

RHODES UNIVERSITY



**Systematic effects and mitigation strategies
in observations of cosmic re-ionisation with
the Hydrogen Epoch of Reionization Array**

by

Ntsikelelo Charles

Submitted in fulfillment of the
academic requirements for the degree of

Doctor

in

Physics

Rhodes University, Makanda

Supervised by Prof. O. M. Smirnov, Prof. G. Bernardi, Dr. L. Bester & Dr.
N. Kern

April 3, 2024

Now unto him that is able to do exceeding abundantly above all that we ask or think, according to the power that worketh in us.

~Ephesians 3:20

For in him we live and move and have our being.

~Acts 17:28

Declaration of Non-Plagiarism

I, **Ntsikelelo Charles** declare that

1. The research reported in this thesis, except where otherwise indicated, is my original research.
2. This thesis has not been submitted for any degree or examination at any other university.
3. This thesis does not contain other persons' data, pictures, graphs or other information, unless specifically acknowledged as being sourced from other persons.
4. This thesis does not contain other persons' writing, unless specifically acknowledged as being sourced from other researchers.

Abstract

The 21 cm transition from neutral Hydrogen promises to be the best observational probe of the Epoch of Reionisation (EoR). It has driven the construction of the new generation of low-frequency radio interferometric arrays, including the Hydrogen Epoch of Reionization Array (HERA). The main difficulty in measuring the 21 cm signal is the presence of bright foregrounds that require very accurate interferometric calibration. However, the non-smooth instrumental response of the antenna as a result of mutual coupling complicates the calibration process by introducing non-smooth calibration errors. Additionally, incomplete sky models are typically used in calibration due to the limited depth and resolution of current source catalogues. Combined with the instrumental response, the use of incomplete sky models during calibration can result in non-smooth calibration errors. These, overall, impart spectral structure on smooth foregrounds, leading to foreground power leakage into the EoR window. In this thesis we explored the use of fringe rate filters ([Parsons et al., 2016](#)) as a mean to mitigate calibration errors resulting from the effects of mutual coupling and the use of an incomplete sky model during calibration. We found that the use of a simple notch filter mitigates calibration errors reducing the foreground power leakage into the EoR window by a factor of $\sim 10^2$.

[Thyagarajan et al. \(2018\)](#) proposed the use of closure phase quantities as a means to detect the 21 cm signal, which has the advantage of being independent (to first order) from calibration

errors and, therefore, bypasses the need for accurate calibration. In this thesis, we explore the impact of primary beam patterns affected by mutual coupling on the closure phase. We found that primary beams affected by mutual coupling lead to a leakage of foreground power into the EoR window, which can be up to $\sim 10^4$ times and is mainly caused by the unsmooth spectral structure primary of primary beam sidelobes affected by mutual coupling. This power leakage was confined to $k < 0.3$ pseudo h Mpc $^{-1}$.

Lastly, we also proposed and demonstrated an analysis technique that can be used to derive a flux scale correction in post-calibrated HERA data. We found that after applying flux scale correction to calibrated HERA data, the bandpass error reduces significantly, with an improvement of 6%. The derived flux scale correction was antenna-independent, and it can be applied to fix the overall visibility spectrum scale of H4C data post-calibration in a fashion similar to [Jacobs et al. \(2013\)](#)

Acknowledgments

“Imfundo ayigugelwa! kwaye indela ubuzwa kwaba phambili!”

This note of thanks is the final touch on my dissertation dedicated to all those whose paths have crossed mine. I want to reflect on the people who supported and helped me throughout my Masters journey.

I thank my supervisors, Prof. Oleg Smirnov, Prof. Gianni Bernardi, Dr Landman Bester and Dr Nickolas Kern for their sound advice, inspiration and encouragement.

The financial assistance of the South African Radio Astronomy Observatory (SARAO) towards this research is hereby acknowledged (www.sarao.ac.za).

I thank the staff of the Rhodes University Physics and Electronics Department. Sincere gratitude goes to my colleagues, members of the RATT team and for all their support, discussions, reviews and engaging arguments, which gave this work a perfect finish.

Finally, I thank my family for all their support and encouragement.

Contents

Declaration of Non Plagiarism	i
Declaration of Publications	ii
Abstract	ii
Acknowlegments	iv
1 INTRODUCTION	1
1.1 The redshifted 21 cm emission	2
1.2 Power spectrum of the 21 cm signal from inteferometric observations	6
1.2.1 The Radio Interferometer Measurement Equation (RIME)	6
1.2.2 Sky model-based calibration	8
1.2.3 Redundant calibration	9
1.2.4 Calibration errors	10
1.2.5 Power spectrum of 21 cm signal from visibility measurements	11
1.2.6 Power spectrum of the bispectrum phase	13
1.3 Foreground Problem: an overview	14

1.4	The EoR Window	16
1.5	The Hydrogen Epoch of Reionization Array	23
1.6	Current status of 21 observations	24
1.7	Thesis motivation	25
2	Simulations of primary beam effects on the cosmic bispectrum phase observed with the Hydrogen Epoch of Reionization Array	28
2.1	Abstract	28
2.2	Introduction	29
2.3	Closure phase formalism	32
2.4	Bispectrum phase simulations	35
2.4.1	Beam models	35
2.4.2	Simple sky models	38
2.4.3	Simulations with realistic sky models	44
2.5	Results	47
2.6	Discussion and conclusions	55
3	On the use of temporal filtering for mitigating galactic synchrotron calibration bias in 21 cm reionization observations	57
3.1	Abstract	57
3.2	Introduction	58
3.3	Calibration overview	61
3.3.1	Sky-model based calibration	61
3.3.2	Redundant calibration	63
3.4	Foreground simulations	65
3.5	Fringe Rate Filters for Improving Calibration	69
3.5.1	Fourier representations	69
3.5.2	Calibration Without Filtering	71
3.5.3	Fringe Rate Filtering Before Calibration	75

3.5.3.1	Notch Filters	75
3.5.3.2	Main Lobe Filters	75
3.5.3.3	Computing the Power Spectrum	76
3.5.3.4	Summary	80
3.5.4	Effects of fringe rate filters on calibration	80
3.5.5	Calibration with the notch filter	81
3.5.6	Calibration with the main lobe filter	90
3.5.7	What does this mean for detecting the EoR?	91
3.6	Conclusions	92
4	Assessment of the absolute calibration in the H4C HERA data release	94
4.1	Imaging of wideband HERA observations	95
5	Mitigating the impact of mutual coupling in the calibration of HERA observations with the use of fringe rate filters	107
5.1	Simulation framework	108
5.2	Application of fringe rate filters	116
5.2.0.1	Notch Filters	119
5.2.0.2	Main Lobe Filters	119
5.2.1	Use of a notch filter to improve calibration	120
5.3	Foreground Power Spectrum	124
5.4	Conclusions	125
6	Conclusions	128

CHAPTER 1

INTRODUCTION

With the ever-improving sensitivity of currently built and upcoming radio interferometers such as the Hydrogen Epoch of Reionisation Array (HERA; [DeBoer et al., 2017](#)), the LowFrequency Array (LOFAR; [van Haarlem et al., 2013a](#)), the Murchison Widefield Array (MWA; [Tingay et al., 2013](#)) and the future Square Kilometre Array (SKA; [Koopmans et al., 2015](#)), the detection of the redshifted 21 cm emission line from neutral Hydrogen during the Epoch of Reionisation is nearing the horizon, with power spectrum limits that are becoming sensitive enough to discriminate between models (e.g., [HERA Collaboration et al., 2022](#)). The Epoch of Reionisation (EoR) is the period marked by the reionisation of the Inter-Galactic Medium (IGM) by radiation coming from the formation of the first stars and galaxies at $z \approx 30$ to the complete ionization of the IGM at $z \approx 6$ (e.g., [Mesinger et al., 2011](#)). Currently, several cosmological models are still predicting different reionisation scenarios, and thus, more sensitive power spectrum limits are required to discriminate between re-ionization models. With improved sensitivity, we will be able to understand how the first galaxies formed, and ultimately, this will improve constraints on cosmological parameters (e.g., [Park et al., 2019](#)).

The 21 cm emission line from the EoR is very weak, about 4-5 orders of magnitude dimmer

compared to local radio emission from the Galactic plane and bright extra-galactic sources (e.g., [Bernardi et al., 2009](#); [Thyagarajan et al., 2016](#)). The sky signal received by the ground radio interferometer will contain both the 21 cm signal and foreground emission; it is, therefore, crucial to separate the 21 cm signal from the foreground emission. This is the inherent challenge faced by all 21 cm experiments. Added to this is the non-ideal instrumental response of radio interferometers (e.g., [Barry et al., 2016](#); [Fagnoni et al., 2021b](#)). In this work, we aim to investigate the impact as well as mitigation strategies that can be employed to deal with these challenges.

This introductory chapter provides the necessary background material and a relevant literature review of the work on which this thesis is based upon.

1.1 The redshifted 21 cm emission

The early Universe comprised sub-atomic particles such as protons, electrons and neutrons with **about 25%** of helium nuclei due to the high-temperature conditions ([Fuller et al., 1988](#)). Electrons in the early Universe constantly scattered photons via *Thompson scattering*, preventing the formation of neutral atoms until the Universe was about 380000 years old. This marked the beginning of *recombination* at $z_{rec} \approx 1100$ ([Kamionkowski, 2007](#); [Bennett, 2007](#)). During recombination, electrons bind to protons, forming neutral hydrogen atoms ([Peebles, 1968](#); [Zeldovich, 1972](#)). As a result, matter and radiation decoupled, and thus, photons could finally propagate, making the Universe transparent. This radiation is observed today in all directions as the *Cosmic Microwave Background (CMB)* with an average temperature today of $T_\gamma \approx 2.73$ K.

The baryonic matter, under the force of gravity, combined with the overall cooling effect of matter, began to collapse within already-formed dark matter halos. Eventually, the first generation of stars began to form as the density of matter within dark matter halos increased. This period is often referred to as the *Cosmic Dawn*, expected to take place at $z \approx 20 - 30$. Radiation from these first stars and galaxies initially heated the IGM. However, as the star formation process continued, the radiation began to ionise the IGM. This marked the beginning of the third phase, *Epoch of Reionisation*.

The Hydrogen hyperfine spectral line promises to be a powerful probe for studying the evo-

lution of the IGM during the EoR (Furlanetto et al., 2006; Mesinger et al., 2011). The Hydrogen hyperfine transition occurs when the electron and proton spins are aligned, and this configuration state is denoted as $1_1S_{1/2}$. This configuration has a higher energy than the state when the spins are anti-parallel $1_0S_{1/2}$ (Storey et al., 1994). The emission of a (rest frame) 21 cm photon occurs when the Hydrogen atom transitions from the $1_1S_{1/2}$ state to $1_0S_{1/2}$, and absorption when hydrogen atoms moves from $1_0S_{1/2}$ to $1_1S_{1/2}$.

The spin temperature T_s quantifies the relative number densities of atoms n_1 in the hyperfine level with respect to the ground state n_0 :

$$\frac{n_1}{n_0} = 3 e^{-E_{10}/k_B T_s}, \quad (1.1)$$

where E_{10} is the energy split between the two energy levels. In the early Universe, the spin temperature is set by three competing mechanisms: scattering of CMB photons, collisions between neutral hydrogen atoms and scattering of Ly α photons. In equilibrium, the spin temperature is given by:

$$T_s^{-1} = \frac{T_\gamma^{-1} + x_a T_c^{-1} + x_c T_k^{-1}}{1 + x_a + x_c} \quad (1.2)$$

where x_c is the collisional coupling constant between hydrogen and hydrogen with free electrons, T_c is the colour temperature and x_a is the coupling constant for the absorption and re-emission of Ly α photons, the Wouthuysen-Field (WF, Wouthuysen, 1952) effect. WF effect occurs when a Ly α photon is absorbed by a neutral Hydrogen atom in the ground state $1S$, the electron can jump from the $1S$ state to either of the $2P$ quantum states, following the quantum selection rules. The electron can then decay from the $2P$ state to the fundamental triplet $1_1S_{1/2}$ rather than the singlet $1_0S_{1/2}$, allowing for a 21 cm photon to be emitted.

The 21 cm emission from an IGM cloud is observable as the contrast against the CMB (Furlanetto et al., 2006):

$$\delta T_b \approx 9 x_{H1} (1 + \delta) (1 + z)^{\frac{1}{2}} \left[1 - \frac{T_\gamma(z)}{T_s} \right] \left[\frac{H(z)/(1+z)}{dv_{||}/dr_{||}} \right] \text{mK}, \quad (1.3)$$

where δ is the baryon over-density, x_{H1} the Hydrogen neutral fraction, $H(z)$ is the Hubble parameter at redshift z , $dv_{||}/dr_{||}$ is the gradient of the proper velocity along the line of sight between

the observer and the patch of IGM. The 21 cm emission line can either appear in absorption if $T_s < T_\gamma$ or in emission if $T_s > T_\gamma$.




Figure 1.1: Fiducial evolution of the spin T_s (red solid line), gas T_k (green dashed line) and CMB temperature T_γ (blue dotted line), respectively (from [Mesinger et al., 2011](#)).

