



UvA-DARE (Digital Academic Repository)

Transient study using LoTSS - framework development and preliminary results

de Ruiter, I.; Meyers, Z.S.; Rowlinson, A.; Shimwell, T.W.; Ruhe, D.; Wijers, R.A.M.J.

DOI

[10.1093/mnras/stae1458](https://doi.org/10.1093/mnras/stae1458)

Publication date

2024

Document Version

Final published version

Published in

Monthly Notices of the Royal Astronomical Society

License

CC BY

[Link to publication](#)

Citation for published version (APA):

de Ruiter, I., Meyers, Z. S., Rowlinson, A., Shimwell, T. W., Ruhe, D., & Wijers, R. A. M. J. (2024). Transient study using LoTSS - framework development and preliminary results. *Monthly Notices of the Royal Astronomical Society*, 531(4), 4805-4822. <https://doi.org/10.1093/mnras/stae1458>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Transient study using LoTSS – framework development and preliminary results

Iris de Ruiter ¹★, Zachary S. Meyers,^{2,3} Antonia Rowlinson ^{1,4}, Timothy W. Shimwell,^{4,5} David Ruhe^{1,6} and Ralph A. M. J. Wijers ¹

¹*Anton Pannekoek Institute for Astronomy, University of Amsterdam, Science Park 904, NL-1098 XH Amsterdam, the Netherlands*

²*Deutsches Elektronen-Synchrotron DESY, Platanenallee 6, D-15738 Zeuthen, Germany*

³*Erlangen Center for Astroparticle Physics (ECAP), Friedrich-Alexander-Universität Erlangen-Nuremberg, D-91058 Erlangen, Germany*

⁴*ASTRON, the Netherlands Institute for Radio Astronomy, Postbus 2, NL-7990 AA Dwingeloo, the Netherlands*

⁵*Leiden Observatory, Leiden University, PO Box 9513, NL-2300 RA Leiden, the Netherlands*

⁶*AI4Science (AMLab), Informatics Institute, University of Amsterdam, Science Park 904, NL-1098 XH Amsterdam, the Netherlands*

Accepted 2024 June 7. Received 2024 June 6; in original form 2023 November 14

ABSTRACT

We present a search for transient radio sources on time-scales of seconds to hours at 144 MHz using the LOFAR Two-metre Sky Survey (LoTSS). This search is conducted by examining short time-scale images derived from the LoTSS data. To allow imaging of LoTSS on short time-scales, a modern imaging procedure and fast filtering strategy are introduced. This includes sky model source subtraction, no cleaning or primary beam correction, a simple source finder, fast filtering schemes, and source catalogue matching. This new strategy is first tested by injecting simulated transients, with a range of flux densities and durations, into the data. We find the limiting sensitivity to be 113 and 6 mJy for 8 s and 1 h transients, respectively. The new imaging and filtering strategies are applied to 58 fields of the LoTSS survey, corresponding to LoTSS-DR1 (2 per cent of the survey). One transient source is identified in the 8 s and 2 min snapshot images. The source shows 1 min duration flare in the 8 h observation. Our method puts the most sensitive constraints on/estimates of the transient surface density at low frequencies at time-scales of seconds to hours; $<4.0 \times 10^{-4} \text{ deg}^{-2}$ at 1 h at a sensitivity of 6.3 mJy; $5.7 \times 10^{-7} \text{ deg}^{-2}$ at 2 min at a sensitivity of 30 mJy; and $3.6 \times 10^{-8} \text{ deg}^{-2}$ at 8 s at a sensitivity of 113 mJy. In the future, we plan to apply the strategies presented in this paper to all LoTSS data.

Key words: techniques: image processing – surveys – software: data analysis – radio continuum: transients.

1 INTRODUCTION

The transient radio sky provides a unique opportunity to study the most extreme events that take place in the Universe. The astrophysical processes that generate transient phenomena are often highly dynamic and explosive, allowing us to study environments that are inaccessible on Earth. Transient or highly variable sources have been observed across all wavelengths, however, radio astronomy offers a different perspective as some astrophysical phenomena are either highly beamed, unique to radio frequencies or obscured by dust at other wavelengths. Over the last two decades, radio transients have been discovered all across the transient phase space, which spans orders of magnitude in transient time-scales, observing frequency, and flux density. There are several astrophysical phenomena that are known to be transient at low radio frequencies. These include events like stellar flares, magnetar flares, novae, X-ray binaries, intermittent pulsars, fast radio bursts (FRBs), and strongly scintillating active galactic nuclei (AGNs). In this study, we focus on searching for low-frequency (144 MHz) radio transients with durations of seconds

to hours. To this end, we use survey data obtained with the Low Frequency Array (LOFAR; van Haarlem et al. 2013). Our search is sensitive to various transient phenomena. The most relevant source classes for this study are giant pulses from pulsars and flare stars (Spangler & Moffett 1976; Callingham et al. 2021; Feeney-Johansson et al. 2021), coronal mass ejections (Crosley & Osten 2018), X-ray binaries (Chandra & Kanekar 2017; Chauhan et al. 2021; Monageng et al. 2021), and possibly Algol-type binaries (Lefevre, Klein & Lestrade 1994; Umana, Triglio & Catalano 1998). Recently, long-period magnetars are confirmed to exist and these are also detectable in the low-frequency image plane (Caleb et al. 2022; Hurley-Walker et al. 2022, 2023). The references here point specifically to low-frequency radio detections of these phenomena within the aforementioned time-scales. We note that Algol-type binaries have mainly been studied at higher radio frequencies (>1 GHz) and the variability time-scale of X-ray binaries might be too long (\sim days) to probe with this study. Finally, strongly scintillating background AGN could be interpreted as variables or even transient sources. An overview of radio transient phenomena at various time-scales can be found in figs 3 and 5 in Pietka, Fender & Keane (2015).

Additionally, there are several theories predicting low-frequency, coherent radio emission from short gamma-ray bursts from com-

* E-mail: i.deruiter@uva.nl

pact binary mergers, see e.g. Rowlinson & Anderson (2019) and references therein. Some of these models suggest that the emission mechanisms of these types of sources are similar to FRBs, which are another target of this study due to the low observing frequency of LOFAR. Image domain searches for FRBs (Tingay et al. 2015; Rowlinson et al. 2016; Andrianjafy et al. 2023; Driessen et al. 2023) utilize the signal delay introduced by the dispersion measure. The signal delay $\Delta t \approx 4.15(\nu_{\text{lo}}^{-2} - \nu_{\text{hi}}^{-2})\text{DM}$ ms is defined in terms of the lower and upper limit of the observation bandwidth in GHz, ν_{lo} and ν_{hi} , respectively, and the dispersion measure, DM (see equation 1 in Petroff, Hessels & Lorimer 2019). For the LoTSS bandwidth of 0.120–0.168 GHz (Shimwell et al. 2019) and a typical DM of 500 cm^{-3}pc (see fig. 6 in Chawla et al. 2022 and fig. 1 in Arcus et al. 2021) one thus expects a signal delay of 70.6×10^3 ms. The FRB signal will be spread out over 70 s, implying that a bright burst could be detected in the image domain.

Furthermore, there is a possibility to probe new source classes. Examples of these are frequent in the low-frequency radio sky. A famous example is the class of Galactic centre Radio Transients (Hyman et al. 2002, 2005, 2009), whose bursts last from minutes to months. Specifically, a source like the ‘Galactic Burper’, which has shown a series of ~ 1 Jy bursts (Hyman et al. 2005) and single bursts years later (Hyman et al. 2006, 2007), would be an ideal target for our study. Jaeger et al. (2012) find a low-frequency radio transient that is variable on a time-scale of hours. The source has no counterpart which makes identification difficult, but several characteristics point towards a stellar flare. Obenberger et al. (2014) find two extremely bright transients, at peak flux densities of ~ 3 kJy, lasting for ~ 100 s at 30 MHz. Finally, Stewart et al. (2016) find a bright (possibly Galactic) transient towards the North Celestial Pole at 60 MHz, lasting for < 10 min. These transients discussed in this paragraph are either difficult to ascribe to any of the known source classes, or could possibly be detections of exciting new source classes.

Traditionally, most detections of coherent emitters have been done with time-series techniques (time-domain), while incoherent emission, which generally evolves over longer time-scales, is observed in the image plane (image-domain). Image-domain studies make snapshot images of a patch of sky and use those to look at the time variability of sources. For coherent emission processes, the accelerated particles can cooperate in phase, resulting in emission that can reach extremely high brightness temperatures. Incoherent emission comes from the summation of the radiation from individual accelerated particles and therefore the brightness temperature is limited to $\sim 10^{12}$ K. The stellar flare mechanisms are usually coherent emission mechanisms at 144 MHz, as well as the emission from magnetars (related to short gamma-ray bursts or FRBs), while most of the other source classes operate through incoherent emission mechanisms, usually synchrotron emission. This work presents an image domain transient study that probes both coherent and incoherent processes.

In order to quantify the number of transients we expect to see in the low-frequency radio sky, we determine the rate of transient sources at various time-scales and flux densities. To obtain the most constraining value of the transient surface density, observations of large sky areas at high sensitivities are required. This paper performs a transient search in the image domain using the LOFAR Two-metre Sky Survey (Shimwell et al. 2019, 2022, LoTSS). LoTSS images the low-frequency northern sky at unprecedented sensitivity and resolution. Section 2 presents a method to create snapshot images at 8 s, 2 min, and 1 h time-scales from the 8 h LoTSS pointings, and search for transients in the snapshot images. Section 3 discussed the sensitivity limits of this method via simulated transient sources.

We apply this method to a preliminary data set of 58 LoTSS-DR2 pointings (corresponding to the coverage of LoTSS-DR1) and present the results in Section 4. The implications of these results are discussed in Section 5 and finally, our concluding remarks and outlook for future application of our methods, are presented in Section 6.

2 METHODS

The next subsections will discuss each step in our method to search for transients on snapshot time-scales of 8 s, 2 min, and 1 h in each 8 h LoTSS pointing. Fig. 1 shows a visual schematic of the process. Each box in Fig. 1 corresponds to one subsection. The method presented here is meant to identify transient candidates that leave no signature in the deep 8-h integration image. Variability of sources in the deep image is currently excluded from this analysis. The method described below is developed to identify transient candidates in an efficient manner, follow-up with more traditional and elaborate imaging approaches is then necessary to fully characterize the transient source.

2.1 LOFAR Two-Metre Sky Survey

LOFAR (van Haarlem et al. 2013) is a low-frequency radio telescope that is comprised of many thousands of dipole antennas arranged in stations. These stations are distributed all over the Netherlands and more sparsely throughout Europe. The LoTSS (Shimwell et al. 2017) aims to image the whole northern sky in 3168 pointings. The survey has had two major data releases so far, DR1 (Shimwell et al. 2019) covering 58 pointings and DR2 (Shimwell et al. 2022) covering 814 pointings. LoTSS observes between 120 and 168 MHz. The flux densities are given at the central frequency of 144 MHz.

Whilst LoTSS data contains the entire international baseline coverage of LOFAR the data releases to-date only contain the Dutch stations, yielding a maximum baseline of 121 km (van Haarlem et al. 2013) resulting in an image resolution of 6 arcsec, due to computational limitations. This resolution combined with a median RMS of $83 \mu\text{Jy beam}^{-1}$ for the low-frequency continuum images allows LoTSS to venture into a realm of the radio sky that has been unexplored up to now.

In this work, we perform our transient study on the 58 pointings that cover the first data release of LoTSS (Shimwell et al. 2019), but we note that we do use LoTSS-DR2 products that have been processed according to Shimwell et al. (2022). To date, the LoTSS data releases have only included regions at high Galactic latitudes. However, the goal of LoTSS is to image the whole Northern sky, including the Galactic Plane (Shimwell et al. 2022). The Galactic Plane is the most promising region for most transient candidates discussed in the introduction. One of the goals of this paper is to establish a framework that can be applied to future LoTSS data releases, including the Galactic Plane.

