density scales of the images have been corrected in the way described in Shimwell et al. (2019). For S190426c, we analyse the inner 21 deg^2 core of the mosaic, where sensitivity is uniform to within 5 per cent. For S200213t, we analyse the full mosaic, with reduced sensitivity towards the outer edges, to demonstrate the strategy to be adopted for O4. Root mean square (RMS) noise values for each image are reported in Table 1.

[4]Second data release version (https://github.com/mhardcastle/ddf-pipeline).
[5]We used the tier1-july2018.cfg pipeline configuration file.

3 TRANSIENT SEARCH AND RESULTS

Image mosaics were run through the LOFAR transients pipeline (TRAP, version r5.0; Swinbank et al. 2015) to extract sources and to identify new ones. Contrary to the default settings, we do not specify a source extraction margin, since the image quality of our mosaics is sound up to the image edges. Apart from that, default TRAP settings were used, unless otherwise stated below. The method used for the transient search analysis is based on that developed in Gourdji et al. (2022). Due to improved image quality thanks to DDC, the transient candidate filters outlined there are no longer necessary and can be removed. For instance, we no longer reject candidates near to bright sources, given that DDC drastically diminishes the presence of (and hence false positives caused by) sidelobes. Furthermore, we can lower the source detection threshold from 7σ to 5σ. The image sensitivities are summarized in Table 1.

3.1 S190426c

The average RMS in the first image of this field is slightly higher than in the second (see Table 1), and we thus consider the RMS properties of that image to determine the sensitivity of our search (Fig. 2). Eighty per cent of the field has a measured RMS below 260 µJy and the median 5σ sensitivity is 1.2 mJy beam^{-1}.

TRAP extracted 4295 unique sources across the two images, of which 125 were automatically flagged as new sources. These correspond to sources that were blindly extracted in the second epoch which were not associated with any sources that were blindly extracted in the first epoch. The positions of these new sources were recorded into a monitoring list that was then passed on to TRAP for a second run. This time, we forced TRAP to extract sources at those positions in both images using the shape of the restoring beam. Following Gourdji et al. (2022), for each transient candidate position, we calculated the difference between the peak flux density measurement in each image normalized by their errors summed in quadrature. To check whether a flux difference is significant, we compared to the distribution of normalized flux differences calculated for all compact persistent sources (integrated to peak flux ratio less than 5) that were blindly extracted in the first TRAP run. The results are shown in Fig. 3 where it is evident that none of the transient candidates are significant.