Figure 1.1 and shows the predicted evolution of spin temperature, the gas temperature and CMB temperature as a function of redshift. At high redshifts $z > 200$, CMB photons scatter the residual, a small fraction of free electrons in the IGM, setting the spin temperature to the CMB temperature T_γ , and so the brightness temperature contrast $\delta T_b = 0$. Eventually, the heating via Compton scattering decreases as the density of free electrons reduces basically to zero, and the

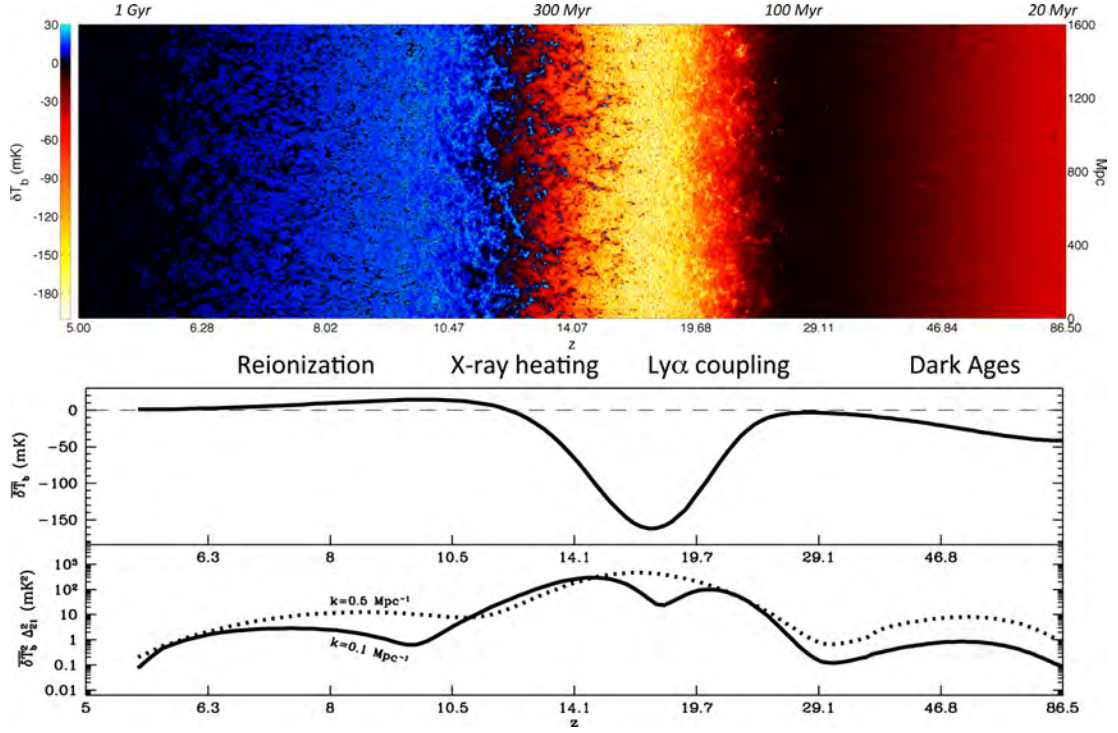


Figure 1.2: Simulated evolution of the 21 cm signal. The top panel shows the evolution of the brightness temperature contrast δT_b with redshift z . The middle panel shows the corresponding sky-averaged signal (i.e. global signal, solid line). The dotted line denotes global signal $\overline{\delta T_b} = 0$. The bottom panel shows the evolution of the power spectrum amplitude at $k = 0.1 \text{ h Mpc}^{-1}$ (solid line) and $k = 0.5 \text{ h Mpc}^{-1}$ (dotted, line; from [Mesinger et al., 2016](#)).

IGM thermally decouples from the CMB at $z \approx 200$.

When the gas thermally decouples from the radiation, the gas begins to cool down faster, and the gas temperature T_k drops below the CMB temperature T_γ . In this phase, the IGM is dense enough that collisions couple the spin temperature to the gas temperature. As the Universe continues to expand, the IGM becomes less dense, and collisions between neutral hydrogen atoms are no longer efficient at coupling the spin temperature to the gas temperature. Thus, the absorption of CMB photons by neutral Hydrogen drives the spin to the CMB temperature, and $\delta T_b \approx 0$.

When the first stars form at $z \approx 30$, they emit background radiation of UV photons that drives the spin temperature to the gas temperature via the WF effect. As star formation proceeds, X-

ray sources form and begin to heat the surrounding gas, eventually above the CMB temperature. This transition is referred to as the *spin temperature saturation*, as the evolution of the 21 cm brightness contrast essentially no longer depends upon the spin temperature T_s (see equation 1.3).

Eventually, as star formation proceeds, ionizing photons from these sources begin to ionise the local IGM, marking the beginning of reionization. Initially, only bubbles of ionised hydrogen form around star-forming regions at $z \approx 10$ and eventually, the entire IGM is completely ionised at $z \approx 6$.

1.2 Power spectrum of the 21 cm signal from inteferometric observations

The simplest radio interferometer is the two-element interferometer, where signals measured from a pair of antennas (i, j) are cross-multiplied and averaged in time. This operation is known as *correlation* and leads to the fundamental quantity measured in radio interferometry, the *visibility function* V_{ij} . The *Van Cittert-Zernike theorem* states that the correlation of signals from the (i, j) pair is related to the sky brightness $I(l, m, \nu)$ by a Fourier transform:

$$V_{ij}(u, v, \nu) = \iint I(l, m, \nu) e^{2\pi i(ul+vm)} dl dm, \quad (1.4)$$

where (u, v) are the components (normally expressed in wavelength units) of the baseline vector \mathbf{b} that connects antenna i to antenna j . The coordinates (l, m) are cosine coordinates, typically centred on the observed source. The units of I are the intensity per solid angle, unit time and frequency, measured in $[\text{Jy Hz}^{-1} \text{s}^{-1} \text{sr}^{-1}]$.

1.2.1 The Radio Interferometer Measurement Equation (RIME)

In ideal conditions and with ideal antennas, the sky brightness is related to the measured visibilities by a Fourier transform (equation 1.4). However, real antennas are not perfect, and the electromagnetic radiation from a radio source passes through the atmosphere/ionosphere before

reaching the antenna. These effects corrupt the signal. The *Radio Interferometry Measurement Equation (RIME)* provides a mathematical model to describe the various propagation effects that corrupt the signal.

Let us consider signals received by two dipole antennas, i and j , for a single polarization. There are two kinds of propagation effects that can corrupt the signal. These are *direction dependent effects*, where the corruptions to the sky signal depend upon the observing direction. An example of a direction-dependent effect, the only direction-dependent effect explored in this study, is the antenna primary beam E . An antenna is not an isotropic receiver, its sensitivity varies across the sky, and the primary beam describes the sensitivity pattern of the antenna across the sky. When the effect of the primary beam is included, the RIME is expressed as (Hamaker et al., 1996; Smirnov, 2011):

$$V_{ij}^m(u, v, \nu) = \iint E_i(l, m, \nu), I(l, m, \nu) E_j^*(l, m, \nu) e^{2\pi i(ul+vm)} dl dm, \quad (1.5)$$

where V^m are the measured visibilities and $*$ denotes complex conjugate, E_i and E_j are antenna primary beams of antenna i and j , respectively. Antenna primary beams are typically considered identical as their differences are small. As a result, an average antenna primary beam $|A|^2$ is used. However, effects such as mutual coupling can lead to antenna-to-antenna primary beam variations (e.g., Neben et al., 2016; Fagnoni et al., 2020).

The other kind of propagation effects are known as *direction-independent effects*, and these are independent of the observing direction. Examples of such are the band-pass response B of the antenna, describing the variation of the receiver response as a function of frequency, and the antenna gain factor G , describing the amplification of the signal by the Low Noise Amplifier. All the direction-independent effects in antenna i can be grouped into a single antenna gain term g_i . When direction-independent effects are included, the RIME takes the form:

$$V_{ij}^d = g_i V_{ij}^m g_j^* + n_{ij}, \quad (1.6)$$

where V_{ij}^d and V_{ij}^m are the corrupted and model visibilities, respectively. For completeness of the visibility product of the RIME from the baseline, we add to the RIME the thermal noise n_{ij} , which is generated by the antennas i and j during an observation.

1.2.2 Sky model-based calibration

Calibration is the process of modelling and correcting the radio signal from corruption introduced by the instrument. Equation 1.6 shows the simplest RIME describing the corruption to a signal received by antenna i and j for single polarization. Gains are solved for by minimizing the chi-square χ^2 :

$$\chi^2 = \sum_{i,j} \frac{|V_{ij}^d - g_i g_j^* V_{ij}^m|^2}{\sigma_{ij}^2}, \quad (1.7)$$

where V_{ij}^m are the constructed model visibilities, and σ_{ij}^2 is the variance of the visibility, and the sum runs over all unique antenna pairs.

Calibration is often divided into three levels: first, second and third-generation calibration. The first and second-generation calibration solves for direction-independent gains and involves the use of a sky model (Smirnov, 2011).

The first-generation calibration generally uses a bright point source with a known spectrum, referred to as a flux calibrator. The bandpass response, which is a frequency-dependent response array, is also solved for during this phase of calibration. Observations of calibration sources are normally interspersed with observations of the target field. The gains from the calibrator source are interpolated and applied to the target field to obtain corrected visibilities V^c :

$$V_{ij}^c = g_i^{-1} V_{ij}^d (g_j^*)^{-1}, \quad (1.8)$$

After performing first-generation calibration, an image of the target field is made. The image quality is assessed or measured by the dynamic range DR :

$$DR = \frac{I_p}{\sigma_I}, \quad (1.9)$$

where I_p is the image peak and σ_I is the image rms noise. A better dynamic range is often a consequence of an improved calibration.

In *second generation* calibration or self-calibration (Smirnov, 2011; Zhao et al., 2019), the image quality, thus calibration, is improved by effectively deconvolving the image from the point spread function and solving for the time-dependent gain terms. This is achieved through an iterative process that begins with an initial sky model, which is used to solve for time-variable

antenna gain terms, and after which, the visibilities are re-calibrated with updated gains and then imaged. The generated image is then deconvolved to further improve the sky model. The process is repeated until a target dynamic range is reached or if further image improvement is not possible.

The last step in the calibration process is to perform *third generation calibration*, where the self-calibrated image is corrected for direction-dependent effects. We will briefly discuss the process of correcting for the various direction-dependent effects. However, for an in-depth discussion of this topic, we recommend the reader consult [Smirnov \(2011\)](#). Direction-dependent calibration is mostly done for sources that are further from the field pointing centre, and these sources can be identified in the self-calibrated image, as they still show artefacts (not present for sources near the pointing centre) that self-calibration could not correct. To correct for these artefacts, a least squares approach similar to the calibration of direction-independent effects is used. However, differential direction-dependent gain terms ΔE are added to the equation, and the sky is effectively broken down into multiple segments/facets s , where each source has its own unique ΔE . The RIME the equation to be solved becomes:

$$\mathbf{V}_{pq}^d = g_i \left[\sum_{s=1}^{N_s} \Delta E_{s,p} X_{s,pq} \Delta E_{s,q}^* \right] g_j \quad (1.10)$$

where X is the the source coherency in the direction s and is given by

$$X_{ij}^s = \iint_{\Omega_s} I(l, m, \nu) e^{2\pi i(ul+vm)} d\Omega \quad (1.11)$$

where Ω_s is the solid angle subtended by source s . The differential gains are then used to correct V_{ij}^d .

1.2.3 Redundant calibration

Interferometric calibration takes a special flavour if an array is deployed in a highly redundant layout, where multiple baselines in the array have the same orientation and length and, therefore, measure the same Fourier mode of the sky. In this case, the knowledge of a sky model is less

necessary, as it becomes a model parameter that can (at least partially) be solved for. For example, assume we have a single baseline type A , uniquely identified via its baseline vector \mathbf{b}_A . For all antenna pairs i, j that share this baseline vector, the calibration equation becomes:

$$V_{ij}^d = g_i V_A g_j^* + n_{ij} \quad (1.12)$$

where V_A is now a parameter of the model, called the “redundant model visibility”. This is repeated for all unique baseline types, eventually, leading to an over-constrained system of equations, which can be solved via a χ^2 minimization (e.g., [Liu et al., 2010](#)).

However, the gain solutions obtained from the system of linear equations are not unique as four degenerate parameters need to be solved for after redundant calibration: the overall gain amplitude, the model visibility amplitude and the phase gradient across the array in the east-west and north-south orientations ([Liu et al., 2010](#); [Zheng et al., 2014](#); [Dillon et al., 2018](#)). These parameters are typically a function of polarization, frequency and time and require a sky model.

1.2.4 Calibration errors

The modelled visibilities are usually simulated using existing source catalogues. Source catalogues, however, do not fully characterise the sky emission due to their limited depth and resolution. The incompleteness of the sky model necessarily introduces errors in interferometric calibration, and these errors impart unsmooth spectral structure in the recovered gains (e.g., [Wijnholds et al., 2016](#); [Barry et al., 2016](#); [Ewall-Wice et al., 2016](#); [Dillon et al., 2020](#); [Kern et al., 2020c](#)).

When incorrect gains are used to calibrate raw data, the unsmooth spectral structure is imparted to the calibrated visibilities, and thus smooth foreground in the calibrated visibilities is pushed to higher delays (Fourier conjugate pair of frequency, discussed in section [1.2.5](#)), thus contaminating the EoR window ([Barry et al., 2016](#); [Dillon et al., 2020](#); [Kern et al., 2020c](#)).

