



UvA-DARE (Digital Academic Repository)

Apparent radio transients mapping the near-Earth plasma environment

Kuiack, M.J.; Wijers, R.A.M.J.; Shulevski, Aleksandar; Rowlinson, A.

DOI

[10.1093/mnras/stab1156](https://doi.org/10.1093/mnras/stab1156)

Publication date

2021

Document Version

Final published version

Published in

Monthly Notices of the Royal Astronomical Society

[Link to publication](#)

Citation for published version (APA):

Kuiack, M. J., Wijers, R. A. M. J., Shulevski, A., & Rowlinson, A. (2021). Apparent radio transients mapping the near-Earth plasma environment. *Monthly Notices of the Royal Astronomical Society*, 504(4), 4706-4715. <https://doi.org/10.1093/mnras/stab1156>

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.

Apparent radio transients mapping the near-Earth plasma environment

Mark J. Kuiack¹,^{*} Ralph A. M. J. Wijers¹, Aleksandar Shulevski¹ and Antonia Rowlinson^{1,2}

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

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

Accepted 2021 April 9. Received 2021 April 9; in original form 2020 March 24

ABSTRACT

We report the discovery of bright, fast, radio flashes lasting tens of seconds with the AARTFAAC high-cadence all-sky survey at 60 MHz. The vast majority of these coincide with known, bright radio sources that brighten by factors of up to 100 during such an event. We attribute them to magnification events induced by plasma near the Earth, most likely in the densest parts of the ionosphere. They can occur both in relative isolation, during otherwise quiescent ionospheric conditions, and in large clusters during more turbulent ionospheric conditions. Using a toy model, we show that the likely origin of the more extreme (up to a factor of 100 or so) magnification events likely originate in the region of peak electron density in the ionosphere, at an altitude of 300–400 km. Distinguishing these events from genuine astrophysical transients is imperative for future surveys searching for low frequency radio transient at time-scales below a minute.

Key words: scattering – turbulence – radio continuum: transients.

1 INTRODUCTION

In recent years, low-frequency (<200 MHz) transient surveys have been completed with the Long Wavelength Demonstrator Array (Lazio et al. 2010), the Murchinson Widefield Array (Bell et al. 2014, 2019), Long Wavelength Array (LWA1; Obenberger et al. 2015), Low-Frequency Array (LOFAR; Carbone et al. 2016), Owens Valley Radio Observatory Long Wavelength Array (OVRO-LWA; Anderson et al. 2019). These focus on the search for bright transient or variable emission, with the hope of discovering new extremely energetic processes taking place in the Universe. Previously, we have reported the discovery of extremely high fluence giant pulses from PSR B0950 + 08 (Kuiack et al. 2020a) during the course of an ongoing blind transients survey with the Amsterdam-ASTRON Radio Transient Facility And Analysis Center (AARTFAAC; Prasad et al. 2014, 2016). We were easily able to associate this emission to the specific astrophysical source because the signal repeated and the source properties and behaviour had been well studied. However, the goal of blind transient surveys such as AARTFAAC is also to search for as-yet unknown phenomena. This necessarily means that any candidate events discovered, may not be so easily associated with known objects, or known processes.

An example is the discovery of an unknown transient source at low frequency in the LOFAR North Celestial Pole survey (Stewart et al. 2016). No strong association to a progenitor, host location, or known astrophysical phenomenon could be made, thus the primary validation is to minimize the chance it could be an imaging artefact. Once we are convinced that an event is not an instrumental or data processing artefact, we can claim detection of a source. However, we then need to check whether the event could not be a noise spike associated with normal measurement error distributions, or variabil-

ity induced by the ionosphere, which has particularly strong effect in the AARTFAAC frequency range (30–80 MHz), (de Gasperin et al. 2018). Fortunately, a high cadence transient survey generates a huge number of measurements (one per second per frequency band in the case of AARTFAAC) and thus allows us to characterize the noise distribution and ionospheric behaviour quite well. We have already studied the variations due to the quiescent ionosphere and noise quite extensively when constructing the AARTFAAC catalogue (Kuiack et al. 2018), and when studying PSR B0950 + 08 (Kuiack et al. 2020a). The bottom line of our findings was that the measurement noise does not play a dominant role in the sources we detect almost by definition because we set the detection threshold at 5–8 times the rms. noise; the ionosphere-induced variations in quiescence have an amplitude of 20–30 per cent, on a time-scale of 10–20 min at 60 MHz.

However, in analysing 545 h for the AARTFAAC transient survey (Kuiack et al. 2020b), more exotic ionospheric behaviours were observed, which confound the detection of transient signals intrinsically emitted from astrophysical sources. Studying these conditions, and their effects on the statistics of measured light curves in detail can inform the criteria required for low frequency transient candidate selection. For example, Varghese et al. (2019) used the kurtosis of the flux measurement distribution to argue that an event that was an outlier among a population of scintillating sources it was more likely a astrophysical transient. Strongly scintillating sources commonly exhibit many scintels over a period of time and therefore the total flux density distribution is skewed towards higher values. Their claimed transient is a single flash, so the total measurement distribution is not skewed. This transient claim does not seem to consider the possibility of single isolated scintels, our work seeks to explore and differentiate the statistical properties of singly and multiply magnified scintillating sources. Further, in attempting to detect transient and variable sources Rowlinson et al. (2019) explore the distributions of two parameters, the weighted reduced χ^2 and the variability index, across all sources observed during

* E-mail: m.j.kuiack@uva.nl

an observation. Simulations illustrated how distinct transient types would appear as outliers in this distribution. We hope to extend the toolkit of transient detection by introducing multifrequency pulse shape and delay analysis to the light curve of each source. Additionally, studying periods of strong scintillation gives insight into the dynamics of the ionosphere (Hill & Frehlich 1996), and due to the coupling of the magnetosphere–ionosphere/thermosphere system (Tu & Song 2016) the broader near-Earth plasma environment as well (Blagoveshchenskii 2013; Danilov & Konstantinova 2019).

Here, we report the discovery of rapid, large-amplitude variations in the 60 MHz flux densities of normally constant sources on short time-scales: sources show apparent flare-ups in brightness of a factor few to over 100 on time-scales of a few tens of seconds. They can occur both in relative isolation, when other nearby sources show their expected quiescent behaviour, or in large groups over a fairly large area of sky, for a period lasting hours. In most cases, careful study reveals that they are due to exceptional effects of relatively nearby plasma. The large regions of strong scintillation are potentially due to, and a new method for observing, a travelling ionospheric disturbance (TID).

In Section 2, we describe the instrument set-up, the set of observations, and the methodology used to analyse light curves and detect isolated bursts. Then, in Section 3 we present our results, both a qualitative and quantitative description of light-curve features that mimic short-duration transients and variables. Next, in Section 4 we explore the potential cause. Lastly, in Section 5 we give our conclusions.

2 INSTRUMENT, OBSERVATION, AND METHODS

AARTFAAC is a project whose aim is to survey the Northern hemisphere for bright (tens of Jy), low-frequency (30–80 MHz) radio transients, with time-scales upwards of 1 s (Prasad et al. 2016). In the AARTFAAC-6 configuration, data are streamed directly from 288 low-band antennas (LBA) in the six central stations of the LOFAR (van Haarlem et al. 2013), the so called ‘superterp,’ providing a maximum baseline of 300 m and resolution of 1° at 60 MHz. The data are processed in parallel during other ongoing LOFAR-LBA observations, or designated ‘filler’ time, employing an independent correlator, and subsequent calibration and imaging pipelines on a dedicated comput storage cluster (Prasad et al. 2014). The data are recorded with a 1 s time integration, in 16 separate sub-bands, each with a bandwidth of 195.3 kHz, and arranged in two contiguous sequences of eight sub-bands covering 57.6–59.0 MHz and 61.1–62.5 MHz.