3.2 S191213g

The RMS around the core of the first and second images, which is centred on AT2019wxt, is 63 and 70 µJy, respectively. This sensitivity is consistent with the typical RMS achieved in LoTSS (Shimwell et al. 2019). We performed a forced source extraction, taking the shape of the restoring beam, at the coordinates of AT2019wxt in our two images. The flux measured in the first and second epochs is 3.8 ± 0.1 and 2.8 ± 0.1 mJy beam^{-1}, respectively. As can be seen in Fig. 4, the host galaxy of the transient is resolved and contaminates the extractions, which complicates accurately characterizing the flux density at the location of the transient. The resolved host galaxy is detected blindly by TRAP in both images and the integrated flux is 12.6 ± 0.3 mJy beam^{-1} in the first epoch and 10.0 ± 0.4 mJy beam^{-1} in the second. The field was observed by LoTSS on 2017 May 23 and the host galaxy radio properties are recorded in the catalogue (Shimwell et al. 2022). The resolved source has a peak flux of 5.7 ± 0.1 mJy beam^{-1}, and integrated flux of 14.7 ± 0.4 mJy beam^{-1} and a major axis of 10.2 ± 0.2 arcsec.
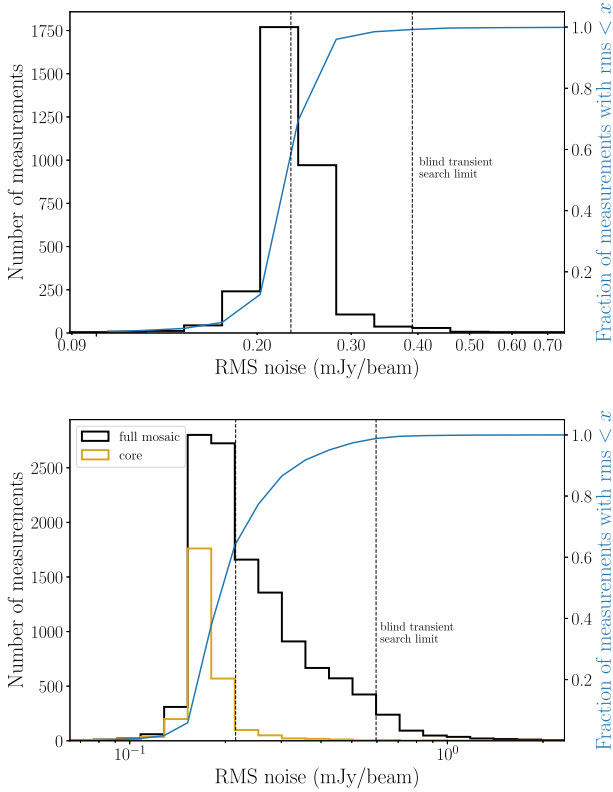
Figure 2. Distribution of RMS measurements from the sensitivity limiting images of S190426c (top) and S200213t (bottom). *Top:* The median RMS is 231 μJy beam^{-1} and is denoted by the left-most vertical line. Twenty one RMS measurements lie beyond the 700 μJy beam^{-1} *x*-axis limit shown here, with a maximum value of 53 mJy beam^{-1}. The right-most vertical line marks the RMS (0.39 mJy beam^{-1}) below which 99 per cent of the measurements lie. This value is used to set the limiting sensitivity of 2.0 mJy beam^{-1} for our blind transient surface density limit. The blue line traces the fraction of the image with RMS $<x$ (right-hand *y*-axis). *Bottom:* Similar to the top panel. The distribution of RMS measurements in the inner 21 deg^2 core is included in yellow, while measurements from the full mosaic are shown in black. The left-most vertical line represents the median RMS of the full mosaic of 216 μJy beam^{-1}. The 99 per cent blind transient limit is at 600 μJy beam^{-1}. There are 23 additional RMS measurements beyond the 2.1 mJy beam^{-1} limit applied to the plot for visualization purposes.

We consider whether the apparent halving of the host galaxy's measured peak flux density in our images (6.1 ± 0.1 to 3.6 ± 0.1 mJy beam^{-1}) is significant by comparing the flux difference normalized by the uncertainties (see Gourdji et al. 2022) to the other persistent compact sources (with major axes less than 12 arcsec) blindly extracted in both of our images. We find that the normalized flux difference is consistent within a standard deviation of the distribution of persistent sources and hence conclude that the host's flux density variation is insignificant. We further note that the source extractor used by TRAP, PYSE, is not designed to handle extended sources (Carbone et al. 2018). Thus, the errors on the integrated flux of the host reported by TRAP are likely underestimated.

While this field turned out not to be actually related to an NS merger, this particular scenario, where the target of interest is contaminated by a nearby radio source, raises the question of whether the blind transient search strategy presented in Gourdji et al. (2022) and used for the other fields in this study would succeed in identifying a merger radio counterpart that is blended with its host.
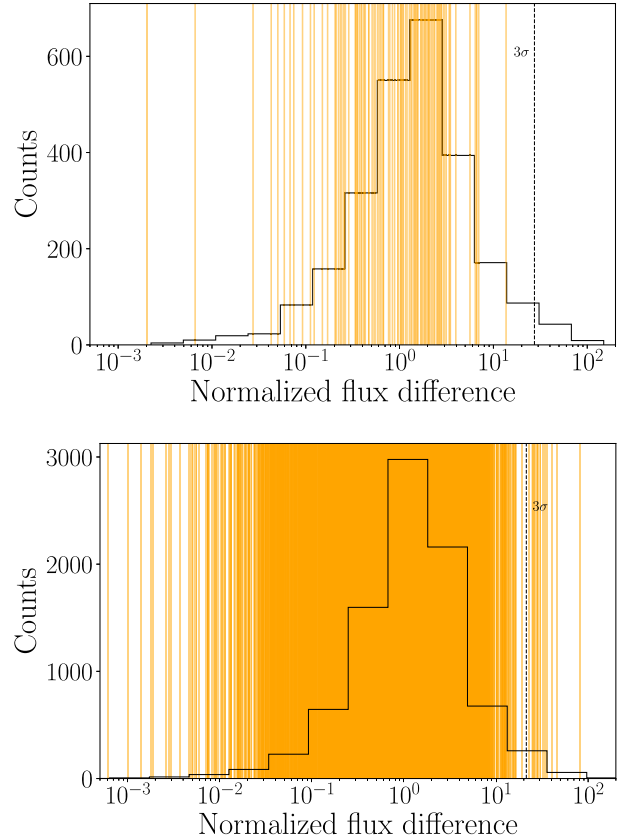