In addition to sky model incompleteness, calibration errors can arise from the inaccurate or incomplete characterisation of the antenna responses. In this thesis, we are particularly interested in the primary beam response of the antenna. The primary beam is spectrally unsmooth on the

sidelobes and is typically not well characterised, and thus, emission falling on the sidelobes is difficult to model and calibrate. Secondly, effects such as mutual coupling further complicate calibration by increasing the spectral variation of primary beam sidelobes, introducing small ripples on sidelobes (Neben et al., 2016; Fagnoni et al., 2020). These effects overall impart unsmooth spectral structure to foreground emission and thus lead to unsmooth calibration errors.

1.2.5 Power spectrum of 21 cm signal from visibility measurements

The spatial fluctuation of the 21 cm signal can be characterised by its power spectrum $P(k)$ (e.g., Furlanetto et al., 2006). The *fractional brightness temperature contrast* is defined as:

$$\delta_{21}(\mathbf{x}) = \frac{\delta T_b(\mathbf{x}) - \overline{\delta T_b}}{\delta T_b}, \quad (1.13)$$

where

$$\overline{\delta T_b} = \frac{\int \delta T_b(\mathbf{x}) d\Omega}{\int d\Omega}, \quad (1.14)$$

is the spatial averaged brightness temperature contrast, i.e. *global signal*. The fractional brightness contrast captures the spatial differences in brightness temperature above or below the average.

The *two point correlation function* of the brightness temperature contrast $\xi(\mathbf{x})$, at a point \mathbf{x} , is defined as an ensemble average product of the brightness temperature fluctuations at two different locations separated by a distance \mathbf{x} :

$$\xi(\mathbf{x}) = \langle \delta_{21}(\mathbf{r}) \delta_{21}(\mathbf{r} + \mathbf{x}) \rangle. \quad (1.15)$$

The two-point auto-correlation function describes the excess probability of finding two over (under) bright regions separated by distance \mathbf{x} relative to the average temperature, i.e., degree of clustering/sparsity of bright/cold regions. The fractional brightness temperature contrast can be described in frequency space by taking a Fourier transform:

$$\delta(\mathbf{x}) = \frac{1}{(2\pi)^3} \int \tilde{\delta}(\mathbf{k}) e^{i\mathbf{k}\mathbf{x}} d^3k, \quad (1.16)$$

where \mathbf{k} is the spatial frequency, the Fourier conjugate of \mathbf{x} and $\tilde{\delta}(\mathbf{k})$ frequency spectrum of the fractional brightness temperature contrast.

The power spectrum of the 21 cm brightness contrast is defined as:

$$P(k)\delta_D(\mathbf{k} + \mathbf{k}')(2\pi)^3 = \langle \tilde{\delta}_{21}(\mathbf{k})\tilde{\delta}_{21}(\mathbf{k}') \rangle, \quad (1.17)$$

where the power spectrum has units of $\text{mK}^2 (\text{h}^{-1} \text{Mpc})^3$. And it can be shown that the power spectrum is related to the correlation function through a Fourier transform:

$$\xi(x) = \frac{1}{(2\pi)^3} \int P(k) e^{i\mathbf{k}\cdot\mathbf{x}} d^3k. \quad (1.18)$$

Thus, the power spectrum describes how much brightness temperature power occurs at the various scales.

Lastly, it can also be shown that the 3-dimensional power spectrum of the 21 cm can be obtained from interferometric measurements. The three-dimensional power spectrum from visibility measurements is expressed as (Parsons et al., 2012; Thyagarajan et al., 2013):

$$P(k) = |\tilde{V}(\tau)|^2 \left(\frac{\lambda^2}{2k_B} \right)^2 \left(\frac{D_c^2 \Delta D_c}{B_{\text{eff}}} \right) \left(\frac{1}{\Omega B_{\text{eff}}} \right), \quad (1.19)$$

where $\tilde{V}(\tau)$ is the Fourier transform of the visibility along its frequency axis, also known as the *delay transform* (Parsons et al., 2012):

$$\tilde{V}(\tau) = \int V(\nu) W(\nu) e^{-2\pi i\nu\tau} d\nu, \quad (1.20)$$

where τ is the delay, the Fourier conjugate variable of ν , $W(\nu)$ is the spectral weighting function, and λ is the centre wavelength of observing bandwidth, B_{eff} is the effective bandwidth, and ΔD_c and D_c are cosmological distances along and perpendicular to the line of sight, respectively, Ω the field of view solid angle and k_B is Boltzmann's constant. The power spectrum is a function of $k = |\mathbf{k}|$ but can also be computed as a function of modes along and perpendicular to the line of sight $\mathbf{k} = (k_{\parallel}, \mathbf{k}_{\perp})$ (Thyagarajan et al., 2013):

$$k_{\parallel} = \frac{2\pi\nu_{12}H_0E(z)}{c(1+z)^2} \tau, \quad (1.21)$$

$$\mathbf{k}_{\perp} = \frac{2\pi\nu}{cD_c} \mathbf{b}. \quad (1.22)$$

Typically the effective bandwidth is chosen to be $B_{\text{eff}} \approx 10 \text{ MHz}$, in the redshift, range $6 < z < 10$. The cosmic evolution limits the largest bandwidth to about 10 MHz since, for larger

bandwidths, the evolution of the Universe begins to impact the result as there are significant changes in IGM morphology over the studied period (Thyagarajan et al., 2013). It should also be noted that the orthogonality of u, v and τ is only valid for short baselines and narrow frequency ranges (Liu et al., 2014a). Ultimately, the power spectrum from visibility measurements can be used to infer the spatial distribution of the 21 cm signal, and thus providing a direct probe of the evolution of the IGM during the EoR.

1.2.6 Power spectrum of the bispectrum phase

Another way of probing spatial fluctuations of the 21 cm signal from visibility measurements was proposed by Thyagarajan et al. (2018). The idea behind the proposal is to detect the 21 cm signal using *closure phases*. The use of closure quantities is beneficial as they are insensitive to direction-independent effects.

The bispectrum C_{ijk} is defined as the triple product of three visibilities from baselines (ij, jk, ki) :

$$C_{ijk} = V_{ij}V_{jk}V_{ki}, \quad (1.23)$$

where indices i, j, k are antenna labels. The closure phase is defined as the complex phase of the bispectrum. Importantly, we assume that gains can be factored into antenna-based gains and that they are not baseline-based. The corrupted visibilities V^d can be written as:

$$V_{ij}^d = g_i V_{ij}^m g_j^* = |g|_i e^{i\phi_i} V_{ij}^m |g|_j e^{-i\phi_j}, \quad (1.24)$$

and thus, the corrupted bispectrum then becomes:

$$\begin{aligned} C_{ijk}^d &= |g|_i e^{i\phi_i} |g|_j e^{i\phi_j} |g|_k e^{i\phi_k} V_{ij}^m V_{jk}^m V_{ki}^m |g|_j e^{-i\phi_j} |g|_j e^{-i\phi_j} |g|_k e^{-i\phi_k} \\ &= |g|_i^2 |g|_j^2 |g|_k^2 V_{ij}^m V_{jk}^m V_{ki}^m, \\ &= |g|_i^2 |g|_j^2 |g|_k^2 C_{ijk}^m, \end{aligned} \quad (1.25)$$

where C_{ijk}^m is the uncorrupted bispectrum. As the antenna-gain phases cancel out in the second step, the complex phase of the corrupted bispectrum C_{ijk}^d is the same as the phase from the uncorrupted bispectrum $C_{ijk}^m = V_{ij}^m V_{jk}^m V_{ki}^m$. This shows that the closure phase is unaffected by antenna-based gains.

The power spectrum of the bispectrum phase P_∇ is defined as (Thyagarajan & Carilli, 2020):

$$P_\nabla(k_{||}) = |\tilde{\Psi}_\nabla|^2 \left(\frac{\lambda^2}{2k_B} \right)^2 \left(\frac{D_c^2 \Delta D_c}{B_{\text{eff}}} \right) \left(\frac{1}{\Omega B_{\text{eff}}} \right), \quad (1.26)$$

where ∇ denotes the baseline triad, $\Delta D = \Delta D(z)$ is the comoving depth along the line of sight corresponding to an effective bandwidth B_{eff} , $\Omega = A_e/\lambda^2$ where A_e is the effective aperture area and:

$$\tilde{\Psi}_\nabla = \tilde{W}(\tau) * \tilde{\Xi}_\nabla(\tau) * V_{\text{eff}} * \delta(\tau), \quad (1.27)$$

where $*$ denotes convolution, \tilde{W} and $\tilde{\Xi}_\nabla$ are the delay transforms of the window function W and the complex closure phase Ξ_∇ respectively:

$$\Xi_\nabla(\nu) = e^{i\phi_\nabla(\nu)}. \quad (1.28)$$

The normalization factor V_{eff} is defined as:

$$(V_{\text{eff}})^{-2} = \sum_{b=1}^3 |V'_b|^{-2}, \quad (1.29)$$

where

$$V'_b = \frac{\int W(\nu) V_b(\nu) d\nu}{\int W(\nu) d\nu}, \quad (1.30)$$

and b denotes baselines in a triad. The closure spectrum is a dimensionless quantity, and the normalisation factor ensures that the power spectra from fields with different brightness can be compared to each other. The power spectrum of the bispectrum is measured in pseudo $\text{mK}^2 (h^{-1} \text{Mpc})^3$ units, and the pseudo is used to differentiate this power spectrum from the power spectrum derived from visibility measurements, and also note that the power spectra are not comparable to each other.

1.3 Foreground Problem: an overview

The challenge faced by all EoR experiments in measuring the redshifted 21 cm signal is the presence of foregrounds that are 3-5 orders of magnitude brighter than the 21 cm signal (e.g., Jelić

et al., 2008; Bernardi et al., 2010; Thyagarajan et al., 2016). The foreground emission is composed of Galactic synchrotron emission as well as radio emission from extra-galactic sources. Synchrotron radiation occurs when relativistic electrons are accelerated in a magnetic field. Accelerated charged particles generate electromagnetic waves (see, Rybicki & Lightman, 1979).

Figure 1.3 shows the brightness temperature of foregrounds relative to the 21 cm signal. Fortunately, the foregrounds have smooth spectra. A power law spectrum can approximate the galactic synchrotron spectrum $I(\nu)$:

$$I(\nu) \propto N_e B_{\perp}^{(\Theta+1)/2} \nu^{-(\Theta-1)/2}, \quad (1.31)$$

where N_e is the electron number density, B_{\perp} is the magnetic field component perpendicular to the line of sight, Θ is the power law index of the electron energy distribution (Pacholczyk, 1970).

The extra-galactic foreground emission mostly comes from Active Galactic Nuclei (AGNs). Relativistic jets created by AGNs can emit synchrotron radiation. Most extra-galactic sources show a power law spectrum, with a median spectral index of $\alpha = 0.8$ (De Breuck et al., 2000):

$$S(\nu) \propto \nu^{-\alpha}, \quad (1.32)$$

where S is the source flux density. The flux density of a radio source is defined to be:

$$S(\nu) = \int_{\Omega} \cos \theta I(\nu)(\Omega') d\Omega', \quad (1.33)$$

where Ω is the solid angle subtended by the source, and θ is the angle between the source centre and the element of solid angle $d\Omega$.

Only a small fraction of the AGN population shows deviations from smooth spectra, mainly caused by *synchrotron self-absorption* (Kellermann & Pauliny-Toth, 1969). Synchrotron self-absorption occurs when the emitting region is opaque to the radiation emitted by the source itself, leading to an inverted spectral index $\alpha = -0.5$. Only 5% of the sources show, however, a turnover between 72 MHz and 1.4 GHz (e.g., Callingham et al., 2017).




Figure 1.3: Illustrations of the various foreground components and their relative brightness temperature in contrast to the 21 cm signal (from Jelić et al., 2008).

1.4 The EoR Window

In this chapter, I will describe the concept of a region of the power spectrum space that may not be contaminated by foregrounds. Let us consider an idealised foreground model composed of a single flat spectrum point source. Its delay transform is a Dirac delta function centred on the geometric delay of the source τ_s .

$$\tau_s = \frac{|\mathbf{b}| \sin \theta}{c}, \quad (1.34)$$

where θ is the angle between the zenith pointing position vector and the position vector pointing to the source. The delay spectrum is thus bound to a maximum delay set by the horizon limit, as the antenna cannot receive the radiation from the source if it is beyond the horizon τ_{\max} (Parsons et al., 2012):

$$\tau_{\max} = \frac{|\mathbf{b}|}{c}. \quad (1.35)$$

If the spectrum of the source is non-flat, its delay transform will have a range of values centred on the geometrical delay, still bound within the maximum delay τ_{\max} . Thus, smooth-spectrum foreground emission is expected to be bounded to a maximum delay set by a baseline-dependent horizon limit. As the baseline length is proportional to k_{\perp} , the foreground power spectrum is expected to be confined to a specific region which assumes a wedge-like shape in k space limited

to (Datta et al., 2010; Parsons et al., 2012; Trott et al., 2012; Thyagarajan et al., 2013; Morales et al., 2019; Mevius et al., 2022):

$$k_{\parallel} \leq \frac{H(z)D(z)}{c(1+z)} k_{\perp}. \quad (1.36)$$

The EoR signal, conversely, fluctuates down MHz scales (Parsons et al., 2012; Thyagarajan et al., 2013; Pober et al., 2013) and is, therefore, expected to extend beyond the horizon limit, and this region beyond the horizon limit is referred to as the *EoR window*.

