



HAL
open science

Transient RFI environment of LOFAR-LBA at 72-75 MHz: Impact on ultra-widefield AARTFAAC Cosmic Explorer observations of the redshifted 21-cm signal

B.K Gehlot, L.V.E Koopmans, S.A Brackenhoff, E Ceccotti, S Ghosh, C Höfer, F.G Mertens, M Mevius, S Munshi, A.R Offringa, et al.

► To cite this version:

B.K Gehlot, L.V.E Koopmans, S.A Brackenhoff, E Ceccotti, S Ghosh, et al.. Transient RFI environment of LOFAR-LBA at 72-75 MHz: Impact on ultra-widefield AARTFAAC Cosmic Explorer observations of the redshifted 21-cm signal. *Astronomy and Astrophysics - A&A*, 2024, 681, pp.A71. 10.1051/0004-6361/202346376 . hal-04297322

HAL Id: hal-04297322

<https://hal.science/hal-04297322>

Submitted on 25 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Transient RFI environment of LOFAR-LBA at 72–75 MHz

Impact on ultra-widefield AARTFAAC Cosmic Explorer observations of the redshifted 21-cm signal[★]

B. K. Gehlot¹, L. V. E. Koopmans¹, S. A. Brackenhoff¹, E. Ceccotti¹, S. Ghosh¹, C. Höfer¹, F. G. Mertens², M. Mevius³, S. Munshi¹, A. R. Offringa^{3,1}, V. N. Pandey^{3,1}, A. Rowlinson^{4,3}, A. Shulevski^{3,1}, R. A. M. J. Wijers⁴, S. Yatawatta³, and S. Zaroubi^{5,1}

¹ Kapteyn Astronomical Institute, University of Groningen, PO Box 800, 9700 AV Groningen, The Netherlands
e-mail: kbharatgehlot@gmail.com

² LERMA, Observatoire de Paris, PSL Research University, CNRS, Sorbonne Université, 75014 Paris, France

³ ASTRON, Netherlands Institute for Radio Astronomy, Oude Hoogeveensedijk 4, 7991 PD Dwingeloo, The Netherlands

⁴ Anton Pannekoek Institute, University of Amsterdam, Postbus 94249, 1090 GE, Amsterdam, The Netherlands

⁵ Department of Natural Sciences, The Open University of Israel, 1 University Road, PO Box 808, Ra'anana 4353701, Israel

Received 10 March 2023 / Accepted 31 October 2023

ABSTRACT

Context. Measurement of the highly redshifted and faint 21-cm signal of neutral hydrogen from the Cosmic Dawn and Epoch of Reionisation promises to unveil a wealth of information about the astrophysical processes that governed the structure formation and evolution of the universe during the first billion years of its evolution.

Aims. The AARTFAAC Cosmic Explorer (ACE) program utilises the AARTFAAC wide-field imager of LOFAR to measure the power spectrum of the intensity fluctuations of the redshifted 21-cm signal from the Cosmic Dawn at $z \sim 18$ corresponding to the global 21-cm absorption feature reported by the EDGES experiment. Radio frequency interference (RFI) from various sources, such as aeroplane communication, contaminates the observed data and it is crucial to exclude the RFI-affected data in the analysis for any reliable detection. In this work, we solely focus on investigating the impact of non-ground-based transient RFI on the analysis of ACE observations.

Methods. We use cross-power spectra and cross-coherence metrics to assess the correlation of RFI over time and investigate the level of impact of transient RFI on the 21-cm signal power spectrum estimation.

Results. We detected moving sky-based transient RFI sources that cross the field of view within a few minutes and appear to be mainly from aeroplane communication beacons at the location of the LOFAR core in the 72–75 MHz band (a part of the EDGES absorption trough), by inspecting filtered images. We find that this transient RFI is mostly uncorrelated over time and is only expected to dominate over the thermal noise for an extremely deep integration time of 3000 h or more with a hypothetical instrument that is sky temperature dominated at 75 MHz. We find no visible correlation over different k -modes in Fourier space in the presence of noise for realistic thermal noise scenarios.

Conclusions. We conclude that the sky-based transient RFI from aeroplanes, satellites and meteorites at present does not pose a significant concern for the ACE analyses at the current level of sensitivity and after integrating over the available ~ 500 h of observed data. However, it is crucial to mitigate or filter such transient RFI for more sensitive experiments aiming for significantly deeper integration.

Key words. dark ages, reionization, first stars – methods: observational – methods: statistical – techniques: interferometric

1. Introduction

Observations of the redshifted 21-cm signal of neutral hydrogen from the Dark Ages (DA), Cosmic Dawn (CD), and Epoch of Reionisation (EoR) are expected to fill the gaps in our understanding of the cosmic structure formation (Furlanetto et al. 2006; Pritchard & Loeb 2012; Zaroubi 2013). Several operational (and upcoming) instruments and associated experiments, such as LOFAR (and AARTFAAC; van Haarlem et al. 2013; Prasad et al. 2016), MWA (Bowman et al. 2013), HERA (DeBoer et al. 2017), NenuFAR (Zarka et al. 2012; Mertens et al. 2021), SKA (Koopmans et al. 2015), LWA (Eastwood et al. 2019; Dilullo et al. 2020), EDGES

(Bowman et al. 2018), SARAS (Singh et al. 2017), LEDA (Bernardi et al. 2016), REACH (de Lera Acedo 2019), are seeking to statistically measure global and spatial properties of the faint redshifted 21-cm signal from CD and EoR. The final frontier is to measure the 21-cm signal from DA, requiring space-based instruments and instruments such as the planned FAR-SIDE (Burns et al. 2021), LuSEE (Bale et al. 2023), DAPPER¹, NCLE (Vecchio et al. 2021), and the ESA-led ALO missions.

However, successfully measuring this highly redshifted and faint 21-cm signal is an extremely challenging task. It requires ultra-precise calibration of the instrument and filtering of bright contaminations with a $10^5:1$ dynamic range

[★] Movies are available at <https://www.aanda.org>

¹ <https://www.colorado.edu/project/dark-ages-polarimeter-pathfinder/>

or higher. These contaminations include (i) spectrally smooth astrophysical foregrounds that are several orders of magnitude brighter than the faint 21-cm signal (Bernardi et al. 2009, 2010; Ghosh et al. 2012); (ii) chromatic effects of the instrument (commonly referred to as mode-mixing) caused by, for instance, the primary beam, uv -coverage, polarisation leakage, cross-talk between receiver elements, and cable reflections that modulate the smooth foregrounds and make it challenging to distinguish them from the signal of interest (Datta et al. 2010; Morales et al. 2012; Vedantham et al. 2012; Trott et al. 2012; Hazelton et al. 2013; Dillon et al. 2014); (iii) contamination effects in the data due to ionospheric phase errors (Mevius et al. 2016; Jordan et al. 2017); (iv) systematic errors due to imperfect processing of the data, for example, calibration and foreground removal (Barry et al. 2016; Patil et al. 2016; Ewall-Wice et al. 2017); and (v) radio frequency interference (RFI) due to human activities (Offringa et al. 2013, 2015, 2019; Whittle et al. 2019; Di Vruno et al. 2023).

Radio frequency interference is a crucial challenge in 21-cm cosmology experiments and comes from sources including, but not limited to, radio broadcasts such as FM, digital TV and audio broadcasts, satellite communication (geostationary and polar), aircraft communication, interference from electricity generation and transmission architecture such as wind and solar farms, and noisy transformers. Additionally, instruments with an ultra-wide field of view (FoV) such as AARTFAAC, LWA, and MWA can be significantly more prone to RFI originating at lower elevations, that is, those on the horizon or in the sky, than instruments with a narrower FoV (for example, LOFAR, HERA, and SKA).

We can broadly classify the RFI signatures from these sources into two categories based on their appearance in the data: (a) Stationary RFI sources: RFI signatures (whether narrowband or broadband) from a majority of sources such as FM, DTV, DAB, and geostationary satellites are usually stationary in time, frequency, and spatial location with relatively well-known frequency bands; (b) transient RFI sources: RFI signatures from polar and low-Earth orbit satellites, aircraft communication channels, drones, and reflections of radio signals off moving objects such as meteorites, etc. These may be considered transient RFI because of their spatially, temporally, and sometimes spectrally transient nature. Because of their transient nature, the impact of these RFI sources on the 21-cm signal has so far, to our knowledge, not yet been studied.

Frequency channels with stationary RFI contamination can be comparatively easily flagged during processing or avoided altogether during observations. Most RFI flagging algorithms use some form of outlier detection in post-correlation dynamic spectra and are designed to detect and flag stationary and transient (bright) RFI signatures (Middelberg 2006; Offringa et al. 2010, 2012; Prasad & Chengalur 2012; Peck & Fenech 2013). However, faint RFI below the noise level is difficult to detect, and techniques that leverage the spectral smoothness of the foregrounds and the Gaussian nature of instrumental noise are required for faint RFI rejection for a given dataset. Offringa et al. (2015) shows an approach to detect and flag faint RFI using baseline-integrated statistics (for example, the standard deviation of dynamic spectra along all baselines) and performing outlier detection. Wilensky et al. (2019) developed another technique, called “SSINS”, that removes the sky component in the visibility spectra by subtracting consecutive time samples and averaging multiple baselines incoherently to lower the noise, and then performing RFI detection on the sky subtracted and averaged spectrum. Other advanced techniques for RFI mitigation include compressive statistical sensing (Cucho-Padin et al.

2019), mitigation using polarisation information of interference signals (Yatawatta 2020), machine learning techniques (Vos et al. 2019), as well as pre-correlation interference mitigation techniques. Readers can refer to An et al. (2017) and Baan (2019), for example, for a review of pre- and post-correlation RFI mitigation techniques.

In this work, we perform an investigative study of the (largely) sky-based transient RFI environment of the LOFAR-AARTFAAC Low-Band Antenna (LBA) system in the frequency range of 72–75 MHz and its impact on the analyses to measure the power spectrum of the redshifted 21-cm signal at redshift $z \sim 18$. More specifically, we focus on the characterisation of the spatially and temporally transient RFI that might require different methods to detect and remove. We do not focus on mitigation techniques for such sky-based transient RFI and defer that to future analyses. For the RFI analysis, we use typical 15-min datasets from eight different nights that were also used to report the first 21-cm signal power spectrum limits from the AARTFAAC Cosmic Explorer (ACE) by Gehlot et al. (2020; referred to as G20 hereafter). We investigate the impact of transient RFI using Fourier space statistics such as their cross-coherence, cylindrical, and spherical power spectra.

