

Search for the Epoch of Reionization with HERA: upper limits on the closure phase delay power spectrum

Pascal M. Keller¹ Bojan Nikolic,¹ Nithyanandan Thyagarajan,² Chris L. Carilli,³ Gianni Bernardi,^{4,5,6} Ntsikelelo Charles⁵ Landman Bester,^{6,5} Oleg M. Smirnov,^{5,6} Nicholas S. Kern,⁷ Joshua S. Dillon⁸, Bryna J. Hazelton,^{9,10} Miguel F. Morales⁹, Daniel C. Jacobs,¹¹ Aaron R. Parsons,⁸ Zara Abdurashidova,⁸ Tyrone Adams,⁶ James E. Aguirre¹², Paul Alexander,¹ Zaki S. Ali,⁸ Rushelle Baartman,⁶ Yanga Balfour,⁶ Adam P. Beardsley,^{11,13} Tashalee S. Billings,¹² Judd D. Bowman,¹¹ Richard F. Bradley,¹⁴ Philip Bull,^{15,16} Jacob Burba,^{15,17} Steven Carey,¹ Carina Cheng,⁸ David R. DeBoer,¹⁸ Eloy de Lera Acedo,¹ Matt Dexter,¹⁸ Nico Eksteen,⁶ John Ely,¹ Aaron Ewall-Wice^{8,19} Nicolas Fagnoni¹, Randall Fritz,⁶ Steven R. Furlanetto¹,²⁰ Kingsley Gale-Sides,¹ Brian Glendenning,³ Deepthi Gorthi,⁸ Bradley Greig^{1,21} Jasper Grobbelaar,⁶ Ziyaad Halday,⁶ Jacqueline N. Hewitt,^{7,22} Jack Hickish,¹⁸ Austin Julius,⁶ MacCalvin Kariseb,⁶ Joshua Kerrigan,¹⁷ Piyanat Kittiwisit¹⁶, Saul A. Kohn¹², Matthew Kolopanis¹¹, Adam Lanman¹⁷, Paul La Plante,^{8,12} Adrian Liu²³, Anita Loots,⁶ Yin-Zhe Ma²⁴, David Harold Edward MacMahon,¹⁸ Lourence Malan,⁶ Cresshim Malgas,⁶ Keith Malgas,⁶ Bradley Marero,⁶ Zachary E. Martinot,¹² Andrei Mesinger²⁵, Mathakane Molewa,⁶ Tshegofalang Mosiane,⁶ Steven G. Murray¹¹, Abraham R. Neben,⁷ Hans Nuwegeld,⁶ Robert Pascua,^{8,23} Nipanjana Patra,⁸ Samantha Pieterse,⁶ Jonathan C. Pober,¹⁷ Nima Razavi-Ghods,¹ James Robnett,³ Kathryn Rosie,⁶ Mario G. Santos,^{6,16} Peter Sims²³, Craig Smith,⁶ Hilton Swarts,⁶ Pieter Van Wyngaarden,⁶ Peter K. G. Williams^{1,26,27} and Haoxuan Zheng⁷

Affiliations are listed at the end of the paper

Accepted 2023 January 27. Received 2023 January 23; in original form 2022 November 3

ABSTRACT

Radio interferometers aiming to measure the power spectrum of the redshifted 21 cm line during the Epoch of Reionization (EoR) need to achieve an unprecedented dynamic range to separate the weak signal from overwhelming foreground emissions. Calibration inaccuracies can compromise the sensitivity of these measurements to the effect that a detection of the EoR is precluded. An alternative to standard analysis techniques makes use of the closure phase, which allows one to bypass antenna-based direction-independent calibration. Similarly to standard approaches, we use a delay spectrum technique to search for the EoR signal. Using 94 nights of data observed with Phase I of the Hydrogen Epoch of Reionization Array (HERA), we place approximate constraints on the 21 cm power spectrum at $z = 7.7$. We find at 95 per cent confidence that the 21 cm EoR brightness temperature is $\leq (372)^2$ ‘pseudo’ mK 2 at 1.14 ‘pseudo’ h Mpc $^{-1}$, where the ‘pseudo’ emphasizes that these limits are to be interpreted as approximations to the actual distance scales and brightness temperatures. Using a fiducial EoR model, we demonstrate the feasibility of detecting the EoR with the full array. Compared to standard methods, the closure phase processing is relatively simple, thereby providing an important independent check on results derived using visibility intensities, or related.

Key words: methods: data analysis – methods: statistical – techniques: interferometric – intergalactic medium – dark ages, reionization, first stars.

1 INTRODUCTION

The Epoch of Reionization (EoR) is a period in cosmic history during which the neutral intergalactic medium (IGM) was ionized by the

first luminous sources. Current constraints on cosmic reionization are derived from the scattering of the Cosmic Microwave Background (CMB) from the ionized IGM (Planck Collaboration VI 2020) and from absorption effects observed in high-redshift quasars and galaxy surveys. Combined, these observations point towards a reionization

* E-mail: pmk46@cam.ac.uk

midpoint at a redshift of $z \sim 7$ and a completion of reionization by $z \sim 5.5$ (Greig & Mesinger 2017).

The 21 cm spin-flip emission of neutral Hydrogen (H1) will be a powerful probe of cosmic reionization (see e.g. Morales & Wyithe 2010; Pritchard & Loeb 2012; Furlanetto 2016 for reviews). Once observed, the redshifted 21 cm signal will provide spatially resolved information about the timing and duration of reionization as well as the physical properties of the neutral IGM. In return, this will give insight into the nature of the sources providing the ionizing photons. An advantage of using the 21 cm signal is that its cosmological redshift is an indicator of the line-of-sight distance, making it a tomographic probe of the neutral IGM (Madau, Meiksin & Rees 1997). However, to increase the sensitivity to the 21 cm signal, current interferometric experiments focus on a statistical detection by means of its power spectrum rather than tomographic imaging.

There are several ongoing, past and future radio interferometers aimed at measuring the power spectrum of the cosmological 21 cm signal. These include the Hydrogen Epoch of Reionization Array (HERA¹; DeBoer et al. 2017), the Donald C. Backer Precision Array for Probing the Epoch of Reionization (PAPER; Parsons et al. 2010), The Murchison Widefield Array (MWA; Tingay et al. 2013), the LOw Frequency ARray (LOFAR; van Haarlem et al. 2013), the Long Wavelength Array (LWA; Eastwood et al. 2019), the Giant Metre Wave Radio Telescope (GMRT; Paciga et al. 2013), and the Square Kilometre Array (SKA; Koopmans et al. 2015). While continuously lowering the upper limits on the 21 cm brightness temperature of the IGM, most of these experiments are currently limited by systematic effects rather than thermal noise (Paciga et al. 2013; Dillon et al. 2015; Beardsley et al. 2016; Patil et al. 2017; Cheng et al. 2018; Barry et al. 2019; Kolopanis et al. 2019; Li et al. 2019; Mertens et al. 2020; Trott et al. 2020; HERA Collaboration 2022a, c).

One particular challenge is to calibrate the instrument to the accuracy required for a detection of the cosmological 21 cm signal. Inaccurate sky models and differences between nominally redundant baselines can introduce calibration errors that overwhelm the weak cosmological signal (Barry et al. 2016; Ewall-Wice et al. 2017; Byrne et al. 2019, 2021). This motivates the use of calibration-independent closure quantities (Thompson, Moran & Swenson 2017; Samuel, Nityananda & Thyagarajan 2022; Thyagarajan, Nityananda & Samuel 2022) to search for the cosmological signal. In this work we use the closure phase, which is defined as the sum of the three visibility phases of an antenna triangle (Jennison 1958). It can be shown that antenna-based direction-independent gain phases cancel in the closure phase.

The basic concept of the closure phase approach was first introduced in Thyagarajan, Carilli & Nikolic (2018) and its mathematical foundation is set out in Thyagarajan & Carilli (2020). Using simulations, these papers confirm that the dynamic range required to detect the weak cosmological signal in the closure phase is comparable to that of a visibility-based approach. Looking at HERA commissioning data, Carilli et al. (2018) find that the closure phase agrees well across redundant measurements. The first results of the closure phase analysis performed on 18 nights of HERA phase I observing are presented in Thyagarajan et al. (2020). While the data are partially affected by systematic effects, they also identify large regions in the power spectra that are limited by thermal-like noise.

The closure phase analysis is carried out in parallel with the standard visibility-based analysis (cf. HERA Collaboration 2022a, hereafter H22a). H22a report improved constraints on the 21 cm

EoR power spectrum using visibility intensities from a full season of Phase I HERA data, finding at 95 percent confidence that $\Delta_{21}^2 \leq (21.4)^2 \text{ mK}^2$ at $k = 0.34 \text{ h Mpc}^{-1}$ and $z = 7.9$ and $\Delta_{21}^2 \leq (59.1)^2 \text{ mK}^2$ at $k = 0.36 \text{ h Mpc}^{-1}$ and $z = 10.4$. These results are an update to previously published limits obtained on a subset of the Phase I data (HERA Collaboration 2022c, hereafter H22c) and provide an improvement by factors of 2.1 and 2.6, respectively. H22a use a similar data processing pipeline to that of H22c, which incorporates elaborate methods for preventing RFI and internal instrumental coupling effects from contaminating the power spectrum (Kern et al. 2019, 2020). The signal losses introduced through these non-linear processing steps are characterized in a validation pipeline that makes use of extensive simulations (cf. Aguirre et al. 2022). The upper limits were then used to set constraints on the IGM and galaxies at $z \sim 8$ and 10 (cf. HERA Collaboration 2022b, H22a.). These constraints require heating above the adiabatic cooling threshold prior to $z \sim 10.4$ and disfavour models with low X-ray heating. While the upper limits set by HERA are going to be increasingly important in constraining the EoR, this also calls for independent and alternative approaches to analysing the data. The closure phase analysis is one such approach which will increase the confidence in these results.

In this paper, we report upper limits of the closure phase delay power spectrum obtained from a full season of HERA Phase I observing. The outline of this paper is as follows. In Section 2, we shortly summarize the mathematical foundations underlying the closure phase approach. Section 3 lists the specifications of the HERA array and Section 4 gives an overview of the data used in the analysis. In Section 5, we describe the modelling used to validate the approach. Section 6 details the data selection and delineates the analysis pipeline. Finally, we present our results in Section 7 and summarize in Section 8.

2 THEORY

The mathematical foundations of the closure phase approach are outlined in Thyagarajan & Carilli (2020). For completeness, we summarize some of the mathematical formalism most relevant to the closure phase delay power spectrum analysis, without repeating the involved mathematics in Thyagarajan & Carilli (2020).

In the limit where the cosmological 21 cm signal is weak relative to the foreground continuum visibility amplitudes, we can treat the cosmological signal as a small perturbation to the foreground visibility phase ϕ_p of baseline p . In a first-order approximation, it can be shown that these phase perturbations are (Thyagarajan & Carilli 2020)

$$\delta\phi_p(v) \approx \Im \left\{ \frac{V_p^P(v)}{V_p^F(v)} \right\}, \quad (1)$$

where V_p^P and V_p^F are the perturbing visibility and the foreground visibility of baseline p , respectively, and \Im denotes the imaginary part. The closure phase ϕ is the sum of the three visibility phases of a closed antenna triangle (triad) and its perturbation is simply the sum of the corresponding phase perturbations,

$$\delta\phi(v) \approx \sum_{p=1}^3 \Im \left\{ \frac{V_p^P(v)}{V_p^F(v)} \right\}. \quad (2)$$

The line-of-sight fluctuations of the cosmological signal will cause the phase perturbations to fluctuate across frequency, while the foreground visibility phases are due to broad-band continuum emission, and relatively smooth in frequency. It is this frequency dependence that allows us to separate the cosmological signal from

¹<http://reionization.org/>

the bright continuum foregrounds. More precisely, we perform a Fourier transform of the closure phase along frequency, which is known as a delay transform (cf. Parsons et al. 2012). In the delay spectrum, the spectrally smooth foregrounds are confined to low delay modes, while higher delay modes can be used to set upper limits on the H₁ 21 cm emission from the EoR. The latter is known as the EoR window (Datta, Bowman & Carilli 2010; Liu, Parsons & Trott 2014). Formally, we define the delay spectrum as

$$\tilde{\Psi}_v(\tau) = V_{\text{eff}} \int e^{i\phi(v)} W(v) e^{i2\pi v\tau} dv, \quad (3)$$

where V_{eff} is a scaling factor, W is a spectral tapering function shaped to fit the observed band, v is the frequency, and τ is the delay. In this analysis, we use a Blackman-Harris function (Blackman & Tukey 1958) for W , which is suited to the high dynamic range requirements of the measurement. Note that instead of directly transforming the closure phase ϕ , we transform its complex exponential. Doing this, we avoid the discontinuities that can arise because of the circularity of phase.