Real-time, streaming, complex-gain calibration is performed with custom built software (described fully in Prasad et al. 2014), which employs a multisource self-calibration scheme (Wijnholds & van der Veen 2009). The sky model consists of five sources Cassiopeia A, Cygnus A, Taurus A, Virgo A, and the Sun. These are each modelled with a single Gaussian component, an accurate enough approximation at AARTFAAC-6 resolution. At each time-step, both direction independent and direction-dependent gain solutions are calculated, however, converging gain solutions are propagated forward in time to reduce the required iterations and thus computation time. Once the visibilities are calibrated, the bright sources used in calibration are subtracted.

The calibrated visibilities are then streamed to the archive disc and our custom built imaging pipeline. The imaging pipeline simply applies flags to visibilities with baselines shorter than 10λ to reduce the appearance of bright diffuse Galactic emission, then performs

frequency-dependent gridding and Fourier transform. This creates a 1024×1024 pixel image for each of the two 1.5 MHz bands, centered at 58.3 and 61.8 MHz. Flux scaling is done in the image plane by comparing many flux density measurements from persistent sources in the image, to the AARTFAAC catalogue (Kuiack et al. 2018). The image stream can then be saved to disc for off-line transient analysis, and/or streamed to the AARTFAAC website¹ for live data quality inspection and public outreach. The total duration of data analysed in this work was 545.25 h, collected between 2016-8-31 and 2019-9-14.

With eight subbands (1.56 MHz) and 1 s integration, AARTFAAC-6 is confusion noise limited, the theoretical thermal noise limit is 1 Jy, while with a resolution of 1° the confusion noise is $10.4 \text{ Jy beam}^{-1}$ (equation 6; van Haarlem et al. 2013). The typical background rms noise in the images, averaged across the field of view and the entire survey is 7.5 Jy beam^{-1} , yielding typical 5σ and 8σ detection thresholds of 37.5 and 60 Jy.

Light curves are generated from the images by the LOFAR Transient Pipeline v3.1.1² (TraP; Swinbank et al. 2015, and references therein). TraP is a PYTHON and SQL based transient detection pipeline for image-based surveys. The only non-standard parameters used in TraP are the one that sets the detection threshold as a function of the local rms noise level, `detection_threshold = 5\sigma`, and one that forces the use of a fixed-width 2D Gaussian beam during source extraction. This is set by the `force_beam = True` parameter in the TraP job configuration file, which locks the widths to the semimajor and semiminor axis of the restoring beam, read from the fits header.

2.1 Light-curve feature analysis

When analysing a potential transient event, we are interested in measuring important quantities such as the peak flux in each frequency, the delay between frequencies, and the duration. This is done automatically using the `scipy.signal` (McKinney 2010) modules `find_peaks` and `peak_widths`. The algorithm for detecting burst-like features in light curves is as follows:

(i) First, the data are regularized in time. By default measurements are only recorded when a signal greater than 5σ is detected. We therefore fill the light curve between the min and max time stamp with zeros, then make a copy of the light curve, smoothed with a flat 3-s kernel to cover missing time-steps, and decrease noise.

(ii) Peaks are detected on the smoothed light curve with `find_peaks` using the parameters:

(a) `min_width = 5 [s]`: The minimum FWHM of light-curve features to be considered. Together with the 3-s smoothing, this ensures that single-bin bright noise spikes are not considered.

(b) `max_width = None`: The maximum FWHM of light-curve feature to be considered.

(c) `min_prominence = 30 [Jy]`: The minimum difference in height between a light-curve feature and its surroundings. Correctly detecting isolated bursts requires that the minimum prominence be larger than the typical measurement variability, reducing the detection of noise spikes on top of smooth light curves (false positives), but small enough that true

¹www.aartfaac.org/live, video stream hosted by YouTube.

²<https://github.com/transientsktp>

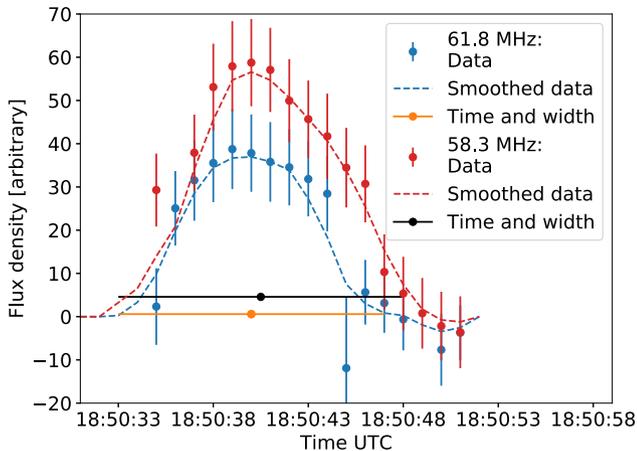


Figure 1. An example light curve from an isolated burst-like magnification, observed on 2016-09-05. The blue and red circles depict the flux density measurements at 58.3 and 61.8 MHz, respectively (in arbitrary units, since the transient finder runs on the data before putting them on the correct flux scale). The dashed lines show the light curve smoothed over 3 s, used during peak analysis. The black and orange dots and the horizontal lines show the peaks’ mid-point, used for interfrequency association, and the width at 10 per cent of peak flux level. The underlying source is 4C 26.55, a galaxy with a quiescent flux density of 15.9 Jy at 74 MHz (Lane et al. 2014).

positives are not ignored. Therefore, this is manually tuned to approximately the minimum height of a true positive.

(d) `max_prominence = None`: The maximum difference in height between a light-curve feature and its surrounding.

(e) `rel_height = 0.9`: The relative height of the feature where the width is measured. This number corresponds to the proportion above therefore 1.0 is the base, and 0.0 the peak.

(iii) Peaks are associated using the central time, matching nearest neighbours within a delay interval of -120 to $+20$ s between the 61.8 and 58.3 MHz light curves (120 s delay would result from a DM of 900 pc cm^{-3}). One-to-many matches are eliminated by choosing the pair with the minimum delay.

(iv) Highly skewed peaks, where the time difference between the peak and mid-point are greater than 5 s, are removed. This ensures that we are selecting only isolated peaks. Consequently, we may also be removing real transients that follow a fast rise-exponential decay behaviour or series of peaks that are overlapping.

(v) Lastly, once the features are matched between frequencies, the effective duration of the burst is measured on the smoothed light curve at `rel_height = 0.9`, as is the delay (the difference between the mid-points in time), and the peak is measured from the original light curve, which consists of measurements over 5σ .

Using this feature-detection algorithm, we searched the candidate data bases, output by TraP, for light curves consisting of only a single burst. Fig. 1 depicts an example transient candidate consisting of one burst, and the various parameters of interest described above.

2.2 Strong scintillation zones

We sometimes find large concentrations of magnification events in sky region and time, which we call ‘strong scintillation zones’. We quantify these by counting the number of sources whose flux density is typically below the confusion noise limit, rendering them undetectable during calmer conditions, but are temporarily magnified above our threshold. This is done by extracting all of the sources

observed in each 5-min interval, and calculating the surface density of sources that are not present in the AARTFAAC catalogue. We model the source distribution in the sky with a Gaussian kernel density estimation (KDE), which yields a 2D probability density function (PDF) of detected sources. By comparing relatively calm observations to those where a zone is plainly visible by eye, in time-lapse videos, we heuristically set the PDF threshold that is characteristic of strong scintillation.