Figure 3. Normalized flux difference of transient candidates measured in each epoch image (orange) plotted over the distribution of normalized flux differences of persistent compact sources blindly extracted in the S190426c (top) and S200213t (bottom) fields. Thirteen transient candidates exceed three standard deviations from the persistent source distribution in the S200213t field and are further considered in the text.

Our transient search does not distinguish overlapping sources by, for example, fitting them with multiple overlapping Gaussians. Rather, as described in Gourdji et al. (2022), using its default settings, TRAP would likely fit a single Gaussian to a multicomponent source, as it would to any extended source. A possible way to prevent this would be to use TRAP's 'deblend' option (Carbone et al. 2018), though using this setting for such GW transient searches would require non-trivial testing and adjustments to our search algorithm and is beyond the scope of this study.

3.3 S200213t

With its 8-h integration time, this field represents the observing strategy to be adopted in O4 and we therefore search the full mosaic rather than the inner core to demonstrate the end-to-end strategy that will be used for events observed with LOFAR in the upcoming observing run. The transient search strategy used for this field is identical to that of S190426c; however, the larger search area and additional epoch result in more transient candidates. The image from the first epoch has the highest RMS and so we use its properties to determine our transient search's median 5σ sensitivity of 1.1 mJy beam^{-1} (870 μJy beam^{-1} in the mosaic core).

TRAP blindly extracted 16 257 unique sources across three images, each encompassing an area of 89 deg^2. TRAP identified 3002 transient
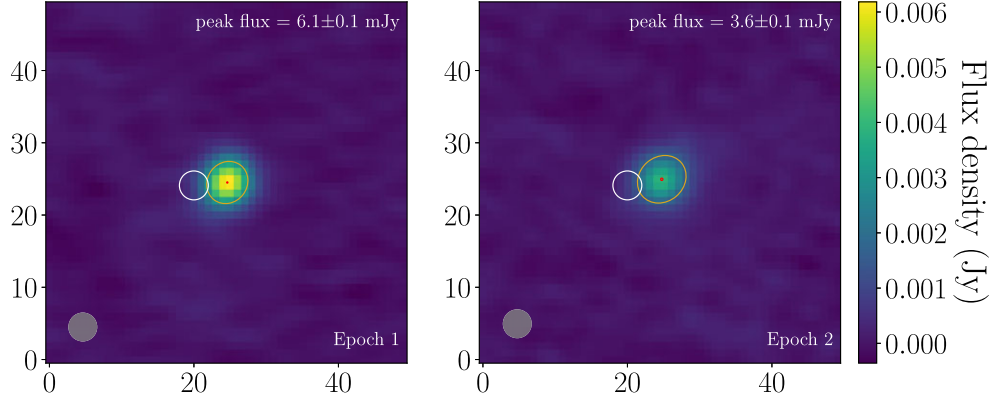
Figure 4. LOFAR images of AT2019wxt and its host galaxy, KUG 0152+311, 108 and 139 d after the transient was identified (which itself was identified 3 d following GW event S191213g). AT2019wxt is denoted by a white circle with size corresponding to the synthesized beam. The host position and 2D Gaussian fit are denoted by the red and yellow markers, respectively, with values as measured by TRAP during blind extraction. The circular synthesize beam is shown in the bottom left of each image.