Figure 1.4 and 1.5 show the visibility spectra and power spectra of simulated foreground emissions for both diffuse and compact sources and the EoR signal. Note the instrumental response of the primary beam is included. Thus, the foreground emission slightly exceeds the horizon limit of the 29 m baseline. However, the foreground visibility spectrum (Figure 1.4) is still relatively smooth, and thus the Fourier transform of the visibility spectra, i.e. delay spectrum (Figure 1.5), is confined to small k -modes. This is seen by the sharp drop in the foreground power at $k_{\parallel} \approx |0.3| h \text{ Mpc}^{-1}$. This is in contrast to the visibility spectra of the 21 cm signal, which has variations down to a few MHz, and thus has power even at high k modes. Notice that the 21 cm power spectrum at $k_{\parallel} = 0 h \text{ Mpc}^{-1}$ is about eight orders of magnitude fainter than the foregrounds. However, the foreground power drops rapidly below the 21 cm power spectrum at high delays $k_{\parallel} > |0.3| h \text{ Mpc}^{-1}$. As most arrays have greater sensitivity at small k modes due to the greater number of short baselines with respect to long baselines, preserving the low k modes of the EoR window free from foreground contamination is important. Thus, as long as the instrumental response is relatively smooth, the 21 cm power spectrum can be cleanly separated from the foreground power spectra.

But, this is never the case in actual observations, as antenna responses are typically non-smooth and complex to model, and thus, this overall leads to calibration errors which impart frequency structure on smooth foreground spectra, and thus, a portion of the foreground emission leaks into the low k modes of the EoR window.

Figure 1.6 shows gains derived from actual observations taken with the Hydrogen Epoch of Reionization Array (HERA, see the next section for a description of the HERA array). The gains show structure at various delays, which can be attributed to errors in calibration from different

sources. Typically, in HERA calibration, the diffuse emission component of the sky is usually excluded. This, in part, is because current diffuse emission models are not accurate enough and usually make calibration errors worse if included (see, [Kern et al., 2020a](#)). However, excluding the diffuse emission component still introduces calibration errors as the recovered gains contain structure at low delays, in the $200 < \tau < 400$ ns range. Thus, these errors have the potential to contaminate k modes in the EoR window. Another source of error in calibration comes from the effects of mutual coupling, where an incident sky signal on an antenna is reflected to nearby antennas. And thus, copies of the sky signal with various strengths are repeated at various delay modes. This introduces structure in recovered gains at delays range $700 < \tau < 1200$ ns.

Similarly to visibility spectra, [Thyagarajan et al. \(2018\)](#) showed that the closure spectrum of foreground emission is smooth compared to that of the 21 cm signal. They performed closure spectra simulations, including foregrounds and a fiducial 21 cm signal for a 14.6 m baseline, including an idealized HERA primary beam. They showed that closure spectra remain relatively smooth across the HERA observing band (right panel of [Figure 1.4](#)). Conversely, the EoR closure spectrum highly fluctuates along frequency. This behaviour is similar to the visibility spectra and suggests that an avoidance-like approach can be used to separate the foreground closure spectra from the 21 cm closure spectra. Indeed, through 21 cm and foreground simulations [Thyagarajan et al. \(2020\)](#) showed the power spectrum of the closure phase has a similar dynamic range (the ratio of the power at $k_{\parallel} = 0$ h Mpc⁻¹ to the power at a specified k_{\parallel} value) to that of the visibility based power spectrum at $k_{\parallel} > 0.5$ h Mpc⁻¹.

Although closure phase quantities are immune to calibration errors under the assumption that antenna gains can be factored into antenna-based terms, they can still be affected by the antenna primary beam response, particularly if the effects of mutual coupling are significant. [Figure 1.7](#) shows the closure spectra from actual HERA observations. The closure spectra are derived from 14 m redundant equilateral triads and show variations with a median absolute deviation $\sim 4^\circ$. The differences in the primary beams of the antenna within the redundant triads may be the cause. These differences, coupled with the unsmooth sidelobe primary beam response, can impart unsmooth frequency structure to the average closure spectra, thus causing some of the foreground emissions to leak into the EoR window.

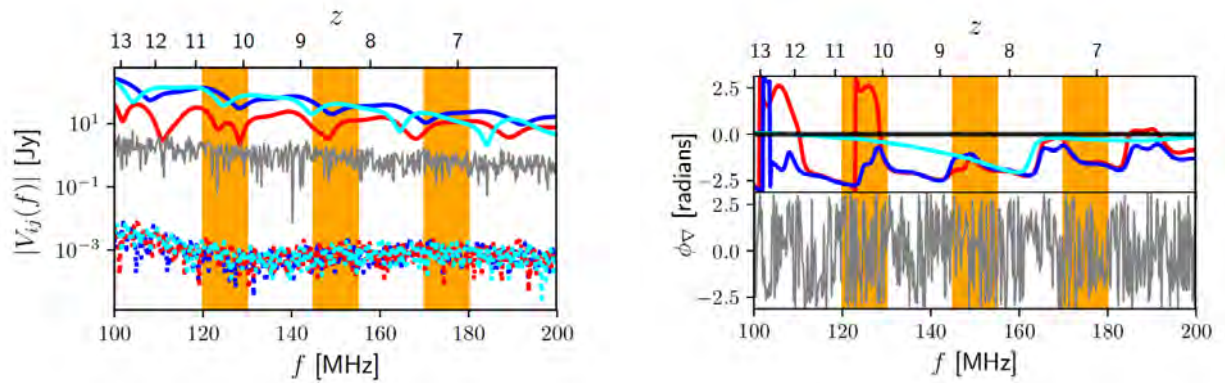


Figure 1.4: Left: Amplitude of visibility spectra for a 14.6 m baseline in the case of foreground emission (solid line) and a fiducial EoR model (dotted lines). Different colours denote visibility spectra from three baselines forming an equilateral triangle. The grey line denotes the simulated noise level obtained after 1-minute integration. The orange denotes sub-bands of 10 MHz centred at 125 MHz, 150 MHz, and 175 MHz, respectively. Right: closure spectra of a point source (black), compact sources (cyan), diffuse foregrounds (blue), diffuse and compact components combined (red), and the 21 signal fluctuations (grey) for a 14.6 m equilateral triad (from, [Thyagarajan et al., 2018](#)).




Figure 1.5: Simulated power spectra for 29.2 m baseline at $z = 7.7$. The black dotted line denotes foreground spectra and the black solid power spectrum of 21 cm, including instrumental effects. The solid grey line denotes the power spectrum of the 21 cm signal without instrumental effects. The black and grey dashed lines denote the noise level obtained after incoherent averaging (averaging different baselines) and coherent averaging, where redundant baselines are averaged together (from, [Thyagarajan et al., 2020](#)).

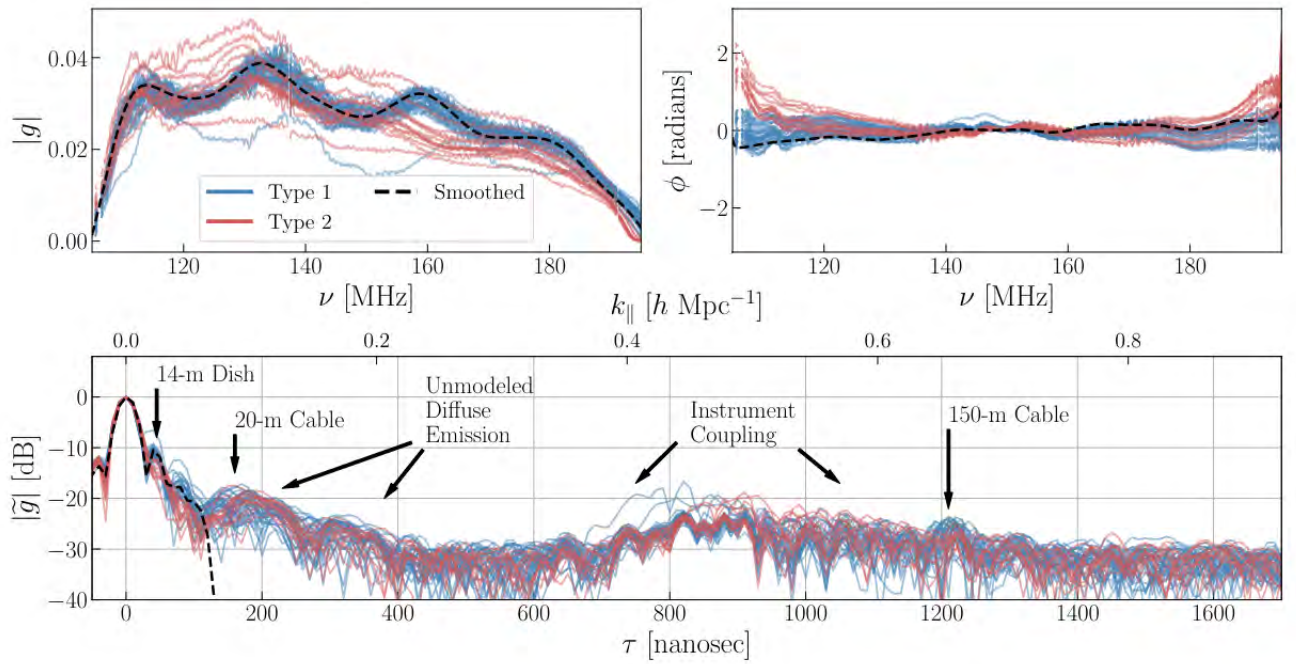


Figure 1.6: Antenna gains derived from actual HERA observations. The peak normalized amplitude response of the gains (top-left), the phase of the gains after cable delay is calibrated out (top-right). The peak normalized delay response of the gains (bottom). Notably here is that the gains show structure at various delay modes, these result from errors in calibration originating from different sources, such as unmodeled diffuse emission, mutual coupling and cable delay. Type 1 (red) and Type 2 (blue) denote data sets from the two feeds that were used for the observation, (for details see, [Kern et al., 2020a](#)).




Figure 1.7: Closure spectra from redundant 14 m equilateral triads averaged over 1 minute. Red shows the median closure spectra, and black denotes the median absolute deviation (from, [Carilli et al., 2018](#)).

1.5 The Hydrogen Epoch of Reionization Array




Figure 1.8: HERA array composed of 14 m parabolic transit dish antennas. credit: *South African Radio Astronomy Observatory (SARAO)*.

The Hydrogen Epoch of Reionisation Array (HERA) is an instrument designed to measure the 21 cm signal in the $6 < z < 30$ range (DeBoer et al., 2017). It is currently under construction in the Karoo radio-quiet area in South Africa. At the time of writing, phase II is complete and includes wide-band feeds sensitive to the 50 – 250 MHz range. The array uses both analog and digital systems (see, Fagnoni et al., 2021a). The array comprises 350 dishes, reaching out to a maximum distance of ~ 1 km, and a compact, hexagonal core that includes 320 antennas

(Figure 1.8), with a minimum spacing of 14.6 m in order to achieve a maximally redundant configuration (DeBoer et al., 2017). Redundant arrays have poor imaging performances but maximum sensitivity on a certain number of k modes and are, therefore, a promising strategy to detect the 21 cm signal (Parsons et al., 2012). However, once the outer-rigger antennas are incorporated into the array, the imaging performance is expected to improve.

1.6 Current status of 21 observations

There are two main ways of detecting the 21 cm signal: spatial fluctuations and sky-averaged (global) observations. Interferometric arrays detect spatial fluctuations (e.g., Koopmans et al., 2015; DeBoer et al., 2017; Thyagarajan et al., 2018), whereas single dipole antennas probe the global signal (e.g., Bernardi et al., 2015; Singh et al., 2017; Bowman et al., 2018; Philip et al., 2019).

Two methods are used to measure spatial fluctuations of the 21 cm signal that we will discuss in this thesis: the power spectrum (Parsons et al., 2012; Thyagarajan et al., 2013) and the power spectrum of the *closure phase* (Thyagarajan et al., 2018). The Giant Metrewave Radio Telescope (GMRT) was the first instrument to provide an upper limit of 4900 mK^2 on the 21 cm power spectrum at $z \approx 8.6$ (Paciga et al., 2011), which was subsequently revised by Paciga et al. (2013), with a more robust upper limit, although less sensitive. The robust upper limit was set to $\approx 61500 \text{ mK}^2$. The Precision Array for Probing the Epoch of Reionisation (Parsons et al., 2010), also provided several upper limits (Parsons et al., 2014; Ali et al., 2015; Kolopanis et al., 2019). The Murchison Widefield Array (MWA), indeed, also provided the best upper limits at $z = 7$ (Barry et al., 2019) and $z = 6.5$ (Trott et al., 2020). The LOw-Frequency ARay (LOFAR) probed the Cosmic dawn at high redshift ($20 < z < 25$) (Gehlot et al., 2019) and also they provided upper limits in the $8 < z < 10.5$ range (Patil et al., 2017; Mertens et al., 2020), these are all shown in Figure 1.9. At the time of writing, the deepest power spectrum upper limits on the power spectrum are 457 mK^2 and 3496 mK^2 , at $z = 7.9$ and $z = 10.4$ respectively (The HERA Collaboration et al., 2022). This represents an improvement by a factor of ~ 17 over a decade-long period. Whilst these upper limits are orders of magnitude far from the expected 21 cm

power spectrum, they are consistent with models where the IGM was heated above the adiabatic cooling limit as early as $z = 10.4$, which rules out a broad set of so-called "cold reionisation models", i.e., EoR models where the IGM is not heated above the adiabatic cooling limit.