2.2 Sky model subtraction

A technique that is often employed in transient studies is the use of *difference imaging*; once you create all snapshot images, subtract images of consecutive time-steps, resulting in a difference that should allow for easier transient identification. However, this technique is not well suited for radio transient studies, because with many facilities it is challenging to create good-quality images for each time-step due to sparse uv-coverage giving an irregular point spread function (PSF). Due to the irregular PSF and structured noise in individual images, a large number of artefacts are created when subtracting

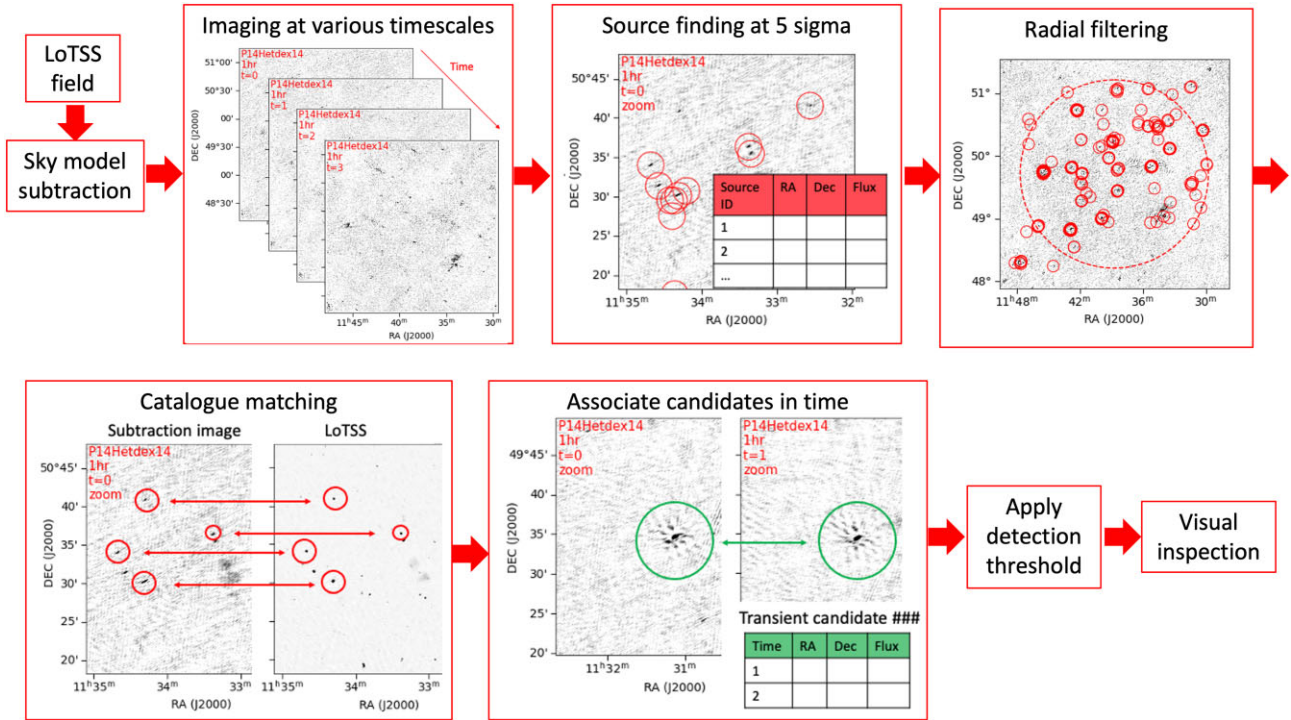


Figure 1. Illustration of our method to search for transients on snapshot time-scales of 8 s, 2 min, and 1 h in each 8 h LoTSS pointing. Each box corresponds to a subsection in the Methods.

subsequent images. Therefore, the difference imaging technique is not well-suited for radio transient searches.

As an alternative, here we use source *subtracted data* to search for transients (see e.g. Fijma et al. 2023; Wang et al. 2023). Source subtracted data are created by subtracting the full-observation sky model from the uv-plane using DDFacet (Tasse et al. 2018, 2021) to apply the direction-dependent calibration solutions during the subtraction. The images at various time-scales created from these data show the difference between the sky during the snapshot time and the full observation sky model. A more mathematical description of this technique can be found in Appendix A. The source subtracted data is imaged at various cadences and because most of the signals are removed by subtracting out the sky model, the adverse effects of the poor uv-coverage and associated PSF in the images are limited.

There are three major benefits to working with source subtracted data, compared to traditional methods. First of all, subtraction of the sky model from the uv-plane greatly reduces the compute time spent on imaging, because no primary beam correction or cleaning is required, as there are typically no sources in the field. A reduction of the imaging time is critical when imaging a full 8-h observation on an 8 s cadence. An additional benefit is that we can create source subtracted images even when the PSF is poorer due to sparse uv-sampling. Secondly, the subtracted images should, in theory, only contain the sources that are not in the sky model, or sources with a variable flux density compared to the sky model, which simplifies a transient search. Lastly, by subtracting the full 8-hour sky model from the shorter time-scale snapshots, one can remove the high confusion noise from the snapshots, which allows for a deeper transient search (see e.g. Fijma et al. 2023). This is because the removal of the sources allows one to image at lower resolution without being limited by confusion noise.

One aspect that we want to point out is that the direction-dependent calibration solutions have time-variable behaviour. Therefore, a

potential problem of this method is that the solutions could absorb an actual transient in order to make the result look like the sky model against which the data is calibrated. To this end, we perform simulations where we inject a transient source before calibration (against a sky model without the transient source) and analyse the result. See Section 3.1 for more details on these simulations.

In conclusion, the subtraction imaging method should be able to probe any variable and transient behaviour on time-scales shorter than the duration of the full sky model observation, which is 8 h in the case of LoTSS. A script¹ to perform the sky model subtraction exists within the pipeline used for LoTSS processing (DDF-pipeline, Tasse et al. 2018, 2021)² and yields a single sky model subtracted measurement set for each LoTSS field (or pointing). Throughout the rest of this work, when referring to *subtracted images*, we refer to the sky model subtracted images.

2.3 Snapshot imaging

After the sky model has been subtracted, images are produced from the subtracted data column in the measurement sets using WSCLEAN (Offringa et al. 2014). The subtracted data is imaged on cadences of 8 s, 2 min, and 1 h, which yields a large number of images. For example, the 8 s cadence will yield 3600 images for the full 8 h observation per field. Therefore, imaging parameters are chosen to minimize the compute time spent on imaging. No CLEAN algorithm iterations need to be applied because our first goal is to find any emission that is not subtracted out (i.e. transients sources). To this end, there is no need to deconvolve the PSF from the dirty image.

¹<https://github.com/mhardcastle/ddf-pipeline/blob/master/scripts/sub-sources-outside-region.py>

²<https://github.com/mhardcastle/ddf-pipeline>




Figure 2. Imaging artefact around a 2.1 Jy source in an 8 s subtracted image snapshot. The left-hand panels show the subtracted image, while the right-hand panels show the corresponding sky area in the LoTSS survey.

Table 1. The WSCLEAN settings used to image all the observations presented in this analysis. All other settings are default settings.

Setting	Value
Number of clean iterations	0
Size of image (pixels)	2200
Size of one pixel (arcsec)	6
(Min, max) uv-l	(50,60000)
Briggs weighting	-0.25
Intervals out	8, 225, 3600
Padding	1.6

Additionally, no primary beam correction is performed because we do not need accurate flux densities to identify transient sources. Furthermore, transient sources should be easier to identify against a more uniform noise background. An example of a subtracted snapshot image is given in Fig. 2. The left-hand panel shows an 8 s integration subtracted image and the right-hand panel shows the corresponding part of the sky in the full integration 8 h LoTSS data. Most sources have been subtracted out nicely, except for the bright 2.1 Jy sources in the centre of the image, which shows an artefact of inaccurate source subtraction in the subtracted image. Section 2.6 elaborates on how we mitigate these particular artefacts.

The most important WSCLEAN parameters include a pixel size of 6 arcsec, implying that the LOFAR Dutch station baseline resolution is mapped to one pixel. The LoTSS pointings are typically separated by 2.58° (Shimwell et al. 2019), and to ensure substantial overlap we image $3.67 \times 3.67^\circ$ for each pointing. Finally, an important parameter is the padding factor, which specifies the factor by which the image size is increased beyond the field of interest to avoid edge issues. We found it was crucial to increase this parameter from its default value of 1.2–1.6 to reduce some of the fake transient source introduced by aliasing effects. Details on this issue are given in Appendix B. The WSCLEAN imaging parameters are summarized in Table 1.

2.4 Source finding

We use the ‘Live Pulse Finder’ software presented in Ruhe et al. (2021, 2022) to perform source finding in the subtracted images. Traditional source detection algorithms divide the image into a grid (see e.g. Spreeuw et al. 2018). In every grid cell, sigma-clipping is performed and the local statistics are used to determine a detection

threshold in each grid cell. A downside to this approach is that isolated faint sources might not be detected at the edge of a grid cell that is filled with bright sources. To counteract this, one can fold-in the sigma-clipping statistics of neighboring grid cells, to prevent the grid statistics in an individual cell from being biased, if there are relatively more sources in the grid cell. The LPF source finder used in this work takes this idea a step further by calculating grid statistics at every pixel coordinate. This effectively becomes a convolution of the grid statistics with the image, which is implemented in an effective manner using Fourier transform, accelerated on GPUs (Ruhe et al. 2022). The LPF source finder accurately models the noise structures that are present in subtracted images. After constructing a noise model the source finder looks for peaks in signal-to-noise ratio (SNR) above the prespecified threshold and saves the locations. Since this source finder simply identifies pixels that significantly stand out from the noise it is fast compared to methods that perform Gaussian fits to the source shape (e.g. PYSE; Carbone et al. 2018)).

We perform a blind search for sources in our subtracted images with a detection threshold of 5σ . This detection threshold is rather low for the enormous number of pixels that is searched. In Section 2.8 this value is updated to a more stringent detection threshold. We decide to perform the source finding for 5σ sources as this will allow us to characterize transient candidates in more detail later on. For example, a marginal detection above the stricter detection threshold in one snapshot could be accompanied by several subthreshold detections in adjacent snapshots at the same location. This information helps to decide whether or not a transient candidate is a potential imaging artefact (see Section 2.8).

2.5 Radial filtering

After finding sources in the full $3.67 \times 3.67^\circ$ square subtracted image, a radial filter is applied. As the image quality decreases with radial distance to the centre of the pointing Shimwell et al. (2022), we filter out all sources that lie more than 1.5° away from the beam/pointing centre. This still ensures significant overlap between pointings (see Section 2.3) but allows us to disregard noisier parts of the image.

2.6 Source catalogue matching

Ideally, the subtracted image only shows sources that have a significantly different flux density at the time-step considered, compared to the full image. In reality, the subtracted image also contains artefacts that arise mainly around bright and/or extended constant sources.

These subtraction artefacts are not the transient candidates that are targeted in this study and should be filtered out. Fig. 2 shows the importance of this particular step in the analysis. The left-hand panel shows an 8 s snapshot subtracted image and the right panel shows the LoTSS image of the corresponding part of the sky. The red circle indicates the location of a source identified by the source finder. In fact, all separate parts of the artefact in the subtracted image are identified as separate sources by the source finder. However, looking at the LoTSS image (right-hand panel Fig. 2) it is clear that these sources are a result of improper subtraction of the central, bright 2.1 Jy source. We therefore want to disregard all these sources as transient candidates, which is done via the catalogue matching scheme detailed below.