candidates. These were treated following the same procedure as described in Section 3.1 and results from our candidate significance check are shown in Fig. 3. Thirteen candidates have significant normalized flux differences and were inspected manually. All but one were clearly extended sources and visible in all three epochs and thus are rejected from further consideration. The remaining candidate is slightly resolved and detected in the second and third epochs (25.0σ and 22.4σ, respectively), but not present in the first (the forced fit, using the shape of the restoring beam, gives a peak flux of $-2.4 \pm 0.9\,\text{mJy beam}^{-1}$). However, large sidelobe image artefacts are present in the first image (as in the other two), which indicates that the source may have been erroneously cleaned during the direction-dependent data reduction process. We confirm this by checking the intermediate calibration products and find that the source is present in the data up until the final DDF-PIPELINE calibration step. We thus rule out this candidate and conclude that no physical transients are present in this field. Diagnostic images of the transient candidate are included in the appendix (Fig. A1).

3.4 Transient surface density

To set a limit on the blind transient surface density, we consider a minimum sensitivity corresponding to 99 per cent completeness in each field considered. This corresponds to $2.0\,\text{mJy beam}^{-1}$ and $20.7\,\text{deg}^2$ of sky surveyed for S190426c, $3.0\,\text{mJy beam}^{-1}$ and $85.5\,\text{deg}^2$ for S200213t, and $1.5\,\text{mJy beam}^{-1}$ and $20.6\,\text{deg}^2$ for the latter's inner core. Following the same strategy as in Gourdji et al. (2022), we calculate the corresponding 95 per cent confidence surface density limits. We find 0.15, 0.018, and $0.073\,\text{deg}^{-2}$, respectively. It is, however, more constraining than the former two values to consider the surface density limit of the combined area surveyed above $2.0\,\text{mJy beam}^{-1}$ across both S190426c and S200213t ($20.7 + 76.4\,\text{deg}^2$), which yields $0.017\,\text{deg}^{-2}$. These constraints are summarized and compared to other relevant transient searches in Fig. 5, where it is evident that this analysis constitutes the most sensitive blind search for transients on time-scales longer than 1 min at frequencies between 60 and 340 MHz.

4 CONCLUSIONS AND O4 STRATEGY

We have presented a LOFAR search at 144 MHz for radio transients related to three GW events detected in O3. We have demonstrated
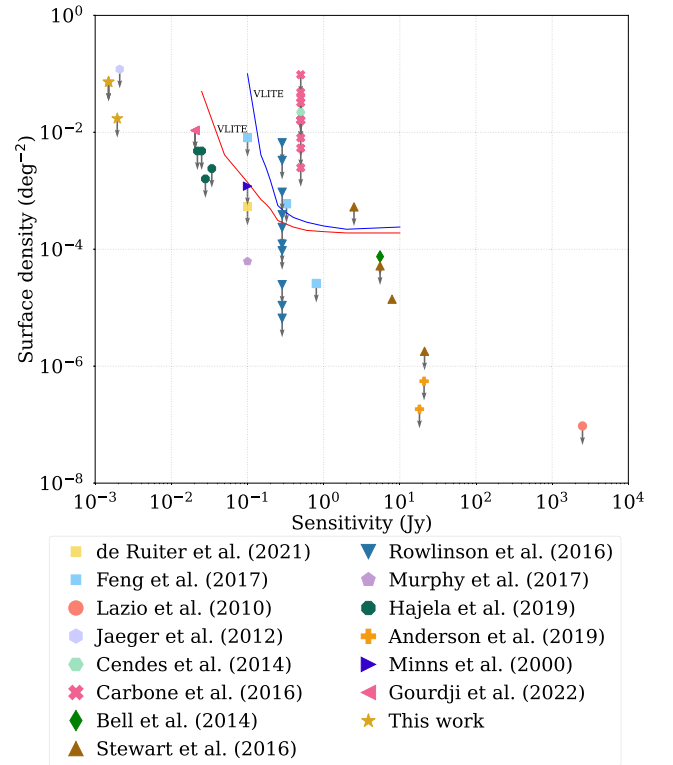