2. The LOFAR-AARTFAAC Cosmic Explorer Programme

The EDGES collaboration (Bowman et al. 2018) reported a putative detection around 78 MHz of a 21-cm absorption feature, during CD, in the otherwise smooth frequency spectrum of the low-frequency radio sky. Such an absorption feature can result from the first generation of stars that appeared during the CD ($z \sim 18$), coupling the spin temperature of neutral hydrogen to the gas temperature. Since the latter is below the Cosmic Microwave Background (CMB) temperature, the 21-cm signal is observed in absorption against the CMB. This putative detection is unusual in the sense that it is unexpectedly deep and also wider than the most extreme scenarios predicted by standard astrophysical simulations. Hence, non-standard (“exotic”) models are required to explain the unusual nature of the reported detection. Several concerns were raised by Hills et al. (2018), Bradley et al. (2019) regarding the validity of the signal. Furthermore, the claimed detection was recently contested by the SARAS collaboration (Singh et al. 2022) whose observations are inconsistent with the EDGES global 21-cm absorption feature at a marginally significant level. However, if the EDGES detection is confirmed, it may unveil the presence of previously unknown fundamental physical processes such as the interaction between dark matter and baryonic matter, or the presence of a strong radio radiation background in addition to the CMB (Barkana 2018; Ewall-Wice et al. 2018; Dowell & Taylor 2018; Feng & Holder 2018; Fialkov & Barkana 2019; Reis et al. 2021). Additionally, a deeper global 21-cm absorption feature is expected to enhance the brightness-temperature fluctuations of the 21-cm signal from redshifts corresponding to the absorption trough (Bowman et al. 2018). This makes it possible to detect these brightness temperature fluctuations in a significantly shorter integration time. To test this hypothesis, in 2018 we commenced the ACE programme using the AARTFAAC-12 wide-field imager of the LOFAR telescope to observe the radio sky (for up to a thousand hours) between 72 and 75 MHz, within the absorption feature. The key science goal of the ACE programme is to measure or set the first limits on the power spectrum of the redshifted 21-cm signal from the Cosmic Dawn at $z \sim 18$ (G20)

and constrain/exclude non-standard models that could explain the unusual nature of the EDGES absorption feature.

2.1. The wide-field imaging mode

The Amsterdam Astron Radio Transients Facility and Analysis Center (AARTFAAC) is an independent system and observing mode that utilises the LOFAR in-field hardware (receivers and the signal chain) to perform full-sky imaging with the LOFAR-LBA dipoles² using an independent correlator. AARTFAAC can piggyback on ongoing LOFAR observations by tapping off the voltage streams from individual LBA dipoles (or HBA tiles) before beam-forming, which are then transported to a GPU-based correlator located at the Center for Information Technology of the University of Groningen. The AARTFAAC system has two operating modes, A6 where 288 receiver elements³ from the six innermost stations (“superterp”) of the LOFAR core are correlated, and A12 where 576 receiver elements from the 12 innermost stations are correlated. Correlating signals from individual LBA dipoles provide an instantaneous full-sky FoV which is optimal for large angular-scale 21-cm cosmology experiments. Although AARTFAAC can observe in the full LBA frequency range of 10–90 MHz, the instantaneous bandwidth is currently limited to 3.1 MHz⁴ due to the limited bandwidth of the signal transport chain. The correlator has an integration time of 1 s and produces XX , XY , YX , YY auto and cross-correlation products, where X and Y correspond to the orthogonal polarisations of the dual-polarisation receivers. Unlike standard tracking and beam-formed observations with LOFAR, AARTFAAC observes the sky in the drift-scan mode and stores the correlation products as raw binary files. Readers are referred to [van Haarlem et al. \(2013\)](#), [Prasad et al. \(2016\)](#), [Kuiack et al. \(2019\)](#), [Shulevski et al. \(2022\)](#) for more information about the LOFAR and AARTFAAC observing modes.

2.2. Observational setup

For the ACE observing campaign, we observed the northern sky in drift scan mode targeting the 72.36–75.09 MHz frequency range ($z = 17.9$ – 18.6) by placing 14 subbands (with 3 channels of 65 kHz width each) together. We use the remaining two subbands as frequency outriggers placed about 4 MHz apart on either side of the spectral window. The outrigger subbands can be used for quantification of frequency smooth foregrounds and systematics. Due to the large number of antenna elements, the AARTFAAC data volumes in A12 mode are enormous. Therefore, we limit ourselves to a frequency resolution of 65 kHz (3 channels per subband) and 1 s correlator integration time as a compromise between large data volume and the total number of frequency channels (42 channels for the targeted spectral window + 6 outrigger channels) that are sufficient for 21-cm power spectrum analyses. The data rate is about 0.84 TB per hour of observation with the A12 mode, for the correlator time and frequency resolution of 1 s and 65 kHz respectively. The ACE observing campaign lasted for three LOFAR observing cycles (10, 11 and 12) and around 500 h of data (between sunset and sunrise) were obtained. The recorded drift-scan data are split into

² For the HBA observations, the FoV $\sim 25^\circ$ at 150 MHz.

³ The receiver element can be either LBA dipoles or HBA tiles, with different user selectable configurations. LBA_OUTER is the default configuration for AARTFAAC observations.

⁴ The bandwidth is divided into 16 subbands of 195.3 kHz width and each subband can have a maximum of 64 channels of 3 kHz width.

Table 1. ACE observing campaign setup details.

Parameter	Value
Telescope	AARTFAAC (A12 mode)
Number of antennas	576 (dual polarisation)
Baseline range	3–1290 m
Frequency range	72.36–75.09 MHz
Outrigger frequencies	68.4 MHz and 79.1 MHz
Number of subbands	14 (target) + 2 (outrigger)
Channels per subband	3 (Total channels: 48)
Channel width	65 KHz
Correlator integration time	1 s
Observing mode	“Semi drift-scan”
Total data recorded	~ 500 h
Recorded data volume	~ 420 TB (uncompressed)

local sidereal time (LST) “slices” of 15 min duration and each LST bin is phased to the direction that transits the zenith halfway through the selected LST slice. We refer to this observing strategy as “semi drift-scan” mode (G20). Table 1 summarises the observing setup of the ACE program.

3. Observations and data reduction

In this work, we limit ourselves to eight observing blocks of 15 min duration spanning the LST range 23.5–23.75 h. This LST range corresponds to the LST-bin1 in the analyses presented in G20 and these observing blocks are referred to as “time-slices”. Each of these time-slices is processed with an updated version of the ACE pipeline as used by G20. The eight time-slices are processed as follows:

Phase tracking. We phased the drift-scan data to the direction $(\alpha, \delta) = (23\text{h } 37\text{m } 30\text{s}, +52\text{d } 37\text{m } 21.44\text{s})$ which transits the zenith halfway through the LST range 23.5–23.75 h. The phased data were then converted to a measurement set (MS) format table and stored. Phasing of the drift-scan data and conversion to MS is done using the AARTFAAC2MS⁵ package (Offringa et al. 2015).

Initial RFI flagging and time averaging. We used AOFLAGGER⁶ (Offringa et al. 2010, 2012) to detect and flag the RFI affected data. We have devised a customised flagging strategy for AARTFAAC data instead of using a generic flagging strategy. The strategy is specifically tuned to detect transient RFI in noisy data (key parameters: `base_threshold = 3.0`, `iterations = 5`, `transient_threshold_factor = 0.5`). This step flags the bright RFI per baseline that appears above the noise level. After flagging, the data are averaged to a lower time resolution of 4 s. Any averaged time-frequency blocks with missing samples are also flagged to avoid low-level artefacts due to flagging and averaging (Offringa et al. 2019). Any bright RFI that severely impacts further processing and analysis is flagged during this step. We note that the analysis of transient RFI in this paper is therefore the RFI that is not detected during this customised per baseline visibility-based flagging step.

Initial calibration. This step performs the initial direction-independent calibration of the flagged and averaged data to set the flux scale (for every frequency channel) and direction-independent phase. We used a sky model consisting of about

⁵ <https://github.com/aroffringa/aartfaac2ms>

⁶ <https://gitlab.com/aroffringa/aoflagger>

9990 compact sources extracted from VLSS (Kassim et al. 2007) and Cambridge radio catalogues (Laing et al. 1983) with a flux-density greater than 2 Jy at 74 MHz after extrapolating with the corresponding spectral models. The flux-density threshold we have used for the sky model is four times the classical confusion limit for the A12 system at 74 MHz. We also added the bright “A-team” sources, Cas A, Cyg A, and Tau A⁷ (with an angular extent of several arcminutes) to the sky model in addition to the 9990 sources. We used the `gaincal` task of the DP3 package⁸ (van Diepen et al. 2018) to perform the calibration. We used a solution interval of 65 kHz and 4 s (the same as the frequency and time resolution of the datasets after flagging and averaging) and excluded baselines shorter than 20λ from the calibration. The lower baseline cut helps mitigate two issues: first, our calibration model does not include large-scale diffuse emission, which is partly polarised and dominates these short baselines, and excluding short baselines from calibration thus avoids biases due to the missing diffuse component in the sky model. Second, the excluded baselines are dominated by intrastation baseline pairs which share the same electronic cabinets and might be affected by low-level cross-talk between dipoles. We note that the expected noise level for the 65 kHz and 4 s solution interval is about 2500 Jy (for an LBA-dipole SEFD of ~ 1.8 MJy), and a signal-to-noise ratio of about 36 per solution and baseline for the given calibration model. This step enables the detection of outlier antennas, baselines, and times, as discussed in the next step, and the data are re-calibrated (discussed in the following steps) after flagging these bad antennas/baselines to improve the overall quality of the calibration.

Statistic based visibility flagging. The AOFLAGGER software also calculates various quantities: RFI percentage, mean, and standard deviation of visibilities (and their frequency difference) along different axes such as time, frequency, and baselines. These statistics were calculated using the data after the initial calibration in the previous step and can be used to identify outlier antennas, baselines, times, and frequency channels and flag them. We use a robust standard deviation estimate, $\sigma \approx 1.4826$ MAD (median absolute deviation), to flag outlier baselines (with 6σ threshold for RFI percentage and 5σ for the rest of the statistics) and times (4σ threshold for all statistics).

Bandpass calibration. After flagging visibility outliers, we performed a bandpass calibration on the uncalibrated data after flagging and averaging with additional flags from the previous step applied. We used the same sky model as in the initial calibration for the bandpass calibration but with a 15 min solution interval instead of 4 s and a single gain solution was obtained per 65 kHz frequency channel. This step allows us to capture rapid direction-independent (thus independent of the 21-cm signal) spectral variations in the global bandpass (time-averaged) of the antenna elements and correct for these variations.

Direction-independent (DI) calibration. After correcting the visibilities for the bandpass shape, a direction-independent calibration step was performed. This step is similar to initial calibration and uses the same calibration model and settings, that is, a solution interval of 4 s and 65 kHz, and the same baseline cut of 20λ . Repeating the DI-calibration step after flagging bad anten-

nas improves the overall quality of the calibration. This step corrects for the direction-independent effects on both short time and frequency scales.

Direction-dependent (DD) calibration. The next data-processing step was to remove the dominant sources Cas A and Cyg A, each with a flux-density of about 17 kJy (at 74 MHz), from the calibrated data. We used the `DDECal`⁹ task of DP3 to perform the direction-dependent calibration, and obtain the direction-dependent gains towards these two sources. These gains are then used to subtract the sky model components of the two sources from the calibrated visibilities. For this step, the sky model (the same as in the DI calibration) is split in 3 directions, that is, Cas A, Cyg A, and “main”, where the main direction contains all other sources in the sky model. Only Cas A and Cyg A are subtracted from the visibilities. We did not calibrate towards the direction of Tau A because of the insufficient signal-to-noise. The time and frequency solution intervals were kept the same as during the DI calibration. We excluded baselines shorter than 60λ in the calibration and performed a constrained optimisation in which the direction-dependent gains are constrained to be smooth on a 3 MHz scale. We used a different baseline cut in direction-dependent calibration to avoid overlap between the baselines used for calibration and power spectrum analysis (20 – 60λ). These settings are used to avoid over-fitting during the direction-dependent calibration and to perform an accurate subtraction of the bright sources without affecting the rest of the emission (whether compact or diffuse) on baselines shorter than 60λ (Mouri Sardarabadi & Koopmans 2019; Mevius et al. 2022).