In order to define regions around known sources that are disregarded in the transient search, we first investigate up to which radius to filter around known sources in the subtracted image. We perform this investigation separately on the three snapshot time-scales used in this study. To this end, we gather known sources in

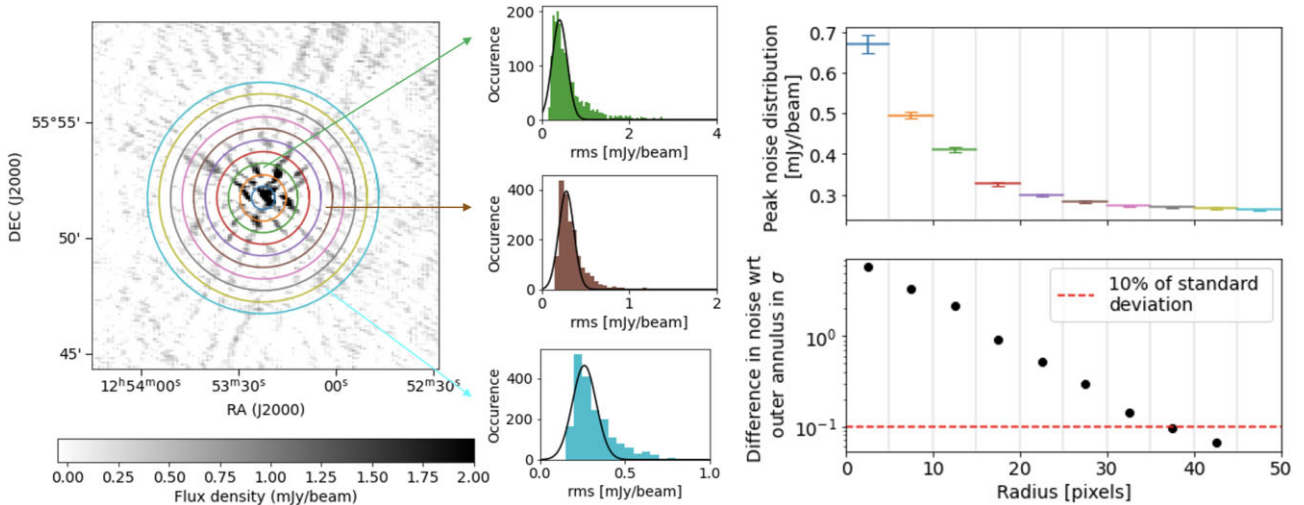


Figure 3. Example of determining a filter radius around artefacts of known sources with flux 100–500 mJy in 1 h subtract snapshots. The left-hand panel shows an example 1 h subtracted image with a significant structure due to improper subtraction of a 314 mJy source at that location. We define different annuli around this source and investigate the noise properties; examples of the pixel distributions in three of these annuli are shown in the plots in the middle of this figure. The top right panel shows the noise in each annulus for a sample of artefacts around 100–500 mJy in 1 h subtract snapshots. The bottom right panel shows the difference between the peak of the noise distribution in each respective annulus compared to the outer annulus. This difference is shown in standard deviations of the noise distribution of the outer annulus. The filter radius is defined at the point where the noise distribution starts to deviate by more than 10 per cent of the standard deviation of the outer noise distribution, indicated by the red dashed line.

the following integrated flux density intervals [10, 50], [50,100], [100,500], [500,1000], and >1000 mJy and define a ‘filter’ radius for each of these flux categories. An example of this investigation for 1 h subtracted snapshots targeting 100–500 mJy known sources is shown in Fig. 3. The left-hand panel shows an example of a 1 h subtracted image. The image shows significant structure, which is due to a known 314 mJy source at this location. Different coloured annuli around this structure define different trial filter radii. The histograms in the middle show examples of pixel distributions in these annuli, where the noise goes down for the annuli with larger radii. The top plot in the right-hand panel shows the mode of the pixel histogram distribution in each annulus. This represents the noise level in each annulus. We now assume that the noise distribution in the outer annulus is representative of the local noise distribution and compare the noise distribution in each annulus to the outer annulus. The bottom panel shows the difference between the peak of the noise distribution for each respective annulus compared to the outer annulus, defined in terms of the standard deviation of noise distribution in the outer annulus. Once the noise difference becomes more than 10 per cent of the standard deviation of the noise distribution in the outer annulus, we set a filtering limit. This is the point where the subtraction artefacts start to significantly impact the noise distribution. The red dashed line in the bottom right panel shows this cutoff. In this case setting a filter radius of 35 pixels, corresponding to 3.5 arcmin. Note that although the left-hand panel shows a single instance of an artefact around a 100–500 mJy known source in a single 1 h snapshot, the histograms and plots on the right side are created from a representative sample of such sources.

The process outlined in this section is repeated for each of the flux density intervals and each snapshot time-scale. This results in the filtering radii summarized in Table 2. Catalogue sources with flux density between 10 and 50 mJy did not give significant improper subtraction artefacts in the 2 min and 8 s subtracted snapshots, therefore no filtering is necessary there. The same holds 50–100 mJy source in the 8 s subtracted snapshots. In summary, transient

Table 2. Filter radii around catalogue sources determined for various flux density intervals and the three imaging time-scales considered in this study. The filter radii are in arcmin around the location of the source in the LoTSS catalogue.

Lower limit flux [mJy]	Upper limit flux [mJy]	Filter radius [arcmin]		
		1 h	2 min	8 s
1000	∞	4	4	3
500	1000	3.5	3	2.5
100	500	3.5	2	1.5
50	100	2	1.5	–
10	50	1	–	–

candidates that lie within the filter radius of catalogue sources (as defined in Table 2), are disregarded in the further steps of the transient search.

2.7 Associate candidates in time

All steps detailed above were performed on individual snapshots. The next step is to group together transient candidates that are found in multiple images of the same snapshot time-scale. For example, if the first 100 images of the 8 s snapshots contain a transient candidate at roughly the same location (within 5 arcmin), these candidates are grouped together as one individual candidate. This way, a source only has to be visually inspected in a few of these images before deciding whether it is an artefact or an actual interesting transient candidate. A radius of 5 arcmin is used because the different components of a single source can be quite (spatially) dispersed in a subtracted image, see e.g. Fig. 3.

2.8 Detection threshold

So far we have not done any filtering on the detection significance of transient candidates. The source finder (Section 2.4) identifies

Table 3. Detection thresholds for transient candidates to be considered for visual inspection.

Time scale	Number of pixels (N)	Detection threshold
1 h	1.2×10^9	6.02σ
2 min	3.3×10^{10}	6.54σ
8 s	5.3×10^{11}	6.94σ


all sources with an $\text{SNR} \geq 5$. This value is low and will yield false positive transient identifications because a large number of pixels is trialed. A detection threshold is calculated based on the probability that one pixel is encountered that exceeds the detection threshold for the total number of pixels per imaging time-scale. In other words, we calculate the probability $P(X \leq x) = 1 - \frac{1}{N}$ with N the total of number of pixels sampled per time-scale.

We calculate the value of x expressed in σ (i.e. the threshold) for this probability using the percent point function (or inverse cumulative distribution function) for the Gaussian distributions of the images per time-scale using the `scipy stats norm ppf` functionality (Virtanen et al. 2020). The results of this calculation are shown in Table 3. Note that for we perform this calculation for a random subset of all images, but repeating this procedure for a different subset gives similar results implying that the subset is representative of the full sample. Finally, the areas around known catalogue sources are not searched in for transients but those cuts are not taken into account here. The full inner region of the images ($\pi * 900^2 = 2544\ 690$ pixels per image) is assumed to be searched in this calculation. Accounting for the cut areas does not make much difference, for the 1 hour subtracted images up to about 20 per cent of the image is cut out, but this results in the detection threshold being lowered to 5.99σ (as opposed to 6.02σ , see Table 3).

We apply these thresholds to all light curves that come out of the previous step. A light curve is put forward for visual inspection if at least one source in the time series has an SNR above the threshold calculated here. This is done such that other subthreshold detections of the same transient candidate can be considered in visual inspection.

2.9 Visual inspection

After the filtering process, there are $\mathcal{O}(10)$ transient candidates left per field for all time-scales. There are fields, for example P173+55, that have one very bright source (3.7 Jy) that affects the sky model subtraction over a large fraction of the field and produces many spurious candidates. The remaining sources are visually inspected. There are two main categories of false positive transient candidates that still come through the pipeline. First of all, there are subtraction artefacts due to faint sources that generally do not give significant subtraction artefacts. Secondly, there are bright known sources that show extended subtraction artefacts, that are not filtered out by the filter radii presented in Table 2. Additional filtering, around fainter sources or extending the filter radii around bright sources, would solve both issues but comes at the cost of a further reduced sky area that is searched for transients. The final visual inspection step thus removes transient candidates that are associated with faint sources in the deep field, for which we have not applied any filtering radius, or transient candidates that are associated with artefacts of bright sources that extend beyond the filtering radius.

**Figure 4.** Simulated transient source of 250 mJy with a duration of 8 s, captured by one 8 s subtracted image snapshot. The left-hand panel shows the subtracted image, while the right-hand panel shows the corresponding sky area in the LoTSS survey. The coloured dots and legend indicate the flux densities of the sources in the LoTSS source catalogue.

3 SENSITIVITY

3.1 Simulations

To assess the sensitivity of our method we perform simulations of transient sources. To this end, transient sources are injected into a field, and the data is recalibrated. Here, the full direction-dependent calibration is repeated to investigate whether or not the calibration strategy might absorb faint short-duration transients. Afterwards the subtraction imaging and filtering methods described in Section 2 are applied. By injecting simulated transient sources with different integrated flux densities throughout the field, the sensitivity of the method as a function of distance from the beam centre can be determined. For these simulations, the P23Hetdex20 field is used, which is a typical field containing a variety of diffuse, moderately bright sources and normal levels of calibration artefacts.

A total of 8 s transients are injected with integrated flux densities of [70.0, 80.2, 96.5, 113.3, 133.0] mJy. One hour transients are injected with integrated flux densities of [0.2, 0.6, 2.0, 6.3, 20.0] mJy. Each flux density value is injected at 5 different radii across the whole field, because we expect a dependence of sensitivity as a function of radius, where the search is less sensitive away from the field centre, due to the primary beam. We inject simulated transients such that they do not overlap with bright LoTSS sources, as those would be filtered out via the steps described in Section 2.6.

Fig. 4 shows an example of a simulated transient source. The left-hand panel shows the injected transient in the subtracted image. On the right, the LoTSS image is shown for this part of the sky for comparison. The simulated transient is injected into a relatively empty part of the sky. Furthermore, the subtracted image shows that in this case the subtraction has worked quite well and the injected transient dominates everything else in the nearby region. This particular simulated transient has a duration of 8 s and an integrated flux density of 250 mJy.