The effective visibility V_{eff} provides the delay spectrum with units of Jy Hz, which are the units of a standard visibility delay spectrum. Furthermore, V_{eff} should be designed to gauge the strength of the closure phase fluctuations to the strength of the perturbing signal. That is, it should counteract the inverse proportionality to the foreground visibility in equation (2) so that the spectral fluctuations in $\delta\phi$ are of similar strength to the fluctuations in V_p^F . This allows one to combine measurements of regions on the sky with different foreground visibility amplitudes. Motivated by this, we define

$$V_{\text{eff}}^{-2} = \sum_{p=1}^3 \left(\hat{V}_p^F \right)^{-2}, \quad (4)$$

where \hat{V}_p^F is an estimate of the foreground visibility amplitude weighted by the window function W and averaged over the observed subband (cf. Thyagarajan & Carilli 2020). The summation in inverse quadrature gives weight to the baselines with the smallest foreground visibility amplitudes, where the spectral fluctuations of the 21 cm signal are expected to be strongest. In this analysis, we use calibrated data to estimate V_p^F , as it is readily available from the visibility processing pipeline. However, in principle, we could use an accurate sky model instead, making this approach completely independent of calibration. Note that V_{eff} is deliberately chosen to be frequency independent to avoid further foreground contamination into the EoR window. This assumption should be reasonable over the relative narrow bands considered herein (Thyagarajan & Carilli 2020). Regardless, any claimed limits or detection using this technique comes down to comparison with physical models for the cosmological and foreground signals, which are treated identically to the real data.

Ultimately, we are interested in the power spectrum $|\tilde{\Psi}_v(\tau)|^2$, which, statistically, is independent of direction and polarization. This is due to the assumed isotropy and polarization-independence of the 21 cm signal and allows for the incoherent averaging of different portions of the data. Section 6.5 describes how we do this in practice.

3 HERA

HERA is a low-frequency radio interferometer designed to measure the 21 cm emissions of neutral Hydrogen during the EoR. Located in a radio quiet zone in the Karoo desert in South Africa, it is minimally affected by radio frequency interference (RFI). In its complete form, HERA will consist of 320 closely packed 14-m dishes arranged in a split-hexagonal core and complemented with 30

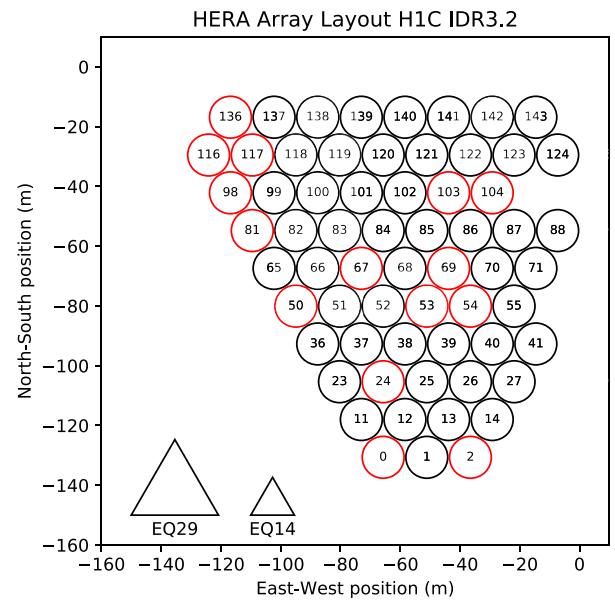


Figure 1. The layout of the antennas used in this analysis. These antennas form a part of the south-west section of the hexagonal array. Their diameter is 14.0 m and the separation between two neighbours is 14.6 m. Antennas that were flagged completely using flags from H22a are shown in red (cf. Section 6.2). The equilateral triad shapes (EQ14 and EQ29) used in this analysis are shown in the bottom left-hand corner of this plot.

outrigger antennas (DeBoer et al. 2017). The highly redundant layout allows for high-precision redundant calibration and is optimized for the delay spectrum approach (Dillon & Parsons 2016). Furthermore, it is ideally suited to the closure phase analysis presented herein, due to the many redundant closure triads in the array.

The data used in this analysis was taken with the HERA Phase I system, which re-used several system elements from its precursor instrument PAPER. These elements include dipole feeds, parts of the analogue signal chain, and the correlator. The Phase I system operated at frequencies between 100 and 200 MHz and observed 1024 channels simultaneously, resulting in a spectral resolution of 97.7 kHz. The temporal resolution is 10.7 s (DeBoer et al. 2017, H22c).

4 DATA

The observations used for this paper cover 94 nights between the Julian Dates 2 458 041 and 2 458 208 (2017 October 15 to 2018 April 1), during which the array was under active construction. As a result, the number of antennas changes throughout the observing season. All together, we use 48 antennas, which constitute a part of the south-west segment of the final array (see Fig. 1). These antennas are flagged on a nightly basis using the flags from H22a, resulting in 35 to 41 unflagged antennas at any one night. The per antenna and per night flags are further detailed in Section 6.2.

In this paper, we only consider measurements on closed antenna triads, rather than baselines or single antennas. We employ closed triad classes that are equal or point-symmetric at the origin in uv -space. For example, a north facing equilateral triad and a south facing equilateral triad belong to the same class, as their uv -geometry can be matched by inversion through the origin (i.e. they are conjugates of one another). As a result their closure phases differ only by sign. We

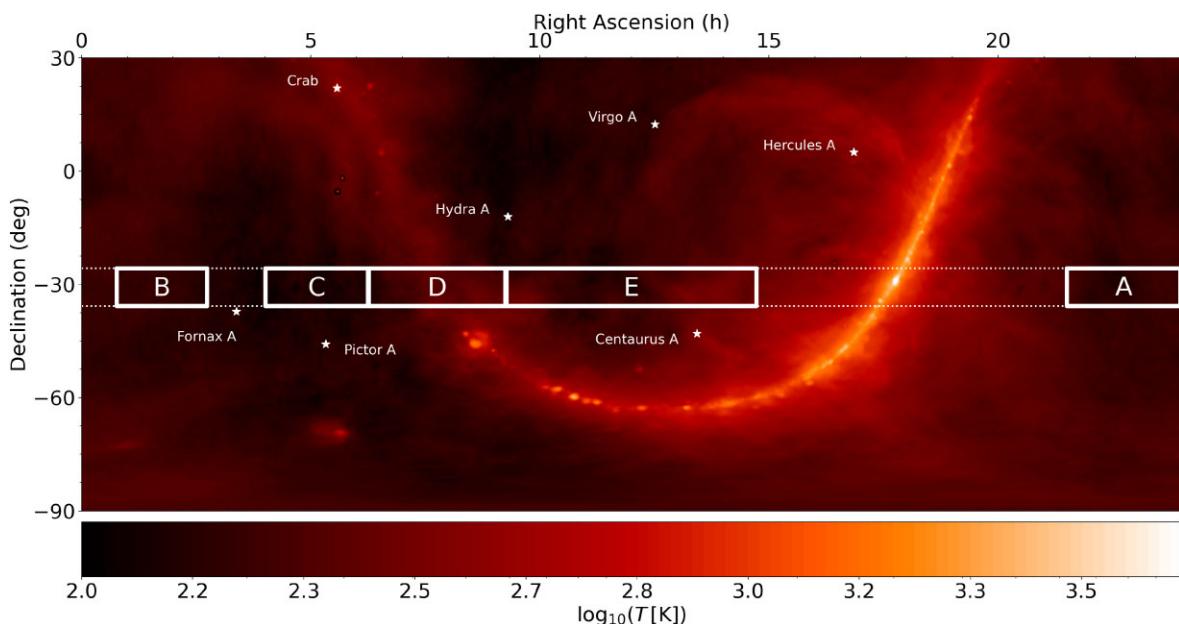


Figure 2. The observed fields A, B, C, D, and E plotted on the global sky model (GSM) at 160 MHz from Zheng et al. (2017). The locations of other bright radio sources are also shown on the map. The dashed lines indicate the $\sim 10^\circ$ field-of-view of HERA at 150 MHz. The fields are chosen to avoid the Galactic plane and the bright radio galaxy Fornax A, which is located only $\sim 6^\circ$ from the centre of the HERA strip.

Table 1. The LST ranges of the fields used in this work and their total observation time in hours.

Field	A	B	C	D	E
LST (h)	21.5–0.0	0.75–2.75	4.0–6.25	6.25–9.25	9.25–14.75
Total (h)	55	89	148	214	210

perform the analysis on two triad classes individually, equilateral 14.6 and 29.2-m triads, which we refer to as EQ14 and EQ29, respectively. This will allow us to see if the two triad classes are affected differently by systematic effects and to what extent their different responses to the foregrounds influence the final power spectra. Ultimately, they probe different spatial scales, leading to more stringent constraints on cosmic reionization.

As a zenith pointing array, the observable portion of the sky is limited to a strip centred at a declination of -30.7° . The width of this strip is defined by the full-width at half maximum (FWHM) of an antenna beam. For a HERA dish, the FWHM is approximately 10° at 150 MHz (Fagnoni et al. 2021). As shown in Fig. 2 the Galactic centre transits overhead at Local Sidereal Time (LST) ~ 18 h. Moreover, the bright radio galaxy Fornax A at ~ 3.3 h RA and -37° DEC is located close enough to the HERA strip to produce a sizable fraction of the total power received by an antenna. We chose to analyse observations from five fields, denoted A, B, C, D, and E, that avoid these bright regions on the sky and hence limit the dynamic range required to measure the cosmological 21 cm signal. The LST-ranges of the observed fields and their total observation time are listed in Table 1.

In this paper, we report results for a frequency band ranging from 160.59 to 167.97 MHz corresponding to a central redshift of about 7.7. This band overlaps with Band 2 (152.25–167.97 MHz) used in the visibility processing (H22a). The reason for the trimming of Band 2 is that we found residual RFI in the lower part of the band after averaging the closure phases. The inclusion of such RFI can lead to excess power in the power spectrum, making it indistinguishable

from a sky based signal. The evidence by which we decided to cut the band is further discussed in Appendix A. Note that the initial bandwidth of 7.4 MHz of our band is reduced to an effective bandwidth of 3.7 MHz after applying the Blackman-Harris function in equation (3).

5 MODELLING

We use data simulations to validate the closure phase approach and ultimately to compare the data and the expected signal of a given EoR model. As our measurements are currently limited by noise and systematics, we use this comparison to estimate the additional sensitivity needed to achieve a detection of a fiducial EoR model (Mesinger, Greig & Sobacchi 2016). In the following two sections, we describe the sky models used here, consisting of foregrounds, a 21 cm signal component and noise.

5.1 Foregrounds

We use the GaLactic and Extragalactic All-Sky MWA survey catalogue (GLEAM; Hurley-Walker et al. 2017) as a basis for our foreground models. Since the GLEAM catalogue does not cover the Galactic Plane, we restrict our simulations to fields A, B, and C of this analysis. For a given LST, we select GLEAM point sources that have an integrated flux density exceeding 50 mJy at 151 MHz and fall within a radius of 15° of the pointing centre (i.e. they lie within the main lobe of the antenna beam). Flux densities are then interpolated in frequency using the fitted spectral indices and fluxes provided in the catalogue. Where no spectral index is provided, we fit it using the integrated flux densities at 122, 130, 143, 151, 158, 166, and 174 MHz, respectively, assuming that the flux is described by a power law in that frequency range. Applying a discrete version of the van Cittert-Zernike theorem (van Cittert 1934; Zernike 1938),

we compute the visibilities V_p for two antennas separated by \mathbf{b}_p as

$$V_p = \sum_i A(\mathbf{s}_i, \nu) I_i(\nu) e^{2\pi i \nu \frac{\mathbf{b}_p \cdot \mathbf{s}_i}{c}}, \quad (5)$$

where c is the speed of light, $I_i(\nu)$ is the flux density of source i at a location given by the unit vector \mathbf{s}_i , and at frequency ν and $A(\mathbf{s}_i, \nu)$ is the response of a simulated HERA-beam (Fagnoni et al. 2021).

Several strong radio sources have been peeled from the GLEAM catalogue (see table 2 in Hurley-Walker et al. 2017). We simulate visibilities for these sources whenever they are above the horizon and assume that they are point-like, which is justified for the short baselines used in this analysis. For Fornax A we use a three component model informed by McKinley et al. (2015). The model consists of point sources representing the two lobes and the core with flux densities of 478 Jy (west), 260 Jy (east), and 12 Jy, respectively, at 154 MHz. The spectral indices of the lobes are -0.77 while that of the core is -0.88 . For the other peeled sources, we use the spectral indices and total flux densities provided in table 2 of Hurley-Walker et al. (2017).