Post-calibration outlier detection and flagging. After subtracting Cas A and Cyg A, we ran AOFLAGGER once more on the subtracted data and also repeated the statistics-based flagging (with the same settings used previously) to identify and flag the outlier visibilities, baseline pairs, and times that are badly calibrated or behave differently from the rest. After the data-processing steps, each night have typically around 20–30% flagged data in the baseline range of 20 – 60λ , except two nights for which about 35 and 40% of data, respectively, are flagged.

Imaging. After the data are calibrated and visibility outliers are flagged, they are ready to be imaged. We used WSCLEAN¹⁰ (Offringa et al. 2014; Offringa & Smirnov 2017) to image the visibilities. We used the `w-griddler` algorithm (Arras et al. 2021; Ye et al. 2022) implementation in WSCLEAN to grid and image the data with the griddler-accuracy of 10^{-4} . The `w-griddler` algorithm dynamically sets gridding kernel parameters (kernel size, oversampling, number of w-layers) for a given griddler-accuracy to achieve optimal and high-fidelity imaging. We do not perform any deconvolution on the images. We used only the 20 – 60λ baselines for imaging because the shortest baselines may be affected by low-level cross-talk and the uv -sampling becomes relatively sparse for baselines longer than 60λ (about 250 m). The longer baselines are also used during the DD-calibration step, making the excluded baselines non-optimal for 21-cm cosmology analyses. Every 1-min in the LST interval and 65 kHz channel was imaged separately such that 15 image cubes (with 42 channels each) were produced for each of the eight time-slices. The key imaging parameters are listed in Table 2.

⁷ The models of these sources consist of delta functions and Gaussians to describe compact and extended emission, respectively. These models were obtained using LOFAR-LBA observations and are a part of the list of calibrator sources in the LOFAR Initial Calibration (LINC) pipeline <https://git.astron.nl/RD/LINC>

⁸ <https://git.astron.nl/RD/DP3>

⁹ Readers are referred to Gan et al. (2023) for a description of the algorithm.

¹⁰ <https://gitlab.com/aroffringa/wsclean>

Table 2. Imaging parameters for WSCLEAN.

Parameter	Value
Baselines	20–60 λ
Pixel size	12 arcmin
Number of pixels	576 \times 576
Gridder	<i>w</i> -gridder
Gridding kernel	Modified ES kernel
Gridder accuracy	10 ⁻⁴ (default)
Weighting	Natural
<i>uv</i> -pixel size	0.5 λ \times 0.5 λ

4. Detection of transient RFI

Although our flagging strategy is very effective in removing RFI that is stationary and above the thermal noise, the impact of transient RFI has never been studied in detail to the best of our knowledge. To investigate the impact of transient RFI on 21-cm cosmology analyses, the RFI in each spectral image cube needs to be highlighted and extracted. To achieve this, we produce a 3D mask for the pixels that are dominated by spectrally smooth emission associated with foregrounds, sidelobe structure due to instrumental chromaticity, and spectrally uncorrelated noise. This 3D mask is then applied to the corresponding image cube to highlight the transient RFI and mask the pixels dominated by the foregrounds and the noise. We follow the steps listed below to produce the mask and residual cubes:

First, we subtracted the continuum emission (multi-frequency synthesis or MFS image produced with all channels combined) in a given 1-min cadence image cube. Since the A12 system has a uniform and dense *uv*-sampling in the 20–60 λ baseline range, the PSF does not change significantly over the 2.7 MHz frequency band. Thus, the residual cube, to the first order, contains any emission significantly deviating from the pixel-average continuum emission, frequency-dependent PSF sidelobes, and noise. This emission is expected to be dominated by non-stationary sources.

For every line of sight in the residual image cube, we selected voxels with the highest (F_{\max}) and the lowest flux density (F_{\min}) along the frequency direction and added them in quadrature, that is, $F_{\text{total}} = \sqrt{0.5(F_{\max}^2 + F_{\min}^2)}$, to obtain a spatial map of outliers. This method highlights the bright positive and negative outliers without enhancing the noise-dominated voxels. Based on qualitative inspection of RFI spatial maps, we found that this metric performs better in terms of separating residual foregrounds and bright (narrow-band) outliers compared to the traditional sigma thresholding method using standard deviation which seems to falsely detect foregrounds as transient RFI. We excluded the directions more than 85° away from the phase centre, to avoid false positives due to projection-based effects and proximity to the horizon in the mask creation process. We note that this mask excludes transient RFI from very close to the horizon from the analysis. We also masked the 3° region around the bright sources Cas A, Cyg A, and Tau A to avoid any false detection of RFI, due to calibration and imaging artefacts. An example of a spatial map is shown in Fig. 1. In essence, this map highlights the combined level of deviations of the highest peaks and deepest valleys along the frequency axis for every line of sight. We also provide [online](#) movies of the spatial maps looping over 1-min duration timesteps, for all eight LST time-slices used in the analysis.

Next, we used the spatial map produced in the previous step and selected the pixels that are 12 σ (robust standard deviation estimator based on median absolute deviation) higher than the median of all pixels of the spatial map. The selected pixels correspond to the location of the RFI source(s) in the sky. For every such pixel, we performed an outlier detection with more than 6 σ deviation from the median along the frequency direction and masked all pixels with a non-detection (including the region <5° above the horizon). This provides a 3D mask for a given image cube where everything is masked except for the transient, narrow-band RFI peaks and tracks localised in image space. We note that this step only selects narrow-band transient RFI that remains undetected by any RFI detection and flagging steps. Although the RFI, irrespective of its nature, could impact the 21-cm power spectrum measurement, we expect the bright RFI (whether narrow-band or broadband) also to affect the calibration process, resulting in bad calibration. Such badly calibrated data are likely to be rejected by flagging steps that are performed post-calibration.

We used the final 3D mask to set all the masked pixels that were dominated by the sky emission and noise to zero in the spectral image cube, leaving only the transient RFI structures. The RFI detection process was performed on Stokes *I* image cubes at a 1 min time cadence for all 8 observations.

Figure 2 shows Stokes *I* images of extracted RFI for the observations corresponding to the eight nights. For the sake of visualisation, we have added the RFI images for all timesteps and frequency channels to produce a combined Stokes *I* RFI map for each night. For a single timestep, the majority of the transient RFI peaks that we observe have a flux density between 10 and 100 Jy/PSF. The filtered Stokes *I* RFI cubes are subsequently used to investigate the impact of this transient RFI on the ultra-widefield 21-cm cosmology experiment observations as described in Sect. 2.

5. Impact of transient RFI

In this section, we discuss the impact of the transient RFI that we observed in the data on 21-cm power spectrum analysis. We use the `ps_eor`¹¹ package to estimate the power spectra and cross-coherence of the RFI image cubes produced in the previous section. This package is extensively used for LOFAR-EoR (Mertens et al. 2020) and ACE power spectrum analyses (G20). Power spectrum calculations involve the following steps: trimming image cubes to a user-provided FoV, applying a spatial taper to the trimmed image cubes, converting image flux to temperature units, spatial Fourier transforming the trimmed and tapered image cubes to obtain gridded visibility cubes, Fourier transforming the gridded visibilities along frequency, and averaging the data in Fourier space to obtain cylindrical/spherical power spectra. We follow the above-mentioned process for most of the analysis unless specified otherwise.

5.1. RFI correlation over time

The RFI contamination, regardless of its properties is of serious concern in 21-cm experiments. If the RFI is bright enough, it is mostly flagged by the RFI detection algorithms. On the other hand, the impact of residual transient and non-stationary RFI that stays below the thermal noise level and remains undetected by traditional RFI detection methods is less understood. First, we investigate whether the transient RFI that we observed

¹¹ https://gitlab.com/flomertens/ps_eor

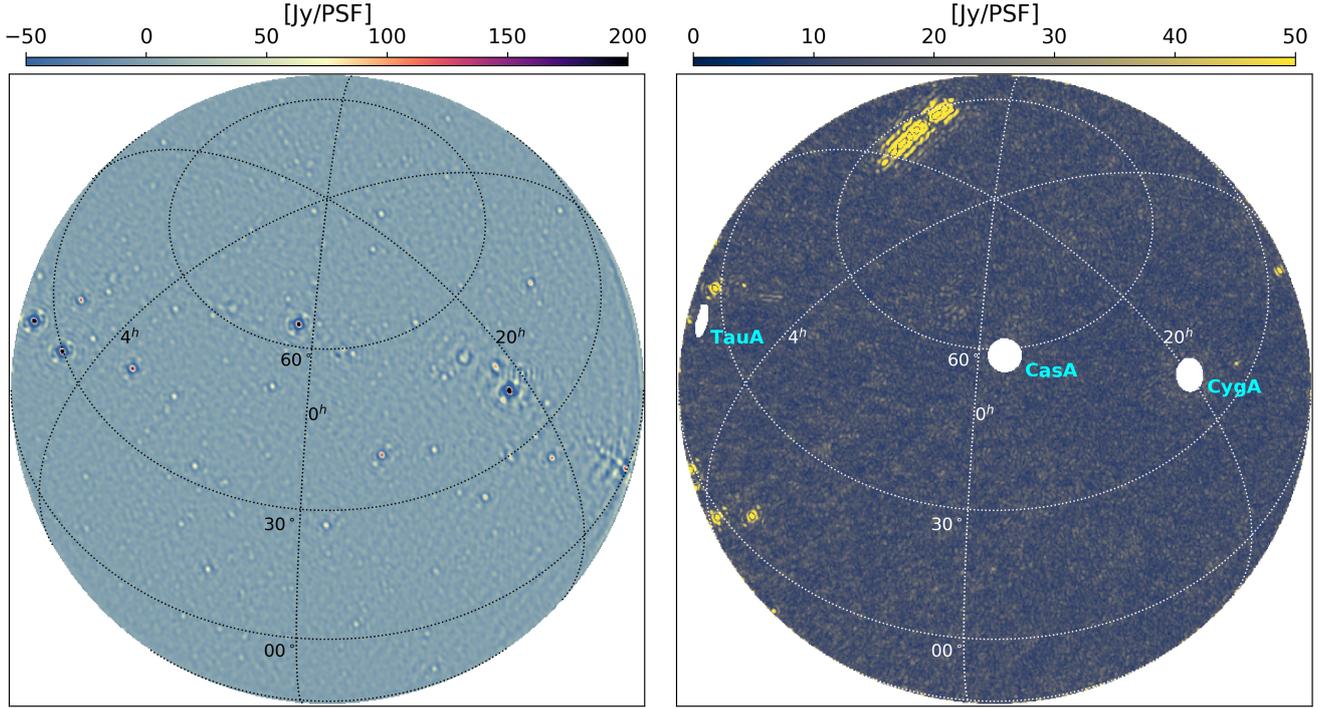


Fig. 1. Images produced with a subset of ACE data. The left panel shows a dirty continuum image of the sky produced with a one-minute subset of data from one of the nights used in the analysis. The colour scale of the image has been saturated to highlight a wide range of fluxes. The right panel shows a spatial map highlighting the transient RFI over the one-minute duration, extracted from the same subset of data used to produce the continuum image in the left panel. Blank pixels in the panel correspond to the bright A-team sources, Cas A, Cyg A, and Tau A. Additionally, the 5° region above the horizon has been excluded from both images. The bright pixels are narrow-band emissions deviating from the continuum emission shown in the left panel and are likely to be from the transient RFI. The stretched structure in the top part of the spatial map is likely an aeroplane and moves slowly over the sky (several minutes across the field).