We inject the transient sources in the data in such a way that they will be captured by exactly one snapshot, i.e. the transient is injected at a time where it will fully fall within one snapshot image and is not captured partly by two consecutive snapshot images. After injecting the transients in the field, we recalibrate the data against the LoTSS sky model that does not contain the transient. This is to check whether

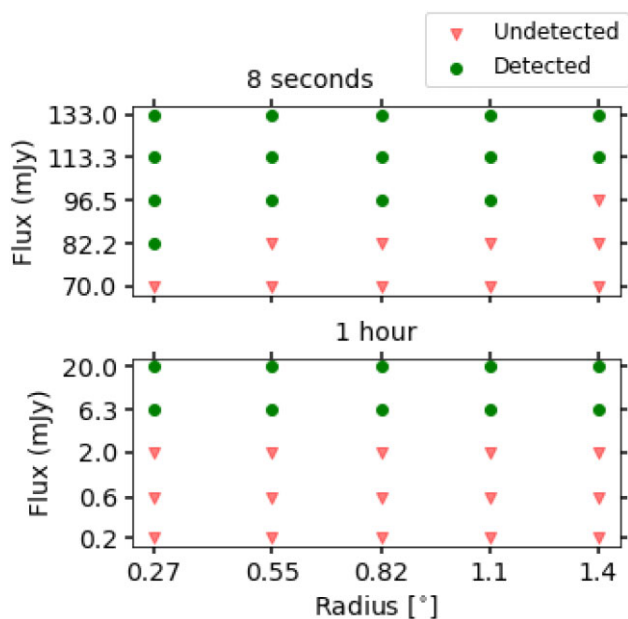


Figure 5. Results of the injecting simulated transients with various flux densities (y-axis) at different radii throughout the field (x-axis). The top panel shows the 8 s duration simulated transients, the bottom panel shows the 1 h duration simulated transients.




Figure 6. The green shaded area shows the flux density of potential transient sources as a function of time-scale that the transient search method presented in this work is sensitive to. The upper limits indicated with blue triangles are based on the filtering of known catalogue sources outlined in Section 2.6 and Table 2. The lower limits are determined using simulations, as outlined in Section 3.1. The black crosses are the lower limits of our search as determined by multiplying the rms noise in the subtracted images (Section 4.1) by the detection thresholds (Section 2.8).

or not a transient can be absorbed into the calibration solutions, and therefore might appear differently than initially expected. The recalibrated data are then processed via the exact same steps outlined in Fig. 1 and Section 2. Fig. 5 shows which of the injected simulated transients were recovered by our method. The plot shows the flux densities of the injected transients injected at different radii, where the green dots indicate that the particular transient was recovered at the end of the pipeline, while the red triangles indicate that the transient was not found by our pipeline. The sensitivity is fairly constant throughout the beam for this set of simulations for the 1 h transients. For the 8 s transients a relation between the flux of the

recovered simulation and radius is visible. We draw the conservative conclusion that the sensitivity of the 1 hour transient search is around 6 mJy and around 113 mJy for the 8 transient search. Extrapolating the sensitivity of 6.3 mJy at 1 h and 113 mJy at 8 s to 2 min, using that the detection sensitivity scales with $1/\sqrt{t}$, we find a sensitivity of roughly 30 mJy for the 2-min snapshots.

We want to note that these sensitivity estimates are lower limits because we assume that the transient is exactly covered by one snapshot. In reality, an 8 s transient will not exactly line up with the binning of our snapshots, and therefore we will not reach the aforementioned sensitivity of 113 mJy, as the source is split up over two snapshots. The worst case limits for detections is when the fluence of a transient is split equally between two bins, in which case the limits should be multiplied by a factor $\sqrt{2}$, which yields detection thresholds of [160, 42, 8.9] for the 8 s, 2 min, and 1 h snapshots, respectively.

3.2 Upper limits

The previous section determines the lower sensitivity limit to the transient search. In this section the upper limit is explored. When transient sources get extremely bright, they might get into the sky model and show up on the deep image, despite being active only a fraction of the time of the full observation. As explained in Section 2.6, the sources in the deep image are used to filter out subtraction artefacts in the subtracted images, but this implies that also these extremely bright transients will be filtered out. A simple estimate of the brightest source this method would find per time-step is given by multiplying the lowest flux density value included in our filtering scheme as outlined in Table 2 by the number of snapshots that are made for a particular time-scale. For example, for the 1 h snapshots filters are applied around deep field sources of 10 mJy. A transient source that is active for 1 h would therefore have to emit at $10 \times 8 = 80$ mJy to make it into the deep image and be disregarded in the transient search. Fig. 6 shows the upper and lower limits calculated via the methods described in this and the previous section. The blue triangles show the upper limits of flux density values of transient sources we would be sensitive to at various time-scales, based on the filters described in Table 2. The red triangles show the lower limits on the transient search based on the simulations (the 2 min snapshot value is inferred from the 8 s and 1 h subtracted images). The green-shaded region between the blue and red markers indicates the flux density values the method presented in this paper is sensitive to. Furthermore, Fig. 6 shows black stars that, per time-scale, indicate the rms noise level per image (see Section 4.1) times the detection threshold (defined in Section 2.8). This gives the theoretical lower limit of sources we could detect in the subtracted images per time-scale. Finally, Fig. 6 shows the $1/\sqrt{t}$ relation that is used to derive the lower limit to the search in the 2 min snapshots, based on the 8 s and 1 hour subtraction images.

Finally, we want to point out that in some cases the subtraction images could slightly underestimate the flux of a transient candidate. For example, a 40 mJy transient that is ‘on’ for one hour throughout the 8 h observation will show up as a $(40/8 =) 5$ mJy source in the deep image. This means that a 5 mJy source will be subtracted out at the location of the transient during the source subtraction. This will lead to a transient flux density that is slightly lower than the true flux density, and additionally, it will create negative subtraction artefacts at times when the transient is off. We do not consider this effect in detail in our search, as the main goal of the subtracted images is to identify transient sources, characterization will follow with additional imaging.




Figure 7. Rms noise in a sample of 500 subtracted images per time-scale. The rms distributions for the 8 s snapshot images, 2 min timeslices and 1 h images are shown from left to right, respectively. Only the innermost circle with radius 1.5° is used to construct these distributions.

4 RESULTS

We apply the methods outlined above to 58 pointings of LoTSS. These pointings correspond to the sky area covered by LoTSS-DR1 (Shimwell et al. 2019), but we note that the data is reprocessed using the LoTSS-DR2 approach (Shimwell et al. 2022).

4.1 Subtracted images quality

Fig. 7 shows the rms noise for the subtracted images for a subset of 500 images per time-scale. The rms distributions for the 8 s images, 2 min timeslices, and 1 h time-scale are shown from left to right, respectively. No sigma clipping is performed on these distributions, so they also contain source pixels. The rms noise is around $4.6 \text{ mJy beam}^{-1}$ for the 8 s snapshots, $1.3 \text{ mJy beam}^{-1}$ for the 2 min snapshots and $0.3 \text{ mJy beam}^{-1}$ for the 1 h snapshots. Multiplying these numbers with the detection thresholds applied at the various time-scales (see Table 3) gives a lower limit on the flux density of the faintest transient that we could detect at each time-scale. In Fig. 6 these numbers are indicated by the black stars and they are in good agreement with the detection thresholds determined with the simulations in Section 3.1. These lower limits seem to probe slightly deeper than the simulations, but that is mainly an effect of the quite crude steps in flux density in our simulations, see Fig. 5. Furthermore, these numbers are in good agreement with the expected scaling of the sensitivity with time as $1/\sqrt{t}$.

4.2 Method efficiency

Fig. 8 shows the number of sources left after each step in the analysis for a single field for different snapshot time-scales. The numbers above the final step show the percentage of sources left for visual inspection compared to the number of sources found by the source finder in the first step. The steps described in the methods section are able to filter out $\sim 0.003 - 0.06$ per cent of sources as potential transient candidates. This corresponds to in total $\mathcal{O}(10)$ candidates per field, which is feasible but tedious for visual inspection.

4.3 Transient candidates

After processing 58 fields that correspond to the fields presented in the first data release of LoTSS (Shimwell et al. 2019), 11 transient candidates are identified. The figures in Appendix C show

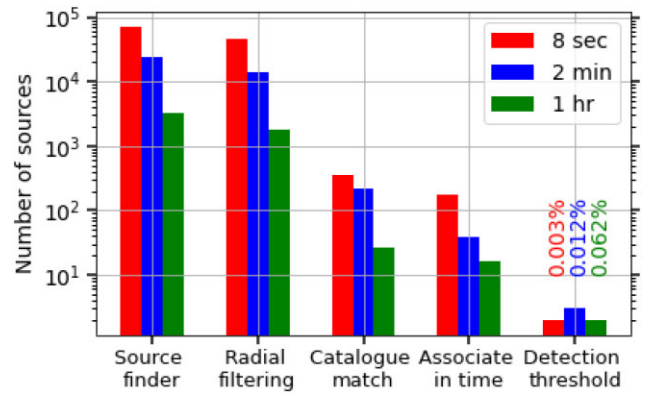


Figure 8. Number of sources left after each step in the analysis for the P164+55 field for different snapshot time-scales. The numbers above the final step show the percentage of sources left for visual inspection compared to the number of sources found by the source finder.

all transient candidate sources that are left after a first round of visual inspection. These sources did not immediately fall within one of the two categories of sources that are normally vetted by visual inspection (subtraction artefacts of bright sources, or sources that are associated with a faint deep fields source).

In the following sections, we decide to consider the transient candidates shown in Fig. C1, C2, and C3. This is because these are the brightest candidates and the candidates in the subtracted images have a shape similar to the dirty beam, similar to the simulations (See Fig. 4). As the candidates in Figs C1 and C2 are two instances of the same candidate detected in different time-scale, these two detections are discussed jointly in more detail in Section 5.1, we will refer to this source as transient candidate C1\C2. The candidate in Fig. C3 is discussed in more detail in Section 5.2, we will refer to this source as transient candidate C3.

We decide to not follow-up on the candidates presented in Figs C4, C5, and C6 because the source in the subtracted image is not dirty-beam shaped, unlike expected (see Section 3.1). Finally, the candidates shown in Figs C7–C11 are disregarded because upon closer inspection they are associated with faint sources in the deep field.

5 DISCUSSION

In this Section, we discuss the implications of our results. First, in Sections 5.1 and 5.2 we discuss in more detail the origin of the two transient candidates we identified. To study the candidates in more detail the sources were imaged with primary beam correction and cleaning using the direction-independent data products. In the next sections, we discuss the upper limits we can place on the transient surface density. Finally, we discuss how this work could be extended to look for variable sources.

5.1 Transient candidate C1\C2

The first transient candidate found in this study is presented in Figs C1 and C2. There is a bright dirty-beam shaped source in the subtracted image that cannot be associated with any source in the deep image. The source is detected in three consecutive 8 s subtracted images with SNR increasing from 6.1 to 7.7 and 8.8. Additionally, the source is identified in the 2 min snapshot that encompasses this 24 s interval with an SNR of 7.3.




Figure 9. RMS noise in the 8 s subtracted images of field where the transient candidate C1\C2 was found. The coloured vertical lines show the rms noise in the particular subtracted images where the transient candidate was identified.

5.1.1 Validity checks

To make sure that this transient source is real and not an imaging artefact, we perform some additional checks. In contrast to the artefacts discussed in Appendix B and transient candidate C3 discussed in Section 5.2, the location of the source does not change in the subtracted images, as the brightest pixel is identified to be the exact same pixel in each of the three images. This is what we expect to be the case for an astrophysical transient. The distance from the transient location to the centre of the beam 0.82° , most artefacts as discussed in Appendix B were found close to the edge of the search radius (1.5°). Furthermore, recreating the subtracted images with additional padding and slightly different imaging settings does not make the transient candidate disappear. This assures us that transient candidate C1\C2 is not one of the PSF artefacts discussed in Appendix B.

Additionally, Fig. 9 shows the noise distribution of the particular subtracted images where the transient is identified compared to the rms noise in all 8 s subtracted images of this pointing. Fig. 9 shows that these subtracted images do not show increased noise compared to the full set of images. Comparing Fig. 9 to Fig. 7 reveals that the noise in this particular field is on the high side compared to 8 s subtracted images of other fields. This is due to poorer ionospheric conditions during this observations, which we will shortly discuss in the next section.