3 RESULTS

Here, we report on three different types of transient or variable events observed in the seconds-time-scale, low-frequency radio sky. The first are isolated, spontaneous flashes, which coincide in position with 10–20 Jy radio sources, that are otherwise undetectable by AARTFAAC. Although they are coincident in direction, the flashes do not have a sufficient time-frequency delay to originate from an extragalactic distance. The second, is strong variability characterized in some sources by many magnification spikes in a relatively short time, and in others by singular flashes similar to the first, both clearly induced by a localized structure that persists over many hours. And third, we observed a few potential cosmic transients that, although similar to the other transients in peak shape, duration, and magnitude, have a time-frequency delay that is consistent with an extragalactic origin, these sources are analysed further in the AARTFAAC low-frequency transient survey paper Kuiack et al. (2020b).

3.1 A detailed look at light curves

3.1.1 Singular magnification events

Using the method described in Section 2.1, a population of 76 transient candidates, each of which appears as a single isolated burst, were discovered. One example observed on 2016-09-05, is shown in Fig. 1. This light curve shows a magnification of the flux density of 4C 26.55, which has a quiescent flux density of 15.9 Jy at 74 MHz (Lane et al. 2014). This burst was analysed more deeply by re-imaging all 16 individual sub-bands in a 30-min interval around the candidate. The resulting 16 sub-band light curve is shown in Fig. 2. Here, we can see the relative stability of the background in the 10-min prior, and 20 min after that burst, albeit with a few sub-detection-threshold correlated noise events around 18:55. This indicates that this region of sky, at this time, is not subject to the type of strong scintillation described in Section 3.2.

An isolated, bright outburst, with a time-scale on the order of seconds is exactly the kind of transient signal we are searching for with AARTFAAC. This could potentially indicate the detection of a low frequency fast radio burst, the prompt emission of a neutron star merger, or another as of yet unobserved phenomenon dumping a tremendous amount of energy in to coherent emission in a short time, an exciting prospect. However, given the lack of any noticeable time delay between 62.5 and 57.6 MHz, it is apparent that the signal can not be of extragalactic origin. In Fig. 2, we see that there is no appreciable delay between the sub-bands.

In order to estimate an upper limit on the distance, we attempted to measure any dispersion delay in the signal. The time delays are measured by the lag in the cross-correlation functions between the light curve at 62.5 MHz and each successive frequency. Additionally, in order to estimate the uncertainty in the inferred DM, we use a simple Monte Carlo method: We randomly generate 10 000 sets of 16 light curves, according to the measured flux densities and uncertainty, at each frequency and time. We then fit the delays for each set. The

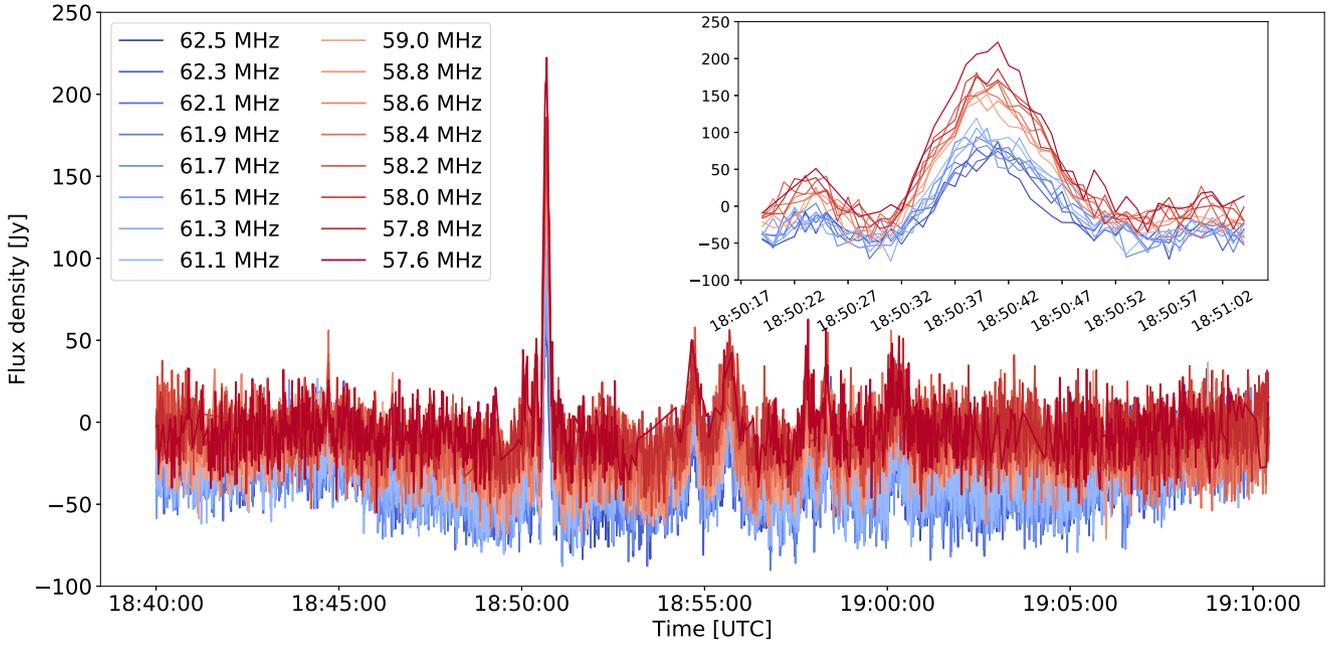


Figure 2. Singular non dispersed burst-like profile from the 16 frequency channels sample with a 1 s time resolution, observed on 2016-09-05.

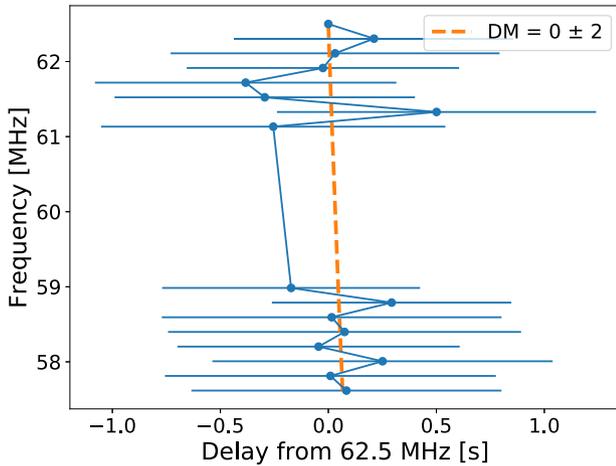


Figure 3. The time lag between the 62.5 MHz light curve, the highest sub-band frequency, and each of the 16 sub-bands. The delay uncertainties were calculated by randomly generating 10000 sets of light curves using the measured flux densities and their uncertainties. The orange line shows the best-fitting dispersion delay, and is clearly consistent with 0 DM. Additionally, with the data quality during this event, we are sensitive to DMs of even a few.

mean and standard deviation for each frequency are shown in Fig. 3 and allow us to conclude that for this transient candidate, the DM is 0 ± 2 , indeed consistent with zero.