over different nights correlate with the LST, and if so, at what level? This is a vital test because any contamination that correlates over time is expected to impact the power spectrum and it needs to be mitigated to avoid any biases in the final deep integration power spectra and their interpretation. We use the cross-coherence metric in Fourier space to study the temporal correlation of the transient RFI. For a given pair of gridded visibility cubes $\tilde{T}_i(\mathbf{k})$ and $\tilde{T}_j(\mathbf{k})$ in Fourier space, we define the 2D cross-coherence, $C_{ij}(k_\perp, k_\parallel)$, as:

$$C_{ij}(k_\perp, k_\parallel) = \frac{\langle |\tilde{T}_i^*(\mathbf{k})\tilde{T}_j(\mathbf{k})| \rangle^2}{\langle |\tilde{T}_i(\mathbf{k})|^2 \rangle \langle |\tilde{T}_j(\mathbf{k})|^2 \rangle}. \quad (1)$$

The $C_{ij}(k_\perp, k_\parallel)$ can have a value between 0 and 1 that corresponds to no correlation (zero) or maximum correlation (one) because it is a normalised quantity. This definition is similar to the one used by Mertens et al. (2020). We used only the real component of the cross-power spectrum to calculate the cross-coherence instead of using the complex values. The averaging was performed over all (k_\perp, k_\parallel) modes in the range $0.14 \lesssim k \lesssim 2.83$ and a single C_{ij} value was obtained for each pair of image cubes. We used all 1-min image cubes for all 8 nights and calculated their cross-coherences. We trimmed the image cubes to an FoV angle diameter of 120° (or π sr¹² in terms of solid angle footprint), and applied a spatial Hann taper (Blackman & Tukey 1958) before Fourier transforming the images, to avoid stationary RFI on the ground and horizon (which is easier to mitigate and not a topic of this study on transient sky-based RFI).

¹² We calculate the solid angle as $\Omega = 2\pi(1 - \cos \theta/2)$, where θ is the angle diameter.

We note that the Fourier Transform of full-sky (curved) images under the flat-sky approximation can lead to geometric effects, but we since perform a relative comparison of nights with and without noise, this assumption does not affect our findings and conclusions. In forthcoming analyses, we plan to investigate the application of spherical harmonics-based inversion methods (Ghosh et al. 2018) for the ACE 21-cm power spectrum measurement.

Figure 3 shows the average cross-coherence of transient RFI in the absence of noise. We find that for most image-cube pairs the coherence is below 0.01 and only a tiny fraction of pairs, about 7.5% show coherence higher than 0.01 (and about 1.7% of pairs show coherence greater than 0.02). The image cube pairs with high coherence and short temporal scale are likely to be associated with transient RFI that is largely incoherent over time. However, we do observe slightly increased coherence (albeit lower than 0.02) of 1-min image cube pairs corresponding to the same LST. This manifests itself as diagonal structures in Fig. 3 in some pairs of nights. We suspect that this behaviour could be due to some low-level correlated foreground (and their side-lobes) and noise components that remain in the filtered images (overlapping with the pixels with transient RFI) being treated as transient RFI.

Next, we introduced three different thermal noise levels to the gridded visibilities of every 1 min image cube to assess the correlation of transient RFI in the presence of thermal noise. The thermal noise rms per gridded visibility cell (σ_{grid}) can be calculated as (Thompson et al. 2001):

$$\sigma_{\text{grid}} = \frac{\text{SEFD}}{N_{\text{vis}} \sqrt{2\Delta\nu\Delta t}}, \quad (2)$$

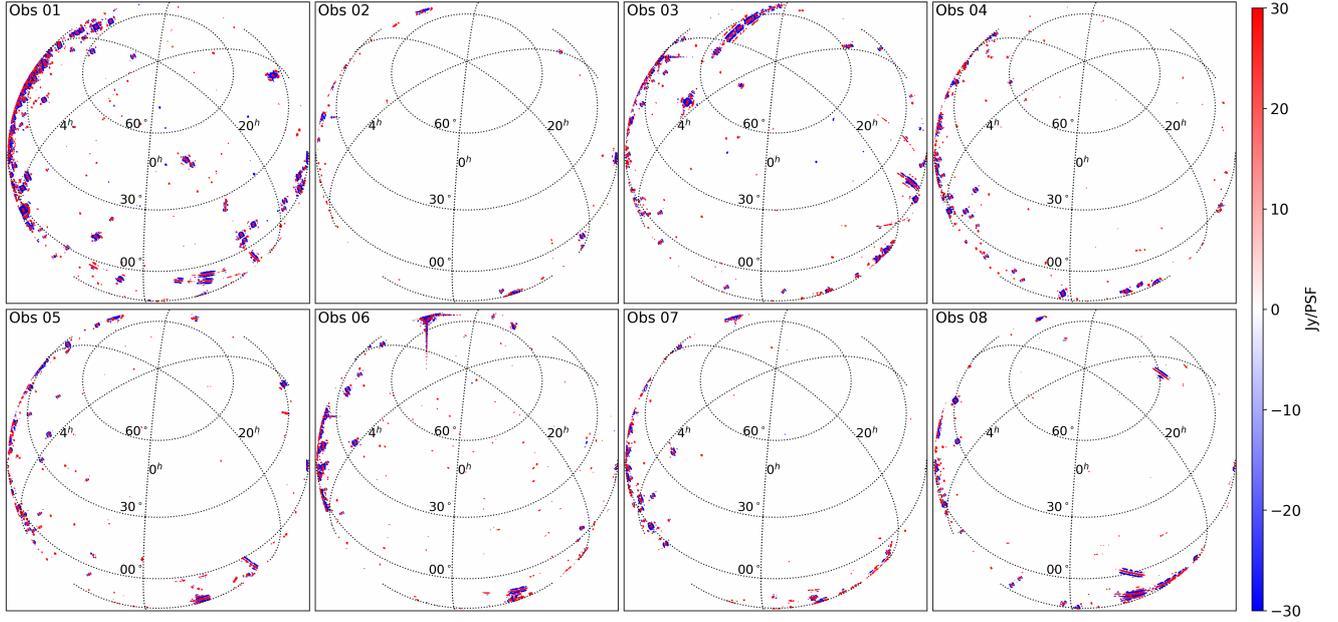


Fig. 2. Stokes I images of the transient RFI detected using the method described in Sect. 4. For visualisation purposes, RFI from each channel and time step (1 min) is plotted on the same panel for every night. Each panel corresponds to a different night used in the analysis. Similar to Fig. 1, the 5° region above the horizon has been excluded from the images.

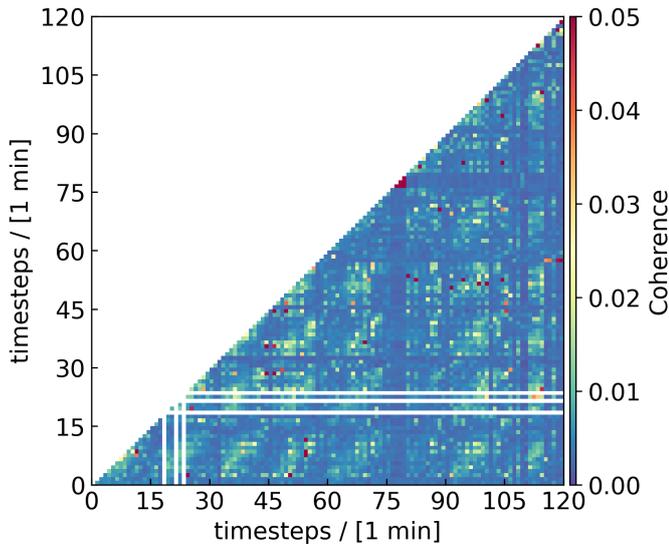


Fig. 3. Average cross-coherence of all image-cube pairs in the absence of thermal noise. The grid lines separate timesteps corresponding to different nights. The blank pixels (between 15 and 30) belong to timesteps that do not have any detected RFI.

where SEFD is the system equivalent flux density, N_{vis} is the number of visibilities in a uv -cell, $\Delta\nu = 65$ kHz is the channel width, $\Delta t = 4$ s is the integration time of a single visibility. To obtain the thermal noise level for deep integrations, we scaled down σ_{grid} by multiplying it by the factor $(t_{\text{image}}/t_{\text{obs}})^{1/2}$, where $t_{\text{image}} = 1$ min (observing time of a single data-cube), and t_{obs} is the total observing time per data-cube for deep integration. We used three different observing times to calculate the thermal noise, $t_{\text{obs}} = 1$ min, 4.17 h, and 25 h per data cube which is equivalent to the total deep integration (t_{int}) of 2 h, 500 h, and 3000 h, respectively, for the full dataset containing 120 such data cubes. Independent noise realisations were generated for each

Table 3. Thermal noise calculation parameters.

SEFD	Integration time	
	Single data-cube (t_{obs})	Full dataset ($120 \times t_{\text{obs}}$)
1.8 MJy	1 min	2 h
1.8 MJy	4.17 h	500 h
900 kJy	25 h	3000 h

data cube and integration time and added to the corresponding gridded visibilities. For the first two cases, we chose the SEFD of 1.8 MJy which is approximately equal to the observed sensitivity of AARTFAAC-12 at 74 MHz. However, for the third case, we chose a hypothetical instrument twice as sensitive as AARTFAAC-12 and a SEFD of 900 kJy¹³. Table 3 summarises the parameters used to calculate the thermal noise.

Figure 4 shows the cross coherence in the presence of the thermal noise with different levels as described earlier. We observe that the average cross-coherence of transient RFI in the presence of thermal noise virtually disappears for the realistic noise level (1-min integration time for a 1-min datacube) and it is visually identical to the cross-coherence of the thermal noise itself. We start to observe some coherence between a small fraction of image cube pairs for the low thermal noise scenarios. We note that the low-level structures that we observe in Fig. 3 disappear in the presence of thermal noise because the unmasked region is a tiny fraction of the sky and the low-level correlated component in the unmasked region is subdominant in the cross-coherence compared to the simulated noise.

We also perform a two-sample Kolmogorov–Smirnov test to quantify how the transient RFI coherence distribution, in the presence of thermal noise, compares with the coherence

¹³ Assuming $T_{\text{sys}} \approx T_{\text{sky}} = 50\lambda^{2.56}$ K, this is the minimum achievable SEFD for a sky temperature dominated dipole receiver with the effective area ($A_{\text{eff}} \approx 5.33$ m²) same as of a LOFAR-LBA dipole at 75 MHz.