5.1.2 Reimaging

Now that a viable transient candidate is identified, additional imaging is performed to fully characterize the source. Using the publicly available products presented in Shimwell et al. (2022), we reimage the direction-independent calibrated visibilities with particular interest around the time when transient candidate C1\C2 is found. In this process, the size of a pixel is decreased in order to get a more accurate position measurement, some of the shortest baselines are cut-out to get rid of extended emission and deep cleaning is performed. The most important imaging parameters are presented in Table 4. Now that cleaning and a primary beam correction are applied, a more accurate estimate of the flux of the transient can be made.

Fig. 10 shows the peak flux density as a function of time based on the reimaged 8-s snapshots for transient candidate C1\C2. The spectral index at the peak of the flare is negative $\alpha \approx -1.0$ (for $S \propto \nu^\alpha$).

Table 4. The WSCLEAN settings used to reimage transient candidate C1\C2 presented in Section 5.1.

Setting	Value
Number of clean iterations	150 000
Auto-mask	2.5
Auto-threshold	0.5
Minuv-l	80
Channels-out	3
Size of image (pixels)	4400
Size of one pixel (arcsec)	3
Briggs weighting	-0.5
Interval	2612–2621
Intervals out	9
Padding	1.4
Apply primary beam	True




Figure 10. Peak flux density of transient candidate C1\C2 as a function of time. Different markers indicate different levels of SNR of the detection.

Separate images were made for the Stokes I and V components of the signal, and no circularly polarized flux was detected. The data quality is insufficient to perform the frequency slicing necessary for QU fitting. It would be better to perform an in-depth study of the polarization properties using the direction-dependent self-cal solutions. Unfortunately, this pointing has particularly bad ionospheric conditions and we were unable to improve the calibration compared to the LoTSS-DR2 images. Improved calibration would not only be useful in studying the polarization properties but also in understanding the spectrum of the flare, as this procedure would increase the SNR. To this end additional observations under more favourable ionospheric conditions are necessary.

5.1.3 Future work

We leave it to a future paper (de Ruiter et al. in preparation) to determine the true nature of this transient source. To this end, we will analyse additional observations of this field. The 8 s integration time of the data and relatively low SNR (preventing us from making many frequency slices), hampers a dispersion analysis for the flare. If an intrinsically short duration signal is dispersed to $6 \times 8 = 48$ s (taking only the detections with $\text{SNR} > 5$), this points to a DM of roughly

$$\text{DM} = \frac{1.2 \times 10^5 \times \tau_{\text{DM}} \times \nu^3}{B} = \frac{1.2 \times 10^5 \times 48 \times 0.144^3}{48} \quad (1)$$

$$= 358 \text{pc cm}^{-3}$$

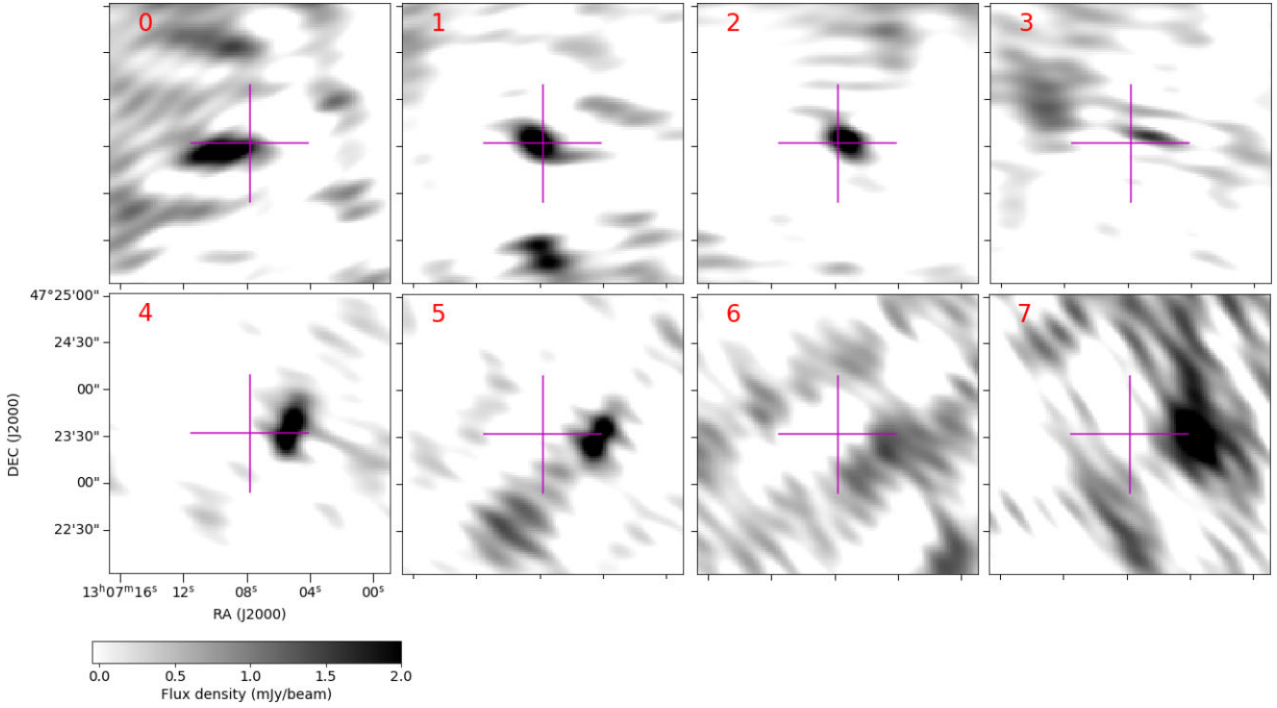


Figure 11. Insets at the location of transient candidate C3 in the snapshot images, which are created with primary beam correction and cleaning using the direction dependently calibrated data. The crosshair is fixed at the same location in all panels. The panels show the 1 h snapshots throughout the 8 h observation. The transient candidate moves throughout the observation.

with τ_{DM} the dispersion measure smearing in seconds, ν the central frequency in GHz, and B the bandwidth in MHz. A DM of 358 pc cm^{-3} would make this an extragalactic source (Cook et al. 2023).

Another possibility is that the radio pulse produced by this source intrinsically has a duration of half a minute to a minute. Recently, galactic sources of this type have been discovered by Hurley-Walker et al. (2022, 2023). These sources also show similarities with extremely bright minute duration bursts discovered by Hyman et al. (2005). For all aforementioned sources, multiple bursts have been identified, thus warranting follow-up and archival searches for additional flares from transient C1\C2. The progenitor scenario for the aforementioned transient sources is still under debate. Isolated neutron star and white dwarf rotating dipole scenarios have been explored to explain these long-period radio pulsars, but those do not seem to explain all observed features (Rea et al. 2022, 2024). Another progenitor scenario could be a compact object in a binary with a main-sequence star, as in Pelisoli et al. (2023). There is a star within 5 arcsec of the position of transient C1\C2, but whether this radio signal is associated with that star remains to be seen. A more in-depth follow-up study on this transient source will be presented in de Ruiter et al. (in preparation).

5.2 Transient candidate C3

The same reimaging process that was outlined in Section 5.1 is repeated for the transient candidate C3 (presented in Fig. C3).

The new images reveal that the location of transient candidate C3 changes throughout the observation. Fig. 11 shows insets of the transient candidate in the new images. The crosshair is fixed at the same location in all panels. The panels show the 1 h snapshots throughout the 8 h observation. From these images, it is clear that the




Figure 12. The right ascension and declination of transient candidate C3 in degrees minus its average location. The numbers indicate the snapshot number. This plot only contains the snapshots where the source was detected at $>5\sigma$ above the rms noise.

transient candidate moves throughout the observation. This effect is quantified in Fig. 12, which shows the right ascension and declination of the source minus its average location. The numbers indicate the snapshot number. This plot only contains the snapshots where the source was detected with $>5\sigma$. The average position of the source is $(196.778^\circ, 47.392^\circ)$. Studying the deep LoTSS image (right-hand panel in Fig. C3) carefully, there is an arc-shaped artefact visible in the 8 h average image. This arc roughly corresponds to the path of the source, see Fig. 12.

Including the marginal detection of the last hour snapshot, this implies that our candidate sources moves about 1 arcmin on the sky over 8 h. This movement is clearly different from the slight shift in position that other sources experience from snapshot to snapshot.

Especially in right ascension the movement of this source is 8σ away from the jitter that other sources in the field experience. Furthermore, most sources seem to move in a random direction from image to image, but transient candidate C3 seems to follow a trajectory, it does not move back toward its original position. If we assume this is a real astrophysical source, this movement implies a high proper motion and/or an extremely nearby object. The displacement places the source within roughly one parsec, where it would travel with the speed of light. This suggests that the source, if astrophysical, is most likely within our Solar system. At an average position of $(196.778^\circ, 47.392^\circ)$ this source lies approximately 70° above the Galactic Plane and about 50° above the ecliptic.

Additional checks and tests that were performed do not provide any further insight. The emission is broad-band, but hints to be brighter towards lower frequency. This seems to exclude most reflected signals and satellites. There is no significant Stokes Q, U, or V detection. Due to the position shift it is difficult to match the source to other catalogues. No minor planets were found at the location of our transient using <https://minorplanetcenter.net/cgi-bin/checkneo.cgi> around the observation time of this field (2014-07-14 14:00:00 UTC). A highly speculative origin for this type of emission is an asteroid or comet reflecting emission from the Sun. Radio reflection tomography has been proposed as a technique to study the interior of asteroids and comets (see eg. Safaeinili et al. 2002) and the Sun has been used as a radio source to use a similar reflection technique to probe the thickness of glaciers on Earth (Peters et al. 2021). During the observation of this field, there was no extreme solar activity.³

A final note we would like to make is that the source trajectory (Fig. 12) is quite similar to the trajectories that we observe as a result of artefacts close to the edge of the image, such as for example shown in Fig. B2. However, all the ‘fake’ transient sources that were found to show such a trajectory were detected quite close to the edge of the image, and the sources disappeared altogether when reimaging with additional padding. Following this procedure did not make the transient source discussed here disappear.

The most likely explanation is that this source is an imaging artefact. The arc-shaped artefact present in the LoTSS deep image is similar to artefacts introduced by facet calibration due to aliasing or flagging. We note that neither increased padding nor turning off flagging removed this source from the images. An extensive investigation into the origin of this artefact is beyond the scope of this work.

In conclusion, we find a source in the data that passes all tests we do in our analysis, and proper reimaging shows that there is indeed a transient point source in the data. However, we are currently unaware of an astrophysical process that would explain a source that shows broad-band transient radio emission and shows a high proper motion, which implies it is located in the Solar system. The most likely explanation for this source is that is an imaging artefact. If this is an astrophysical source we expect to find more similar sources in our follow-up study (~ 13), where we plan to repeat this study for a larger sky area.

5.3 Transient rates

Following Rowlinson et al. (2016) we calculate the transient surface density limit using Poisson statistics via

$$P(X = n) = \frac{(\rho \times \Omega_{\text{total}})^n e^{-\rho \times \Omega_{\text{total}}}}{n!}, \quad (2)$$

³https://solen.info/solar/old_reports/2014/july/20140715.html

Table 5. The transient surface density for various time-scales.

Time-scale	N	Ω_{total} [deg ²]	Ω_{corr} [deg ²]	ρ [deg ⁻²]
1 h	8	3280	3210	$<4.0 \times 10^{-4}$
2 min	225	92 245	92 218	5.7×10^{-7}
8 s	3600	1475 920	1 475 908	3.6×10^{-8}

where Ω_{total} is the total area surveyed at a certain time-scale, ρ is the surface density limit, and P is the confidence interval. In case no transient candidates are detected this equation reduces to $P(X = 0) = e^{-\rho \Omega_{\text{total}}}$ which allows one to define an upper limit on the transient surface density. Following Bell et al. (2014), we use $P = 0.05$ to give a 95 per cent confidence limit. The total sky area surveyed in this work is summarized in Table 5. Ω_{total} shows the naive calculation of the surveyed sky area, ie. the number of fields, 58, times the number of images per time-scale, N . However, in this work we perform cuts around bright sources that reduce the effective sky area we search. Ω_{corr} shows the total sky area surveyed while taking into account the area we cut out. As we do not find any transients in the 1 h snapshots, an upper limit is calculated for this time-scale. One transient candidate (C1\C2) is found in 2 min and 8 s snapshots, and the corresponding transient surface densities are shown in the final column of Table 5.