In Fig. 4, we illustrate another example of an isolated, non-dispersed burst-like event. Here, the magnified source is 3C 234. With a persistent flux density of 90.06 ± 0.15 Jy during regular conditions (Kuiack et al. 2018), the light curve surrounding the central magnification remains above the detection threshold. This reveals how the flux density is de-magnified, prior to, and following the apparent burst, before returning to the steady state. Together with some secondary peaks, it somewhat resembles an Airy function,

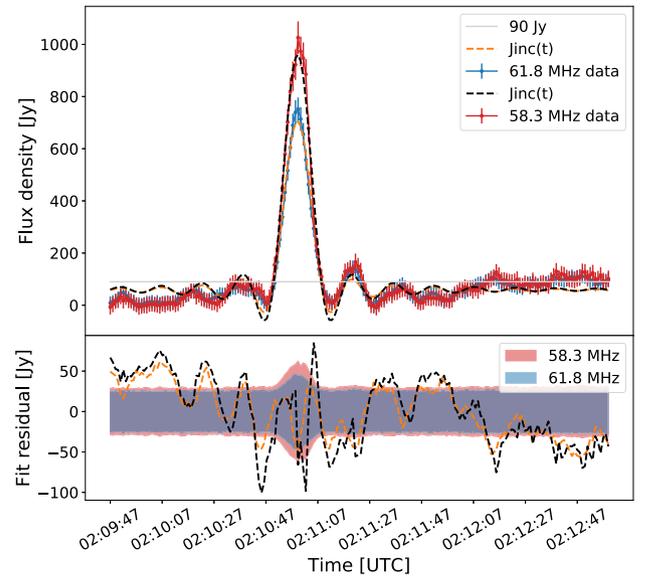


Figure 4. Isolated magnification of the signal from 3C 234, captured during 2019-01-13. We fit a Jinc function (top), and plot the residuals (model – data, bottom), where the shaded regions are defined by the flux density uncertainty. The horizontal grey line in the top panel indicates the 90 Jy persistent flux density as measured in the AARTFAAC catalogue.

except that the secondary peaks are a bit high relative to the central peak. We fit a Jinc function, which fits both the central peak and the positions and amplitudes of the satellite peaks fairly well:

$$\text{Jinc}(t) = a \frac{2J_1(t)}{t} + d, \quad (1)$$

where $J_1(t)$ is the Bessel function of the first kind, and $t = \frac{x}{b} - c$, with b and c used to scale and translate the time x of the light curve during fitting. This pattern could be caused by a cone-shaped

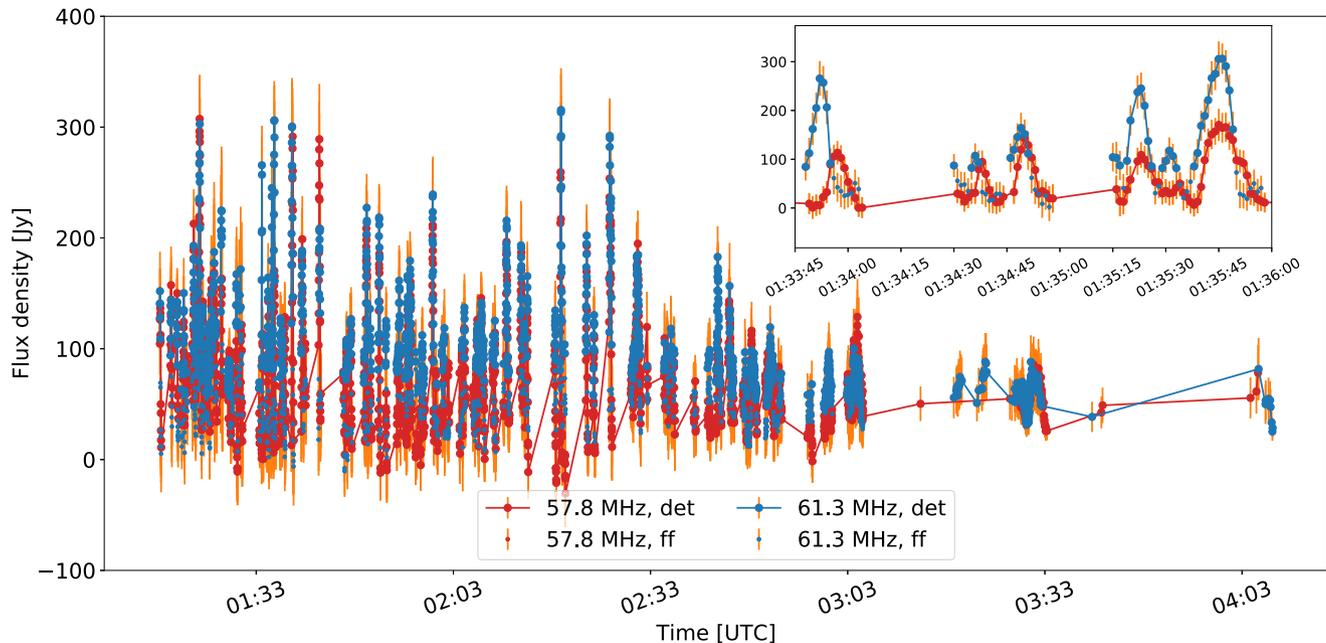


Figure 5. Example light curve illustrating strong scintillation of 4C 14.27 during an observation on 2019-01-01, with the inset expanding a typical region of the light curve, illustrating the scintel shape and height variability. The quiescent flux is 15.07 Jy at 74 MHz, typically below our detection threshold. Large connected dots labelled ‘det’ denote measurements where the signal is blindly detected, i.e. above 5σ . smaller dots labelled ‘ff’ are forced-fit fluxes, which are done automatically for a few time-steps after a detection at the same location.

notch in a plasma sheet, but we shall return to that in the discussion (Section 4.1).

3.1.2 Strong variability due to scintillation

Sometimes events similar to those described in the previous section appear in large numbers in a region of the sky, affecting a given source multiple times during some period and affecting many neighbouring sources as well. Fig. 5 is an example light curve, recorded on 2019-01-01, when such strong scintillation was observed across nearly half of the visible sky. The source is 4C 14.27, a Seyfert 1 galaxy, with a flux density of 15.07 ± 0.16 Jy in the VLSSr catalogue, and very compact, <10 arcsec according to the FIRST survey, which has a synthesized PSF FWHM of 5 arcsec (Helfand, White & Becker 2015). We will call the individual events ‘scintels’, and during such periods of strong scintillation, many tens of scintels, each similar in character to the isolated case in Fig. 2, can be produced. This is clearly a persistent background source temporarily exhibiting strong variability through scintillation. In Section 3.2, we will measure the properties of these scintels and compare their distributions with the distributions of properties of the isolated events.

3.2 Statistical description of strong scintillation zones

In attempting to disentangle the different scintillation regimes discussed earlier, we have observed zones of strong scintillation in opposition to the Sun. Their locations are not generally fixed either in celestial or terrestrial coordinates, so it is not immediately clear from their motion what to associate them with. An example of these zones is shown in Fig. 6.³ The zones vary in size from 20° to 40°

³For an animated version visit https://aartfaac.org/images/scint_20170225.gif.

ecliptic latitudinal extent and from 20° to 90° ecliptic longitudinal extent. They can appear stationary in ecliptic coordinates and in size, or evolve along the latitude lines, decreasing from 60° to 40° latitude, or remain stable between 0 and 20. Of the 28 nights where we observed between 22 and 05 UTC, strong scintillation regions are clearly defined in seven. The total lifetime of the zone can be hours, as seen in the observation on 2018-03-09, where the scintillation zone begins around 22:00 UTC, peaks in intensity at 00:00 UTC, then diminishes until disappearing at 04:00 UTC, lasting a total of 6 h. However, on 2017-02-25, Fig. 6, the zone is present at the start of the observation at 1:30 UTC, increases in intensity until 4:20 and then begins to decrease until the end of the observation at 5:00. Then in the next night at 1:30, when we start observing again, the zone is present at 40° latitude, but restricted in longitudinal extent compared with the previous night, indicating that this may be a single case lasting more than 24 h.