5.2 EoR signal

We use the public ‘Faint Galaxies’ simulation output of 21cmFast as an EoR model (Mesinger et al. 2016) and follow a similar procedure to the one used in Thyagarajan et al. (2020). The EoR light cone consists of two transverse axes and one line-of-sight axis which each have a length of 1.6 cMpc and a resolution of 1024 voxels. In our simplified model we identify the transverse direction with orthographically projected angular extent and the line-of-sight direction with frequency. At $z \sim 7.8$ the corresponding angular resolution is thus ~ 28 arcmin and the frequency resolution ~ 93 kHz, where the conversion was done using the cosmological parameter from Planck Collaboration VI (2020). Subtending an angle of $\sim 10^\circ$, the EoR light cube does not cover the entire HERA field of view. For this reason, we tile the model three times along the transverse axes such that each tile is a reflection of its neighbouring tiles, thus avoiding discontinuities at the tiling boundaries. To reduce subsequent computing time, we then smooth and downsample the model to an effective angular resolution of ~ 7 arcmin corresponding to 256 pixels per side. The resolution thus obtained is still well below the FWHM of $\sim 5^\circ$ of the synthesized HERA beam. Treating each pixel as a point source, we compute visibilities in the same way as for the foregrounds described in Section 5.1.

5.3 Noise

We use calibrated autocorrelation visibilities, V_p^{auto} , to model the system temperature of the array, T_{sys} , as a function of time t , and frequency ν (cf. Tan et al. 2021). The system temperature of a single antenna, indexed by p , is calculated as

$$T_{\text{sys},p}(t, \nu) = \frac{c^2 V_p^{\text{auto}}(t, \nu)}{2k_b \nu^2 \Omega}, \quad (6)$$

where Ω is the integrated beam area and k_b is the Boltzmann constant and V_p is taken from LST-binned and systematics filtered autocorrelations of the visibility processing pipeline (H22a). The values of $T_{\text{sys},p}$ vary strongly across antennas, with an average relative standard deviation of 17 per cent. To obtain a model of the system temperature that characterizes the whole array, we calculate the quadratic average of $T_{\text{sys},p}$ over all antennas. Fig. 3 shows the

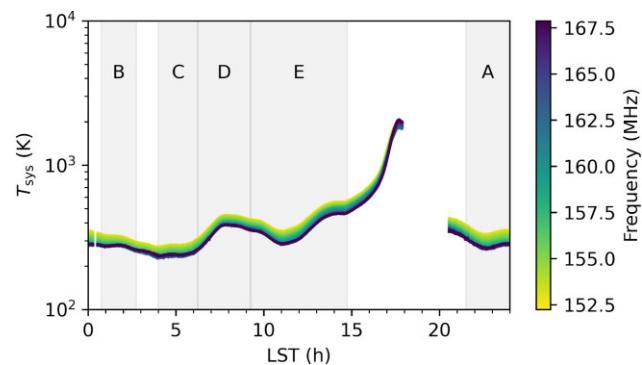


Figure 3. Our model of the system temperature T_{sys} as a function of LST and frequency. The observed fields are shaded in grey.

model of T_{sys} plotted against LST and coloured by frequency. As expected, T_{sys} rises around the transit of the Galactic Plane and is at its lowest in the colder fields, A, B, and C.

The standard deviation of the visibility noise is related to T_{sys} through the radiometer equation

$$\sigma(t, \nu) = \frac{\sqrt{2}k_b \nu^2 \Omega T_{\text{sys}}(t, \nu)}{c^2 \sqrt{\Delta t \Delta \nu}}, \quad (7)$$

where Δt and $\Delta \nu$ are the integration time and frequency resolution, respectively. In our model, we draw the noise from a Gaussian distribution with standard deviation σ and add it to the real and imaginary parts of the simulated visibilities individually. These visibilities are then propagated through the data analysis pipeline described below (Section 6) using the same flags as for the data.

6 ANALYSIS

This section provides an overview of our analysis pipeline. We start by explaining how we obtain closure phases from raw HERA data and how we apply flags to the data. We then continue with our averaging techniques as well as the power spectrum and error estimation. Lastly, we describe how we estimate upper limits. The processing steps and the data products of this analysis are delineated in Fig. 4.

6.1 Computing closure phases

Since closure phases are independent from antenna-based calibration, we compute them directly from raw visibility data. For a given triad, we do this by taking the phase of the triple product of the three visibilities. We then generate a uniformly spaced LST-grid with a resolution given by the integration time of a measurement and replace the time stamps of the closure phases with their nearest points on the grid. This allows us to combine the closure phases to a five dimensional data array of shape $N_p \times N_j \times N_t \times N_l \times N_f$. In this analysis the number of polarizations is $N_p = 2$, as we only keep the parallel ‘East–West’ and ‘North–South’ polarization products and the number of frequencies is $N_f = 76$ for our chosen band. The number of nights N_j , triads N_t , and LST-integrations N_l depend on the observed field (see Section 4). Note that the closure phases will later be averaged across nights, so the effect of LST-gridding on the power spectrum will be similar to that of coherent time averaging described in Appendix B. The introduced signal loss is expected to be negligible, since the gridding interval is equal to the time interval of a single integration.

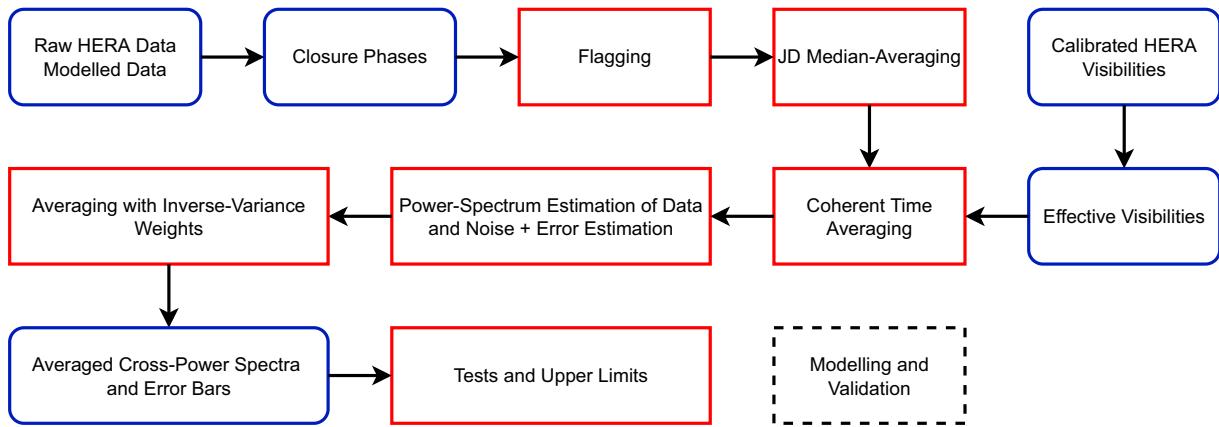


Figure 4. A flowchart showing the processing steps (red) and data products (blue) of the closure phase analysis. The closure phases are computed directly from raw HERA data, i.e. there is no calibration step in the pipeline. The modelling and validation steps (black dashed) are described in Section 5 and Appendix B and are performed independently from the data processing steps.

6.2 Data flagging

We take the nightly antenna flags from H22a to flag triads formed by such antennas. These flags are informed by specially designed metrics for detecting malfunctioning antennas (Storer et al. 2022) and by the redundant-baseline calibration process of H22a which is able to identify particularly non-redundant antennas (cf. Dillon et al. 2020). The exact reason for flagging individual antennas can be found in Dillon (2021) and the notebooks referenced therein. We further adopt the time flags listed in tables 2 and 3 in H22a that are due to broad-band RFI and digital system failures. Time flags that are due to calibration issues are not applied, as these may not be relevant to the closure phase analysis. We also flag times during which the sun is above the horizon and LST's at which $V_{\text{eff}} < 5 \text{ Jy}$. Setting this condition on V_{eff} , we avoid instances where one of the visibility amplitudes happens to be close to zero or is of the order of the thermal noise because of the coincidence of the orientation of foreground sources. First, the closure phase is poorly defined if one of the visibility amplitudes is zero. We observed that this can cause poles in the closure phase spectrogram, which in return can lead to excess power in the power spectrum. Secondly, the noise variance of the scaled closure phase will tend to zero together with V_{eff} . This has a strong effect further down the analysis pipeline (see Section 6.5), where we compute an inverse variance weighted average across LST. That is, LST's at which V_{eff} is small relative to the visibility noise variance will be weighted disproportionately high relative to other LST's. We find that a threshold of $V_{\text{eff}} \sim 5 \text{ Jy}$ eliminates these unwanted effects. Across all times, both polarization products and both triad classes the data is flagged 7 per cent of the time because of this condition. For model data, we find this flagging condition to be inadequate, since the visibility amplitudes are generally lower due to unmodelled diffuse emission. We therefore flag the model data by hand in regions where one of the visibility amplitudes crosses zero, resulting in a flagging frequency of 6 per cent. We use the model flags on the data and vice versa to prevent a bias between the two.

Other than that, we do not flag data on a per-integration and per-frequency basis and rely instead on robust averaging to filter out remaining RFI (see Section 6.3). In particular, we do not use the final ‘by-hand’ flags of H22a.

Fig. 5 shows the amount of data used after flagging. As can be seen in the left plot, the observing season can be divided into four epochs, where, except for the fourth epoch, each successive

epoch has an increased number of triads in use. The reason for the decreased number of triads in the fourth epoch can be traced to an increased number of malfunctioning antennas, despite there being more antennas connected (cf. Dillon 2021, and notebooks referenced therein). The number of triads ranges from 13 (EQ29, Epoch 4) to 36 (EQ14, Epoch 4). Note that some of these triads share a baseline and hence do not have independent noise. The right-hand plot in Fig. 5 shows the number of nights covering a given LST, ranging from 14 night at $\sim 22 \text{ h LST}$ to 82 unflagged nights at $\sim 7 \text{ h LST}$.

6.3 Averaging closure phases

Data containing the same sky-based signal can be averaged coherently. Under the assumption of perfect redundancy, the closure phase data has two axes suited for coherent averaging, namely, the repeated LST-integration of different nights and triads from the same class. Indeed, the night-to-night variations of the closure phase can be attributed to noise and RFI. The variations among nominally redundant triads, on the other hand, are clearly non-random and are caused by non-redundancies (see Fig. 6). For visibilities, the loss of sensitivity to the cosmological signal due to non-redundancy has been found to be at a level of 1–2 per cent (Choudhuri, Bull & Garsden 2021, H22a). It has yet to be established, if these estimates also apply to the closure phase, but, for the time being, we assume that the loss of sensitivity under averaging of nominally redundant triads is at a similar level.

In our processing pipeline, we first average the complex exponentials of the closure phases across nights. This is done using a geometric median, which is defined as the data point that minimizes its Euclidean distance to all other data points on the complex plane, i.e.

$$\hat{\phi} = \arg \min_{\phi_n} \sum_m |e^{i\phi_n} - e^{i\phi_m}|, \quad (8)$$

where the subscripts stand for the Julian Dates. This estimator has the advantage of being robust to outliers, such as RFI, while also respecting the circularity of the phase. The application of the geometric median is the only measure we take to mitigate the imprint of RFI on the final power spectrum. A persistent or repeated signal will not be eliminated by this treatment, therefore necessitating a band that is free of such effects (cf. Section 4 and Appendix A).

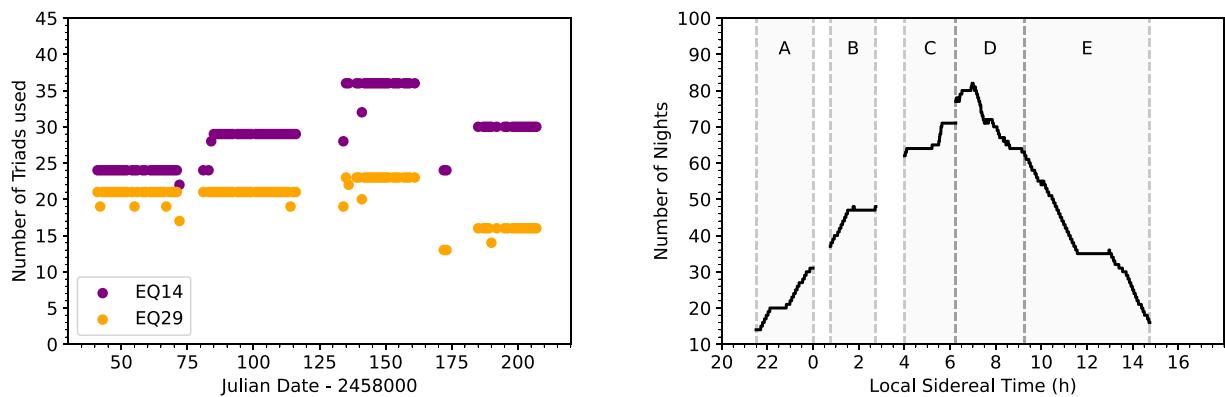


Figure 5. The number of triads after flagging for a given triad class and parallel polarization product (East–West, North–South) as a function of Julian Date (left-hand panel) and the number of nights covering a given Local Sidereal Time for fields A, B, C, D, and E (right-hand panel). In the first three epochs, the number of triads increases as more antennas were put into operation.