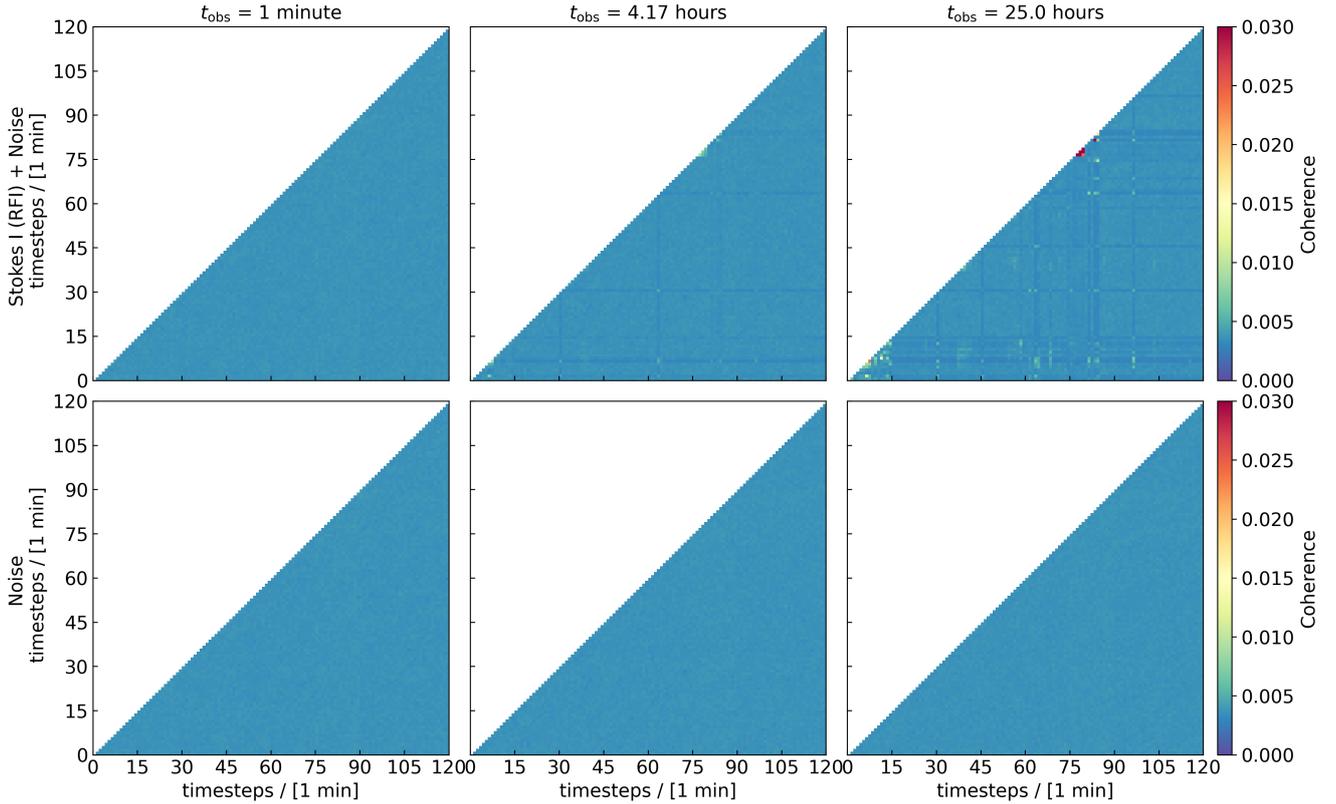


Fig. 4. Average cross-coherence of all image-cube pairs. The top row shows cross-coherence of transient RFI in Stokes I images in the presence of thermal noise with different levels (left to right columns: $t_{\text{obs}} = 1$ min, 4.17 h, 25 h), and the bottom row shows the cross-coherence of the corresponding thermal noise levels. The transient RFI coherence in the presence of thermal noise is visually indiscernible from the noise coherence itself.

distribution of the thermal noise. Figure 5 shows the coherence distributions of the transient RFI in the presence of thermal noise and the thermal noise itself. The p -values for the three thermal noise scenarios are 1.0, 0.0, and 0.0, respectively. This suggests that in the realistic thermal noise scenario with an integration time of 1 min per image-cube, transient RFI coherence is indiscernible from the thermal noise coherence for the AARTFAAC-12 instrument. For the integration time of 4.17 h per image cube (equivalent to the current total integration time in hand per 1 out of 120 image cubes), the transient RFI coherence distribution becomes slightly skewed, where a tiny fraction of image-cube pairs show slightly higher coherence but still lower than 0.01. However, for a hypothetical instrument twice as sensitive as AARTFAAC-12, the transient RFI starts to appear above the thermal noise for 25 h of integration per image cube. Still, the transient RFI observed is slow enough that it is detectable (directly or via the PSF sidelobes) by integrating over 1-min time scales. Therefore, the sky-based masking could also be used to filter this transient RFI from the gridded visibilities before foreground removal and power spectrum estimation using Gaussian Process Regression (Mertens et al. 2018, 2024) to avoid any low-level biases in the 21-cm signal measurement with extremely sensitive instruments. However, such a strategy to mitigate the transient RFI will be investigated in detail in future work.

5.2. Effect of a spatial taper

The next part of the study is to investigate the effect of a spatial taper and trimming in the image domain on the transient RFI power in Fourier space. For this test, we use two types

of spatial tapering functions, a Hann window and a Tukey or a tapered cosine window (Harris 1978) with taper parameter $\alpha = 0.2$ which means 80% inner region of the image is covered by flat-top of the window and the remaining 20% region is cosine tapered. Applying a Hann window reduces the effective fraction of the image to 25%, whereas the Tukey window with $\alpha = 0.2$ reduces the effective fraction to 81% of the image. We also use three different FoV trims, full-sky (2π sr), 150° (1.48π sr), and 120° (π sr) diameter (same as in G20). We use the cross-coherence metric (Eq. (1)) to quantify the impact of the taper and trimming of images on the behaviour of transient RFI in Fourier space.

We select two nights that show significant transient RFI levels, that is, nights corresponding to the first and third panels of the top row in Fig. 2 and calculate cross-coherence for all combinations of 1 min image-cubes from the two nights. These two nights represent the worst-case scenarios in the set of nights we have used for this work. Figure 6 shows the transient RFI cross-coherence for Hann (top row) and Tukey tapers (bottom row) with different image trimming levels (different columns). We observe that the Tukey taper with $\alpha = 0.2$ has an overall worse performance compared to the Hann window regardless of the trimming used. This is mainly due to the decreased sensitivity towards lower elevation caused by the Hann window but with a trade-off of sky sensitivity compared to a Tukey ($\alpha = 0.2$) taper. For both Hann and Tukey windows, decreasing the FoV leads to a decrease in overall coherence. This is expected as excluding the regions near the horizon showing high transient RFI reduces the overall contamination in the field. A low-level leakage is still expected due to side lobes.

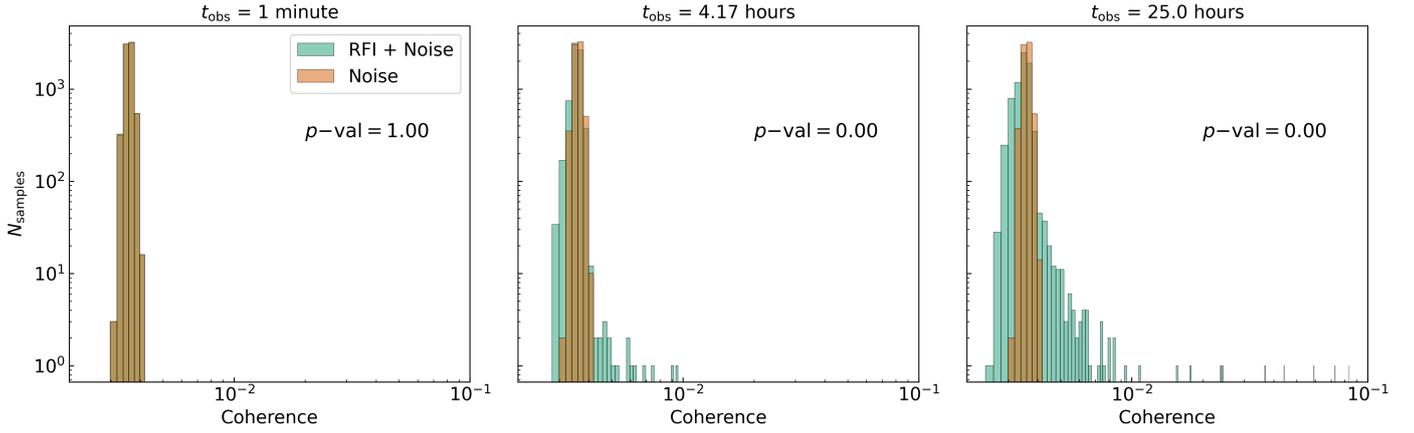


Fig. 5. Distribution of average cross-coherences of transient RFI in the presence of thermal noise (shown in green) and cross-coherences of thermal noise (shown in orange). The left, middle, and right panels correspond to different integration times per image-cube, $t_{\text{obs}} = 1$ min, 4.17 h, and 25 h, respectively.

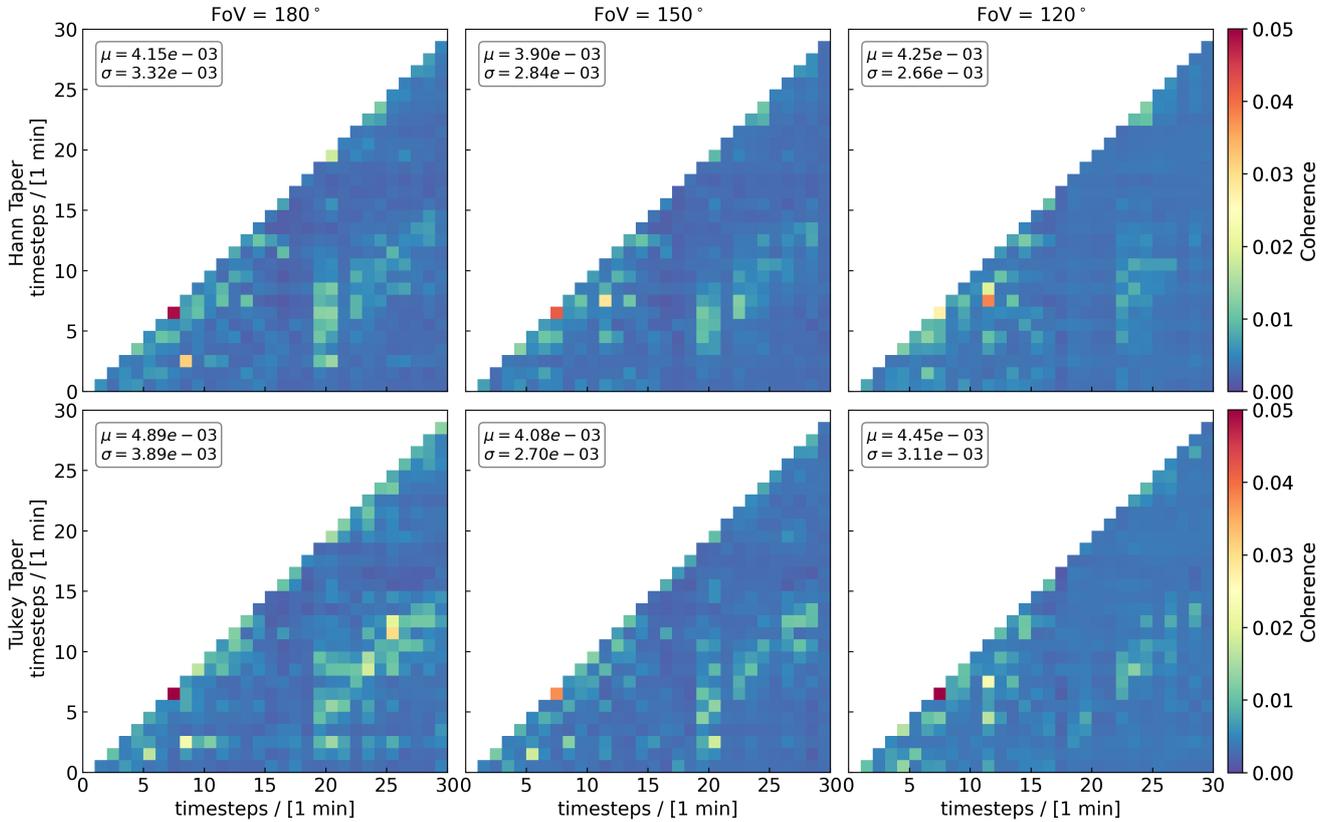


Fig. 6. Cross-coherence for all image-cube pairs for the two nights described in Sect. 5.2. The top and bottom rows correspond to Hann and Tukey taper, respectively. Columns from left to right correspond to FoVs of 180° (no trimming), 150°, and 120° (same as in G20), respectively. The dotted horizontal and vertical lines separate the two nights. The legend in each panel shows the mean and standard deviation of the coherences.