Using the HERA first data release, [Thyagarajan et al. \(2020\)](#) first reported upper limits of the power spectrum of the closure phase at $z = 7.7$ to be 316 pseudo mK^2 at $k = 0.33 \text{ pseudo h Mpc}^{-1}$, where "pseudo" is used to distinguish the closure power spectrum from the visibility-based power spectrum. [Keller et al. \(2023\)](#) also reported upper limit of 372 pseudo mK^2 at $k = 1.14 \text{ pseudo h Mpc}^{-1}$.

Progress has also been made in global signal experiments. The Shaped Antenna measurement of background RADio Spectrum (SARAS), designed to measure the global (sky averaged) signal, constrained the duration of the EoR to $6 < z < 10$ ([Singh et al., 2017, 2018](#)). SARAS observations also disfavored a class of models with weak X-ray heating and rapid reionisation. [Bowman et al. \(2018\)](#) were the first to report the detection of a 21 cm absorption profile centred at 78 MHz ($z \approx 18$), with a width at half-maximum of approximately 19 MHz and an amplitude of about 500 mK. The result is very controversial as the amplitude is a factor two larger than the brightest models (e.g., [Cohen et al., 2017](#)). Various solutions have been explored to explain this discrepancy (e.g., [Fialkov & Barkana, 2019](#); [Natwariya & Bhatt, 2020](#)), some have also questioned the results themselves (e.g., [Singh & Subrahmanyam, 2019](#); [Hills et al., 2018](#)). Competing experiments like the Large Aperture to Detect the Dark Ages ([Bernardi et al., 2015](#); [Price et al., 2018](#)) and Probing radio Intensity at high-Z from Marion ([Philip et al., 2019](#)) will need to collaborate these results.

1.7 Thesis motivation

As we emphasised in the previous sections, foreground separation by the avoidance method is hindered by instrumental effects and calibration errors. Both effects may impart spectral structure on smooth foreground spectra, thus challenging the detection of the 21 cm signal by contaminating small k modes of the EoR with foreground emission. These errors can be in the form of direction-independent calibration errors ([Trott & Wayth, 2016](#); [Barry et al., 2016](#); [Procopio et al., 2017](#)), or direction-dependent errors from primary beam differences ([Neben et al., 2016](#);




Figure 1.9: Current upper limits on the EoR power spectrum for values in the $0.1 < k < 1 \text{ Mpc}^{-1}$ range. The dashed line is a simulated fiducial model (from Liu & Shaw, 2020).

Ewall-Wice et al., 2017; Dillon et al., 2018).

Closure phases are independent of calibration errors resulting from direction-independent calibration effects, however, they are not immune to errors resulting from direction-dependent effects. In the HERA case, a prominent source of direction-dependent errors may be due to primary beams that can be different from dish to dish due to mutual coupling effects generated by the closely packed configuration.

In this thesis, we investigate the impact that systematic errors have on EoR observations, with specific regard to the HERA case. Our goal is to model systematic effects in order to study their impact on the EoR detection and, ultimately, provide mitigation strategies - if/where possible.

In Chapter 2 we simulated the effect that different primary beams have on the power spectrum of the closure phase, i.e. what is the level of contamination to the EoR signal observed with HERA if closure quantities are used.

In Chapter 3, we turn our attention to the impact that model inaccuracies have on HERA calibration, in particular, we carry out simulations to validate the effectiveness of fringe rate filters (see Parsons et al., 2012) as a technique to filter out diffuse Galactic emission, a component difficult to model and, therefore, always neglected in the calibration process.

In Chapter 4, we assess the calibration accuracy of the first observing season taken with the HERA wide-band feeds.

In Chapter 5, we further explore the use of fringe rate filters to mitigate non-redundancy effects in HERA observations, and we conclude in Chapter 6.

Simulations of primary beam effects on the cosmic bispectrum phase observed with the Hydrogen Epoch of Reionization Array

Published as: N Charles, G Bernardi, H L Bester, O M Smirnov, C Carilli, P M Keller, N Kern, B Nikolic, N Thyagarajan, E de Lera Acedo, N Fagnoni, M G Santos, *Monthly Notices of the Royal Astronomical Society*, Volume 512, Issue 2, May 2022, Pages 2716–2727

2.1 Abstract

The 21 cm transition from neutral Hydrogen promises to be the best observational probe of the Epoch of Reionisation. The main difficulty in measuring the 21 cm signal is the presence of bright foregrounds that require very accurate interferometric calibration. Closure quantities may circumvent the calibration requirements but may be, however, affected by direction dependent effects, particularly antenna primary beam responses. This work investigates the impact of antenna primary beams affected by mutual coupling on the closure phase and its power spectrum. Our

simulations show that primary beams affected by mutual coupling lead to a leakage of foreground power into the EoR window, which can be up to $\sim 10^4$ times higher than the case where no mutual coupling is considered. This leakage is, however, essentially confined at $k < 0.3 h \text{ Mpc}^{-1}$ for triads that include 29 m baselines. The leakage magnitude is more pronounced when bright foregrounds appear in the antenna sidelobes, as expected. Finally, we find that triads that include mutual coupling beams different from each other have power spectra similar to triads that include the same type of mutual coupling beam, indicating that beam-to-beam variation within triads (or visibility pairs) is not the major source of foreground leakage in the EoR window.

2.2 Introduction

The detection of the redshifted 21 cm emission line from neutral Hydrogen during the Epoch of Reionisation (EoR) is one of the main goals of (current and upcoming) low frequency radio telescopes like the Low-Frequency Array (LOFAR, [van Haarlem et al., 2013b](#)), the Murchison Widefield Array (MWA, [Tingay et al., 2013](#)), the Giant Metrewave Radio Telescope EoR experiment (GMRT, [Paciga et al., 2013](#)), the Hydrogen Epoch of Reionisation Array (HERA, [DeBoer et al., 2017](#)), and the Square Kilometre Array (SKA, [Koopmans et al., 2015](#)). The EoR is one of the least known areas of cosmology, from an observational point of view. Advancing our understanding of reionisation will enable us to understand how the first galaxies formed and, ultimately, improve constraints on cosmological parameters (e.g., [Furlanetto et al., 2006](#); [Mesinger et al., 2016](#); [Park et al., 2019](#); [The HERA Collaboration et al., 2021](#)).

The hyperfine transition from neutral Hydrogen (21 cm emission) is one of the most promising probes of structure formation, imprinted in the intergalactic medium. Measurements of the 21 cm signal are challenged by the presence of foreground emission from the Galaxy and extragalactic sources, which are orders of magnitude brighter (e.g., [Santos et al., 2005](#); [Bernardi et al., 2009](#); [Ali et al., 2015](#)). Foregrounds are spectrally smooth, unlike the 21 cm emission line which fluctuates rapidly (e.g., [Santos et al., 2005](#)). If foreground spectral and spatial characteristics are known, they can be subtracted to isolate the 21 cm emission. The process generally begins

with the subtraction of bright, compact sources, After bright source subtraction, the sky brightness is dominated by the diffuse foreground emission (i.e., [Bernardi et al., 2010](#); [Pober et al., 2013](#); [Dillon et al., 2014](#)), which can be subtracted leveraging, again, on the foreground spectral smoothness (e.g., [Mertens et al., 2018](#); [Ghosh et al., 2020](#); [Kern & Liu, 2021](#)). In practice, though, smooth spectrum foregrounds are corrupted by instrumental effects. The calibration process attempts to correct for these corruptions. High accuracy in calibration is therefore required in order not to compromise the foreground spectral smoothness (e.g., [Wang et al., 2013](#); [Chapman et al., 2015](#); [Sims et al., 2016](#); [Datta et al., 2017](#); [Kern & Liu, 2021](#)). Additionally, some of the brightest sources can have complicated, extended morphologies: failing to model and subtract them accurately can leave residual foreground contamination that may prevent the 21 cm detection. Over-fitting diffuse emission may equally lead to 21 cm signal loss (e.g., [Wang et al., 2013](#); [Cheng et al., 2018](#)).

The calibration process makes use of a sky model to correct for instrumental effects ([Smirnov, 2011](#)). Sky models are built from catalogues of compact sources with known properties, and often cover an area larger than the field of view of the observation ([Yatawatta et al., 2013](#); [Pober et al., 2016](#)). The sky model ideally should contain the entire sky emission but, inevitably, remains incomplete due to the limited angular resolution and depth of existing catalogues ([Grobler et al., 2014](#); [Wijnholds et al., 2016](#); [Trott & Wayth, 2016](#); [Procopio et al., 2017](#); [Barry et al., 2021](#)). Calibration errors due to incomplete sky models lead to leakage of foreground power into the EoR window ([Barry et al., 2016](#); [Ewall-Wice et al., 2017](#)).

The need for highly accurate calibration required for foreground subtraction has led to alternate methods, known as *foreground avoidance methods*. As the name suggests, the idea is to avoid the foreground emission rather than subtracting it (e.g., [Parsons et al., 2012](#); [Thyagarajan et al., 2013](#)). The *delay spectrum* is one such method; it makes use of interferometric delays to isolate the power spectrum of the 21 cm emission ([Parsons et al., 2012](#)). Due to the spectral nature of the 21 cm signal, its power spectrum appears at all k modes, whereas the foreground emission is limited to a wedge like region in k -space ([Datta et al., 2010](#); [Parsons et al., 2012](#); [Vedantham et al., 2012](#); [Trott et al., 2012](#); [Thyagarajan et al., 2013](#); [Hazelton et al., 2013](#); [Pober et al., 2013](#); [Liu et al., 2016](#); [Morales et al., 2019](#)). Foreground avoidance methods remain,

however, prone to calibration errors.

Yet another alternative method to detect the 21 cm signal was proposed by [Thyagarajan et al. \(2018\)](#) and takes advantage of closure quantities. The use of closure phases mitigates calibration requirements as closure quantities are independent (to first order) of antenna based corruptions. In terms of foreground separation, [Thyagarajan et al. \(2020\)](#) showed that the dynamic range required to detect the 21 cm signal is similar to the standard power spectrum approach (e.g., [Parsons et al., 2012](#); [HERA collaboration et al., 2022](#)). For a massively redundant array like HERA ([Dillon et al., 2015](#); [DeBoer et al., 2017](#)) closure quantities may, therefore, represent an appealing alternative to the mainstream power spectrum analysis.

Closure phase quantities, however, are affected by direction-dependent effects such as varying antenna primary beams due to mutual coupling induced by a very compact configuration (e.g., [Fagnoni et al., 2021b](#); [Josaitis et al., 2022](#)). Variations of primary beams across the array invalidates the assumption of redundancy (i.e., that baselines with the same length and orientation measure exactly the same signal from the sky) which is the core of the HERA calibration strategy ([Dillon et al., 2020](#)). Several authors have empirically simulated the impact that deviations from redundancy have on the calibration and proposed possible mitigation schemes ([Ewall-Wice et al., 2017](#); [Joseph et al., 2018](#); [Orosz et al., 2019](#); [Choudhuri et al., 2021](#)). It has been established that variations from redundancy due to enhanced structure in primary beams couples foreground structure observed through sidelobes into spectral structure in calibration solutions ([Orosz et al., 2019](#); [Kern et al., 2020a](#); [Choudhuri et al., 2021](#)).

The analysis of HERA data using closure phase also showed some evidence of deviation from redundancy ([Carilli et al., 2018](#)), in particular the presence of a baseline-dependent systematic effect appearing at $k_{\parallel} \sim 0.5 h \text{ Mpc}^{-1}$ ([Thyagarajan et al., 2020](#)). At higher k modes, however, the closure phase analysis of ~ 2 h of HERA observations shows no evidence of systematic effects, suggesting that longer integrations may reduce the thermal noise ([Thyagarajan et al., 2020](#)).

In this paper we simulate the impact that different primary beams have on closure phase in the case of HERA observations, specifically investigating the case when two beams are different within a baseline pair. We quantify the effect that such deviations from redundancies have on

the power spectrum of the bispectrum phase of foreground emission and the impact on the EoR window.

The paper is organized as follows: Section 2.3 summarizes the closure phase formalism, Section 2.4 describes our simulations, Section 2.5 presents the simulated closure spectra and power spectra and we conclude in Section 2.6.

2.3 Closure phase formalism

The simplest radio interferometer is the two-element interferometer, where signals measured from a pair of antennas (p, q) are cross-multiplied and averaged in time. This operation is known as *correlation* and leads to the fundamental quantity measured in radio interferometry, the *visibility function* V_{pq} . The *Van Cittert-Zernike theorem* states that the correlation of signals from the (p, q) pair is related to the sky brightness $I(\mathbf{s}, \nu)$ by a Fourier transform like relation:

$$V_{pq}(\nu) = \iint I(\mathbf{s}, \nu) \exp(-2\pi i \frac{\nu}{c} \mathbf{b}_{pq} \cdot \mathbf{s}) \frac{dldm}{n(\mathbf{s})}, \quad (2.1)$$

where \mathbf{b}_{pq} is the baseline vector connecting antenna p and q , $\mathbf{s} = [l, m, n]^T$ is a unit vector (so that $n = \sqrt{1 - l^2 - m^2}$ with (l, m, n) the direction cosines of \mathbf{s}) representing a direction on the celestial sphere, ν is the observing frequency and c the speed of light.