Fig. 13 shows our new result (red cross) compared to other results in the literature. The figure consists of three panels probing three different transient duration times. We compare our results to the most constraining studies below 340 MHz and find that our results are probing the lowest sensitivities to date at all transient time-scales. The structure of this plot was taken from Murphy et al. (2017), but a more up-to-date sample of the most constraining studies was compiled using an overview⁴ from Mooley et al. (2016). Markers with a downward pointing arrow represent upper limits. We choose to show the results of our study as datapoints at a fixed sensitivity, instead of as a curve showing the transient surface density as a function of sensitivity by including a larger portion of images with higher rms noise values (see e.g. Rowlinson et al. 2022). We opt not to do this as we show that the sensitivity is quite uniform across the image (see Fig. 5). From Fig. 13, it is clear that the transient source presented in Section 5.1 could not have been detected by previous studies, as those did not probe sufficient sky area and/or lacked sensitivity compared to our search.

5.4 Transient rates at higher frequencies

Many radio telescopes have typical observing frequencies of around 1.4 GHz. Therefore, in this section, we compare the transient surface density values from this work (at 144 MHz) to some of the most constraining studies at 1.28–1.4 GHz. From the papers below, we distill the information necessary to calculate the transient surface density following equation (2).

Chastain et al. (2023) use MeerKAT to perform a commensal search for transients within the ThunderKAT program (Fender et al. 2017). Images are created with integration times of 4 h, 15 min, and 8 s, with median noise levels of 10, 30, and 176 μJy , respectively (table 2 in Chastain et al. 2023). Furthermore, detection thresholds of 5.3σ , 5.7σ , and 6.4σ are used. These observations are taken at 1.28 GHz and each image has a size of at least 2.8 by 2.8° . In total 28: 4-h images, 406: 15-min images, and 43:964 8-s images were

⁴<http://www.tauceti.caltech.edu/kunal/radio-transient-surveys/index.html>




Figure 13. Limits on the transient surface density from this study compared to previously published results below 340 MHz. The three panels show studies probing different transient time-scales. The result presented in this paper is shown as a red cross. Other results are from Lazio et al. (2010), Obenberger et al. (2015), Polisensky et al. (2016), Rowlinson et al. (2016), Stewart et al. (2016), Feng et al. (2017), Anderson et al. (2019), Hajela et al. (2019), Tingay & Hancock (2019), Varghese et al. (2019), Kuiack et al. (2021), and Sokolowski et al. (2021).

Table 6. Comparison of the transient surface density as found in this study, compared to transient searches at higher frequency. The sensitivity at the frequency of the original study is extrapolated to 144 MHz assuming a spectral index $\alpha = -1.2$, where $S_\nu \propto \nu^\alpha$.

S (mJy)	$S_{144\text{MHz}}$ (mJy)	Transient surface density (deg^{-2})	Time-scale	Frequency (GHz)	Ref.
0.053	0.73	$<1.4 \times 10^{-2}$	4 h	1.28	[1]
3.9	59.8	$<3.2 \times 10^{-2}$	1 h	1.4	[2]
–	6.3	$<4.0 \times 10^{-4}$	1 h	0.144	
–	6.3	$<1.6 \times 10^{-3}$	4 h	0.144	[**]
0.17	2.34	$<9.4 \times 10^{-4}$	15 min	1.28	[1]
1.2	18.39	3.2×10^{-4}	15 min	1.4	[3]
19.2	294.2	$<1.1 \times 10^{-3}$	2 min	1.4	[2]
–	30	5.7×10^{-7}	2 min	0.144	
–	30	4.3×10^{-6}	15 min	0.144	[**]
1.13	15.55	$<8.7 \times 10^{-7}$	8 s	1.28	[1]
56.4	864	$<6.7 \times 10^{-5}$	8 s	1.4	[2]
–	113	3.6×10^{-8}	8 s	0.144	

Note. References [1]:Chastain et al. (2023), [2]:Fijma et al. (2023), [3]:Wang et al. (2023), [**]: results from this study that have been extrapolated to longer time-scales to allow a direct comparison with other studies.

created. The transient surface density resulting from this study is calculated using equation (2) and shown in Table 6. We also note the work by Anderson et al. (2022), who find a flare star in a commensal search for transients in MeerKAT observations, but this search does not go as deep (in the number of epochs) as the work by Chastain et al. (2023).

Additionally, we compare our results to the work by Fijma et al. (2023), as their imaging strategy is similar to ours. They perform a transient search in MeerKAT data around the NGC 5068 field, creating snapshot images using the subtraction imaging methods as explained in Section 2.2. Furthermore, the time-scales used for

imaging are identical to what was used in this study. The transient surface density limits from Fijma et al. (2023) are summarized in Table 6.

Finally, Wang et al. (2023) perform a transient search on 15-min time-scales using the Variable and Slow Transient survey using the Australian SKA Pathfinder (ASKAP; Hotan et al. (2021)). They find 3 single flare transients in 754 images at 1.4 GHz. Again, this study is of particular interest as a similar source subtraction method is used to create images. The size of an individual image is $2.1 \times 2.1^\circ$ and the median rms noise in an individual image is around $0.2 \text{ mJy beam}^{-1}$. A detection threshold of 6σ was used. The transient surface density resulting from this study is calculated using equation (2) and shown in Table 6. We also note the work by Dobie et al. (2023) but instead, choose to include just the work by Wang et al. (2023), as the latter provides a more constraining result.

Table 6 shows an overview of the transient surface density values as presented in the studies detailed above. The first column shows the detection sensitivity as originally mentioned in the study, whereas the second column shows how that sensitivity would translate to 144 MHz assuming a spectral index $\alpha \approx -1.2$, where $S_\nu \propto \nu^\alpha$. We chose this spectral index as most of the transient phenomena we are interested in here have steep spectral indices. Flare stars show either narrow-band (Callingham et al. 2021) frequency behaviour or steep negative spectral indices (Spangler & Moffett 1976); pulsars have spectral indices of $\alpha \approx -1.6$ (Lorimer et al. 1995), and giant pulses might show even steeper spectral indices of $\alpha \approx -2.7$ (Popov et al. 2008); the spectral indices of fast radio bursts are estimated to be around $\alpha \approx -1.5$ (Macquart et al. 2019); and the spectral index of the Galactic Center Transient (Zhao et al. 1992) is steep ($\alpha \approx -1.2$), as well as the spectral indices of the long-period radio sources that have been found recently (Hurley-Walker et al. 2022, 2023). There are radio transient sources, such as X-ray binaries that do not follow this trend. Generally, in X-ray binary radio flares the

spectral turnover can be observed, which shows as an evolving spectral index. For example, in Monageng et al. (2021) α evolves from -0.89 to -0.25 , and in Chauhan et al. (2021) and Chandra & Kanekar (2017) the spectral index evolves from negative to positive. Note that although our choice of spectral index is appropriate for most transient phenomena we are interested in here, it might not be appropriate for all sources. The other columns in Table 6 show the transient surface density, time-scale, observing frequency and reference to the study respectively. Each block of Table 6 compares studies that have been performed at a similar time-scale. The final rows in each block show the results from our work (as calculated in Table 5). For the hour and minute time-scale we additionally extrapolate our results to the exact time-scales probed in Chastain et al. (2023) and Wang et al. (2023) to make a fair comparison. This is done by multiplying the transient surface density with a factor of 4 and 7.5 for hours and minutes time-scales, respectively.

From Table 6, it is evident that even when assuming a flat spectral index the detection sensitivity in our work is of the same order of magnitude as in Fijma et al. (2023). Extrapolating the detection sensitivity in Fijma et al. (2023) to 144 MHz shows that our study probes over an order of magnitude deeper at hours and minutes time-scales, and roughly three times deeper at second time-scales. Additionally, our work provides much deeper limits on the transient surface density at all time-scales.

The last row in the minutes time-scale block of Table 6 shows our 2 min time-scale results extrapolated to 15 min. Our results are of similar sensitivity compared to Wang et al. (2023), but the transient surface density is much more constraining. Finally, our study is less sensitive than the extrapolated sensitivity presented in Chastain et al. (2023) for the hours, minutes, and seconds time-scales, but in contrast to Chastain et al. (2023) we do find a transient source in the 2 min and 8 s snapshots, most likely due to our much larger observing time. Therefore, our transient surface density rates are more constraining.

A curiosity we would like to point out regarding Table 6 is that Wang et al. (2023) do find 3 interesting flaring transient sources at a sensitivity in between ours and Chastain et al. (2023), which is surprising. The discrepancy between Wang et al. (2023) and Chastain et al. (2023) could be due to the larger sky area probed by Wang et al. (2023), but there is also almost an order of magnitude sensitivity difference, which is expected to play a role. Additionally, the fact that Wang et al. (2023) find three sources at minutes time-scales, compared to our one source implies that we might probe different transient source populations at 144 MHz and 1.4 GHz. This could for example be due to a synchrotron self-absorption break in the spectrum that causes some of the transients found at 1.4 GHz to be fainter at 144 MHz, causing us to miss them in our transient searches. Finally, it would be good in the future thoroughly compare methods with Wang et al. (2023) as there might be some intrinsic differences to our search, leading to different results.

5.5 Stewart et al. transient

To date, only a handful of transient surveys at low radio frequencies (<1 GHz) have detected transient sources. Examples include Hyman et al. (2009), Bannister et al. (2011), and Jaeger et al. (2012) at time-scales of days to months, which have low detection probability in our survey because the largest snapshot time we use in 1 h. More relevant is the bright (15–25 Jy) transient identified at 60 MHz by Stewart et al. (2016), with a duration of a few minutes. In the next section, we calculate the expected number of observed Stewart-like transients in our survey, as a function of spectral index. Using a conservative flux of 15 Jy and survey frequency of 60 MHz for the

Table 7. The expected number of transients on the minutes time-scale, as a function of spectral index, scaled from the Stewart et al. (2016) transient detection.

Spectral index	Number predicted	Null detection probability
−4	16	8.8×10^{-8}
−5	4.4	1.3×10^{-2}
−6	1.17	0.31
−7	0.32	0.73
−8	0.085	0.92

Stewart et al. (2016) survey, we can calculate the flux we expect to observe at our survey frequency (144 MHz), assuming some spectral index α , where $S \propto \nu^\alpha$.

The sensitivity of this survey can be scaled to Stewart et al. (2016) survey using $N \propto S^{-3/2} \cdot \Omega$ as expected for an isotropic homogeneous distribution of sources throughout a flat space, where S is the sensitivity of the survey and Ω is the total surveyed area. The expected number of Stewart-like transients in our survey is then equal to the ratio of the surveyed area in our survey and the Stewart survey times the detection sensitivity at minutes time-scale in our survey divided by the expected flux of the Stewart-like transient at the observing frequency of our survey, to the power $-3/2$:

$$N = \frac{\Omega_{\text{LoTSS}}}{\Omega_{\text{Stewart}}} \left(\frac{S_{\text{survey, minutes Jy}}}{15 \text{ Jy} \times \left(\frac{144 \text{ MHz}}{60 \text{ MHz}} \right)^\alpha} \right)^{-3/2}. \quad (3)$$

For the detection sensitivity of our survey at minutes time-scale we use a value of ~ 30 mJy, by extrapolating our simulation (see Section 3.1). Note that we use a flux density of 15 Jy for the Stewart et al. transient, in contrast to Rowlinson et al. (2016) where the detection limit of the survey is used. Table 7 gives the expected number of transients on the minutes time-scale, as a function of spectral index, scaled from the Stewart et al. (2016) transient detection.