3.2.1 Magnification factor and peak flux of the underlying source

In an effort to understand whether the singular magnification events are different in cause and character to the scintels observed during periods of strong scintillation, we compare the magnification distributions from multiply and singly magnified sources. In Fig. 7, we show the distribution of magnification factors of 725 scintels described in Section 3.1.2. These were measured from 16 sources, during one period of strong scintillation. The typical flux density of each is below our detection threshold, thus the individual scintels are well isolated from each other. Additionally, 76 singular scintels, instances where the flux density has been magnified only once during the observation (Section 3.1.1), for that specific background source, are also plotted. In both populations, the distribution of magnification factors is bounded on the lower side by our sensitivity limit. Therefore, the bottom of the distribution follows a power

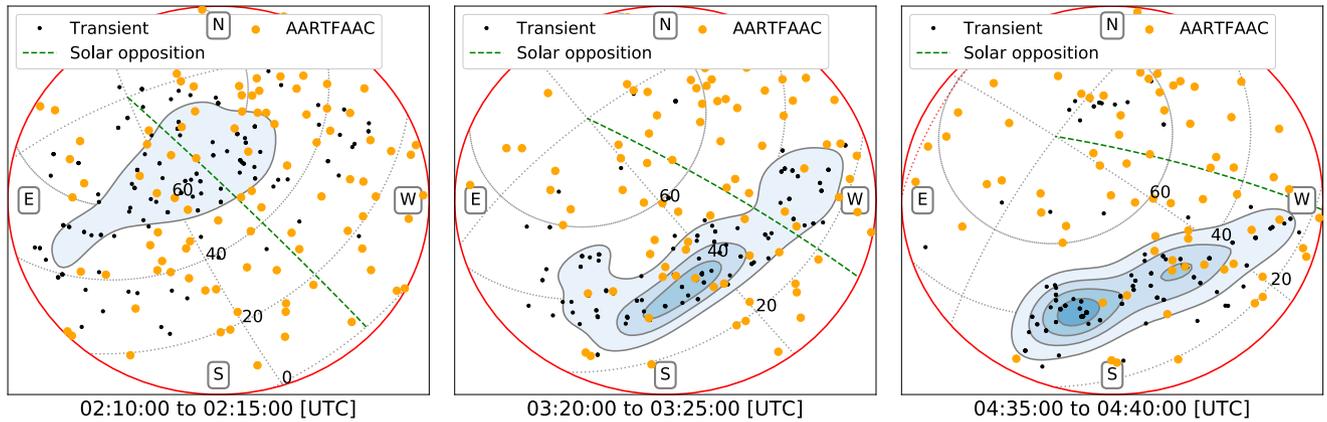


Figure 6. Evolution of the scintillation zone, during three h on 2017-02-25. Each frame illustrates the source detection region, down to 55° from zenith (the red line). Ecliptic coordinates (the grey-dotted lines), with the latitudinal parallels labelled at 180° longitude. The green-dashed line defines geomagnetic midnight; the longitude of solar opposition. The AARTFAAC catalogue sources (the orange circles), and other sources detected (the black dots) during each time interval are shown. The contours outline the KDE of the scintels, with levels from 0.15–0.35, in steps of 0.05, scintels $\text{deg}^{-2} 5 \text{ min}^{-1}$ and reveal the extent of the zone of strong scintillation.

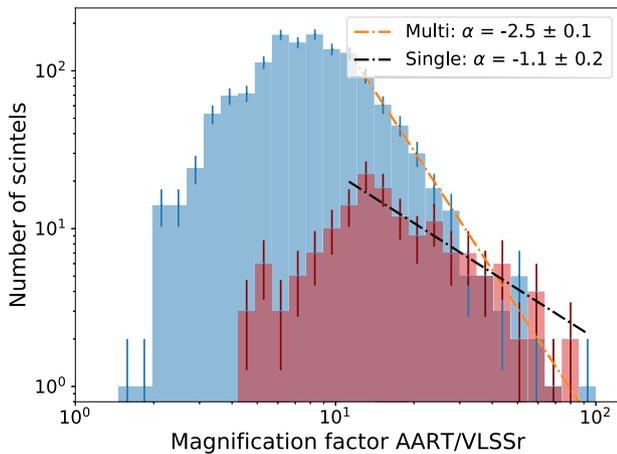


Figure 7. Magnification factor distributions for the 725 multiply magnified scintels (blue), and the 76 singly magnified scintels (red).

law, with index -1 . The distribution is bounded on the top by the decreasing likelihood of observing a scintel with greater and greater magnification, and on the right by the decreasing number of sources at higher unmagnified brightness. The distribution of magnification factors, illustrated in Fig. 7, drops according to a power-law distribution with an index of -2.5 ± 0.1 for the multiply magnified scintels. For the singly magnified scintels the distribution is flatter, with a power-law index is -1.3 ± 0.2 .

3.2.2 Duration distribution

Next, in an effort to infer physical properties of the turbulent medium producing the scintillation, we investigated the distribution in the temporal widths of the scintels, as well as the relationship between the scintels' width and the peak magnification factor. The width distribution is approximately exponential with a characteristic time of $\tau = 1/\lambda = 8.5 \pm 1.5 \text{ s}$, from the peak of 10 s (Fig. 8). For comparison, the isolated magnification events have a somewhat longer characteristic duration, with a characteristic time of a fitted exponential distribution

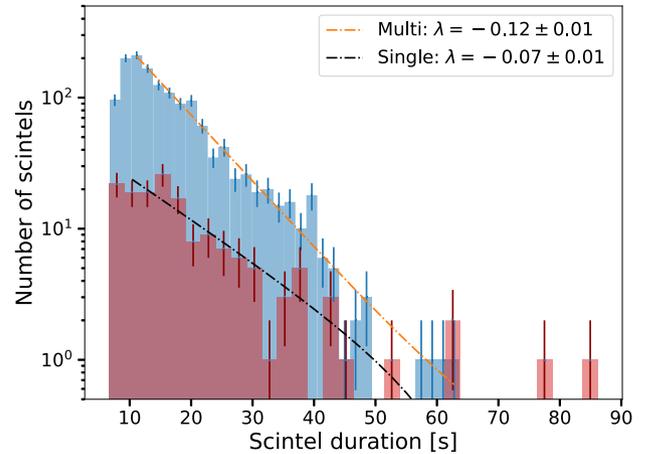


Figure 8. Distribution of scintel duration for multiply magnified sources (blue) and single events (red). Closely follows an exponential distribution.

of $\tau = 15 \pm 2 \text{ s}$. Additionally, the scintel duration and magnification factor are almost uncorrelated (correlation coefficient 0.2, Fig. 9).

3.2.3 Scintel delay distribution

Lastly, in order to deduce to what extent strong scintillation confounds the search for extragalactic transients, we investigated the distribution of delay times between peaks measured at 61.8 and 58.5 MHz for each scintel or isolated magnification event. The distributions are shown in Fig. 10 for both types separately. The distribution of scintel delay times is fairly narrow and symmetric around zero, and its width is consistent with noise errors. Hence, we conclude that they have no measurable delays and thus are consistent with $\text{DM} = 0$. For the isolated events, the majority fall into a similarly narrow and symmetric distribution, but here we find a significant number of outliers, and more so on the negative side where true dispersion delays should lie. The potential extragalactic event described in detail in Kuiack et al. (2020b) is indicated with the red arrow.