However, it does not correlate over time because the transient RFI is expected to be uncorrelated. Additionally, the average coherences for the different cases are similar ($\sim 4 \times 10^{-3}$). Whereas, the standard deviation of the coherence values decreases as the FoV is decreased and is slightly lower for the Hann taper compared to that of the Tukey taper. From the findings above, we conclude that using a Hann taper and a FoV of 120° is a better choice for ACE analysis compared to the Tukey ($\alpha = 0.2$) taper to reduce the contamination from transient RFI at low elevations. This was our choice of the taper used for the analysis presented in G20.

5.3. Power spectra of transient RFI

In 21-cm cosmology, cylindrically and spherically averaged power spectra are the most preferred and widely used statistics for diagnostic purposes, investigating the impact of various contaminations, as well as the measurement of the redshifted 21-cm signal from the CD and EoR. Therefore, it is useful to assess the impact of the transient RFI on these two power spectra. We use the standard definition of the cosmological power spectra (Morales & Hewitt 2004; Morales et al. 2019) to compute cylindrically and spherically averaged power spectra for each

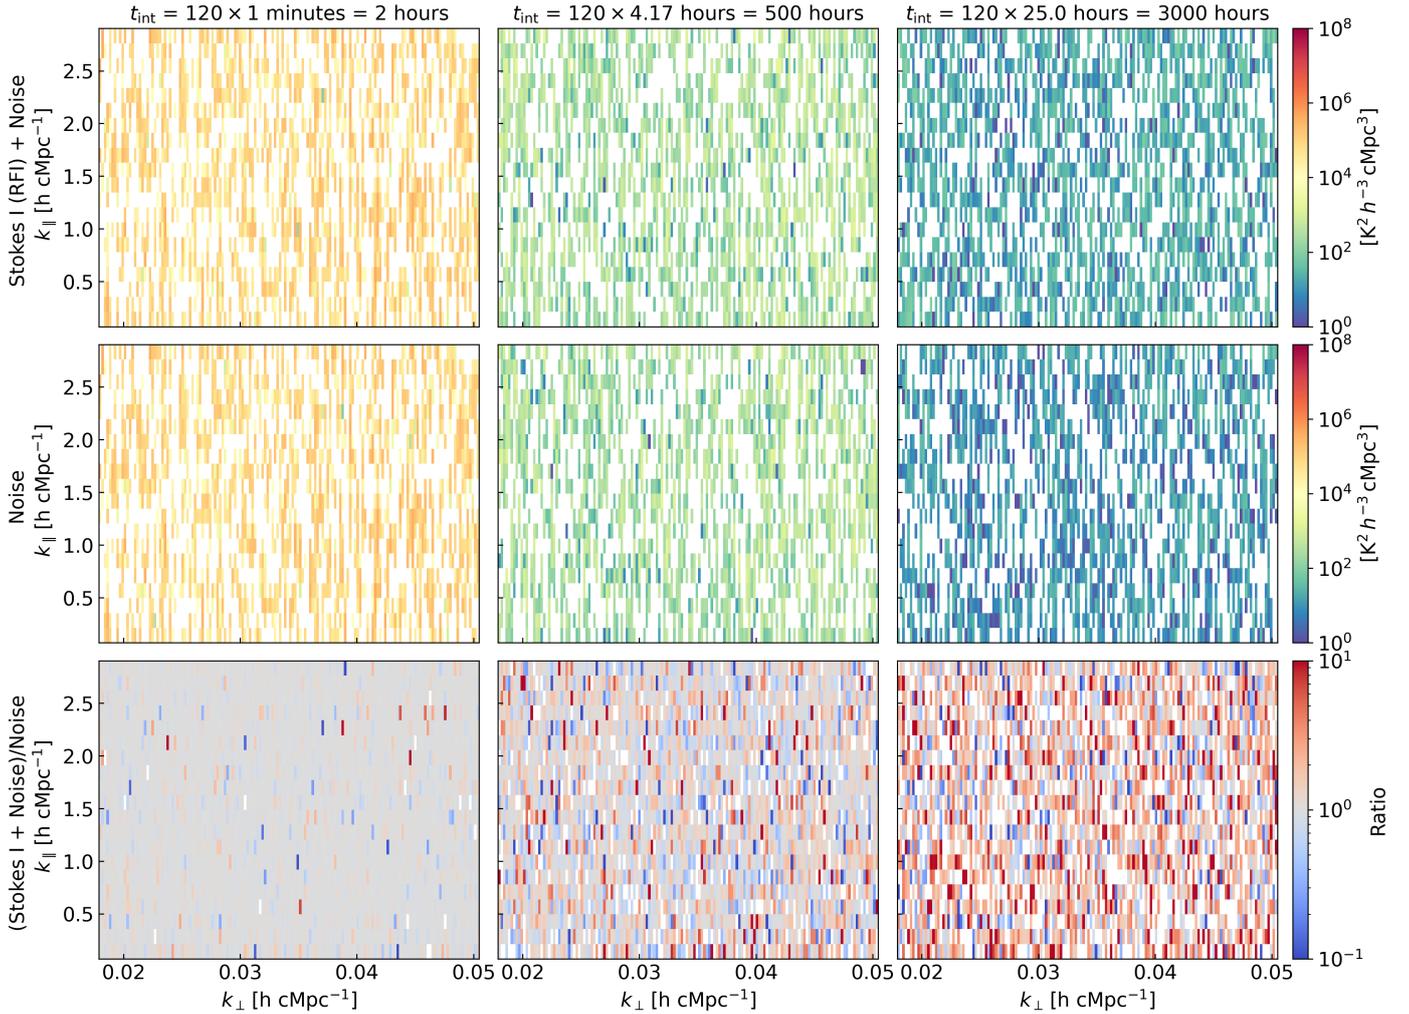


Fig. 7. Inverse variance weighted average of the cylindrical cross-power spectra of all image-cube pairs. The top row shows the power spectra of transient RFI in the presence of thermal noise for different deep integration times (left to right columns), the middle row shows the cross-power spectra of the thermal noise itself, and the bottom row shows the ratio of the top row and the middle row. The blank pixels correspond to negative power which occurs due to the presence of the uncorrelated nature of data.

transient RFI image cube in the dataset. However, we choose to calculate cross-power spectra by cross-multiplying two different gridded visibility cubes, instead of calculating the auto-power spectrum by multiplying a gridded visibility cube by itself, to avoid the bias due to the noise and incoherent RFI. This allows us to inspect the correlation of transient RFI. We then average all cross-power spectrum pairs using inverse variance weighting to obtain a single cylindrical and spherical power spectrum. For this test, we use gridded visibilities from the transient RFI image cubes with thermal noise (the three scenarios used in Sect. 5.1) and corresponding thermal noise realisations to calculate cylindrically and spherically averaged power spectrum. In addition to cross-power spectra, we also compute the cylindrically averaged cross-coherence which is essentially a normalised cross-power spectrum. This is done using Eq. (1), but cylindrical averaging is performed instead of averaging all k -modes together.

Figure 7 shows the inverse variance weighted average of the cross-power spectra calculated from all image-cube pairs corresponding to the eight nights. We do not include auto-power spectra in the averaging to avoid incoherent RFI and noise bias in the final spectrum. We observe that the cylindrically averaged power spectra for transient RFI behave like uncorrelated noise regard-

less of the integration time. Although it is not very straightforward to interpret the power spectra with higher integration time, we can think of them as if a similar RFI level is observed over the given integration time or the instrument has higher sensitivity. This suggests that the transient RFI we observe is uncorrelated when different nights are combined and leaves little to no impact on k -modes where we expect to measure the faint 21-cm signal from Cosmic Dawn. The ratio of the power spectra of transient RFI and the thermal noise has a median of about unity for the three scenarios. However, the variance of the ratio increases as the thermal noise decreases. This is likely due to the presence of uncorrelated transient RFI that enhances the variance relative to the thermal noise. For the integration time of 500 h, we do not observe a significant difference between the transient RFI power and the thermal noise power spectrum. This is an important result because RFI is considered one of the major hurdles in 21-cm cosmology experiments. However, this conclusion should be treated with some care because it only applies to the transient RFI that is sky-based, varies over time, is above the location of the LOFAR core, is in the frequency range studied here, and does not show any night-to-night correlation at the same LST. Persistent RFI (whether narrow band or broadband)