Table 7 shows that we can rule out spectral indices ≥ -5 at 95 per cent confidence for the Stewart et al. (2016) transient, based on our extrapolated survey sensitivity estimate. Combining this result with Rowlinson et al. (2016), we conclude that either the transient rate derived in Stewart et al. (2016) is too high or that this transient event had an extreme spectral index.

5.6 Variability

In this work, we do not study variable sources, but the subtracted images presented in this work are suitable to study variable sources. A variable source would show up in the subtracted images just like the example in Fig. 2. Either a positive artefact would be present if the source is brighter than the sky model flux at the subtracted image snapshot time, or a negative artefact if the source is dimmer than the sky model flux at the specific snapshot time. In our analysis, we disregard these sources as they are associated with a LoTSS catalogue source. However, one could imagine setting a threshold where very bright variable source artefacts would be kept for further analysis. Highly variable sources can be used to study interstellar scintillation, see e.g. Anderson et al. (2022). Therefore, an interesting future extension to the work presented here would be to include (highly) variable sources in our analysis. Additionally, our work excludes transient sources associated with bright radio sources or galaxies. A transient that lies in close proximity to radio sources with a flux density that falls within the filter limits described in

Table 2 will be excluded from our search. This biases our search against transients occurring in galaxies with high radio flux.

6 CONCLUSION AND OUTLOOK

In this paper, we have presented the results of a search for transient sources at 144 MHz using the LOFAR Two-metre Sky Survey (Shimwell et al. 2022). The search covers mostly extragalactic sky areas and covers 410 deg² of sky. This search was performed on snapshot images with integration time-scales of 8 s, 2 min, and 1 h, splitting the 8 h pointings in roughly 3600, 225, and 8 snapshots. In order to create these snapshot images we use a new approach where we search for transients in the subtracted images. These are created by subtracting the full (8-h) sky model from the visibilities of the snapshot and imaging those visibilities without cleaning or primary beam correction. This process greatly reduces the otherwise computationally very expensive imaging step. Afterwards, source finding, filtering, and source catalogue matching steps are applied to find the transient sources. One transient candidate is identified, but follow-up is necessary to determine its true nature. This work identifies the lowest transient surface density at time-scales of seconds to hours at the highest sensitivity to date. In the future, this method will be applied to the second data release of LoTSS. This will increase the number of processed pointings from 58 in this study to 841.

Another approach that might be explored in the future is the use of a different source finding technique, more suitable for subtracted images. The source finder used in this work already accounts for complicated noise structures throughout the images, but subtracted images introduce other difficulties. Since the subtracted images are not cleaned, each source appears as a blob consisting of multiple components. Using the LPF sourcefinder all these components are identified as individual sources. A source finder that would automatically identify these components to be part of the same uncleaned source would simplify the filtering steps afterwards and speed up the process as there would be fewer sources to consider. Secondly, if the techniques presented in this paper would be applied to future studies where many transients are identified, an automated cross-matching to other catalogues might be useful in order to determine the origin of the transient emission. Finally, future data sets or surveys will potentially need a more stringent catalogue match or filtering scheme to reduce the number of transient candidates left for visual inspection, but we find the current filter scheme gives a good trade-off between not disregarding transient candidates too soon and time needed for visual inspection.

ACKNOWLEDGEMENTS

This research made use of ASTROPY (Astropy Collaboration 2013, 2018) for FITS file handling and coordinate matching and MATPLOTLIB (Hunter 2007) was used to create plots.

LOFAR is the Low Frequency Array designed and constructed by ASTRON. It has observing, data processing, and data storage facilities in several countries, which are owned by various parties (each with their own funding sources), and which are collectively operated by the ILT foundation under a joint scientific policy. The ILT resources have benefited from the following recent major funding sources: CNRS-INSU, Observatoire de Paris and Université d'Orléans, France; BMBF, MIWF-NRW, MPG, Germany; Science Foundation Ireland (SFI), Department of Business, Enterprise and Innovation (DBEI), Ireland; NWO, the Netherlands; The Science and Technology Facilities Council, UK; Ministry of Science and Higher

Education, Poland; The Istituto Nazionale di Astrofisica (INAF), Italy.

This research made use of the Dutch national e-infrastructure with support of the SURF Cooperative (e-infra 180169) and the LOFAR e-infra group. The Jülich LOFAR Long Term Archive and the German LOFAR network are both coordinated and operated by the Jülich Supercomputing Centre (JSC), and computing resources on the supercomputer JUWELS at JSC were provided by the Gauss Centre for Supercomputing e.V. (grant CHTB00) through the John von Neumann Institute for Computing (NIC).

This research made use of the University of Hertfordshire high-performance computing facility and the LOFAR-UK computing facility located at the University of Hertfordshire and supported by STFC [ST/P000096/1], and of the Italian LOFAR IT computing infrastructure supported and operated by INAF, and by the Physics Department of Turin university (under an agreement with Consorzio Interuniversitario per la Fisica Spaziale) at the C3S Supercomputing Centre, Italy.

This research was funded through project CORTEX (NWA.1160.18.316), in research programme NWA-ORC, financed by the Dutch Research Council (NWO).

DATA AVAILABILITY

The data and data products presented in this paper are available in a reproduction package via Zenodo, at <https://dx.doi.org/10.5281/zenodo.10118363>.

REFERENCES

- Anderson M. M. et al., 2019, *ApJ*, 886, 123
 Anderson G. et al., 2023, *MNRAS*, 523, 4992
 Andrianjafy J. et al., 2023, *MNRAS*, 518, 3462
 Arcus W., Macquart J., Sammons M., James C., Ekers R., 2021, *MNRAS*, 501, 5319
 Astropy Collaboration, 2013, *A&A*, 558, A33
 Astropy Collaboration, 2018, *AJ*, 156, 123
 Bannister K., Murphy T., Gaensler B. M., Hunstead R., Chatterjee S., 2011, *MNRAS*, 412, 634
 Bell M. E. et al., 2014, *MNRAS*, 438, 352
 Caleb M. et al., 2022, *Nat. Astron.*, 6, 828
 Callingham J. et al., 2021, *A&A*, 648, A13
 Carbone D. et al., 2018, *Astron. Comput.*, 23, 92
 Chandra P., Kanekar N., 2017, *ApJ*, 846, 111
 Chastain S., van der Horst A., Rowlinson A., Rhodes L., Andersson A., Diretse R., Fender R., Woudt P., 2023, *MNRAS*, 526, 1888
 Chauhan J. et al., 2021, *Publ. Astron. Soc. Aust.*, 38, e054
 Chawla P. et al., 2022, *ApJ*, 927, 35
 Cook A. M. et al., 2023, *ApJ*, 946, 58
 Crosley M., Osten R., 2018, *ApJ*, 862, 113
 Dobie D. et al., 2023, *MNRAS*, 519, 4684
 Driessen L. N. et al., 2024, *MNRAS*, 527, 3659
 Feeney-Johansson A. et al., 2021, *A&A*, 653, A101
 Fender R. et al., 2017, preprint ([arXiv:1711.04132](https://arxiv.org/abs/1711.04132))
 Feng L. et al., 2017, *AJ*, 153, 98
 Fijma S. et al., 2023, *MNRAS*, 528, 6985
 Hajela A., Mooley K., Intema H., Frail D., 2019, *MNRAS*, 490, 4898
 Hotan A. et al., 2021, *Publ. Astron. Soc. Aust.*, 38, e009
 Hunter J. D., 2007, *Comput. Sci. Eng.*, 9, 90
 Hurley-Walker N. et al., 2022, *Nature*, 601, 526
 Hurley-Walker N. et al., 2023, *Nature*, 619, 487
 Hyman S. D., Lazio T. J. W., Kassim N. E., Bartleson A. L., 2002, *AJ*, 123, 1497
 Hyman S. D., Lazio T. J. W., Kassim N. E., Ray P. S., Markwardt C. B., Yusef-Zadeh F., 2005, *Nature*, 434, 50

Hyman S. D., Lazio T. J. W., Roy S., Ray P. S., Kassim N. E., Neureuther J. L., 2006, *ApJ*, 639, 348

Hyman S. D., Roy S., Pal S., Lazio T. J. W., Ray P. S., Kassim N. E., Bhatnagar S., 2007, *ApJ*, 660, L121

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., Hyman S., Kassim N., Lazio T., 2012, *AJ*, 143, 96

Kuiack M., Wijers R. A., Shulevski A., Rowlinson A., Huizinga F., Molenaar G., Prasad P., 2021, *MNRAS*, 505, 2966

Lazio T. J. W. et al., 2010, *AJ*, 140, 1995

Lefevre E., Klein K.-L., Lestrade J.-F., 1994, *A&A*, 283, 483

Lorimer D., Yates J., Lyne A., Gould D., 1995, *MNRAS*, 273, 411

Macquart J.-P., Shannon R., Bannister K., James C., Ekers R., Bunton J., 2019, *ApJ*, 872, L19

Monageng I. M., Motta S. E., Fender R., Yu W., A Woudt P., Tremou E., Miller-Jones J. C., van der Horst A. J., 2021, *MNRAS*, 501, 5776

Mooley K. et al., 2016, *ApJ*, 818, 105

Murphy T. et al., 2017, *MNRAS*, 466, 1944

Obenberger K. S. et al., 2014, *ApJ*, 785, 27

Obenberger K. S. et al., 2015, *J. Astron. Inst.*, 4, 1550004

Offringa A. et al., 2014, *MNRAS*, 444, 606

Pelisolì I. et al., 2023, *Nat. Astron.*, 7, 931

Peters S., Schroeder D., Chu W., Castelletti D., Haynes M., Christoffersen P., Romero-Wolf A., 2021, *Geophys. Res. Lett.*, 48, e2021GL02450

Petroff E., Hessels J., Lorimer D., 2019, *A&AR*, 27, 1

Pietka M., Fender R., Keane E., 2015, *MNRAS*, 446, 3687

Polisensky E. et al., 2016, *ApJ*, 832, 60

Popov M. V. et al., 2008, *Astron. Rep.*, 52, 900

Rea N. et al., 2022, *ApJ*, 940, 72

Rea N. et al., 2024, *ApJ*, 961, 214

Rowlinson A., Anderson G., 2019, *MNRAS*, 489, 3316

Rowlinson A. et al., 2016, *MNRAS*, 458, 3506

Rowlinson A. et al., 2022, *MNRAS*, 517, 2894

Ruhe D., Kuiack M., Rowlinson A., Wijers R., Forré P., 2021, *Astrophysics Source Code Library*. recordascl:2103.015

Ruhe D., Kuiack M., Rowlinson A., Wijers R., Forré P., 2022, *Astron. Comput.*, 38, 100512

Safaenili A., Gulkis S., Hofstadter M., Jordan R., 2002, *Meteorit. Planet. Sci.*, 37, 1953

Shimwell T. et al., 2017, *A&A*, 598, A104

Shimwell T. et al., 2019, *A&A*, 622, A1

Shimwell T. et al., 2022, *A&A*, 659, A1

Smirnov O. M., 2011, *A&A*, 527, A106

Sokolowski M. et al., 2021, *Publ. Astron. Soc. Aust.*, 38, e023

Spangler S. R., Moffett T. J., 1976, *ApJ*, 203, 497

Spreeuw H., Swinbank J., Molenaar G., Staley T., Rol E., Sanders J., Scheers B., Kuiack M., 2018, *Astrophysics Source Code Library*. record ascl:1805.026