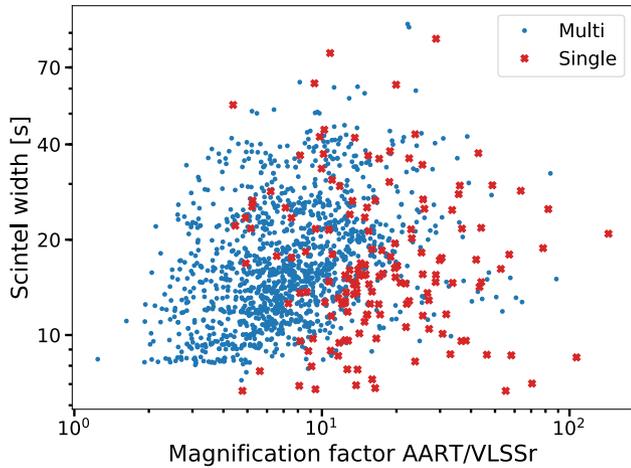


Figure 9. Duration versus magnification. The quantization of the width is due to our one second sampling time.

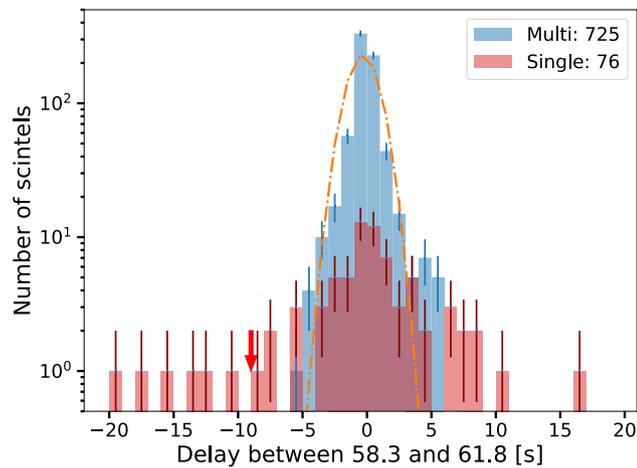


Figure 10. Distribution of peak delay between 61.8 and 58.3 MHz for 725 isolated scintels from 16 different sources on 2019-01-01 (blue), and 76. The red arrow denotes the delay of the transient candidate presented in Kuiack et al. (2020b). The delay is consistent with an extragalactic distance. The orange dash-dotted line shows a Gaussian distribution fit to the population.

4 DISCUSSION

As is clear, most of the phenomena we described above are not consistent with a distant astrophysical origin and must be caused by plasma lensing or other propagation effects relatively near to the Earth. We now discuss what constraints we can place on the medium that causes them.

4.1 Detailed light curves of single magnification events

Among the single magnification events, there are a few with very high signal to noise, and we consider these first in order to see what we can learn from the details of their light curves. The first light curve is that of 4C 26.55, shown in Fig. 2. The total DM along this Galaxy along this line of sight is 65 or 71 pc cm^{-3} according to the YMW16 (Yao, Manchester & Wang 2017) and NE2001 (Cordes & Lazio 2002) models, respectively. Using only this range, which would ignore any contribution from the intergalactic medium or the local environment would result in a dispersion delay of 12.3–13.3 s between 57.6 and

62.5 MHz. However, we measure no significant delay, and our earlier analysis (Fig. 3) showed a formal measured DM of 0 ± 2 . Hence, the brightness increase cannot originate in the source, but must be caused by plasma within a few tens of parsecs.

Our highest signal-to-noise magnification event peaks at 1000 Jy (Fig. 4). Again, the lack of measurable dispersion delay indicates the lensing plasma must be fairly nearby. A Bessel beam diffraction pattern, which matches this light curve, could be caused by a conic depression in a higher density, plane-parallel plasma sheet. (Since the index of refraction due to cold free electrons is close to, but below unity, a converging lens must be underdense and convex relative to its surroundings, or overdense and concave.) The somewhat depressed flux before and after the peak relative to the normal state may also suggest we see the ‘shadow’ of this lensing plasma there.

One type of magnification event known to occur for compact sources is the so-called extreme scattering event (e.g. Clegg, Fey & Lazio 1998), caused by isolated bubbles in the interstellar electron density. We do not think these are exactly the cause here because they would typically be further away and also could not cause the scintillation storms, in which the individual events look similar. They also have time-scales of months to years typically and amplitudes of only a factor few. Furthermore, Dong, Petropoulou & Giannios (2018) show that the lensed-to-unlensed brightness ratio tends strongly to one below 700 MHz, resulting from the increase in background source size.

We also note that these singular magnifications are similar in shape, time-scale, and magnitude (in the case of the burst in Fig. 4) to a recently reported low-frequency cosmic transient observed by LWA1 and LWA-SV (Varghese et al. 2019, esp. figs 3 and 4). While this was not their preferred explanation, they noted that it could be a magnification of the compact radio source 4C+1.06 by about a factor 30, a value on the high side, but within the range we find here; its duration of 15–20 s is also within our range. The observing frequency of 34 MHz is even lower than ours. If we may therefore regard this event as part of the same class the LWA gives important extra information: while they report only one event, and without dispersion information, the fact that they detected it with two stations 75 km apart is important. First, the absence of a noticeable parallax between the two stations indicates a minimum distance of 1400 km to the lens. LWA also report other events detected that they do attribute to scintillation (17 over a period of about 340 d). These are much less frequent than the scintillation storms we find, but that is likely due to the much higher flux threshold of LWA and the fact that they limited their analysis to cases where there was a simultaneous spike in the two stations. Due to this selection, we should probably be careful to attach too much significance to the simultaneity of these events in the two stations (to within about 5 s). Such simultaneity would imply a size of the magnification spot on the ground of more than 75 km, and a transverse speed of any lensing plasma of at least a few tens of km s^{-1} . However, if we look at their Fig. 4, we also see many more scintillation spikes in their example station light curves in the half hour before and after the simultaneous spike that are not coincident between the two stations, so smaller and/or slower magnification spots are perhaps more common. Before discussing a toy model for these magnifications, we will first look into what additional constraints follow from the strong scintillation zones.

4.2 Constraints from strong scintillation zones

To the extent that we may regard the scintillation zones or periods as the occurrence of many events like the isolated ones in a short period of time, they provide at least one extra constraint: we can see

them move in the sky and thus get their angular velocity. We use the example in Fig. 6, where the apparent motion is from 60° to 40° ecliptic latitude in 70 min, i.e. an angular velocity of $8 \times 10^{-5} \text{ rad s}^{-1}$. First of all, this angular velocity allows us to very much reduce the maximum distance of the lensing plasma from the dispersion delay limit of a few tens of pc since the physical transverse speed must be less than the speed of light, the distance is at most 20 au. Other than that, the scintillation zone is clearly not fixed either relative to the horizon or in ecliptic coordinates, so it cannot be obviously associated with either phenomena fixed to the Earth [such as the flux tubes discovered by Loi et al. (2015, 2016) with the MWA], or with something fixed in the Solar system. We do note a preference for the zones to appear in the anti-Sun half of the sky, though, perhaps suggesting a link with downstream effects of the solar wind impacting on the Earth's magnetosphere. To analyse this constraint in more detail, we must first note that the scintillation zones are tens of degrees across in the sky, but as we will show below, the individual lensing plasma bubbles are much smaller. However, scintillation studies done with LOFAR by Fallows et al. (2020) indicate that a fair approximation is that a scattering screen with turbulent cells moves past our line of sight rather faster than that it changes in its own rest frame, as in the case of interstellar scintillation. Therefore we will assume that the angular velocity we measure for the overall region also applies to the individual turbulent cells.