Figure 5. Transient surface density limits for sources longer than 1 min and surveys between 60 and 340 MHz. Upper limits are represented by down-facing arrows. Limits from this study are denoted by gold stars and represent the most sensitive limits yet. The VLITE survey limits come from Polisensky et al. (2016). This figure was adapted from fig. 5 of Gourdji et al. (2022). Galactic centre radio transients are excluded.

a significant improvement to the strategy previously presented in Gourdji et al. (2022) by employing DDC and creating mosaics out of seven-beam pointings. We have achieved uniform $870\,\mu\text{Jy beam}^{-1}$ 5σ sensitivity across the inner $21\,\text{deg}^2$ of the resulting 8-h mosaics, and a median sensitivity of $1.1\,\text{mJy beam}^{-1}$ across the full $89\,\text{deg}^2$ mosaiced field. We have also demonstrated our ability to attain LoTSS sensitivities for single-beam fields (RMS $\sim 70\,\mu\text{Jy beam}^{-1}$). Further, the multi-epoch data sets presented here constitute the

deepest search for transients longer than 1 min at frequencies between 60 and 340 MHz.

The next GW observing run is expected to commence on 2023 May 24 and to last 18 months, with a likely month-long maintenance break part-way.[6] This will be preceded by a month-long engineering run where exceptional events may be reported to the public. The next LOFAR observing cycle will be from 2023 June 1 to 2024 May 31, providing excellent overlap before stations are gradually brought offline thereafter as part of the LOFAR2.0 roll-out.

We plan to use the strategies used in the analyses presented here for LOFAR follow-up observations of events in O4. In particular, our wide-field search strategy will match that presented for S200213t (870 μJy beam^{-1} sensitivity across an instantaneous 21 deg^2 field and further coverage at reduced sensitivity out to 89 deg^2) and we will apply the S191213g strategy for events that have a previously identified EM counterpart. When available, LoTSS data will be used as reference for our transient search, and observations would be taken 1, 3, and 6 months post-merger. The cadence is chosen to properly sample the afterglow light curve and to probe events that are significantly off-axis (e.g. refer to fig. 2 of Broderick et al. 2020). Additionally, we plan to make use of LOFAR's rapid response capability to trigger on significant GW events, enabling us to gather interferometric data across 89 deg^2 within minutes of the merger. With these data sets, we will use the strategy presented in Rowlinson et al. (2019, 2021) to search for coherent emission resulting from the merger remnant. The LVK forecasts 36^{+49}_{-22} BNS and 6^{+11}_{-5} NSBH merger alerts per year[7] during O4, and we can expect about a quarter to have their location probabilities peak in LOFAR's field of view.

ACKNOWLEDGEMENTS

DATA AVAILABILITY

The LOFAR visibility data used in this analysis are publicly available at the LOFAR Long Term Archive (https://lta.lofar.eu) under

project code LT10_013. The image products used for the scientific analysis presented here will be shared on reasonable request to the corresponding author.