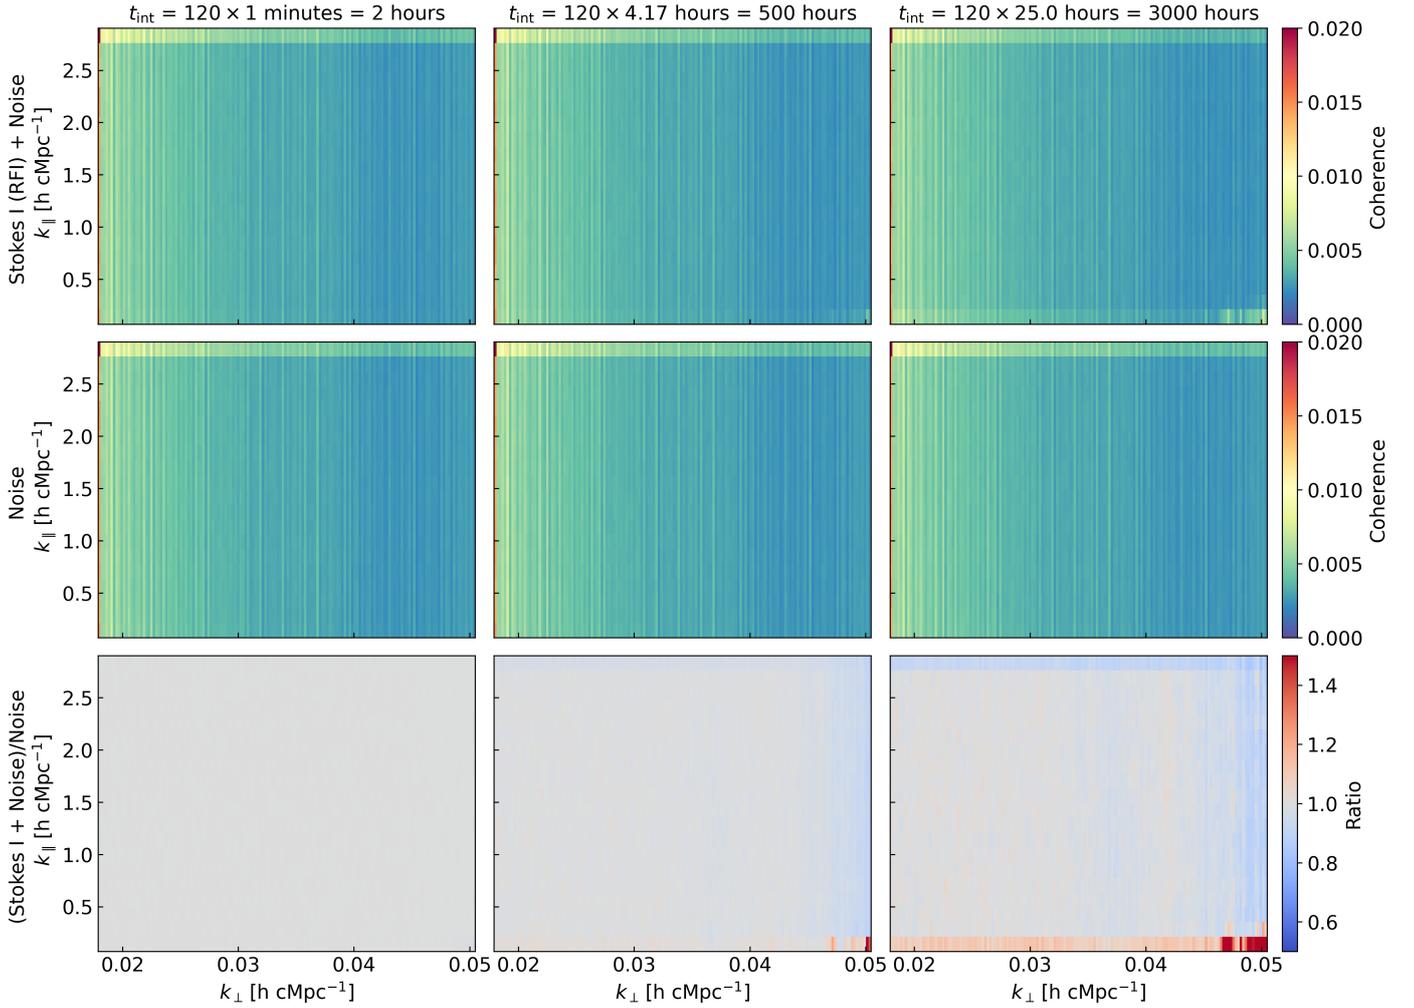


Fig. 8. Average cross-coherence of all image-cube pairs. The top row shows the cross-coherence of the transient RFI in the presence of thermal noise for different deep integration times (left to right columns), the middle row shows the cross-coherence of the thermal noise itself, and the bottom row shows the ratio of the top row and the middle row.

from sources such as geostationary satellites, FM transmissions, digital audio and video broadcasts contaminates the observed data at approximately the same level and location (in frequency spectrum) regardless of the sidereal time of the observation and leaves strongly correlated artefacts in the data. The latter contamination is of grave concern in 21-cm experiments and needs to be carefully mitigated.

We can draw similar conclusions from Fig. 8 that shows the cylindrical cross-coherence averaged over all image-cube pairs. We do not observe any discernible structure due to the transient RFI in addition to the thermal noise at different k -modes for 2 h and 500 h of integration, suggesting that the transient RFI does not correlate over time and combining observations of transient RFI localised differently in image space does not result in correlated structures in Fourier space. We start to see very low-level coherence at the smallest k_{\parallel} mode for the third scenario which we suspect is due to spectrally smooth foregrounds (and side-lobes) that are falsely detected as transient RFI. Finally, we show inverse variance weighted averages of the spherically averaged cross-power spectra for all image-cube pairs in Fig. 9. We see that for all three scenarios, 2 h, 500 h, and 3000 h (and twice the AARTFAAC-12 sensitivity) integration, the transient RFI power in the presence of thermal noise is consistent with the thermal noise floor within 2σ error.

6. Conclusions and discussion

The LOFAR-AARTFAAC Cosmic Explorer project aims to measure the power spectrum of the redshifted 21-cm signal of neutral hydrogen from the Cosmic Dawn ($z \sim 18$) using the AARTFAAC-12 wide-field imager of LOFAR-LBA and to constrain or exclude non-standard (“exotic”) astrophysical models that could explain the unusual nature of the putative 21-cm absorption feature as reported by the EDGES collaboration (Bowman et al. 2018). In G20, we presented the first upper limits on the redshifted 21-cm signal at $z = 18.26$ using a small subset of AARTFAAC data recorded as a part of the ACE observing campaign. However, since the LOFAR low-band receivers that compose AARTFAAC have a full-sky FoV and are situated in a high population density area of Western Europe, they may be especially susceptible to the radio frequency interference at low elevations as well as sources of interference that rapidly move across the FoV, for example, aeroplanes, satellites, meteors etc. In this work, we investigated the impact of transient RFI on the power spectrum which is the main statistic that we use to measure the redshifted 21-cm signal. Our findings are summarised below:

1. Different levels of transient RFI are found for different nights. Most of the RFI shows up closer to the horizon, but

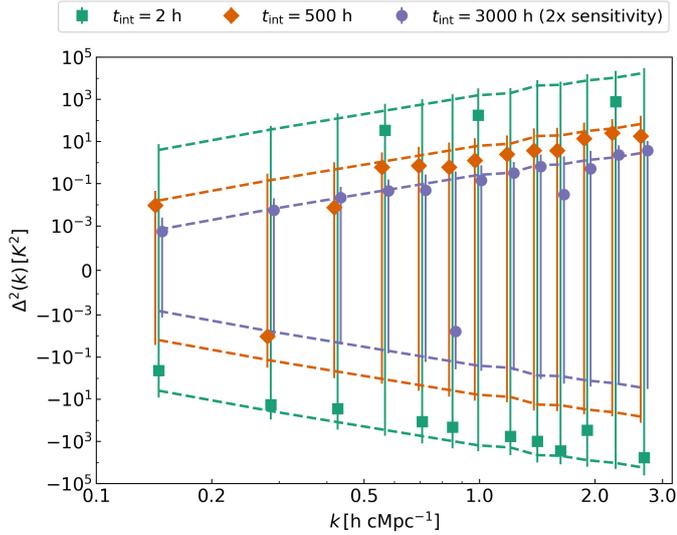


Fig. 9. Inverse variance weighted spherically averaged power spectrum of transient RFI in the presence of thermal noise. Different markers (and colours) show transient RFI cross-power spectra for different thermal noise levels and the dashed lines correspond to the 1σ error due to the thermal noise. We note that the k -values for 500 h and 3000 h have been shifted slightly to better visualise the error bars. The dashed lines also correspond to the power spectrum sensitivity of the given integration time and instrument sensitivity.

we also observe rapidly moving RFI sources that cross over the FoV of AARTFAAC within a few minutes. Most of the transient RFI appears to be narrow-band and localised in image space. They are most likely aeroplane communication beacons in the observed spectral window.

2. A cross-coherence analysis shows that the transient RFI between different points in time (image-cubes) exhibit little to no correlation (2% or below) except for a few outliers that show a low-level correlation of the order of a few percent. Additionally, in the presence of thermal noise, the RFI coherence behaviour is mostly indistinguishable from the thermal noise, except for the extremely low noise regime (equivalent to >500 h of integration) where outliers start to appear.
3. Trimming the images to a certain FoV or applying a spatial taper to exclude the transient RFI at low elevation has a substantial impact on the coherence of the transient RFI in Fourier space. For example, trimming the image to a smaller FoV of 120° reduces the RFI coherence compared to the full FoV. Similarly, applying a Hann spatial taper shows relatively lower transient RFI coherence compared to a Tukey taper, suggesting that most RFI is generated on the ground plane and/or near the horizon.
4. The cylindrically averaged power spectrum of the transient RFI in the presence of thermal noise and after applying a Hann spatial taper behaves similarly to uncorrelated noise regardless of the different thermal noise levels. However, the overall variance of the transient RFI relative to the thermal noise increases for deep integration of 500 h and above. Similarly, cylindrically averaged cross-coherence of transient RFI shows no additional coherence in Fourier space for up to at least 500 h of integration. However, the 3000 h of integration with the hypothetical instrument twice the sensitivity of AARTFAAC shows very low-level coherence at the smallest k_{\parallel} mode which is likely due to residual correlated foregrounds remaining after the filtering process. The

spherically averaged power spectrum of transient RFI in the presence of thermal noise is consistent with instrument sensitivity for all integration times, 2 h, 500 h, as well as 3000 h for a hypothetical instrument twice as sensitive as AARTFAAC.

From the investigation presented in this work, we conclude that the sky-based transient RFI at the location of LOFAR core (located in Exloo, Netherlands) between 72 and 75 MHz for any practical purpose is uncorrelated over time, and indistinguishable from the thermal noise equivalent to 500 h of integration, which is the same as the amount of data recorded during the ACE observing campaign. Additionally, following a more conservative approach in future, we plan to investigate advanced mitigation techniques for transient RFI such as:

1. Image-based filtering of transient RFI: In the analysis, we observed that the transient RFI is localised in image space over short periods. A 4D mask can be used to filter out the voxels (in space, frequency and time) affected by transient RFI. These voxels can then be filled by artificial emission derived from a physically and statistically motivated model of the structure (foregrounds and noise) using a robust approach such as Gaussian Process Regression (Mertens et al. 2018), to avoid artefacts in the power spectrum. However, an in-depth investigation of the impact of such an approach on 21-cm signal measurement is still required.
2. Flagging of RFI in visibility space: Visibilities observed with AARTFAAC have a high noise level because of the small effective collecting area of receiver elements. Traditional algorithms that operate on per baseline dynamic spectra of visibilities to search for outliers, therefore, are sub-optimal for detecting low-level transient RFI. The use of baseline-integrated statistics provides a method to search for faint RFI signatures from aeroplanes and digital broadcasts. Offringa et al. (2015) shows an example of applying their detection method on the baseline-detected standard deviation, thereby detecting RFI that is not detectable in a single baseline. Similarly, Wilensky et al. (2019) performs detection on the incoherently averaging the visibility spectra of the baselines to improve detection. This method is shown to be particularly effective for MWA datasets. However, it may lead to significant over-flagging because of the averaging of multiple baselines, which are all flagged even if only a subset exhibits RFI. We plan to explore a similar approach by grouping baselines based on length and orientation to avoid over-flagging. Additionally, we plan to explore the use of cross-polarisations for faint RFI detection (Yatawatta 2020) in future analyses.

In closing, significantly deeper observing campaigns with highly sensitive instruments may be susceptible to low-level transient RFI and would require a transient RFI mitigation strategy. However, conclusions from this investigation can be extrapolated to the next generation of 21-cm cosmology experiments such as with the SKA and HERA, which are located at more radio-quiet sites and have a narrower FoV. These experiments are therefore less likely to suffer from contamination from the transient RFI that is uncorrelated over time impacting the redshifted 21-cm signal measurement.

Acknowledgements. B.K.G., L.V.E.K., S.A.B., S.G., C.H., and S.M. acknowledge the financial support from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 884760, "CoDEX"). L.V.E.K., A.R.O., and E.C. acknowledge support from the Centre for Data Science and Systems Complexity (DSSC), Faculty of Science and Engineering at the University of Groningen.