In real observations, signals are corrupted by the antenna response. Corruptions are modeled using antenna-based gain terms (the so called measurement equation [Smirnov, 2011](#)) that can include the antenna primary beam pattern E . The antenna primary beam depends upon the observing direction, frequency and time - the latter normally due to the rotation of the sky with respect to the feed orientation. In this paper we investigate the response of a single polarisation feed, for which the measurement equation takes the following form:

$$V_{pq}(\nu) = \iint J_p(\mathbf{s}, \nu) I(\mathbf{s}, \nu) J_q^*(\mathbf{s}, \nu) K_{pq}(\mathbf{s}, \nu) \frac{dldm}{n(\mathbf{s})}, \quad (2.2)$$

where $J_p = G_p E_p$ is a general antenna gain that can include a direction independent contribution G_p as well as the direction dependent antenna primary beam pattern E_p and we have introduced the abbreviated notation $K_{pq}(\mathbf{s}, \nu)$ representing the exponential term introduced in (2.1). It is

normally assumed that the primary beam is the same for the two receiving elements, i.e. $E_p = E_q$, but here we explicitly explore beams that are different from each other, i.e. $E_p \neq E_q$.

In this work we assume that the sky can always be modelled as a collection of point sources so that equation 2.2 can be discretised as:

$$V_{pq\nu} = \sum_s J_{ps\nu} I_{s\nu} J_{qs\nu}^* K_{pq\nu} = \sum_s J_{ps\nu} X_{pq\nu} J_{qs\nu}^*, \quad (2.3)$$

where s labels sources so that $I_{s\nu}$ is the point source flux density $I(\mathbf{s}, \nu)$ in the direction of source s at frequency ν , for example. Here we have also introduced the source coherency $X_{pq\nu}$ corresponding to the visibilities of a specific source. As shown below, closure phases are independent of direction independent antenna based gains. Thus we use equation 2.3 with $J_p = E_p$ to apply differing primary beams to simulated sky models. Equation 2.3 is implemented efficiently using the *codex-africanus* package (Perkins et al., 2021) and written to a measurement set format using *dask-ms* (see, Perkins, 2021).

At this point we introduce the visibility bispectrum C_{pqr} , defined as the triple product of three visibilities from baselines (pq, qr, rp) :

$$C_{pqr} = V_{pq} V_{qr} V_{rp}, \quad (2.4)$$

where indices p, q, r are antenna labels. Assuming that antenna gains consist of purely direction independent gains, we can rewrite equation (2.3) as

$$V_{pq\nu} = G_{p\nu} X_{pq\nu} G_{q\nu}^*, \quad (2.5)$$

where the individual source coherencies have been combined into a single model coherency term i.e. $X_{pq\nu} = \sum_s X_{pq\nu}$. The bispectrum then becomes:

$$C_{pqr} = |G_p|^2 |G_q|^2 |G_r|^2 X_{pq} X_{qr} X_{rp}, \quad (2.6)$$

as the antenna-gain phases cancel out. Splitting the model coherencies into amplitudes and phases

$$X_{pq} = |X_{pq}| \exp(i\phi_{pq}), \quad (2.7)$$

we see that the phase of the bispectrum (also known as the closure phase):

$$\phi_{\nabla} = \phi_{pq} + \phi_{qr} + \phi_{rp}, \quad (2.8)$$

is independent of antenna based direction independent gains. Here we use ϕ_{∇} to denote the bispectrum phase of a closed triad. It is worth noticing that closure quantities are not immune from spectral structure imparted on the sky signal by the instrument - as it will appear from simulations carried out in this work.

Closure phases will contain contributions from both foregrounds and the cosmological signal of interest. [Thyagarajan et al. \(2018\)](#) showed that it is possible to separate the two contributions by leveraging the different frequency behaviour of the foreground and 21 cm signal closure spectra. In analogy with the delay spectrum approach ([Parsons et al., 2012](#); [Pober et al., 2013](#)), they suggested to form a power spectrum P_{∇} by taking the Fourier transform along the frequency axis of the complex bispectrum phase ([Thyagarajan et al., 2020](#)):

$$P_{\nabla}(k_{\parallel}) = |\tilde{\Psi}_{\nabla}|^2 \left(\frac{\lambda^2}{2k_B} \right)^2 \left(\frac{D_c^2 \Delta D_c}{B_{\text{eff}}} \right) \left(\frac{1}{\Omega B_{\text{eff}}} \right), \quad (2.9)$$

where $\Delta D = \Delta D(z)$ is the comoving depth along the line of sight corresponding to an effective bandwidth B_{eff} , $\Omega = A_e/\lambda^2$ where A_e is the effective aperture area and:

$$\tilde{\Psi}_{\nabla} = \tilde{W}(\tau) * \tilde{\Xi}_{\nabla}(\tau) * V_{\text{eff}} * \delta(\tau), \quad (2.10)$$

where $*$ denotes convolution, \tilde{W} and $\tilde{\Xi}_{\nabla}$ are the delay transforms of the window function W and the complex closure phase Ξ_{∇} respectively:

$$\Xi_{\nabla}(\nu) = e^{i\phi_{\nabla}(\nu)}. \quad (2.11)$$

In this paper we used a Blackman Harris window function (e.g., [Parsons et al., 2012](#); [Thyagarajan et al., 2013](#)) an effective bandwidth $B_{\text{eff}} = 9.77$ MHz, centred at 175 MHz ($z = 7.1$) with a 97.66 kHz channel width i.e. the same observing setup as HERA. We also simulate a single snapshot observation. The normalization factor V_{eff} is defined as ([Thyagarajan & Carilli, 2020](#)):

$$(V_{\text{eff}})^{-2} = \sum_{b=1}^3 |V'_b|^{-2}, \quad (2.12)$$

where

$$V'_b = \frac{\int W(\nu) V_b(\nu) d\nu}{\int W(\nu) d\nu}, \quad (2.13)$$

and b denotes baselines in a triad. The closure spectrum is a dimensionless quantity and the normalisation factor calibrates the power spectra of fields with different brightness distribution on the same scale. We note that the power spectrum of the bispectrum phase is not directly comparable with the standard power spectrum (e.g., [HERA collaboration et al., 2022](#)), even if they share the same units.

In this paper we construct different power spectra of the bispectrum phase by simulating visibilities with beams that are different for each **receiver** using equation [2.3](#).

2.4 Bispectrum phase simulations

2.4.1 Beam models

HERA antennas consist of a dipole feed suspended above a parabolic dish with a diameter of 14 m. The dish structure was initially designed by paying specific attention to its spectral response, i.e. keeping the dish reflections and passband sufficiently smooth so that the EoR window would be preserved at $k_{\parallel} > 0.2 h \text{ Mpc}^{-1}$ modes ([Ewall-Wice et al., 2016](#); [Thyagarajan et al., 2016](#)). Further electromagnetic simulations ([Fagnoni et al., 2021b](#)) generated a primary beam model that is routinely used in the analysis and simulations of HERA observations (e.g., [Martinot et al., 2018](#); [Kern et al., 2020a](#)). Due to compactness of the array, however, effects such as cross coupling amongst antennas, i.e. mutual coupling, cause deviations to the ideal antenna model.

[Fagnoni et al. \(2021b\)](#) also carried out simulations of HERA dishes which included the receiving system and the effects of mutual coupling for the two polarisations, i.e. X and Y , for a redundant, compact array layout that included 19, hexagonally-packed HERA dishes that were the first HERA installment ([Kohn et al., 2019](#), Figure [2.1](#)). Due to the interaction with many more dishes, mutual coupling may be subtly different for the full HERA array compared to the models employed in this work ([Dillon et al., 2015](#)). However, it would be surprising if subtle




Figure 2.1: Simulated HERA 19 array layout. In this work, we used the simulated primary beam pattern corresponding to label 22, for a central dish and beam patterns corresponding to two edge antennas, i.e. label 26 and 6. We only make use of Y polarization (North-South) primary beam patterns.




Figure 2.2: Cut through the HERA primary beam models at $\nu = 175$ MHz for the Y polarisation. Whereas the main lobe structure remains essentially the same for all beams, sidelobes have a more prominent structure for beams with mutual coupling and appear asymmetric in the case of antenna 26.

differences in the beam patterns significantly alter the derived power spectra. Thus we believe that our results should be qualitatively correct and should also hold for the full HERA array.

Fagnoni et al. (2021b) showed that mutual coupling introduces extra sidelobe ripples (Figure 2.2) and increases the sidelobe level by 2 – 4 dB. Figure 23 in Fagnoni et al. (2021b) shows that the gain at zenith also varies as a function of frequency, up to ~ 0.3 dB with respect to the ideal beam and for different antenna positions within the array. The beam value at zenith oscillates with a periodicity of about 20 MHz, which corresponds to reflections occurring at 15 m path length, approximately the distance between the centre of two dishes. These effects lead to further deviations from the smooth ideal beam response.

Lastly, antennas experience a varying degree of mutual coupling and, as a consequence, an antenna at the edge array has an asymmetric primary beam pattern since one side of antenna experiences more mutual coupling, i.e. side facing other dishes, than the other side where there are no dishes (Figure 2.2, antenna 26).

2.4.2 Simple sky models

We begin with simulating simplified sky models in order to demonstrate some basic properties of closure spectra and mutual coupling beams. We generate sky models where we randomly place 100 point sources in the main lobe and another 100 across the beam sidelobes (Figure 2.3). All sources on the main lobe of the primary beam have a 1 Jy flux density at 150 MHz. We vary the flux density of the sources on sidelobes to create four different sky models, which we name A0, A1, A2 and A3 respectively. These sky models are meant to illustrate fields with faint and bright emission in the sidelobe area. Flux densities of sources corresponding to the different models are shown in Table 2.1. All sources have a spectral index $\alpha = 0.7$, where α is defined such that flux density S of a source at frequency ν is given by:

$$S(\nu) = S_0 \left(\frac{\nu}{\nu_0} \right)^{-\alpha}, \quad (2.14)$$

and S_0 is the flux density at some reference frequency ν_0 .

We use beams from dishes 22, 26 and 6 (Figure 2.1) and simulate the effect of mutual cou-








Figure 2.3: Position of sources in simulated sky models, the blue dots mark point source positions within the beam main lobe and the red dots mark point source positions in the sidelobe region. Black lines mark nulls of the HERA ideal beam model at $\nu = 175$ MHz.

Sky model ID	Flux density of sources in main lobe (Jy)	Flux density of sources in the sidelobe region (Jy)
A0	1	0
A1	1	0.01
A2	1	1
A3	0.01	1

Table 2.1: Characteristics of the simple sky models used in simulations (Figure 2.3).



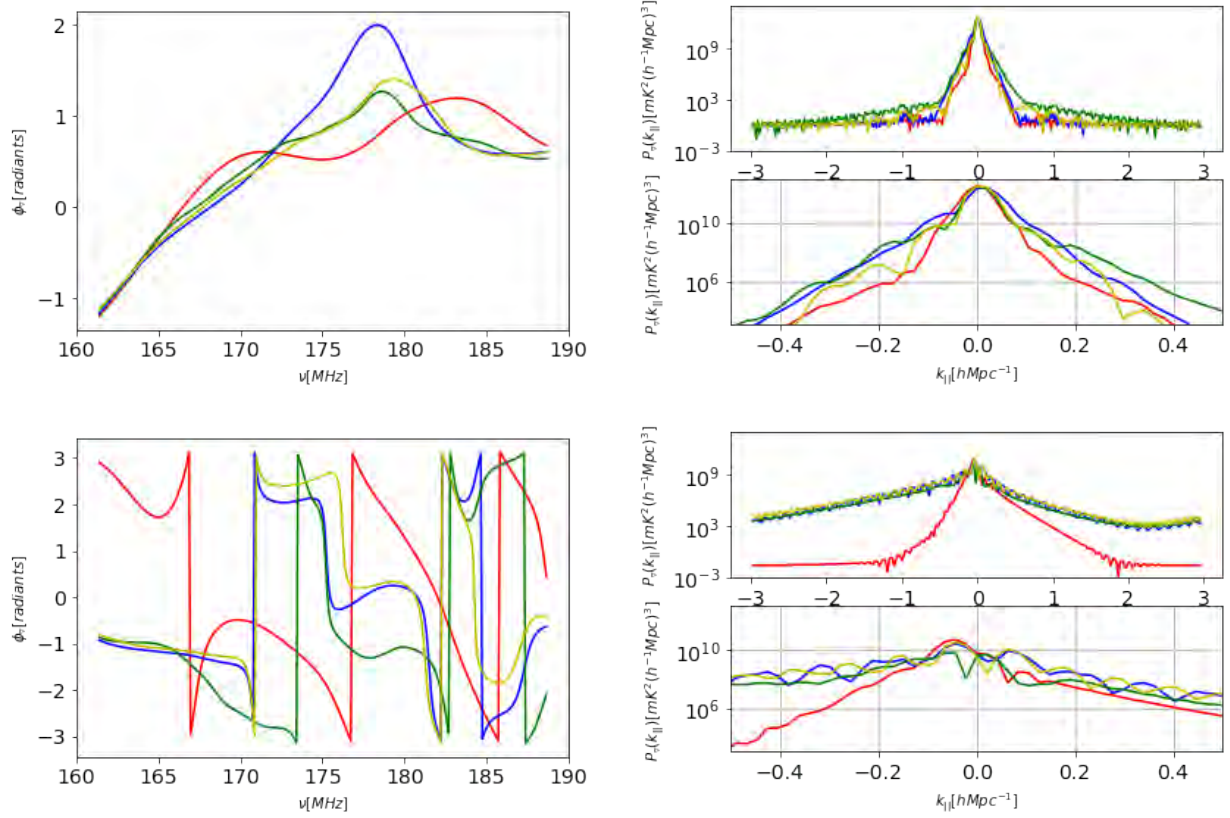


Figure 2.4: Simulated closure spectra (left column) and power spectra of the bispectrum phase (right column) corresponding to sky models A0 (top row), A1 (second row), A2 (third row) and A3 (fourth row) - see text for details. The following triads are shown: ∇_{HHH} (red), ∇_{CCC} (blue), ∇_{CCE} (yellow) and ∇_{CEE} (green). Bottom panels on the right column are zoom into the corresponding upper panels.