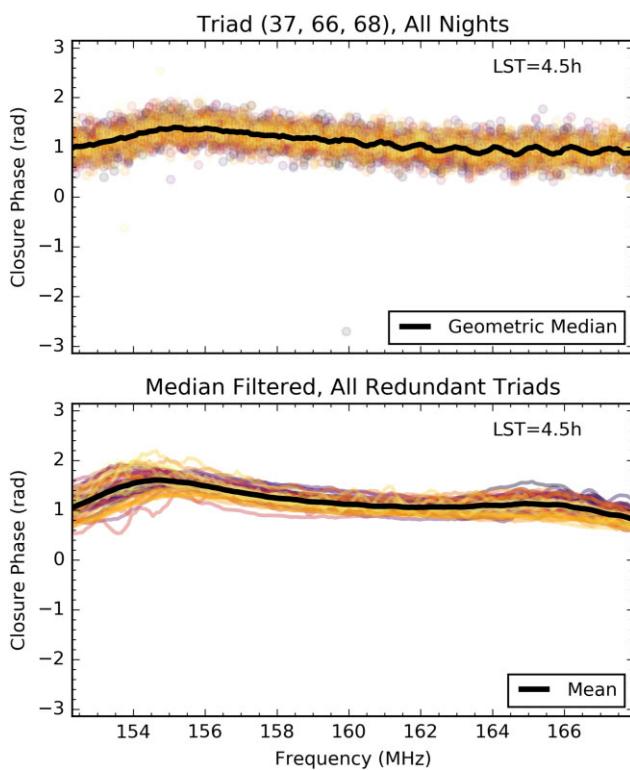


Figure 6. Closure phase spectra of data used in this analysis. The top plot shows the spectra for one triad across all nights and the median-averaged and time-averaged spectrum in black. The night-to-night variations can be attributed to thermal noise and RFI outliers. The averaged spectrum reveals a sinusoidal systematic effect with a period of ~ 1 MHz. The bottom plot shows the filtered and averaged spectra for all triads. Averaging redundant triads reduces the spectral ripple considerably (black line).

Note that in the visibility processing (H22a) the four epochs (see Section 6.2) are at first averaged independently to allow for better systematics mitigation and statistical tests. As these processing steps are not part of this analysis, we include all epochs in the median-average. This increases the efficiency of the median at rejecting outliers.

Next in the pipeline, we average some neighbouring LST integrations. Here, we average in intervals of 171.2 s (16 integrations) using the arithmetic mean. The averaging in time is justified by the

invariance of the closure phase to the translation of the sky. However, the antenna beam breaks this symmetry, as a result of which we expect the averaging to introduce a loss of sensitivity to the cosmological signal. Using our model data, we established the expected scale of this loss to be ~ 2 per cent. The method we used to determine this loss is detailed in Appendix B of this paper.

At this point, we could, in principle, average the closure phases of redundant triads. However, deferring this averaging until after the computation of cross-power spectra will allow us to omit cross-terms between triads that share a baseline (see Section 6.5). This has the advantage of mitigating the effect of baseline based systematics that are coherent across nights (see Thyagarajan et al. 2020).

6.4 Forming cross-power spectra

We divide the data into two bins containing nights with odd and even Julian Dates, respectively, and perform the averaging described in Section 6.3 separately in each bin. This allows us to compute cross-power spectra between the two bins, defined as

$$P_{\nabla}(\kappa_{\parallel}) = \left(\frac{c^2}{2k_B \nu^2} \right)^2 \left(\frac{X^2 Y}{\Omega B_{\text{eff}}} \right) \times \tilde{\Psi}_{\nabla}(\tau) \overline{\tilde{\Psi}'_{\nabla}}(\tau), \quad (9)$$

where $\tilde{\Psi}_{\nabla}$ and $\tilde{\Psi}'_{\nabla}$ are two delay spectra drawn from the first and second bin, respectively, and the bar denotes complex conjugation. The scaling factor is taken from the standard visibility power spectrum, where the first bracket converts flux density to brightness temperature, $X^2 Y$ converts spectral and angular units to cosmological distances (e.g. Liu et al. 2014), and Ω and B_{eff} normalize the power by the integrated squared beam response and the effective bandwidth, respectively. With this definition, the cross-power spectrum has units of ‘pseudo’ $\text{mK} (\text{Mpc } h^{-1})^{-3}$ and is a function of $\kappa_{\parallel} = 2\pi\tau/X$, which has units of ‘pseudo’ $(\text{Mpc } h^{-1})^{-1}$. In line with the practice introduced in Thyagarajan et al. (2020), we use the ‘pseudo’ to emphasize that these units are not physical but are instead used as approximations to the real distance scales and brightness temperatures (see Thyagarajan & Carilli 2020). Note that in computing the scaling factor $X^2 Y$ we require cosmological parameters. We take these parameters from Planck Collaboration VI (2020) and use $H_0 = 100 h (\text{km s}^{-1}) \text{Mpc}^{-1}$ for the Hubble constant.

Ideally, the two delay spectra $\tilde{\Psi}_{\nabla}$ and $\tilde{\Psi}'_{\nabla}$ should contain the same contribution due to sky, but independent thermal noise and transient systematic effects. The sky should therefore occupy the positive real part of $P_{\nabla}(\kappa_{\parallel})$, while the noise can introduce negative

and imaginary components. Hence, we can further reduce the noise by averaging different cross-power spectra for which, statistically, the cosmological signal has the same underlying power spectrum. This averaging is further detailed in the next section.

In Section 7, we also present a form of the power spectrum which was first introduced in Thyagarajan et al. (2020) and is mathematically akin to the widely used cosmological variance. It is computed as $\Delta_{\nabla}^2(\kappa) = \kappa^3 P_{\nabla}(\kappa_{\parallel})/(2\pi^2)$, where the units are mK^2 and $\kappa^2 = \kappa_{\parallel}^2 + \kappa_{\perp}^2$. The perpendicular component is defined in analogy to the perpendicular k -modes in the visibility analysis as $\kappa_{\perp} = 2\pi(|\mathbf{b}|/\lambda)/Y$, where $|\mathbf{b}|$ is the baseline length of an equilateral triad and λ is the wavelength. In the context of the closure phase analysis, this definition of κ is only adequate for equilateral triad classes, where all baselines are sensitive to the same perpendicular k -modes (cf. Thyagarajan & Carilli 2020). However, in this paper it will be useful for obtaining upper limits that correspond approximately to the cosmological variance and, thus, can be compared to the results of standard analysis techniques.

6.5 Averaging cross-power spectra

After forming cross-power spectra the data for a given polarization product has shape $N_t \times N_t \times N_l \times N_{\tau}$, where $N_{\tau} = N_f$ is the number of delays. We proceed with averaging across the two redundant triad axes and the LST axis, noting that the later is an incoherent average making use of the assumed isotropy of the 21 cm signal.

The noise as well as systematic effects vary across triads and by the pointing direction of the telescope, which for a zenith pointing array, such as HERA, coincides with the LST. This motivates a weighted average

$$\tilde{\Psi}_{\nabla}(\tau) \overline{\tilde{\Psi}'_{\nabla}}(\tau) = \sum_{i,j,t} w_{ijt} \tilde{\Psi}_{\nabla}(i, t, \tau) \overline{\tilde{\Psi}'_{\nabla}}(j, t, \tau), \quad (10)$$

where we choose the normalized weights w_{ijt} so that

$$w_{ijt} \propto \begin{cases} w_{it} w_{jt}, & i \text{ and } j \text{ do not share a baseline.} \\ 0, & i \text{ and } j \text{ share a baseline.} \end{cases}, \quad (11)$$

where w_{it} is the inverse variance of triad i at time t . We estimate w_{it} by differencing the scaled complex closure phases from different nightly bins and computing the variance along the frequency axis. Omitting cross-terms between triads with shared baselines prevents baseline based systematic that are coherent across nights from contaminating the power spectrum. This effect has been observed in Thyagarajan et al. (2020), where the cross-power between identical triads produced a positive bias in the high-delay region of the power spectrum.

Assuming an unpolarized and isotropic 21 cm signal, we further average over the two polarization products and in bins of $|\kappa_{\parallel}|$. The cross-power spectra thus obtained are the ones presented in Section 7.

6.6 Estimating uncertainty in the power spectra

We estimate the uncertainty of the real parts of the power spectra from an estimate of their variance. To do this, we divide the data into four bins along the JD axis. Using these four bins, we form six cross-power spectra from which we can compute three independent differences. Any true signal is cancelled in these differences, leaving us with three independent realizations of the noise (see Thyagarajan et al. 2020 and Tan et al. 2021 for similar approaches). We then perform the averaging described in Section 6.5 and obtain error bars by computing the root-mean-square (RMS) of the real parts of the three noise realizations. The errors thus obtained are themselves subject

to uncertainties because of the small sample size. To decrease this uncertainty, we smooth the error bars in quadrature along the delay axis with a flat kernel of width 3. As a result, neighbouring power spectrum errors will be correlated, but instead we have effectively increased the sample size from 3 to 9. Averaging negative and positive $|\kappa_{\parallel}|$ bins further increases the sample size to 18. The relative uncertainty of the resulting error bars is about 17 per cent.

The advantage of this method is that each $|\kappa_{\parallel}|$ bin has its own error, which captures the variance due to thermal noise as well as systematic noise and accounts for cross-terms between the noise and any underlying signal. The subsequent smoothing, however, makes the assumption that neighbouring error bars have similar values. This assumption is accurate for thermal-like white noise, but can lead to biased estimates in regions where the noise-signal cross-terms start to dominate. Fortunately, we are interested in regions where these cross-terms are minimal and the errors are at most overestimated.

6.7 Estimating upper limits

An upper limit x_{UL} on the closure phase power spectrum is implicitly given by the probability $\text{Pr}(0 < \mu < x_{\text{UL}}) = 1 - \alpha$, where μ is the true power. We choose $\alpha = 0.05$ with which the upper limit defines a 95 per cent confidence interval. Assuming that the power spectrum data x is drawn from a normal distribution $N(\mu, \sigma)$ with expectation μ and variance σ , we can apply Bayes theorem:

$$\text{Pr}(0 < \mu < x_{\text{UL}}) = \frac{\int_0^{x_{\text{UL}}} N(x|\mu, \sigma) d\mu}{\int_0^{\infty} N(x|\mu, \sigma) d\mu}, \quad (12)$$

where setting the lower integral limit to zero incorporates our prior knowledge of the true power spectrum and the denominator normalizes the posterior probability. Computing the integrals, we find

$$\text{Pr}(0 < \mu < x_{\text{UL}}) = \frac{\text{erf}\left((x_{\text{UL}} - x)/\sqrt{2}\sigma\right) + \text{erf}\left(x/\sqrt{2}\sigma\right)}{1 + \text{erf}\left(x/\sqrt{2}\sigma\right)}, \quad (13)$$

where erf is the error function. Defining $E := \text{erf}(x/\sqrt{2}\sigma)$, we can solve for the upper limit

$$x_{\text{UL}} = \sqrt{2}\sigma \left(\text{erf}^{-1}(1 - (1 + E)\alpha) + \text{erf}^{-1}(E) \right). \quad (14)$$

Although the closure phase delay power spectra do not initially follow a normal distribution, they will converge to normality as more data is averaged (cf. Tan et al. 2021). This is a consequence of the Central Limit Theorem. The degree to which the noise in our power spectra is normally distributed is investigated in Section 7. A similar derivation of this result can be found in the appendix of Li et al. (2019).