REFERENCES

Abbott B. P. et al., 2016, Phys. Rev. Lett., 116, 061102
Abbott B. P. et al., 2017a, Phys. Rev. Lett., 119, 161101
Abbott B. P. et al., 2017b, ApJ, 848, L12
Balasubramanian A. et al., 2022, ApJ, 938, 12
Bell M. E. et al., 2014, MNRAS, 438, 352
Broderick J. W. et al., 2020, MNRAS, 494, 5110
Carbone D. et al., 2016, MNRAS, 459, 3161
Carbone D. et al., 2018, Astron. Comput., 23, 92
Cendes Y. et al., 2014, preprint (arXiv:1412.3986)
Coulter D. A. et al., 2017, Science, 358, 1556
de Gasperin F. et al., 2019, A&A, 622, A5
de Ruiter I., Leseigneur G., Rowlinson A., Wijers R. A. M. J., Drabent A., Intema H. T., Shimwell T. W., 2021, MNRAS, 508, 2412
Feng L. et al., 2017, AJ, 153, 98
Goldstein A. et al., 2017, ApJ, 848, L14
Gourdji K., Rowlinson A., Wijers R. A. M. J., Goldstein A., 2020, MNRAS, 497, 3131
Gourdji K., Rowlinson A., Wijers R. A. M. J., Broderick J. W., Shulevski A., Jonker P. G., 2022, MNRAS, 509, 5018
Hajela A. et al., 2019, ApJ, 886, L17
Hallinan G. et al., 2017, Science, 358, 1579
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., 2019, Nat. Astron., 3, 940
Hyman S. D., Wijnands R., Lazio T. J. W., Pal S., Starling R., Kassim N. E., Ray P. S., 2009, ApJ, 696, 280
Jaeger T. R., Hyman S. D., Kassim N. E., Lazio T. J. W., 2012, AJ, 143, 96
Kyutoku K., Ioka K., Shibata M., 2014, MNRAS, 437, L6
Lazio T. J. W. et al., 2010, AJ, 140, 1995
LIGO Scientific Collaboration, Virgo Collaboration, 2019, GCN Circ., 24237, 1
LIGO Scientific Collaboration, Virgo Collaboration, 2021, GWTC-2.1: Deep Extended Catalog of Compact Binary Coalescences Observed by LIGO and Virgo During the First Half of the Third Observing Run – Candidate Data Release, https://doi.org/10.5281/zenodo.5148739
LIGO Scientific Collaboration, 2021a, preprint (arXiv:2108.01045)
LIGO Scientific Collaboration, 2021b, preprint (arXiv:2111.03606)
McBrien O. et al., 2019, GCN Circ., 26485, 1
Minns A. R., Riley J. M., 2000, MNRAS, 315, 839
Mooley K. P. et al., 2018a, Nature, 561, 355
Mooley K. P. et al., 2018b, ApJ, 868, L11
Murphy T. et al., 2017, MNRAS, 466, 1944
Nakar E., Piran T., 2011, Nature, 478, 82
Paczyński B., Rhoads J. E., 1993, ApJ, 418, L5
Polisensky E. et al., 2016, ApJ, 832, 60
Rowlinson A., Anderson G. E., 2019, MNRAS, 489, 3316
Rowlinson A. et al., 2016, MNRAS, 458, 3506
Rowlinson A. et al., 2019, MNRAS, 490, 3483
Rowlinson A. et al., 2021, MNRAS, 506, 5268
Sari R., Piran T., Narayan R., 1998, ApJ, 497, L17
Shimwell T. W. et al., 2017, A&A, 598, A104
Shimwell T. W. et al., 2019, A&A, 622, A1
Shimwell T. W. et al., 2022, A&A, 659, A1
Shivkumar H. et al., 2022, preprint (arXiv:2208.09010)
Stewart A. J. et al., 2016, MNRAS, 456, 2321
Swinbank J. D. et al., 2015, Astron. Comput., 11, 25
Tasse C. et al., 2021, A&A, 648, A1
van Haarlem M. P. et al., 2013, A&A, 556, A2

[6] https://observing.docs.ligo.org/plan/index.html
[7] https://emfollow.docs.ligo.org/userguide/capabilities.html