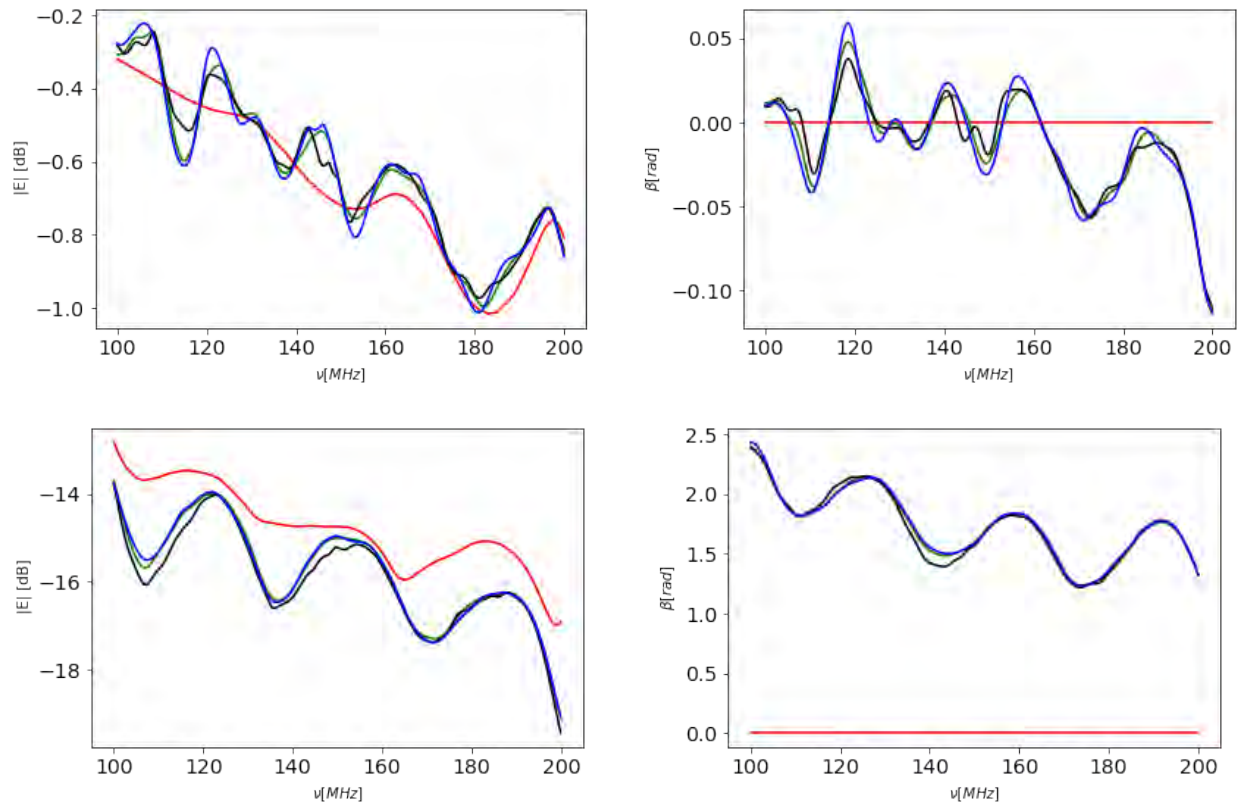


Figure 2.5: Top row: the left panel shows beam response [dB] averaged across all source locations for the sky model with sources only in the main lobe (blue dots in Figure 2.3) as a function of frequency for E_{06} (green), E_{22} (blue), E_{26} (black) and E_H (red). The right panel corresponding beam phase as a function of frequency. Second row: same, but for sources in the sidelobe region (red dots in Figure 2.3).

pling of an antenna placed at the centre of the array and two at the edge, respectively. Hereon, we denote primary beams from dish 22, 26, 6 and the HERA ideal beam, i.e. beam with no mutual coupling, as E_{22} , E_{26} , E_{06} and E_H respectively. We combine different primary beams to simulate four types of 29 m equilateral triads: (1) a triad at the centre of the array, with only E_{22} beams (∇_{CCC}); (2) one at the edge of the array with one centre beam E_{22} and two different edge beams E_{26} and E_{06} (∇_{CEE}); (3) a second triad at the edge of the array with two E_{22} beams and one E_{26} beam (∇_{ECC}) and, finally, (4) a triad unaffected by mutual coupling with just E_H beams (∇_{HHH}).

We simulate noiseless visibilities. We acknowledge that, in practice, each dish has a unique beam as the mutual coupling varies across the array and the primary beam corresponding to dish 8 would be different than the primary beam for dish 6, for example. However, this approximation is acceptable for the scope of our investigation.

Figure 2.4 shows the corresponding closure spectra and power spectra of the bispectrum phase. We first consider sky models with faint or no sources in the sidelobe region, i.e., A0 and A1. Closure spectra of both models are essentially identical, as they are dominated by sources within the primary beam main lobe. Main lobes have a very similar spectral structure for all the beams, yielding very similar closure and power spectra. Among the triads, the one that includes only the ideal beam has the smoothest frequency behaviour - as expected. Similarly, the power spectrum of the ∇_{HHH} triad has a very distinct behaviour: the power is concentrated at small k_{\parallel} modes as it is expected for smooth spectrum foregrounds (Thyagarajan et al., 2018) and falls already by $\sim 10^8$ times at $k_{\parallel} \sim 0.1 h \text{ Mpc}^{-1}$. Power spectra of triads with mutual coupling beams, conversely, show up to $\sim 10^4 \text{ mK}^2 (h^{-1} \text{ Mpc})^3$ higher power starting already at $k_{\parallel} \sim 0.1 h \text{ Mpc}^{-1}$. This is indeed indicative of excess spectral structure in the closure spectra of triads with mutual coupling, likely arising from the gain variation at zenith and the different phases for different beams.

As we increase the brightness of the sources in the sidelobe region, i.e. model A2, closure spectra from all triads show extra frequency structure compared to model A0 and A1, due to sidelobe ripples. This results in an excess power up $\sim 10^4 \text{ mK}^2 (h^{-1} \text{ Mpc})^3$ for triads with mutual coupling, notably at large k_{\parallel} values, $0.5 < |k_{\parallel}| < 2 h \text{ Mpc}^{-1}$. It is worth noting that the

leakage is more pronounced for triads with different primary beam patterns i.e. the edge triads.

When we decrease the brightness of the sources in the main lobe, i.e. sky model A3, the leakage is much worse, with an excess power between 10^4 and 10^8 $\text{mK}^2 (h^{-1} \text{Mpc})^3$ at $|k_{\parallel}| > 0.1 h \text{Mpc}^{-1}$. We also note that in sky model A3, the center triad, where we have no beam variation, shows an excess power leakage comparable to edge triads. This shows that a large portion power leakage in A3 is actually caused by ripples on the sidelobes of mutual coupling beams (see Figure 2.2).

In summary, our simulations show that, in the presence of bright emissions on the sidelobes, we may expect the power to leak at high k_{\parallel} values, and the power leakage increases with brightness of sources on the sidelobes. Indeed, previous work by Choudhuri et al. (2021) shows similar results, as well as analysis of HERA data (e.g., Kern et al., 2020a; Dillon et al., 2020). In addition, our simulations also show that the presence of bright sources on the main lobe mitigates the power leakage by dominating the overall closure spectra. In the case of extremely bright sources on sidelobes, we may expect the sidelobe ripples from mutual coupling beams to contribute a large fraction of the power leakage observed.

2.4.3 Simulations with realistic sky models

After we have treated simplified sky models, we simulate three realistic zenith-pointed observations centred at right ascension $\alpha = (3^{\text{h}} 20^{\text{m}} 6.7^{\text{s}})$, $(5^{\text{h}} 20^{\text{m}} 6.7^{\text{s}})$ and $(10^{\text{h}} 20^{\text{m}} 6.7^{\text{s}})$, that we label field 1, field 2 and field 3 respectively. They are located within the stripe observed by HERA (HERA collaboration et al., 2022). As mentioned in Section 2.3, we only simulate single-snapshot observations. We include three sky model components for each pointing:

- all point sources brighter than 200 mJy at 151 MHz and within a $100^{\circ} \times 70^{\circ}$ region around the centre of each pointing, taken from the GLEAM catalogue (Hurley-Walker et al., 2017);
- Fornax A - which is not included in the GLEAM catalogue. The source is modeled as a core and two lobes, based on observations at 174 MHz (McKinley et al., 2015, Figure 2.6). The core is modelled with a circular Gaussian with a $5'$ axis, a $\alpha = 1$ spectral index and a




Figure 2.6: Fornax A model image at $\nu = 100$ MHz (from [McKinley et al., 2015](#)). Units are Jy pixel^{-1} with a pixel size = 0.75 arcmin.




Figure 2.7: Left column shows sky model images of simulated fields, field 1 (first row), field 2 (second row) and field 3 (third row). Right columns shows the corresponding apparent sky model, i.e. after the primary beam, E_H is applied. Units are $\log_{10} |I(\text{Jy pixel}^{-1})|$.

12 Jy flux density at 154 MHz. The West lobe is modelled with a circular Gaussian with a $20'$ axis, a $\alpha = 0.77$ spectral index and a 260 Jy flux density at 154 MHz. The East lobe is modelled with a circular Gaussian with a $15'$, a $\alpha = 0.77$ spectral index and a 480 Jy flux density at 154 MHz. Visibilities are generated using equation 2.3, treating each image pixel as a point source;

- an all-sky map of Galactic diffuse emission at 408 MHz (Remazeilles et al., 2015) with a $56'$ resolution. The map (in the Healpix format) was extrapolated to 150 MHz using a spatially constant spectral index $\alpha = 0.7$. Like the Fornax A case, each Healpix pixel is treated as a point source in equation 2.3.

2.5 Results

Figure 2.7 shows the model images of field 1, 2 and 3, together with apparent sky model obtained by applying the beam without mutual coupling E_H . Figure 2.8 shows visibility spectra corresponding to our sky models for the 29 m triad ∇_{HHH} .

Fornax A is the dominant source in field 1, and its visibility spectra are essentially the same as the case when the sky model includes both GLEAM sources and the diffuse emission (“full foreground model”). Fornax A is, however, in the primary beam sidelobe region in field 2 and, therefore, largely attenuated, with an apparent flux density up to ~ 8 Jy. Field 2 is a relatively cold patch of the sky, with the Galactic plane on the far sidelobes of the primary beam. As a result, GLEAM sources in the main lobe are the dominant component, largely determining the visibility spectra of the full foreground model. Conversely, field 3 corresponds to an area of relatively bright diffuse emission, particularly at low frequencies ($\nu < 120$ MHz), with the Galactic plane appearing across the second sidelobe of the beam. Beyond this range, GLEAM sources dominate, including the frequency range used for power spectra in this work, i.e. 160 – 190 MHz.