7 RESULTS

The final averaged power spectra for fields A, B, and C are shown in Fig. 7 (EQ14) and Fig. 8 (EQ29), where the left-hand and right-hand columns show the power spectra obtained from data and the corresponding models, respectively. The power spectra obtained on fields D and E are shown separately in Fig. 9. Since fields D and E where not modelled, we show the power spectra obtained from EQ14 triads in the left-hand column and those obtained from EQ29 triads in the right-hand column. Note that the power spectra are scaled by a factor of 2/3 to account for the fact that the fluctuations in the closure phase enter through three visibilities and that, statistically, the phase

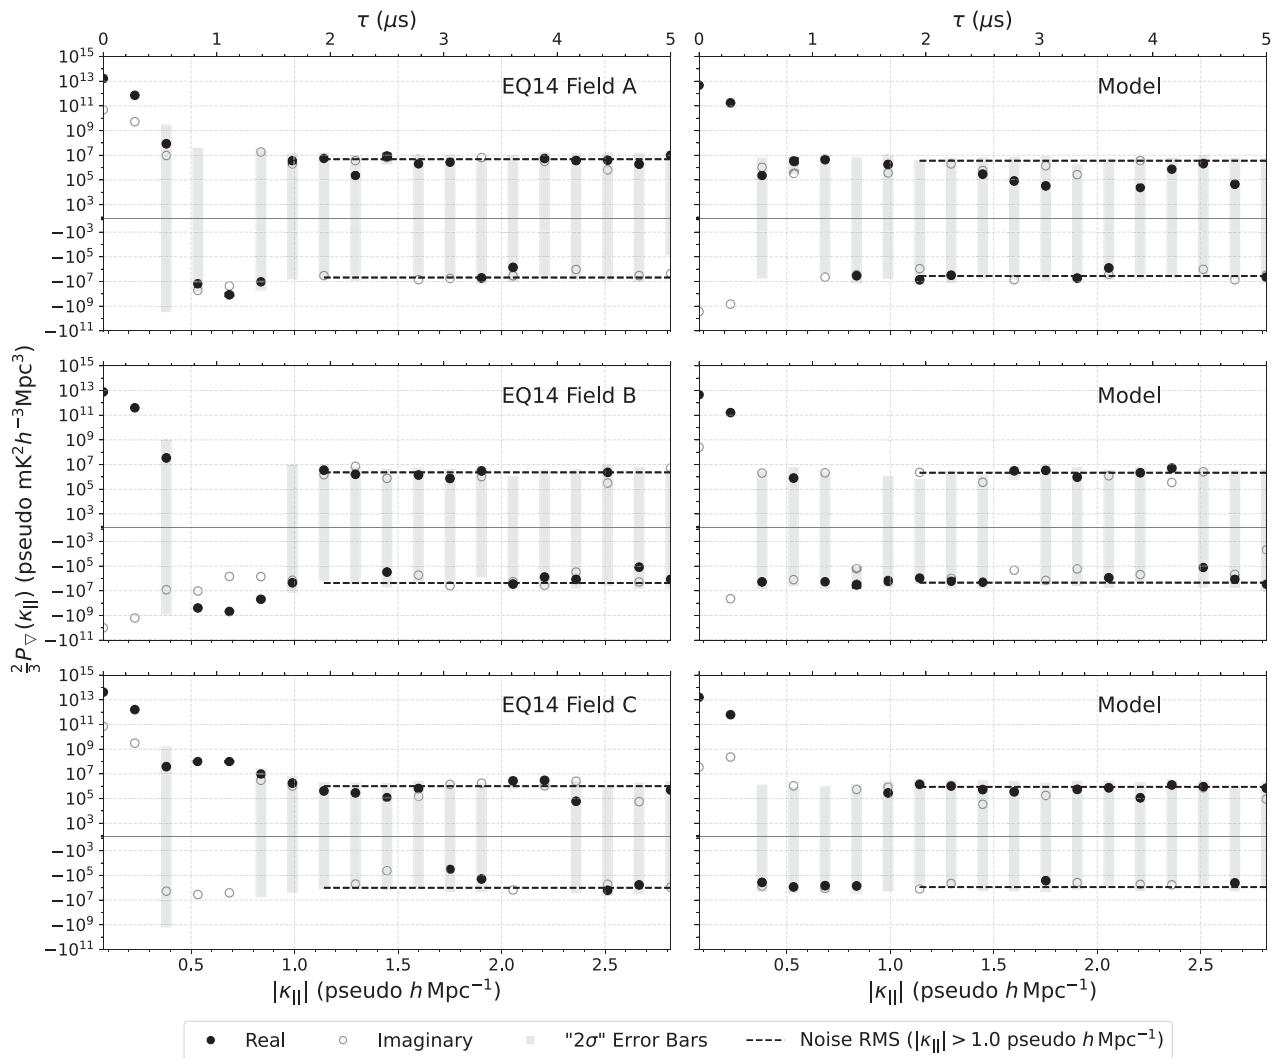


Figure 7. Closure phase delay power spectra for equilateral 14.6 m triads and HERA-fields A, B, and C from data (left-hand panel) and the corresponding models (right-hand panel). The power is scaled by two-thirds to account for the fact that the closure phase is formed by three different phases, which each, statistically, recover half the visibility fluctuations. The observed band is centred at ~ 164 MHz, which corresponds to a redshift of ~ 7.7 . The axes are given in ‘pseudo’ units in order to discriminate between true cosmological scales and the approximate scales used in this analysis. In addition to the real parts of the cross-power spectra (black circles), we also show the imaginary parts (grey circles), the ‘ 2σ ’ error bars (grey boxes) and the root-mean-squares of the power spectrum errors at $|\kappa_{||}| > 1.0$ ‘pseudo’ $h \text{ Mpc}^{-1}$ (black dashed lines). The power axes are linear in the regions between $\pm 10^3$ ‘pseudo’ $\text{mK}^2 \text{h}^{-3} \text{Mpc}^3$ and logarithmic otherwise.

only recovers half of the visibility fluctuations (cf. Thyagarajan & Carilli 2020). The plots show the real part of the power spectra (black filled circles) as well as the imaginary part (grey circles). The latter should be a good proxy of the noise including signal-noise cross-terms and can be compared to the real parts to identify excess power above the noise level. Excess power in the imaginary part, on the other hand, would either indicate significant systematics-noise cross-terms or certain classes of systematics that vary in phase across nights or triads. The plots also show the 2σ error bars (grey boxes) and the RMS of the noise at $|\kappa_{||}| > 1.0$ ‘pseudo’ $h \text{ Mpc}^{-1}$ (black dashed lines). The power axes is in a symmetric log-scale with a linear region between $\pm 10^3$ ‘pseudo’ $\text{mK}^2 \text{h}^{-3} \text{Mpc}^3$, allowing us to show positive and negative powers on the same plot.

We can distinguish between three regions in the power spectra. The first region is at $|\kappa_{||}| \lesssim 0.3$ ‘pseudo’ $h \text{ Mpc}^{-1}$ where the largest fraction of the power is concentrated. This power is attributed to

the spectrally smooth foreground emissions and peaks between 10^{12} and 10^{14} ‘pseudo’ $\text{mK}^2 \text{h}^{-3} \text{Mpc}^3$ depending on the observed field. That is, fields containing strong emissions (e.g. Field E) peak higher than fields with weak emissions (e.g. Field B). Moreover, the power spectra of EQ29 peak lower than those of EQ14, because of the weaker response to large scale emissions on the sky. Since this analysis takes a foreground avoidance approach, we do not set any upper limits in this region of the power spectra.

The second region is characterized by peaks of amplitude 10^8 – 10^9 ‘pseudo’ $\text{mK}^2 \text{h}^{-3} \text{Mpc}^3$, which are centred at delays of $1 \mu\text{s}$ and stretch out to delays of about $1.4 \mu\text{s}$. These peaks are not seen in the models and have uncertainties that are not consistent with thermal-like noise. In closure phase spectra, this effect appears as a spectral ripple with a period of ~ 1 MHz. It is not a multiplicative antenna based effect, as it would otherwise be eliminated in the closure phase. In fact, in a visibility based analysis Kern et al. (2019), Dillon,

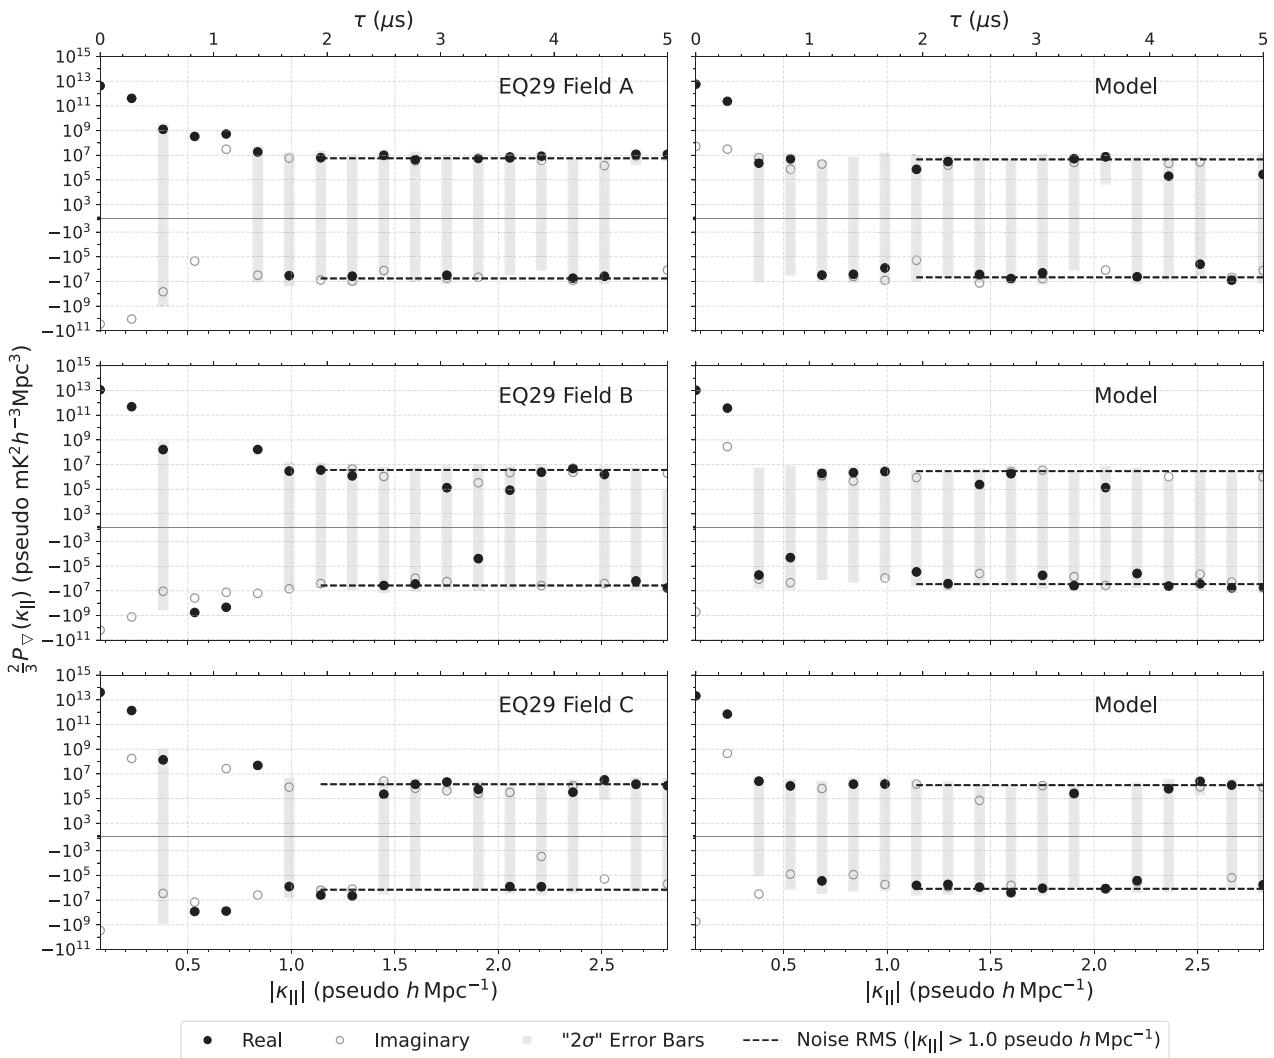


Figure 8. Same as Fig. 7 but for equilateral 29-m triads.

Parsons & Kern (2021) and H22a find evidence that the ripple is a baseline-dependent systematic effect caused by over-the-air coupling between different array elements. Furthermore, the effect is found to vary slowly in time but strongly across baselines (cf. Kern et al. 2019). Both triad classes are equally affected by this systematic effect. In the averaged power spectra, the peaks take on negative as well as positive values, indicating a partial de-correlation between triads. This agrees with the findings of Thyagarajan et al. (2020) that the peaks are considerably suppressed when excluding the cross-power between identical triads (cf. Section 6.5). While the spectral ripple is fitted for and subtracted in the visibility processing (Kern et al. 2019), we do not use filtered data to form closure phases. The subtraction violates closure properties and would have to be performed on raw data rather than averaged data, which would be computationally expensive. Filtering the systematic directly from the closure phase also has its problems. The filtering in Kern et al. (2019) makes use of the fact that the systematic varies slowly in time, therefore having a fringe rate close to zero. However, the same applies to closure quantities, which can be shown to be invariant under the translation of the sky. Filtering the systematic from the closure phase would therefore also remove a large part of the cosmological signal. Hence,

we take the more conservative approach and completely avoid the affected delay modes.