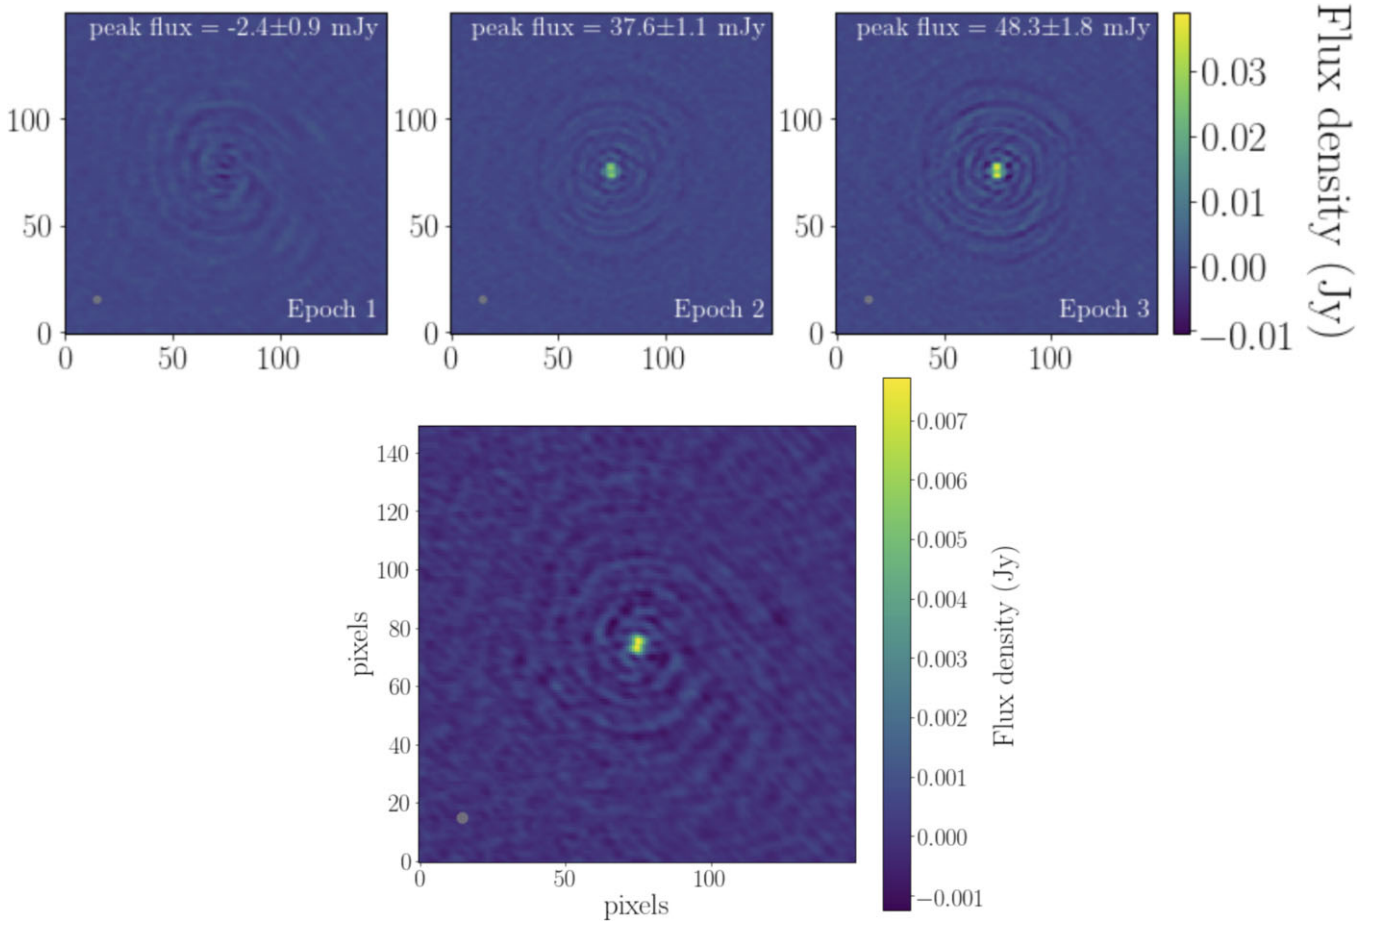## APPENDIX A: TRANSIENT CANDIDATE

Figure A1. Transient candidate detected in the S200213t field. *Top:* Image cutouts at the location of the transient candidate in each epoch. The restoring beam is included in the bottom left corner of each panel. The source is detected blindly in the second and third epochs, and is not apparent in the first epoch. *Bottom:* Cutout at the location of the transient candidate in an intermediate DDF-PIPELINE calibration step image of the first epoch. The source is clearly present and we thus conclude that the source was erroneously removed in the final calibration step and rule out this transient candidate.

This paper has been typeset from a TEX/LATEX file prepared by the author.