F.G.M. acknowledges support from a PSL Fellowship. This work was supported in part by ERC grant 247295 “AARTFAAC” to RAMJW. LOFAR, the Low-Frequency Array designed and constructed by ASTRON, has facilities in several countries, owned by various parties (each with their own funding sources), and collectively operated by the International LOFAR Telescope (ILT) foundation under a joint scientific policy. This research made use of publicly available software developed for LOFAR and AARTFAAC telescopes. Listed below are software packages used in the analysis: AARTFAAC2MS (<https://github.com/aroffringa/aartfaac2ms>), AOFLAGGER (<https://gitlab.com/aroffringa/aoflagger>), DP3 (<https://git.astron.nl/RD/DP3>), WSCLEAN (<https://gitlab.com/aroffringa/wsclean>), and PS_EOR (https://gitlab.com/flomertens/ps_eor). The analysis also relies on the Python programming language (<https://www.python.org>) and several publicly available python software modules: ASTROPY (<https://www.astropy.org/>; Astropy Collaboration 2013), MATPLOTLIB (<https://matplotlib.org/>), SCIPY (<https://www.scipy.org/>), and NUMPY (<https://numpy.org/>).

References

- An, T., Chen, X., Mohan, P., & Lao, B. Q. 2017, *Acta Astron. Sin.*, **58**, 43
- Arras, P., Reinecke, M., Westermann, R., & Enßlin, T. A. 2021, *A&A*, **646**, A58
- Astropy Collaboration (Robitaille, T. P., et al.) 2013, *A&A*, **558**, A33
- Baan, W. A. 2019, *J. Astron. Instrum.*, **8**, 1940010
- Bale, S. D., Bassett, N., Burns, J. O., et al. 2023, ArXiv e-prints [arXiv:2301.10345]
- Barkana, R. 2018, *Nature*, **555**, 71
- Barry, N., Hazelton, B., Sullivan, I., Morales, M. F., & Pober, J. C. 2016, *MNRAS*, **461**, 3135
- Bernardi, G., de Bruyn, A. G., Brentjens, M. A., et al. 2009, *A&A*, **500**, 965
- Bernardi, G., de Bruyn, A. G., Harker, G., et al. 2010, *A&A*, **522**, A67
- Bernardi, G., Zwart, J. T. L., Price, D., et al. 2016, *MNRAS*, **461**, 2847
- Blackman, R. B., & Tukey, J. W. 1958, *Bell Syst. Tech. J.*, **37**, 185
- Bowman, J. D., Cairns, I., Kaplan, D. L., et al. 2013, *PASA*, **30**, e031
- Bowman, J. D., Rogers, A. E. E., Monsalve, R. A., Mozden, T. J., & Mahesh, N. 2018, *Nature*, **555**, 67
- Bradley, R. F., Tauscher, K., Rapetti, D., & Burns, J. O. 2019, *ApJ*, **874**, 153
- Burns, J., Hallinan, G., Chang, T. C., et al. 2021, ArXiv e-prints [arXiv:2103.08623]
- Cucho-Padin, G., Wang, Y., Li, E., et al. 2019, *Radio Sci.*, **54**, 986
- Datta, A., Bowman, J. D., & Carilli, C. L. 2010, *ApJ*, **724**, 526
- DeBoer, D. R., Parsons, A. R., Aguirre, J. E., et al. 2017, *PASP*, **129**, 045001
- de Lera Acedo, E. 2019, *2019 International Conference on Electromagnetics in Advanced Applications (ICEAA)*, 0626
- Dillon, J. S., Liu, A., Williams, C. L., et al. 2014, *Phys. Rev. D*, **89**, 023002
- Dilullo, C., Taylor, G. B., & Dowell, J. 2020, *J. Astron. Instrum.*, **9**, 2050008
- Di Vruono, F., Winkel, B., Bassa, C. G., et al. 2023, *A&A*, **676**, A75
- Dowell, J., & Taylor, G. B. 2018, *ApJ*, **858**, L9
- Eastwood, M. W., Anderson, M. M., Monroe, R. M., et al. 2019, *AJ*, **158**, 84
- Ewall-Wice, A., Dillon, J. S., Liu, A., & Hewitt, J. 2017, *MNRAS*, **470**, 1849
- Ewall-Wice, A., Chang, T.-C., Lazio, J., et al. 2018, *ApJ*, **868**, 63
- Feng, C., & Holder, G. 2018, *ApJ*, **858**, L17
- Fialkov, A., & Barkana, R. 2019, *MNRAS*, **486**, 1763
- Furlanetto, S. R., Oh, S. P., & Briggs, F. H. 2006, *Phys. Rep.*, **433**, 181
- Gan, H., Mertens, F. G., Koopmans, L. V. E., et al. 2023, *A&A*, **669**, A20
- Gehlot, B. K., Mertens, F. G., Koopmans, L. V. E., et al. 2020, *MNRAS*, **499**, 4158
- Ghosh, A., Prasad, J., Bharadwaj, S., Ali, S. S., & Chengalur, J. N. 2012, *MNRAS*, **426**, 3295
- Ghosh, A., Mertens, F. G., & Koopmans, L. V. E. 2018, *MNRAS*, **474**, 4552
- Harris, F. 1978, *Proc. IEEE*, **66**, 51
- Hazelton, B. J., Morales, M. F., & Sullivan, I. S. 2013, *ApJ*, **770**, 156
- Hills, R., Kulkarni, G., Meerburg, P. D., & Puchwein, E. 2018, *Nature*, **564**, E32
- Jordan, C. H., Murray, S., Trott, C. M., et al. 2017, *MNRAS*, **471**, 3974
- Kassim, N. E., Lazio, T. J. W., Erickson, W. C., et al. 2007, *ApJS*, **172**, 686
- Koopmans, L., Pritchard, J., Mellema, G., et al. 2015, *Advancing Astrophysics with the Square Kilometre Array (AASKA14)*, 1
- Kuiack, M., Huizinga, F., Molenaar, G., et al. 2019, *MNRAS*, **482**, 2502
- Laing, R. A., Riley, J. M., & Longair, M. S. 1983, *MNRAS*, **204**, 151
- Mertens, F. G., Ghosh, A., & Koopmans, L. V. E. 2018, *MNRAS*, **478**, 3640
- Mertens, F. G., Mevius, M., Koopmans, L. V. E., et al. 2020, *MNRAS*, **493**, 1662
- Mertens, F. G., Semelin, B., & Koopmans, L. V. E. 2021, SF2A-2021: Proceedings of the Annual meeting of the French Society of Astronomy and Astrophysics, eds. A. Siebert, K. Baillié, E. Lagadec, et al., 211
- Mertens, F. G., Bobin, J., & Carucci, I. P. 2024, *MNRAS*, **527**, 3517
- Mevius, M., van der Tol, S., Pandey, V. N., et al. 2016, *Radio Sci.*, **51**, 927
- Mevius, M., Mertens, F., Koopmans, L. V. E., et al. 2022, *MNRAS*, **509**, 3693
- Middelberg, E. 2006, *PASA*, **23**, 64
- Morales, M. F., & Hewitt, J. 2004, *ApJ*, **615**, 7
- Morales, M. F., Hazelton, B., Sullivan, I., & Beardsley, A. 2012, *ApJ*, **752**, 137
- Morales, M. F., Beardsley, A., Pober, J., et al. 2019, *MNRAS*, **483**, 2207
- Mouri Sardarabadi, A., & Koopmans, L. V. E. 2019, *MNRAS*, **483**, 5480
- Offringa, A. R., & Smirnov, O. 2017, *MNRAS*, **471**, 301
- Offringa, A. R., de Bruyn, A. G., Biehl, M., et al. 2010, *MNRAS*, **405**, 155
- Offringa, A. R., van de Gronde, J. J., & Roerdink, J. B. T. M. 2012, *A&A*, **539**, A95
- Offringa, A. R., de Bruyn, A. G., Zaroubi, S., et al. 2013, *A&A*, **549**, A11
- Offringa, A. R., McKinley, B., Hurley-Walker, N., et al. 2014, *MNRAS*, **444**, 606
- Offringa, A. R., Wayth, R. B., Hurley-Walker, N., et al. 2015, *PASA*, **32**, e008
- Offringa, A. R., Mertens, F., & Koopmans, L. V. E. 2019, *MNRAS*, **484**, 2866
- Patil, A. H., Yatawatta, S., Zaroubi, S., et al. 2016, *MNRAS*, **463**, 4317
- Peck, L. W., & Fenech, D. M. 2013, *Astron. Comput.*, **2**, 54
- Prasad, J., & Chengalur, J. 2012, *Exp. Astron.*, **33**, 157
- Prasad, P., Huizinga, F., Kooistra, E., et al. 2016, *J. Astron. Instrum.*, **5**, 1641008
- Pritchard, J. R., & Loeb, A. 2012, *Rep. Progr. Phys.*, **75**, 086901
- Reis, I., Fialkov, A., & Barkana, R. 2021, *MNRAS*, **506**, 5479
- Shulevski, A., Franzen, T. M. O., Williams, W. L., et al. 2022, *MNRAS*, **513**, 1036
- Singh, S., Subrahmanyan, R., Udaya Shankar, N., et al. 2017, *ApJ*, **845**, L12
- Singh, S., Jishnu, N. T., Subrahmanyan, R., et al. 2022, *Nat. Astron.*, **6**, 607
- Thompson, A. R., Moran, J. M., & Swenson, G. W., Jr. 2001, *Interferometry and Synthesis in Radio Astronomy*, 2nd edn. (Wiley-VCH)
- Trott, C. M., Wayth, R. B., & Tingay, S. J. 2012, *ApJ*, **757**, 101
- van Diepen, G., Dijkema, T. J., & Offringa, A. 2018, *Astrophysics Source Code Library* [record ascl:1804.003]
- van Haarlem, M. P., Wise, M. W., Gunst, A. W., et al. 2013, *A&A*, **556**, A2
- Vecchio, A., Bentum, M., Falcke, H., et al. 2021, *43rd COSPAR Scientific Assembly, held 28 January – 4 February*, 43, 1525
- Vedantham, H., Udaya Shankar, N., & Subrahmanyan, R. 2012, *ApJ*, **745**, 176
- Vos, E. E., Francois Luus, P. S., Finlay, C. J., Bassett, B. A. 2019, *IEEE 29th International Workshop on Machine Learning for Signal Processing (MLSP)*, 1
- Whitler, L. R., Beardsley, A., & Jacobs, D. 2019, *Am. Astron. Soc. Meet. Abstr.*, **233**, 349.17
- Wilensky, M. J., Morales, M. F., Hazelton, B. J., et al. 2019, *PASP*, **131**, 114507
- Yatawatta, S. 2020, ArXiv e-prints [arXiv:2006.00062]
- Ye, H., Gull, S. F., Tan, S. M., & Nikolic, B. 2022, *MNRAS*, **510**, 4110
- Zarka, P., Girard, J. N., Tagger, M., & Denis, L. 2012, in *SF2A-2012: Proceedings of the Annual meeting of the French Society of Astronomy and Astrophysics*, eds. S. Boissier, P. de Laverny, N. Nardetto, et al., 687
- Zaroubi, S. 2013, *Astrophys. Space Sci. Lib.*, **396**, 45