The third region at $\tau \gtrsim 1.4 \mu\text{s}$ or equivalently $|\kappa_{||}| \gtrsim 0.8 \text{ 'pseudo' } h \text{ Mpc}^{-1}$ is dominated by noise, meaning that the power is comparable to its overall variance. In other words, the error bars mostly cross $P_{\nabla} = 0$. Note that by ‘noise’ we mean thermal-like noise as well as non-thermal-like noise such as RFI and instrumental effects. For thermal-like noise the high-delay region of the power spectrum should fluctuate randomly around zero. However, some of the power spectra have a tendency towards positive values in their real parts. To see this more clearly, we investigate the cumulative distribution functions (CDF) of the real and imaginary parts of the power spectrum at $|\kappa_{||}| > 1 \text{ 'pseudo' } h \text{ Mpc}^{-1}$ as well as the combined real and imaginary parts of the differenced power spectra P_{Diff} used to obtain error bars (see Section 6.6). We use the latter as a proxy for the noise. The CDFs are plotted in Fig. 10 together with the CDF of a Gaussian distribution (grey solid line) with the same variance as P_{Diff} . Note that the CDFs of the differenced power spectra (magenta dashed line) are in good agreement with the Gaussian CDFs, which is a consequence of averaging many power spectra together (Central Limit Theorem). To quantify the consistency between the different CDFs, we use two statistical tests, the Shapiro-Wilk (SW, Shapiro &

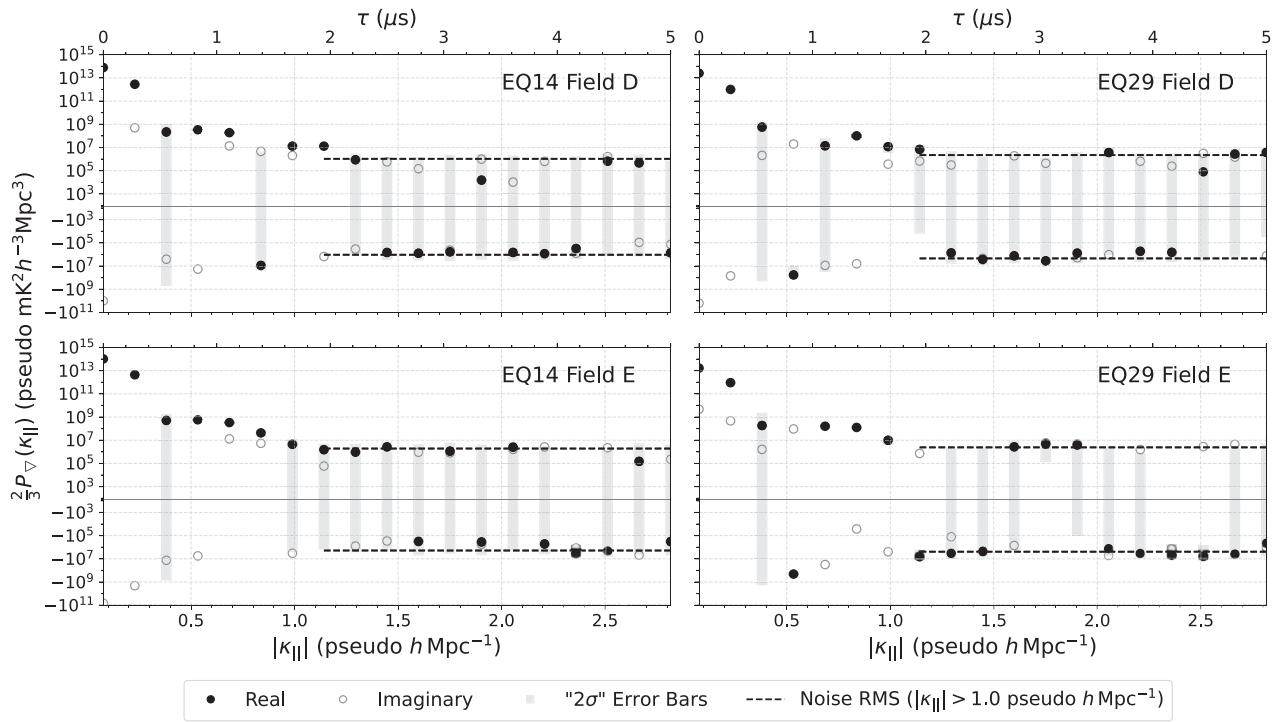


Figure 9. Closure phase delay power spectra of HERA-fields D and E. Unlike the other fields, these two fields cannot be modelled adequately. The left-hand and right-hand columns show the power spectra for equilateral 14.6 and 29.2 m triads, respectively.

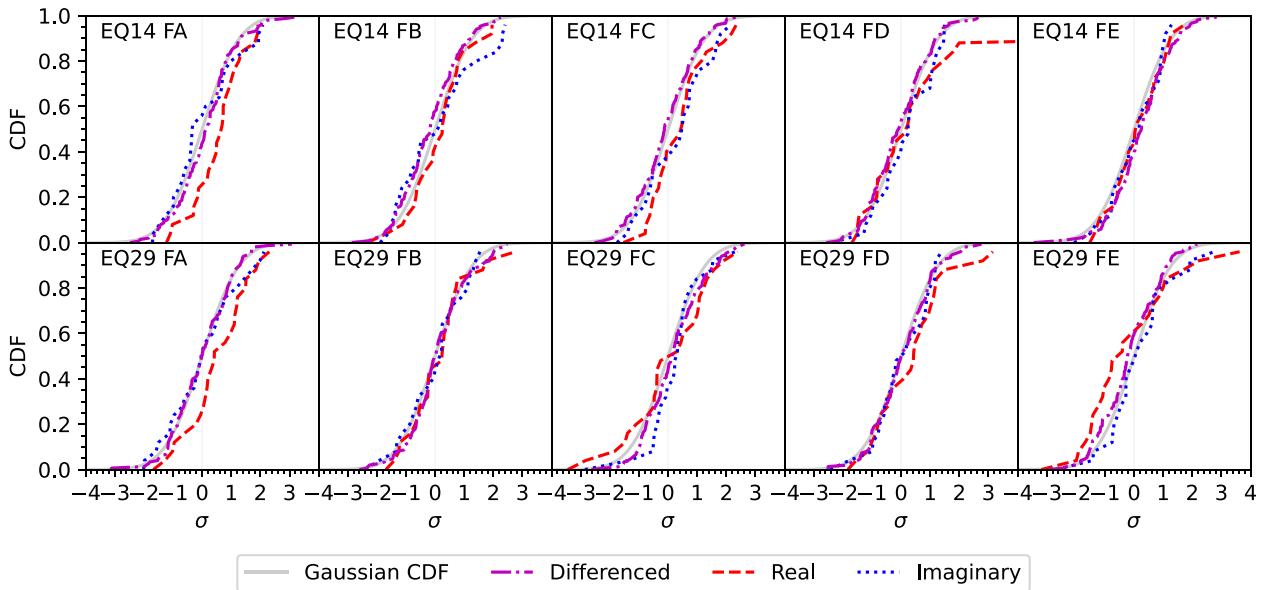


Figure 10. The cumulative distribution functions of the cross-power spectra at $\kappa_{||} > 1.0$ 'pseudo' $h \text{ Mpc}^{-1}$. The red dashed and the blue dotted lines show the CDFs of the real and imaginary parts of the power spectra, respectively, while the magenta dash-dotted line shows the CDF of the noise realization obtained by differencing closure phases with the same underlying sky signal. As a reference, we also plot a Gaussian CDF with the same variance as the noise (grey solid line). The abscissa is given in units of standard deviations of the differenced closure phases.

Wilks 1965) and the Anderson-Darling (AD; Anderson & Darling 1952) test, the results of which are shown in Table 2. We use the former to test the null-hypothesis H_N that the differenced power spectra are drawn from a normal distribution. In all cases the SW test fails to reject H_N at the 5 per cent level. We use the AD test to test

the null-hypothesis H_N that the noise and the real or imaginary parts of the power spectra at $|\kappa_{||}| > 1$ 'pseudo' $h \text{ Mpc}^{-1}$ follow the same distribution. For fields A, B, and C, we also use the AD test to test the data against the corresponding model at $|\kappa_{||}| > 1$ 'pseudo' $h \text{ Mpc}^{-1}$. Unlike the SW test, the AD test does not make any assumptions about

Table 2. Summary of the results from statistical tests performed on the cross-power spectra at $\kappa_{||} > 1.0$ ‘pseudo’ $h \text{ Mpc}^{-1}$. The variables W and A^2 denote the test statistics of the Shapiro-Wilk and Anderson-Darling test, respectively. We set our critical value for rejecting the null-hypothesis H_N at 5 per cent. See the table notes for further elaboration.

Test Samples	Shapiro-Wilk ¹			Anderson-Darling ²		Anderson-Darling ²		Anderson-Darling ²	
		P_{Diff}		P_{Diff} and $\Re\{P_{\nabla}\}$		P_{Diff} and $\Im\{P_{\nabla}\}$		$\Re\{P_{\nabla}\}$ model and data	
Name	W	p -value	reject H_N	A^2	reject H_N	A^2	reject H_N	A^2	reject H_N
EQ14 Field A	0.992	0.605	False	1.889	False	-0.043	False	4.877	True
EQ14 Field B	0.994	0.762	False	-0.062	False	0.425	False	-0.976	False
EQ14 Field C	0.994	0.842	False	0.971	False	0.560	False	0.140	False
EQ14 Field D	0.993	0.627	False	-0.051	False	-0.474	False	-	-
EQ14 Field E	0.989	0.283	False	-0.627	False	-0.413	False	-	-
EQ28 Field A	0.990	0.342	False	2.319	True	-0.523	False	1.787	False
EQ28 Field B	0.985	0.098	False	-0.904	False	-0.605	False	0.883	False
EQ28 Field C	0.991	0.474	False	0.450	False	-0.399	False	-0.106	False
EQ28 Field D	0.993	0.684	False	-0.172	False	-0.912	False	-	-
EQ28 Field E	0.993	0.661	False	0.936	False	-0.185	False	-	-

Notes. ¹ The null hypothesis H_N of the Shapiro-Wilk test is that P_{Diff} is drawn from a normal distribution. ² The null hypothesis H_N of the two sample Anderson-Darling test is that two samples are drawn from the same distributions. Here, we test P_{Diff} against $\Re\{P_{\nabla}\}$ and $\Im\{P_{\nabla}\}$, and the models against the data. The critical values of the test statistic A^2 are 0.325, 1.226, 1.961, 2.718, 3.752, 4.592 and 6.546 at 25, 10, 5, 2.5, 1, 0.5, and 0.1 per cent significant levels, respectively.

the shape of the underlying distributions. For the imaginary parts, the AD test fails in all cases to reject H_N at the 5 per cent level. For the real parts, H_N is rejected for the power spectrum of EQ29 on Field A, and not rejected otherwise. The test of the data against the models, on the other hand, fails to reject H_N for all power spectra except that of EQ14 on Field A. The CDFs of the real parts of Field A show a shift towards positive values, which is consistent with a ‘detection’ of a signal. This signal is of unknown origin and could be due to a variety of different effects such as RFI or digital artefacts. Since Field A is the least sensitive field, it should be less affected by low-level RFI. On the other hand, it is covered by fewer nights, thus making the median-averaging across nights less effective at rejecting RFI.

Although not rejected by the AD test at a 5 per cent level, the CDFs of the real parts of Field D extend towards high positive values. Looking at the power spectrum, we see a peak of excess power at delays of about 2 μs . H22c identify a similar feature in the visibility delay spectrum and trace it to the polarized emissions of the pulsar PSR J0742-2822 (Lenc et al. 2017). The polarization direction is rotated as the radiation passes through magnetic fields. This effect, known as Faraday rotation, is frequency dependent, thus leaving an imprint on certain delay modes in the power spectrum. In the visibility analysis, the Faraday effect can be suppressed below the current noise level by forming pseudo Stokes I visibilities (H22c). This is not possible for the closure phase approach, as each polarization will have independent gains, meaning that the closure phase of a pseudo Stokes I visibility would lose its desirable properties. Consequently, we cannot reliably interpret fields with strong highly polarized sources at high rotation measure.

The real parts of the power spectra on Field C and of EQ29 on Field E also have a somewhat higher AD statistic A^2 compared with other power spectra, albeit not high enough to be rejected at the 5 per cent level. In the CDFs of EQ29 Fields C and E, we see that both have an extended tail towards negative values, which could indicate the presence of a systematic effect that is uncorrelated across triads or nights. Moreover, since we do not see the same effect in the power spectra of the EQ14 triads, this demonstrates that the two triad classes are affected differently by systematic effects.