Stewart A. J. et al., 2016, *MNRAS*, 456, 2321

Tasse C. et al., 2018, *A&A*, 611, A87

Tasse C. et al., 2021, *A&A*, 648, A1

Tingay S., Hancock P., 2019, *AJ*, 158, 31

Tingay S. et al., 2015, *AJ*, 150, 199

Umata G., Tringilio C., Catalano S., 1998, *A&A*, 329, 1010

van Haarlem M. P. et al., 2013, *A&A*, 556, A2

Varghese S. S., Obenberger K. S., Dowell J., Taylor G. B., 2019, *ApJ*, 874, 151

Virtanen P. et al., 2020, *Nat. Methods*, 17, 261

Wang Y. et al., 2023, *MNRAS*, 523, 5661

Zhao J.-H. et al., 1992, *Science*, 255, 1538

APPENDIX A: MATHEMATICAL DESCRIPTION SOURCE SUBTRACTION

Below we describe the sky model subtraction process in a more mathematical manner. An interferometer measures an approximate Fourier transform of the sky intensity, instead of directly measuring

the sky intensity $I(l, m)$. The intensity and visibility are defined as a function of viewing direction cosines (l, m) . For a baseline consisting of antennas i and j , the perfect response to all visible sky emission for a single time instance and frequency is given by the idealized Radio Interferometric Measurement Equation (RIME), see eg. Smirnov (2011). The RIME below does not include the Jones matrices that describe the direction-dependent and direction-independent calibration effects. The visibilities for a baseline consisting of antennas i and j are defined as

$$V_{ij} = \int \int I(l, m) e^{-2\pi i [u_{ij}l + v_{ij}m + w_{ij}(n-1)]} \frac{dl dm}{n}, \quad (\text{A1})$$

where $i = \sqrt{-1}$, $n = \sqrt{1 - m^2 - l^2}$, u_{ij} , and v_{ij} are baseline coordinates in the UV plane-parallel to l and m , respectively, and w_{ij} is the baseline coordinate along the line of sight.

In practice, the visibilities are affected by predominantly antenna-based complex gain factors which may vary with time, frequency, viewing direction, and antenna location. Therefore, the observed visibility is defined as

$$V_{ij,obs} = \int \int a_i(l, m) a_j^\dagger(l, m) I(l, m) e^{-2\pi i [u_{ij}l + v_{ij}m + w_{ij}(n-1)]} \frac{dl dm}{n}, \quad (\text{A2})$$

where \dagger denotes a complex conjugate. The process of calibration tunes these antenna gains $(a_i(l, m)$ and $a_j^\dagger(l, m))$ such that the calibrated or antenna gain corrected visibilities best match the sky model, i.e. $V_{ij,obs} \approx V_{ij}^{\text{sky model}}$. The sky model gives a coarse model of the (brightest) radio sources in the sky.

One can rewrite the RIME (equation A1) as a linear combination of all individual sources k . To this end, we approximate the radio sky by a discrete number of isolated, invariant sources of finite angular extent.

$$V_{ij} = \sum_k V_{ij,k} = \sum_k \int \int I_k(l, m) e^{-2\pi i [u_{ij}l + v_{ij}m + w_{ij}(n-1)]} \frac{dl dm}{n}. \quad (\text{A3})$$

The ‘subtracted’ visibilities refer to the visibilities where we have subtracted the sky model for each source. Mathematically this can be expressed as

$$V_{ij,sub} = V_{ij,obs} - \sum_k \left(g_{ik} g_{jk}^\dagger \right)^{-1} V_{ij,k}^{\text{model}} \quad (\text{A4})$$

with $g_{ik} = g_i(l_k, m_k) \approx a_{ik}^{-1}$, the inverse of the antenna gains. In case of a perfect calibration, this subtracted visibility is zero.

APPENDIX B: POINT SPREAD FUNCTION ARTEFACTS

In the methods section, we indicated the necessity to image 2200 by 2200 pixels, even though we are only using sources from the inner 1800 by 1800 pixels. Fig. B1 shows an example of a fake transient source that is likely introduced by a bright source just at the boundary of the image. The left-hand panel shows a subtracted image of 1800 by 1800 pixels, while the right image shows a subtracted image of the same sky area, imaged with the same imaging settings except for increasing the image size to 2200 by 2200 pixels. The red circle indicates the location of the fake transient source present in the left image. The dashed circle show the 1.5° radius that we use for filtering. The coloured dots indicate the locations bright sources in the LoTSS source catalogue. We think that for example the 1.7 Jy

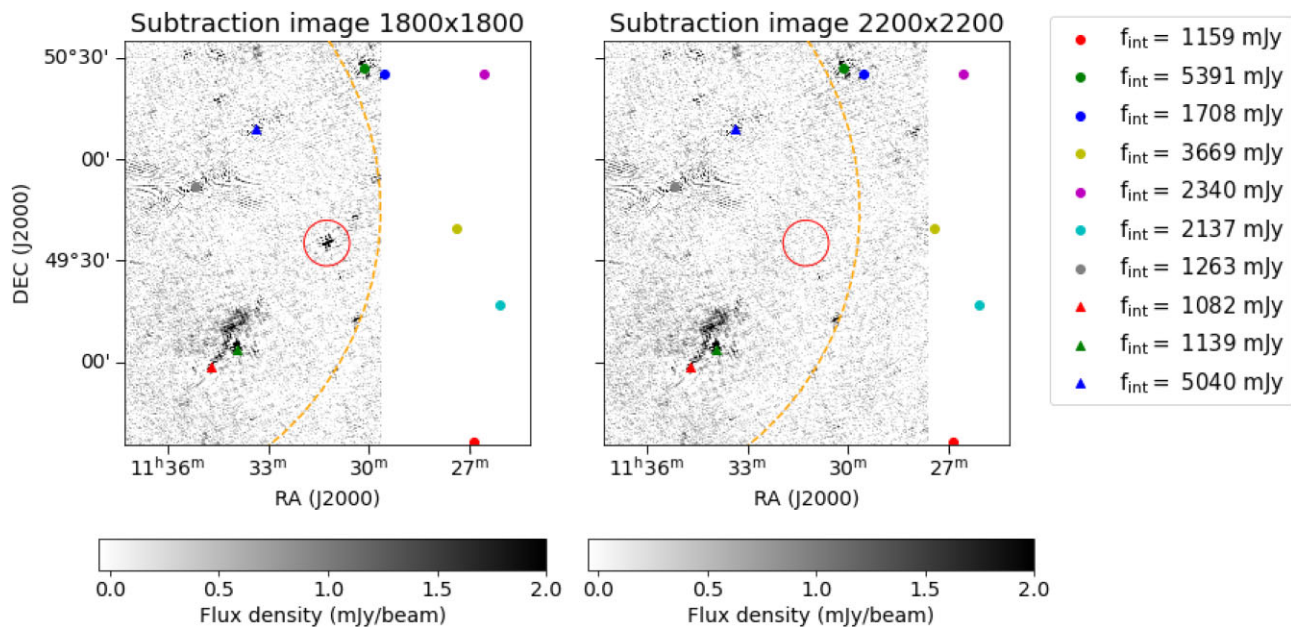


Figure B1. Example of a ‘fake’ transient source that is likely introduced by a bright source just at the boundary of the image. The left-hand panel shows the a subtracted image of 1800 by 1800 pixels, while the right image shows a subtracted image of the same sky area, imaged with the same imaging settings except for increasing the image size to 2200 by 2200 pixels. The red circle indicates the location of the fake transient source present in the left image. The dashed circle show the 1.5° radius that we use for filtering. The coloured dots indicate the locations bright sources in the LoTSS source catalogue.




Figure B2. Example trajectory of two transient candidates, that were excluded after reimaging on a 2200 by 2200 pixel image. The right ascension and declination minus the average location of the source are shown for a source found in 2 min subtracted image snapshots in blue, and a source found in 1 h snapshots, shown in orange.

source (indicated by the navy blue dot) just beyond the 1800 by 1800 limit in the left image could introduce point spread function (PSF) artefacts quite far down the rest of the image. Simply increasing the image size a bit shows that the bright blob in the red circle in the left image is not a real transient candidate.

Initially, it was unclear what the origin of these transient candidates was. There was a strong suspicion that these types of sources were

artefacts, because the source location seemed to move throughout the observation. In Fig. B2, we show the trajectory of two transient candidates, that were excluded after reimaging on a 2200 by 2200 pixel image. The figure shows the right ascension and declination minus the average location of the source. It is clear that for both the example shown in blue, found in 2 min subtracted image snapshots and the example shown in orange, found in 1 h snapshots, the source seems to follow a particular trajectory.

APPENDIX C: TRANSIENT CANDIDATES AFTER VISUAL INSPECTION

The figures in this Appendix show all transient candidate sources that are left after a first round of visual inspection. These sources did not immediately fall within one of the two categories of sources that are normally vetted by visual inspection (subtraction artefacts of bright sources, or sources that are associated with a faint deep fields source). The transient candidate shown in Figs C1 and C2 is discussed in detail in Section 5.1. The transient candidate shown in Fig. C3 is discussed in more detail in Section 5.2.

We decide to not follow-up on the candidates presented in Figs C4, C5, and C6 because the source in the subtracted image is not dirty-beam shaped, unlike expected (see Section 3.1). Finally, the candidates shown in Figs C7–C11 are disregarded because upon closer inspection they are associated with faint sources in the deep field.




Figure C1. Transient candidate C1 identified in two 8 s snapshots. Discussed in more detail in Section 5.1.




Figure C2. Transient candidate C2 identified in one 2 min snapshots. Discussed in more detail in Section 5.1.




Figure C3. Transient candidate C3 identified in all 1 h snapshots in the P35Hetdex10 field. Discussed in more detail in Section 5.2.




Figure C4. Transient candidate in the P205+55 field identified in one 1 h snapshot with an SNR of 6.6.




Figure C5. Transient candidate in the P14Hetdex04 field identified in the first and second 1 h snapshot with an SNR of 5.5 and 6.7, respectively.




Figure C6. Transient candidate in the P19Hetdex17 field identified in one 1 h snapshot with an SNR of 6.1.

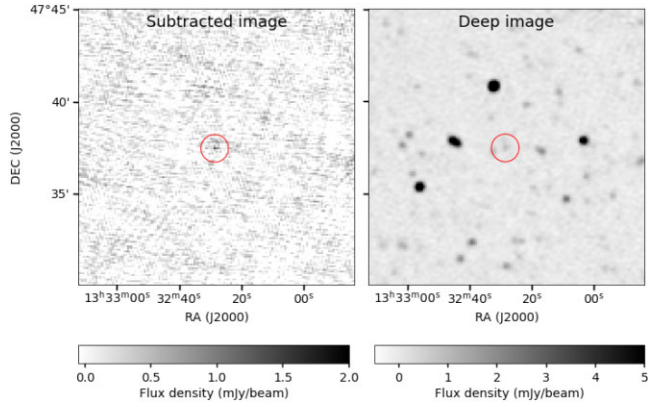


Figure C7. Transient candidate in the P39Hedex19 field identified in the third 1 h snapshot with an SNR of 6.6.




Figure C8. Transient candidate in the P182+55 field identified in the third 1 h snapshot with an SNR of 6.1.




Figure C9. Transient candidate in the P225+47 field identified in the fourth and seventh 1 h snapshot with an SNR of 5.5 and 6.4, respectively.




Figure C10. Transient candidate in the P225+47 field identified in the sixth and seventh 1 h snapshot with an SNR of 5.1 and 6.2, respectively.




Figure C11. Transient candidate in the P191+55 field identified in the third 1 h snapshot with an SNR of 6.3.

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.