In a recently published study of a strong scintillation episode in Cas A with LOFAR, Fallows et al. (2020) perform a detailed analysis of dynamic spectra and of the time sequence in which the scintillation hits different LOFAR stations. They find that this episode is caused primarily by two layers of the ionosphere, the D region at around 80 km, where the plasma speed was about 110 m s^{-1} and the F region at 300–400 km altitude, where that speed was $20\text{--}40 \text{ m s}^{-1}$ at the time. Furthermore, during a 31-h survey with OVRO-LWA between 27 and 84 MHz Anderson et al. (2019) observed source flux increases up to 100s of Jy with time-scales on the order of 13 s (their highest time resolution), which is explained by a scattering screen at a distance of 300 Km, again in the F region of the ionosphere. The typical quiescent electron densities in these regions are $0.5\text{--}1 \times 10^6 \text{ cm}^{-3}$. For the angular velocity we measure, the implied speeds would be 6 m s^{-1} in the D region and $25\text{--}35 \text{ m s}^{-1}$ in the F region; so in case we see the F region, we find about the same speed, but were it the D region then the disturbance we see is considerably slower than the one found by Fallows et al. (2020). For even higher altitudes, as suggested by the lower limit for the LWA event, the required speeds go up: to 200 m s^{-1} at an altitude of 2000 km or even 1 km s^{-1} at 1.5 Earth radii in the geomagnetic tail.

The range of velocities, $100\text{--}1000 \text{ m s}^{-1}$ and altitudes ranging from 300 to 2000 km matches well with TIDs (e.g. Thome 1968; Huang, Andre & Sofko 1998), as does the trajectory southward from high latitude (Hocke & Schlegel 1996). The north-west to south-east orientation, and southwesterly drift is also characteristic of TIDs observed at mid-latitudes (Garcia et al. 2000). However, in the following hour the region remains fixed on the sky, along the 40° parallel. Thome (1968) also report a decrease in velocity from 700 to 125 m s^{-1} over 11 h. Lastly, direct air-glow images of TIDs tend to show multiple waves spanning the entire sky, passing with motion uncorrelated to the background sidereal drift.

Therefore, we conclude that the information we get from the motion of the scintillation regions provides the severest distance restriction and confines the location of the scattering or lensing screen to within 20 au. The motion is consistent with a variety of plausible ionospheric or geomagnetospheric origins and does not strongly favour one among these.

4.3 A toy plasma bubble model

We now try to constrain the properties of the plasma that causes both the isolated events and scintillation storms. To this end, let us first assume a spherical bubble of radius R with a difference in index of refraction δn relative to its surroundings, which is hit by a plane wave from our source, which then converges on AARTFAAC. In order to achieve a large magnification factor, the distance of the bubble to our telescope must be similar to the focal length f of the bubble. For $n \simeq 1$, this is well approximated by $f = R/|2\delta n|$, irrespective of whether the lens is thick or thin in the lens makers' sense. (If we think of it as a spherical bubble, it should be underdense, but we can also imagine an overdense concave body with both size and surface curvature radius about equal to R). Using the standard expressions for the index of refraction and plasma frequency, we derive the following numerical expressions:

$$|\delta n| = 2.24 \times 10^{-2} \delta n_{e,6} \nu_{60} \quad (2)$$

$$f = 223 R_6 \delta n_{e,6}^{-1} \nu_{60}^2 \text{ km}, \quad (3)$$

where $\delta n_{e,6}$ is the electron density in units of 10^6 cm^{-3} , R_6 the lens radius in units of 10 km, and $\nu_{60} = \nu/(60 \text{ MHz})$. Here, we have put some prejudice in the numbers by scaling them to values relevant for the quiescent ionosphere. For the ideal ray optics case, the magnification becomes infinite in the focus, but of course there it will always meet the diffraction limit. More realistically, AARTFAAC will cut through the beam some distance before or after the focus as the bubble drifts through the line of sight and throws a spot of radius r on to the ground. The magnification is then simply $(R/r)^2$. The formal diffraction limit on the spot size is then of course $r_{\min} = f\theta_D/2 = 0.305\lambda f/R = 0.61\lambda/\delta n$. Using the same numerical normalizations again, we find

$$r_{\min} = 0.136 \delta n_{e,6}^{-1} \nu_{60} \text{ km} \quad (4)$$

$$M_{\max} = 5.4 \times 10^3 R_6^2 \delta n_{e,6}^2 \nu_{60}^{-2} \quad (5)$$

$$\frac{\Delta f}{f} = 3.4 \times 10^{-2} R_6^{-2} \delta n_{e,6}^{-2} \nu_{60}^{-2}, \quad (6)$$

where M_{\max} gives the diffraction limit to the peak magnification, and the depth of field, Δf , is the range of focal length over which the spot size and magnification are close to the diffraction limit. This shows that these extreme numbers will be very rare even in the ideal case, and of course one should note that realistic plasma bubbles are quite unlikely to be so smooth as to act as ideal lenses, and so rather larger spots and more modest magnifications are likely to be the rule. But the observed magnification factors (Fig. 7) can at least be reassuringly far below that theoretical maximum. Of course the duration T of an event is related to its actual spot size as $T = 2r/\nu_L$, but unfortunately via the unknown transverse speed ν_L of the lens.

Another quantity that is accessible to observational testing is the angular size of the lenses on the sky, since large lenses would affect nearby sources in similar ways. The angular diameter on the sky is of course simply $\theta_L = 2R/f$, again approximating that AARTFAAC's distance to the lens is similar to the focal length, which is needed for a significant magnification. Numerically this works out to

$$\theta_L = \frac{2R}{f} = 4\delta n = 5.1^\circ \delta n_{e,6} \nu_{60}^{-2}. \quad (7)$$

This is quite sizeable and less dependent on parameters than some of the other quantities. Unfortunately, it is smaller than the typical distance between AARTFAAC catalogue sources on the sky, and so the test is non-trivial.

Now, we can use these equations with the properties of the observed events and of AARTFAAC to derive some general constraints on possible lenses. First, we note that a detectable magnification will require a spot that covers the AARTFAAC superterp, i.e. $r_{\min} > 300$ m. Via equation (4), this places a lower limit on the electron density perturbation: $\delta n_{e,6} < 0.45$. If we now use equation (3) and the requirement that $f \simeq h$, where h is the altitude of the lens, we get that

$$R_6 = \frac{h}{223 \text{ km}} \delta n_{e,6}. \quad (8)$$

Here and below, we set $\nu_{60} = 1$, all observations reported here are near 60 MHz. We can now set a constraint on the magnification by eliminating R_6 from it and using the electron density constraint:

$$M_{\max} = 5.4 \times 10^3 \left(\frac{h}{223 \text{ km}} \right)^2 \delta n_{e,6}^4 < 230 \left(\frac{h}{223 \text{ km}} \right)^2. \quad (9)$$

Since we observe magnifications up to 120, this implies that some of them must originate from a height at least 160 km, excluding the D region we mentioned previously. This exclusion is even stronger if we remind ourselves that we calculated the spot sizes and magnifications for ideal lenses, and realistically we therefore expect higher altitudes. Near the density peak of the F region, at perhaps 400 km, one could theoretically get magnifications of 750 and perhaps that makes the observed values possible even with somewhat non-ideal lenses. In this layer, the quiescent electron density is also highest, so that the maximum density perturbation we derive is only about half the total value. If one goes up to one Earth radius, the density has dropped by a factor 1000, more than undoing any favourable effect of a larger distance on the possible magnification. This is valid as long as we consider waves and turbulence that perturb the normal plasma values around the mean. If one considers events that change the electron density by orders of magnitude there would be more freedom.

We also note that for the spot sizes needed for large magnifications and the speeds we measure, the expected durations of the magnification events are a few seconds to a few minutes, which nicely encompasses the range we observed. From equation (7) we also find that the angular size of the lenses must be less than 2° . We find very few cases of sources closer than that on the sky that both scintillate, so with our current data we cannot test whether sources closer than that vary in sync because they are lensed by the same bubbles.