Comparing the noise RMS at $|\kappa_{||}| > 1$ ‘pseudo’ $h \text{ Mpc}^{-1}$ between the data and the models, we find that the RMS’s of the models are on average a factor of ~ 1.2 lower than those of the data. Possible

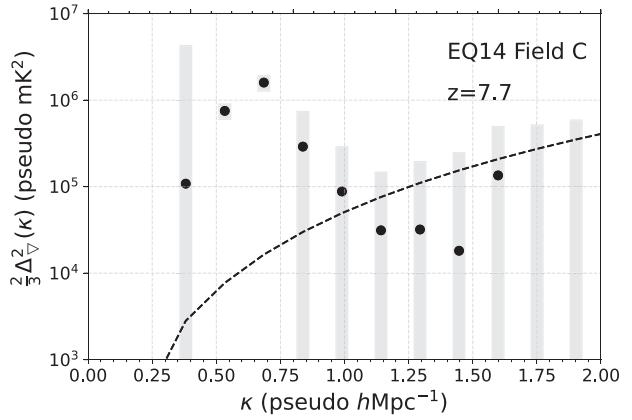


Figure 11. The closure phase power spectrum of EQ14 Field C in units of mK^2 (cf. Fig. 7). Data points with negative powers are only shown by their error bars (grey boxes) on this plot. The dashed line indicates the RMS of the power spectrum errors at $|\kappa_{||}| > 1.0$ ‘pseudo’ $h \text{ Mpc}^{-1}$.

explanations of this discrepancy are the presence of non-thermal effects (e.g. RFI), an underestimate of the system temperature, or other model inaccuracies. As expected, we also find that the noise RMS of EQ14 power spectra are lower than those of EQ29 power spectra, since there are fewer triads in EQ29 that can be averaged.

Despite the presence of non-thermal like effects in some of the power spectra, we can use the Bayesian framework described in Section 6.7 to set upper limits on the cosmological 21 cm signal. Here, we provide the limits for our deepest field and triad class, Field C and EQ14. Fig. 11 shows the power spectrum in units of ‘pseudo’ mK^2 and Table 3 shows the associated upper limits at different κ -modes. The strongest noise-limited upper limit is $(372)^2$ ‘pseudo’ mK^2 at 1.14 ‘pseudo’ $h \text{ Mpc}^{-1}$. We re-emphasize that this limit should only be interpreted as approximations to the physical distance and brightness scales of the conventional 21 cm power spectrum (cf. Thyagarajan & Carilli 2020).

Table 3. The upper limits $(2/3)\Delta_{\nabla}^2 \text{UL}$ at 95 per cent confidence and the standard deviation σ of the closure phase delay power spectrum obtained from EQ14 triads on Field C between 0.38 and 1.45 h Mpc^{-1} .

κ 'pseudo' h Mpc^{-1}	$(2/3)\Delta_{\nabla}^2 \text{UL}$ 'pseudo' mK^2	σ 'pseudo' mK^2
0.38	$(2058)^2$	$(1443)^2$
0.53	$(937)^2$	$(274)^2$
0.69	$(1367)^2$	$(401)^2$
0.84	$(824)^2$	$(473)^2$
0.99	$(518)^2$	$(317)^2$
1.14	$(372)^2$	$(241)^2$
1.29	$(431)^2$	$(285)^2$
1.45	$(491)^2$	$(338)^2$

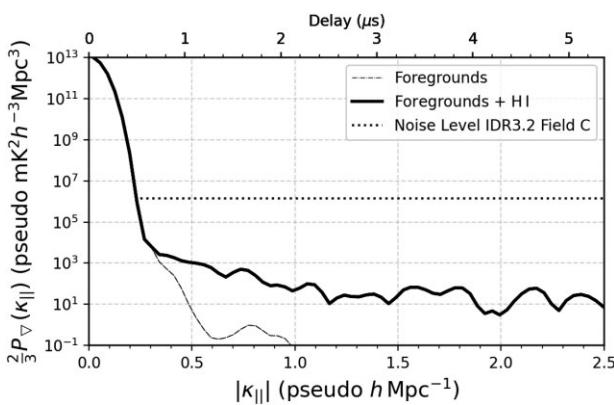


Figure 12. A simulated delay power spectrum of the frequency band 152.25–167.97 MHz ($z \sim 7.9$). The dashed line is the power spectrum of a sky with foregrounds only (GLEAM sources at 5 h LST) and the solid line is that of a sky including a cosmological H I signal (Faint Galaxies, 21cmFast). While the foreground dominate the low $|\kappa_{||}|$ -modes, the H I signal dominates at $|\kappa_{||}| > 0.25$. Also shown is the noise level at $|\kappa_{||}| > 1.0$ of our most sensitive field (horizontal dotted line). An improvement in sensitivity of at least two orders of magnitude is required for a hypothetical detection of this commonly used EoR model.

8 SUMMARY

We present closure phase delay power spectra using data from a full season of HERA Phase I observing. The data was observed over 94 unflagged nights using 48 antennas. We show power spectra for two triad classes, equilateral 14.6 and 29.2-m triads, and five separate LST ranges. Using our most sensitive field, Field C, which covers 2.25 h of LST centred at 5.125 h, we provide upper limits on the closure phase power spectrum of equilateral 14.2-m triads and find a noise-limited upper limit of $(372)^2$ 'pseudo' mK^2 at 1.14 h Mpc^{-1} . Comparing the power spectrum RMS at $\kappa > 1.0 \text{ h Mpc}^{-1}$ with that of Thyagarajan et al. (2020) at $\kappa > 0.85 \text{ h Mpc}^{-1}$, we find an improvement in sensitivity by a factor of ~ 26 . The limits reported here are only approximately related to the true distance and brightness scales of the redshifted 21 cm power spectrum, which is why they are given in 'pseudo' units. A more careful interpretation of results obtained with the closure phase requires detailed forward-modelling of the sky signals. However, making use of the approximation with 'pseudo' units, we find that our most sensitive noise-limited upper limit is a factor of ~ 5 above that of H22a at $\kappa = 1.16 \text{ h Mpc}^{-1}$.

As can be seen in Fig. 12, the sensitivity of the power spectrum

needs to be improved by at least two orders of magnitude to achieve a hypothetical detection of the EoR model used here (see Section 5). When fully operational, HERA will have 320 antennas which can be used for the closure phase analysis. This amounts to ~ 7 times more nominally redundant triads than are used in this analysis, which can be combined coherently to improve the sensitivity by the same factor. Moreover, HERA will observe all year round for ~ 12 h per night. Combining different fields incoherently could potentially double the sensitivity (~ 4 more data in LST). Repeated nights can also be combined coherently. Taken together, it should therefore be possible to improve the power spectrum sensitivity by a factor greater than 10^2 within the observing horizon of HERA (~ 5 yr). This should be sufficient for a detection of commonly assumed 21 cm signals at $\kappa \sim 0.4 \text{ h Mpc}^{-1}$ provided that we are not limited by systematic effects.

There are some limitations to the analysis as it is presented here. Unlike the visibility analysis, we are not able to model and subtract any baseline-dependent systematic effects such as the 1 μs spectral ripple without losing the desirable properties of the closure phase. Consequently, we do not have access to the lowest κ -modes which, otherwise, are expected to be most sensitive to the cosmological signal. The new HERA system uses fibre optic transmission lines, which may eliminate the 1 μs ripple and allow us to access lower κ -modes.

A further limitation is that we do not combine different triad classes in our analysis. While in the visibility analysis many different baselines can be averaged by the method of the so-called spherical averaging, there is no analogous method to do this for different triad classes. This is especially true for non-equilateral triangle shapes such as isosceles, scalene, or linear triads that are formed by baselines of different lengths. Including all these triad classes would lead to a considerable improvement in sensitivity, but as long as there is no physically motivated way to combine them to a single power spectrum, the classes will have to be analysed separately from one another. In other words, the power spectra of the different triad classes need to be considered as separate measurements, which, combined and in conjunction with forward-modelling, can be used to infer the presence of a signal. This combined analysis may therefore open up a pathway to exploiting the full sensitivity of the array.

Lastly, we find that Faraday rotated emissions may contaminate EoR window of the closure phase delay power spectrum. This effect could be mitigated by forming the closure phase from pseudo Stokes I visibilities. However, the antenna based gains are generally independent between different polarizations, meaning that we would lose some of the advantages of the closure phase approach. To completely bypass the problem of Faraday rotation, one would either have to directly measure Stokes I (i.e. by using circular antenna feeds if the emissions are unpolarized) or resort to polarization independent closure quantities such as the closure trace (Broderick & Pesce 2020). The implications of using the latter have yet to be investigated in detail. For HERA, a closure phase-based approach will need to avoid regions containing strong pulsars such as the Galactic Plane (e.g. Field D).

Despite these limitations, the closure phase analysis still retains advantages over the standard approach. Most importantly, our analysis is independent of multiplicative antenna-based effects, which allows us to bypass conventional calibration. As a result, we require considerably fewer analysis steps and expect fewer errors and systematic effects to be introduced in the data processing.

The closure phase analysis provides an alternative and independent method by which the 21 cm signal during the EoR can be searched for. Initially, the prime objective is a first detection of the signal. The interpretation of a detected signal will require extensive forward-

modelling of the sky, which includes the 21 cm signal as well as foregrounds, as the closure phase is a higher order (non-linear) interferometric quantity. Future work will explore the possibilities of inferring astrophysical properties of the IGM from the closure phase delay power spectrum.

ACKNOWLEDGEMENTS

The software developed for this analysis uses Python and the publicly accessible and open-sourced Python packages NUMPY (van der Walt, Colbert & Varoquaux 2011), SCIPY (Virtanen et al. 2020), ASTROPY (Astropy Collaboration 2022), and MATPLOTLIB (Hunter 2007).

This material is based upon work supported by the National Science Foundation under grants #1636646 and #1836019 and institutional support from the HERA collaboration partners. This research is funded in part by the Gordon and Betty Moore Foundation through Grant GBMF5212 to the Massachusetts Institute of Technology. HERA is hosted by the South African Radio Astronomy Observatory, which is a facility of the National Research Foundation, an agency of the Department of Science and Innovation.

PM Keller is funded by the Institute of Astronomy and Physics Department of the University of Cambridge via the Isaac Newton Studentship. NK gratefully acknowledges support from the MIT Pappalardo fellowship. AL acknowledges support from the New Frontiers in Research Fund Exploration grant program, the Canadian Institute for Advanced Research (CIFAR) Azrieli Global Scholars program, a Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery Grant and a Discovery Launch Supplement, the Sloan Research Fellowship, and the William Dawson Scholarship at McGill.

DATA AVAILABILITY

The code used for the analysis is publicly available at https://github.com/pm-keller/hera_closure and https://github.com/pm-keller/closure_sim. The former repository also contains the data shown in Figs 7, 8, 9, and 11 and Table 3 in a machine readable form.