Hence indeed, taking all constraints together, we conclude that the layers around the ionospheric density peak at a height of about 300–400 km are a plausible location for the plasma that causes the spiky magnification events that we have found. A more common description of the radio scintillation caused by these layers considers the layer as a turbulent scattering screen that causes a Fresnel-zone near-field pattern of intensity variations on the ground. This approach is used in the study of scintillation of Cas A with LOFAR (Fallows et al. 2020). Quite possibly both descriptions apply simultaneously: Cas A is very bright, and a single LOFAR station can detect it, and its variations, continuously. Therefore, monitoring it can reveal the continuous pattern cast on the ground by the turbulent layer. We here see only the tip of the iceberg because all we detect is the occasional very large magnification due a fortuitously shaped turbulent cell and favourable geometry. We find that such events are much more common when the ionosphere is more active, but they can also occur in an otherwise quiescent ionosphere, and we find that the distributions of their properties are somewhat different then.

5 CONCLUSIONS

We have identified a new phenomenon in the low-frequency radio sky, namely strong and well-separated radio flashes, typically lasting 20 s and mostly coincident with bright radio sources that become 3–30 times brighter during the flash, but occasionally up to 100 times brighter. They typically show no significant dispersion delays and thus cannot originate from intrinsic variations in the extragalactic sources they coincide with, so we interpret them as magnification events caused by near-Earth plasma. Using a naive toy model of a plasma lens, we show that they are consistent with originating in the ionosphere around the layer of peak electron density, 300–400 km in altitude, so they may be the extreme end of ionospheric scintillation. They do occur both in times when the ionosphere is otherwise quiescent, albeit rarely, and they become quite frequent when the ionosphere is very active. At those times, which last from a few hours to more than a day, there are so many events that we can trace the development and motion of the active regions across the sky. In the future, with somewhat higher sensitivity and therefore a higher source density on the sky, we should be able to detect or exclude correlated variability of source that are less than 2° apart in the sky and thereby constrain the sizes and motion of the individual lensing plasma bubbles.

ACKNOWLEDGEMENTS

AARTFAAC development and construction was funded by the ERC under the Advanced Investigator grant no. 247295 awarded to RW, University of Amsterdam; This work was funded by the Netherlands Organisation for Scientific Research under grant no. 184.033.109. We thank The Netherlands Institute for Radio Astronomy (ASTRON) for support provided in carrying out the commissioning observations. AARTFAAC is maintained and operated jointly by ASTRON and the University of Amsterdam.

We would also like to thank the LOFAR science support for their assistance in obtaining and processing the data used in this work. We used data obtained from LOFAR, the LOFAR designed and constructed by ASTRON, which has facilities in several countries, that are owned by various parties (each with their own funding sources), and that are collectively operated by the International LOFAR Telescope (ILT) foundation under a joint scientific policy. We thank Harish Vedantham for useful discussions on the lensing mechanism and properties.

This research has used ASTROPY,⁴ a community-developed core PYTHON package for Astronomy (Astropy Collaboration 2013; Price-Whelan et al. 2018), as well as the following: KERN (Molenaar & Smirnov 2018), PANDAS (McKinney 2010), NUMPY (van der Walt, Colbert & Varoquaux 2011), and SCIPY (Jones et al. 2001). Accordingly, we would like to thank the scientific software development community, without whom this work would not be possible.

DATA AVAILABILITY

The data underlying this article can be shared on reasonable request to the corresponding author.

REFERENCES

- Anderson M. M. et al., 2019, *ApJ*, 886, 123
Astropy Collaboration, 2013, *A&A*, 558, A33

⁴<http://www.astropy.org>

- Bell M. E. et al., 2014, *MNRAS*, 438, 352
- Bell M. E. et al., 2019, *MNRAS*, 482, 2484
- Blagoveshchenskii D. V., 2013, *Geomagn. Aeronomy*, 53, 275
- Carbone D. et al., 2016, *MNRAS*, 459, 3161
- Clegg A. W., Fey A. L., Lazio T. J. W., 1998, *ApJ*, 496, 253
- Cordes J. M., Lazio T. J. W., 2002, preprint ([astro-ph/0207156](https://arxiv.org/abs/astro-ph/0207156))
- Danilov A. D., Konstantinova A. V., 2019, *Geomagn. Aeronomy*, 59, 554
- de Gasperin F., Mevius M., Rafferty D. A., Intema H. T., Fallows R. A., 2018, *A&A*, 615, A179
- Dong L., Petropoulou M., Giannios D., 2018, *MNRAS*, 481, 2685
- Fallows R. A. et al., 2020, *J. Space Weather Space Clim.*, 10, 10
- Garcia F. J., Kelley M. C., Makela J. J., Huang C. S., 2000, *J. Geophys. Res.*, 105, 18,407
- Helfand D. J., White R. L., Becker R. H., 2015, *ApJ*, 801, 26
- Hill R. J., Frehlich R. G., 1996, *Appl. Opt.*, 35, 986
- Hocke K., Schlegel K., 1996, *Ann. Geophys.*, 14, 917
- Huang C.-S., Andre D. A., Sofko G. J., 1998, *J. Geophys. Res.*, 103, 2131
- Jones E., Oliphant T., Peterson P. et al., 2001, SciPy: Open source scientific tools for Python. <http://www.scipy.org/>
- Kuiack M., Huizinga F., Molenaar G., Prasad P., Rowlinson A., Wijers R. A. M. J., 2018, *MNRAS*, 482, 2502
- Kuiack M. J., Wijers R. A. M. J., Rowlinson A., Shulevski A., Huizinga F., Molenaar G., Prasad P., 2020a, *MNRAS*, 497, 846
- Kuiack M., Wijers R. A. M. J., Shulevski A., Rowlinson A., Huizinga F., Molenaar G., Prasad P., 2020b, preprint ([arXiv:2003.13289](https://arxiv.org/abs/2003.13289))
- Lane W. M., Cotton W. D., van Velzen S., Clarke T. E., Kassim N. E., Helmboldt J. F., Lazio T. J. W., Cohen A. S., 2014, *MNRAS*, 440, 327
- Lazio T. J. W. et al., 2010, *AJ*, 140, 1995
- Loi S. T. et al., 2015, *Geophys. Res. Lett.*, 42, 3707
- Loi S. T. et al., 2016, *J. Geophys. Res.*, 121, 1569
- McKinney W., 2010, in van der Walt S., Millman J., eds, Proceedings of the 9th Python in Science Conference. p. 51
- Molenaar G., Smirnov O., 2018, *Astron. Comput.*, 24, 45
- Obenberger K. S. et al., 2015, *J. Astron. Instrum.*, 4, 1550004
- Prasad P. et al., 2016, *J. Astron. Instrum.*, 5, 1641008
- Prasad P., Wijnholds S. J., Huizinga F., Wijers R. A. M. J., 2014, *A&A*, 568, A48
- Price-Whelan A. M. et al., 2018, *AJ*, 156, 123
- Rowlinson A. et al., 2019, *Astron. Comput.*, 27, 111
- Stewart A. J. et al., 2016, *MNRAS*, 456, 2321
- Swinbank J. D. et al., 2015, *Astron. Comput.*, 11, 25
- Thome G., 1968, *J. Geophys. Res.*, 73, 6319
- Tu J., Song P., 2016, *J. Geophys. Res.*, 121, 11,861
- van der Walt S., Colbert S. C., Varoquaux G., 2011, *Comput. Sci. Eng.*, 13, 22
- 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
- Wijnholds S. J., van der Veen A.-J., 2009, *IEEE Trans. Signal Process.*, 57, 3512
- Yao J. M., Manchester R. N., Wang N., 2017, *ApJ*, 835, 29

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