REFERENCES

Aguirre J. E. et al., 2022, *ApJ*, 924, 85
 Anderson T. W., Darling D. A., 1952, *Ann. Math. Stat.*, 23, 193
 Astropy Collaboration 2022, *ApJ*, 935, 167
 Barry N., Hazelton B., Sullivan I., Morales M. F., Pober J. C., 2016, *MNRAS*, 461, 3135
 Barry N. et al., 2019, *ApJ*, 884, 1
 Beardsley A. P. et al., 2016, *ApJ*, 833, 102
 Blackman R. B., Tukey J. W., 1958, *Bell Syst. Tech. J.*, 37, 185
 Broderick A. E., Pesce D. W., 2020, *ApJ*, 904, 126
 Byrne R. et al., 2019, *ApJ*, 875, 70
 Byrne R., Morales M. F., Hazelton B. J., Wilensky M., 2021, *MNRAS*, 503, 2457
 Carilli C. L., Nikolic B., Thyagarajan N., Gale-Sides K., 2018, *Radio Sci.*, 53, 845
 Cheng C. et al., 2018, *ApJ*, 868, 26
 Choudhuri S., Bull P., Garsden H., 2021, *MNRAS*, 506, 2066
 Datta A., Bowman J. D., Carilli C. L., 2010, *ApJ*, 724, 526
 Daubechies I., 1988, *Commun. Pure Appl. Math.*, 41, 909
 DeBoer D. R. et al., 2017, *PASP*, 129, 045001
 Dillon J. S., 2021, HERA Memorandum
 Dillon J. S., Parsons A. R., 2016, *ApJ*, 826, 181
 Dillon J. S. et al., 2015, *Phys. Rev. D*, 91, 123011
 Dillon J. S. et al., 2020, *MNRAS*, 499, 5840
 Dillon J. S., Parsons A., Kern N., 2021, HERA Memorandum
 Eastwood M. W. et al., 2019, *AJ*, 158, 84
 Ewall-Wice A., Dillon J. S., Liu A., Hewitt J., 2017, *MNRAS*, 470, 1849
 Fagnoni N. et al., 2021, *MNRAS*, 500, 1232
 Furlanetto S. R., 2016, in Mesinger A. ed., *Astrophysics and Space Science Library* Vol. 423, *Understanding the Epoch of Cosmic Reionization: Challenges and Progress*. Springer, Cham, p. 247
 Greig B., Mesinger A., 2017, *MNRAS*, 465, 4838
 HERA Collaboration, 2022a, preprint ([arXiv:2210.04912](https://arxiv.org/abs/2210.04912))
 HERA Collaboration, 2022b, *ApJ*, 924, 51
 HERA Collaboration, 2022c, *ApJ*, 925, 221
 Hunter J. D., 2007, *Comput. Sci. Eng.*, 9, 90
 Hurley-Walker N. et al., 2017, *MNRAS*, 464, 1146
 Jennison R. C., 1958, *MNRAS*, 118, 276
 Kern N. S., Parsons A. R., Dillon J. S., Lanman A. E., Fagnoni N., de Lera Acedo E., 2019, *ApJ*, 884, 105
 Kern N. S. et al., 2020, *ApJ*, 888, 70
 Kolopanis M. et al., 2019, *ApJ*, 883, 133
 Koopmans L. et al., 2015, *The Cosmic Dawn and Epoch of Reionisation with SKA*, SISSA, Trieste, PoS(AASKA14)001, preprint ([arXiv:1505.07568](https://arxiv.org/abs/1505.07568))
 Lenc E. et al., 2017, *Publ. Astron. Soc. Austr.*, 34, e040
 Li W. et al., 2019, *ApJ*, 887, 141
 Liu A., Parsons A. R., Trott C. M., 2014, *Phys. Rev. D*, 90, 023018
 Madau P., Meiksin A., Rees M. J., 1997, *ApJ*, 475, 429
 Mallat S., 2009, *A Wavelet Tour of Signal Processing* (Third Edition), 3rd edn. Academic Press, Boston
 McKinley B. et al., 2015, *MNRAS*, 446, 3478
 Mertens F. G. et al., 2020, *MNRAS*, 493, 1662
 Mesinger A., Greig B., Sobacchi E., 2016, *MNRAS*, 459, 2342
 Morales M. F., Wyithe J. S. B., 2010, *ARA&A*, 48, 127
 Paciga G. et al., 2013, *MNRAS*, 433, 639
 Parsons A. R. et al., 2010, *AJ*, 139, 1468
 Parsons A. R., Pober J. C., Aguirre J. E., Carilli C. L., Jacobs D. C., Moore D. F., 2012, *ApJ*, 756, 165
 Patil A. H. et al., 2017, *ApJ*, 838, 65
 Planck Collaboration VI 2020, *A&A*, 641, A6
 Pritchard J. R., Loeb A., 2012, *Rep. Prog. Phys.*, 75, 086901
 Rousseeuw P. J., Croux C., 1993, *J. Am. Statist. Assoc.*, 88, 1273
 Samuel J., Nityananda R., Thyagarajan N., 2022, *Phys. Rev. Lett.*, 128, 091101
 Shapiro S. S., Wilk M. B., 1965, *Biometrika*, 52, 591
 Storer D. et al., 2022, *Radio Sci.*, 57, e2021RS007376
 Tan J. et al., 2021, *ApJS*, 255, 26
 Thompson A. R., Moran J. M., Swenson George W. J., 2017, *Interferometry and Synthesis in Radio Astronomy*, 3rd edn. Springer, Cham
 Thyagarajan N., Carilli C. L., 2020, *Phys. Rev. D*, 102, 022001
 Thyagarajan N., Carilli C. L., Nikolic B., 2018, *Phys. Rev. Lett.*, 120, 251301
 Thyagarajan N. et al., 2020, *Phys. Rev. D*, 102, 022002
 Thyagarajan N., Nityananda R., Samuel J., 2022, *Phys. Rev. D*, 105, 043019
 Tingay S. J. et al., 2013, *Publ. Astron. Soc. Austr.*, 30, e007
 Trott C. M. et al., 2020, *MNRAS*, 493, 4711
 van Cittert P. H., 1934, *Physica*, 1, 201
 van Haarlem M. P. et al., 2013, *A&A*, 556, A2
 van der Walt S., Colbert S. C., Varoquaux G., 2011, *Comput. Sci. Eng.*, 13, 22
 Virtanen P. et al., 2020, *Nat. Methods*, 17, 261
 Zernike F., 1938, *Physica*, 5, 785
 Zheng H. et al., 2017, *MNRAS*, 464, 3486

APPENDIX A: RFI AVOIDANCE IN AVERAGED CLOSURE PHASE DATA

We briefly describe how we selected our frequency band for this analysis. Unlike the visibility processing, our RFI treatment relies on robust averaging rather than flagging and inpainting. However, this strategy is only effective if the RFI is transient and does not repeat on a nightly basis (cf. Section 6.3). Hence, before proceeding with

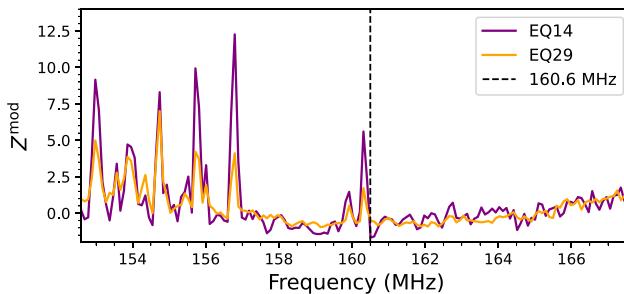


Figure A1. Modified Z-score spectrum of time averaged, high-pass filtered closure phase data. EQ14 is shown in purple and EQ29 in orange. The dashed line shows the lower bound of our selected band.

the computation of the cross-power spectra, we search for residual RFI in the averaged closure phase data.

First, we average the median-averaged data across triads and polarization products using inverse variance weights (computed as in Section 6.5). The resulting data are two dimensional, consisting of a time and a frequency axis. The aim is to identify features in this spectrogram that, if included, would leak into the high-delay region of the power spectrum. This can be achieved using a wavelet high-pass filter. A complete treatment of wavelets would go beyond the scope of this paper, but we give a brief delineation of the concept.

Wavelets are functions that are localized in frequency and delay (or traditionally time and frequency) and are generated by scaling and translating a common function called the mother wavelet. A wavelet transform decomposes a signal into a set of wavelets of different scales and shifts. The concept of the wavelet high-pass filter is to approximate the high-delay components of the original signal using only the finest scale ‘detail’ coefficients. The filter will then be particularly sensitive to features that have a similar shape to the wavelet. Here, we perform a stationary wavelet transform using a Symlet with two vanishing moments (Daubechies 1988) and use the resulting first-level detail coefficients to construct the high-pass filtered data.

The advantage of using wavelets is that they offer sparse representations of a signal, meaning that a signal can be approximated by a small number of wavelets. The smooth components of a signal are represented by large scale wavelets while singularities are represented by highly localized fine scale wavelets. Ringing effects due to discontinuities are highly localized in wavelet approximations. In contrast, Fourier approximations suffer from Gibbs oscillations, which can affect large parts of the approximated signal (Mallat 2009).

To get a picture of the RFI situation across the spectrum of Band 2, we average the absolute values of the high-pass filtered data over time. The spectrum thus obtained will mainly consist of noise fluctuations around zero and peaks corresponding to spectral discontinuities (e.g. RFI). To identify peaks exceeding the noise level, we compute a modified Z-score (cf. HERA Collaboration 2022c), defined as

$$Z_i^{\text{mod}} = \frac{|x_i| - \text{med}(|x|)}{\sigma_{\text{MAD}}} \quad (\text{A1})$$

$$\sigma_{\text{MAD}} = 1.4826 \times \text{med} \{ |x| - \text{med}(|x|) \}, \quad (\text{A2})$$

where σ_{MAD} is the Median Absolute Deviation (MAD) calculated across frequency, x_i is an element of the averaged high-pass filtered data x , and the factor of 1.4826 ensures the equivalence between the standard deviation and the MAD for normally distributed data (Rousseeuw & Croux 1993). Fig. A1 shows the modified Z-score

of the high-pass filtered and time-averaged spectrum. From this it is clear that the lower part of Band 2 is significantly contaminated by residual RFI. We choose to set the lower frequency of our band at 160.59 MHz above which there is little to no significant RFI contamination.

APPENDIX B: COHERENT TIME AVERAGING AND SIGNAL LOSS

As explained in Section 6.3, the averaging of closure phases in time will lead to a small loss in sensitivity to the cosmological signal. To determine the scale of this loss, we use the EoR model described in Section 5. At a given LST, we then place GLEAM point sources on the grid defined by the EoR model and weight by the beam response of a HERA dish. Note that it is important to include foregrounds in these simulations since the closure phase is a higher order quantity. The closure phase will hence consist of cross-terms between the foregrounds and the EoR signal (see Section 2) and the fluctuations due to the latter depend on the foreground structure through equation (2). Different foreground structures could therefore cause different degrees of sensitivity-loss under coherent time-averaging.

To emulate the change of the apparent sky, we shift the beam with respect to the sky in the interval of a pixel corresponding to ~ 7 arcmin or an observing interval of about 28.6 s. We compute closure phases and their delay spectra for these sky models. Using these, we can define the fractional loss as:

$$1 - \eta = \frac{\langle |\Psi(t, \tau)|^2 \rangle_t - |\langle \Psi(t, \tau) \rangle_t|^2}{\langle |\Psi(t, \tau)|^2 \rangle_t}, \quad (\text{B1})$$

where η is the factor by which the EoR signal is reduced in the power spectrum and the angular brackets denote an average over a time interval Δt . We compute fractional losses for 1000 pointings within Field C and average them together. The total fractional loss thus obtained is plotted in Fig. B1 as a function of Δt . It can be seen that the overall loss at $\Delta t = 171.2$ s as used in the data analysis is below 2 per cent. In a similar analysis Aguirre et al. (2022) find that averaging visibilities over an interval of 240 s produces a loss of ~ 1 per cent. The observing interval in our model is greater than the true interval of 10.7 s, which results in a slight underestimation of the loss. However, the loss at short intervals is considerably lower than the overall loss and should therefore be considered a small effect.

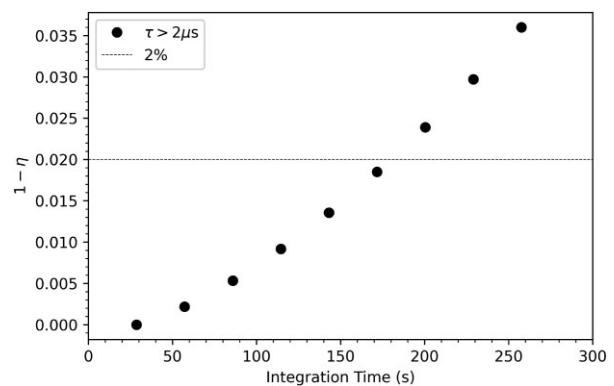


Figure B1. The average fractional loss at $\tau > 2 \mu\text{s}$ as a function of integration time. Here, we average in intervals of 171.2 s, which corresponds to a loss smaller than 2 per cent.

¹Cavendish Astrophysics, University of Cambridge, Cambridge CB3 0HE, UK

²Commonwealth Scientific and Industrial Research Organisation (CSIRO), Space & Astronomy, P. O. Box 1130, Bentley, WA 6102, Australia

³National Radio Astronomy Observatory, Socorro, NM 87801, USA

⁴INAF-Istituto di Radioastronomia, via Gobetti 101, I-40129 Bologna, Italy

⁵Department of Physics and Electronics, Rhodes University, PO Box 94, Grahamstown 6140, South Africa

⁶South African Radio Astronomy Observatory, Black River Park, 2 Fir Street, Observatory, Cape Town 7925, South Africa

⁷Department of Physics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

⁸Department of Astronomy, University of California, Berkeley, CA 94720-7300, USA

⁹Department of Physics, University of Washington, Seattle, WA 98195-1560, USA

¹⁰eScience Institute, University of Washington, Seattle, WA 98195-1560, USA

¹¹School of Earth and Space Exploration, Arizona State University, Tempe, AZ 85287-1404, USA

¹²Department of Physics and Astronomy, University of Pennsylvania, Philadelphia, PA 19104-6396, USA

¹³Department of Physics, Winona State University, Winona, MN 55987, USA

¹⁴National Radio Astronomy Observatory, Charlottesville, VA 22903-4608, USA

¹⁵Jodrell Bank Centre for Astrophysics, University of Manchester, Manchester M13 9PL, UK

¹⁶Department of Physics and Astronomy, University of Western Cape, Cape Town 7535, South Africa

¹⁷Department of Physics, Brown University, Providence, RI 02912, USA

¹⁸Radio Astronomy Lab, University of California, Berkeley, CA 94720-3411, USA

¹⁹Department of Physics, University of California, Berkeley, CA 94720-7300, USA

²⁰Department of Physics and Astronomy, University of California, Los Angeles, CA 90095-1547, USA

²¹School of Physics, University of Melbourne, Parkville, VIC 3010, Australia

²²MIT Kavli Institute, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

²³Department of Physics and McGill Space Institute, McGill University, 3600 University Street, Montreal, QC H3A 2T8, Canada

²⁴School of Chemistry and Physics, University of KwaZulu-Natal, Scottsville, 3209, South Africa

²⁵Scuola Normale Superiore, PI, I-56126 Pisa, Italy

²⁶Center for Astrophysics, Harvard & Smithsonian, Cambridge, MA 02138, USA

²⁷American Astronomical Society, Washington, DC 20006, USA

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