It is worth noting the striking difference between visibility spectra for field 3 in the case of the HERA ideal beam and the mutual coupling beam: although point source visibility spectra

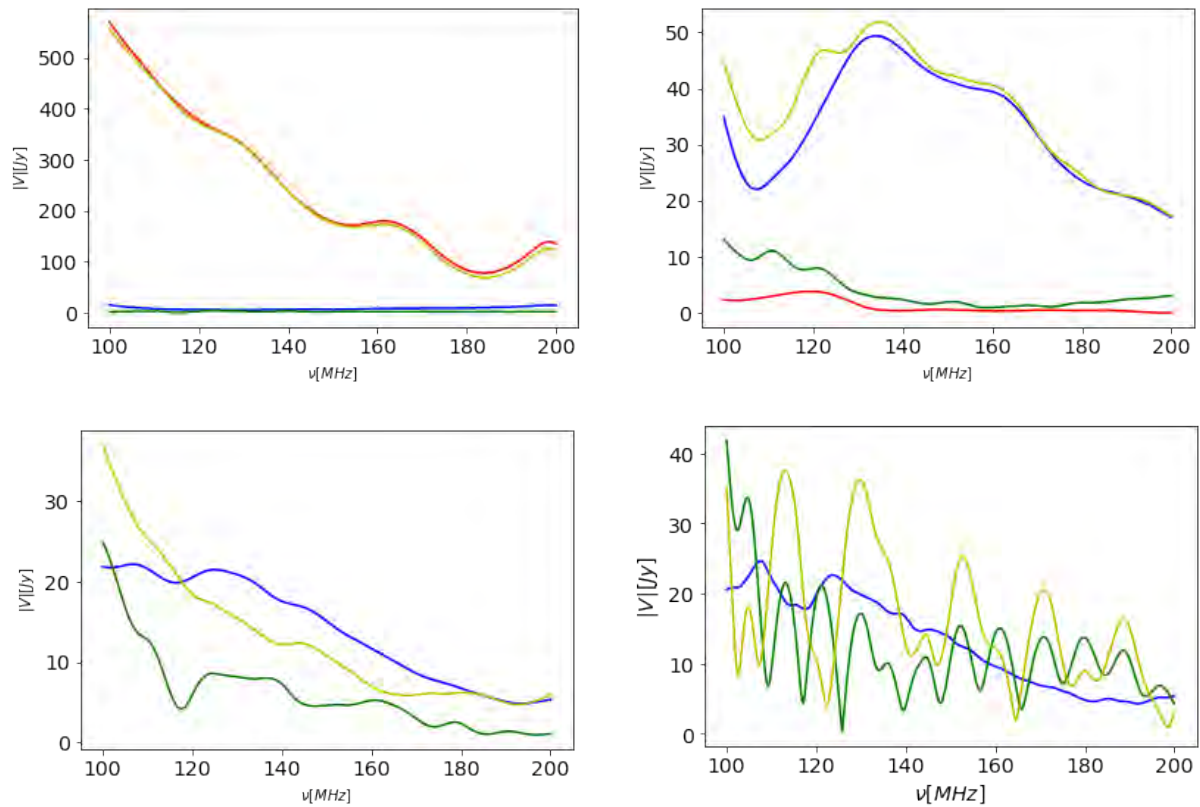


Figure 2.8: Simulated visibility spectra corresponding to sky models for field 1 (top left panel), field 2 (top right panel) and field 3 (bottom left) for a baseline b_{HH} in triad ∇_{HHH} . Colours indicate sky model components: Fornax A (red), GLEAM sources (blue), diffuse emission (green) and full sky model (Fornax A + GLEAM sources + diffuse emission; yellow). Note that Fornax A is not included in field 3. The bottom right panel shows simulated visibility spectra for triad ∇_{CCC} corresponding to the sky model for field 3.

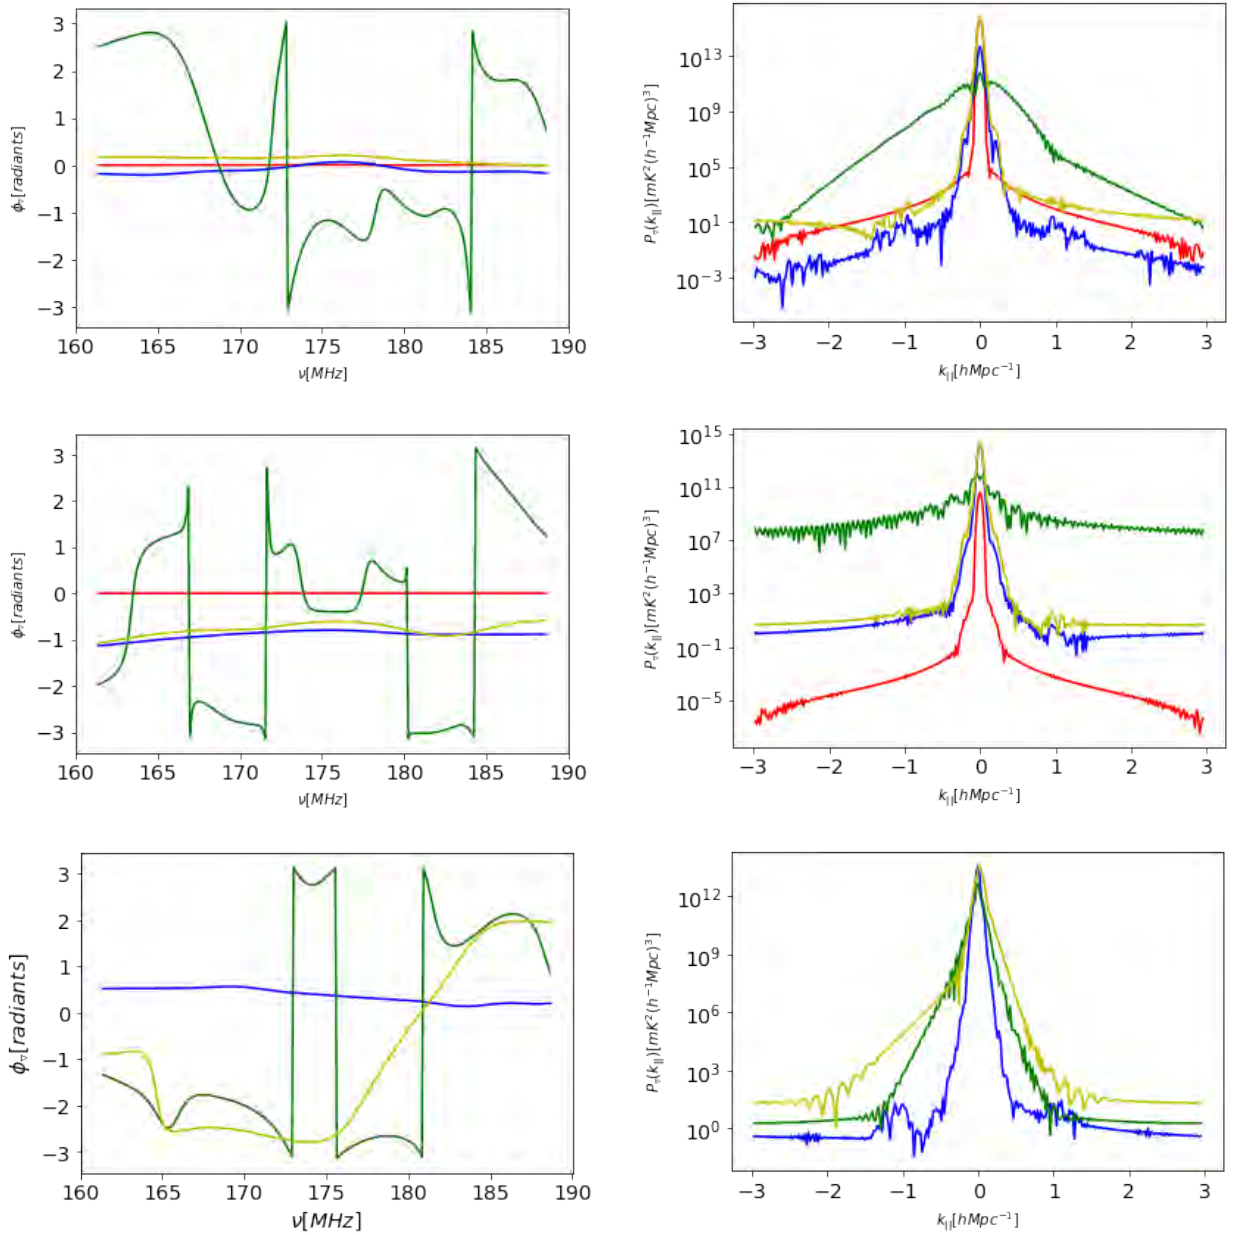


Figure 2.9: Simulated closure spectra (left column) and power spectra (right column) corresponding to sky models field 1 (first row), field 2 (second row) and field 3 (third row) from triad ∇_{CCC} . Colours indicate sky model components: Fornax A (red), GLEAM sources (blue), diffuse emission (green) and full sky model (Fornax A, GLEAM sources and diffuse emission; yellow). Note that Fornax A is not included in field 3.

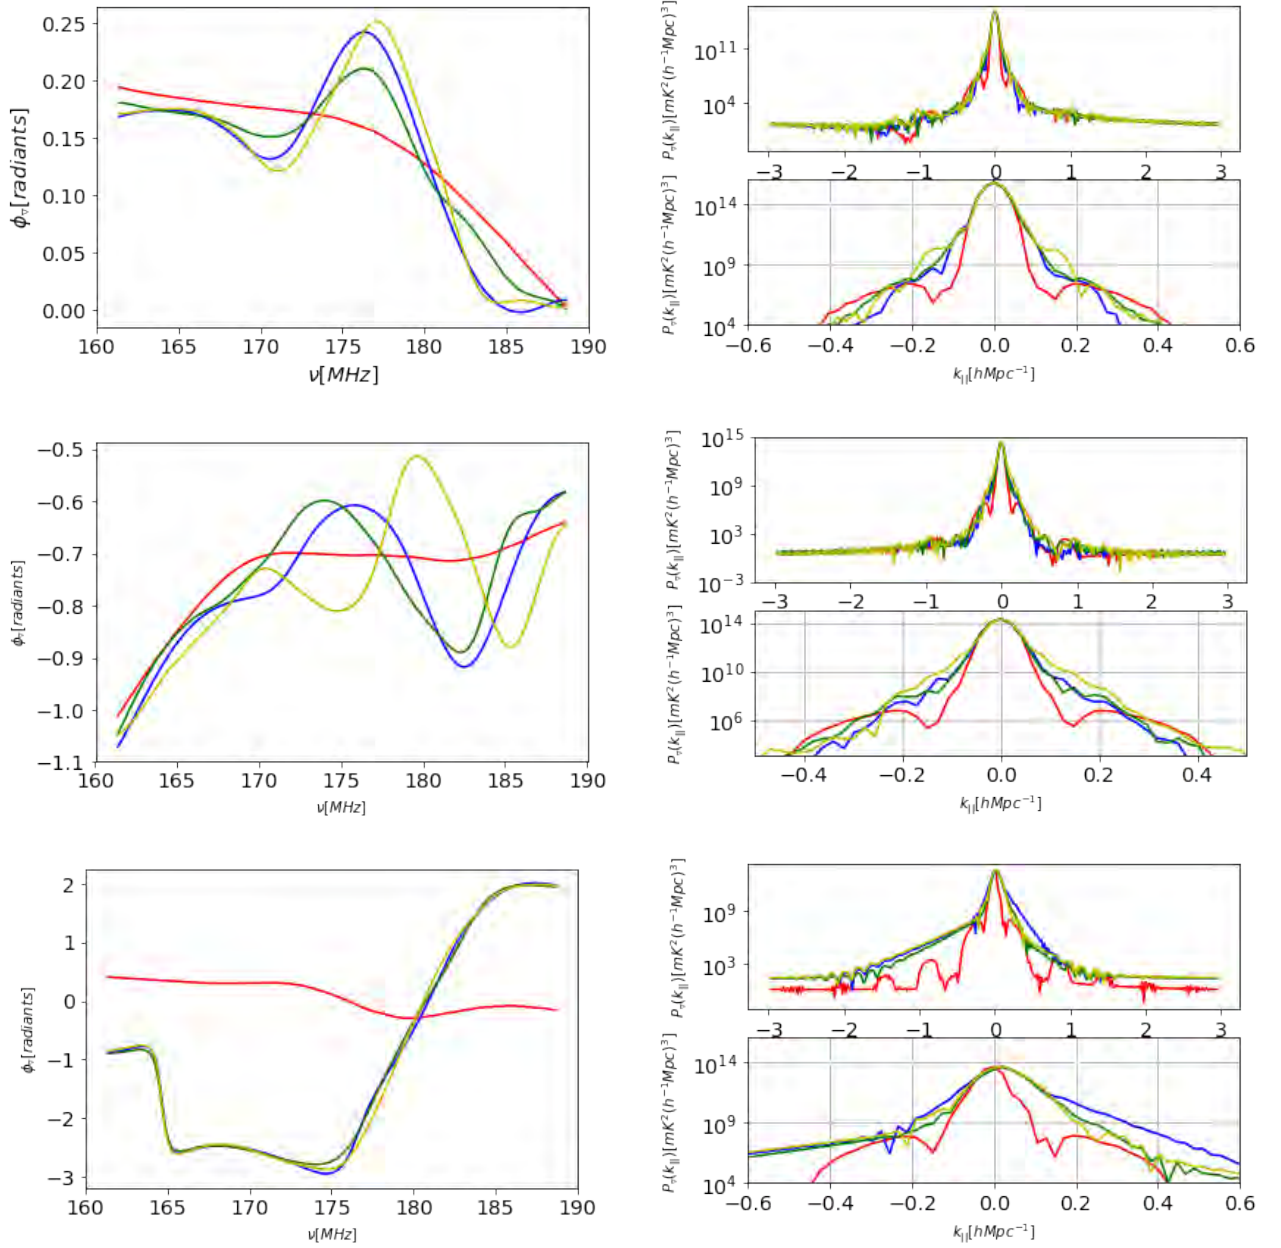


Figure 2.10: Left column: closure spectra corresponding to sky models field 1 (first row), field 2 (second row) and field 3 (third row). The colours indicate here, the closure spectrum from different simulated triads: ∇_{HHH} (red), ∇_{CCC} (blue), ∇_{CEE} (green) and ∇_{ECC} (yellow). Right column: corresponding power spectra.

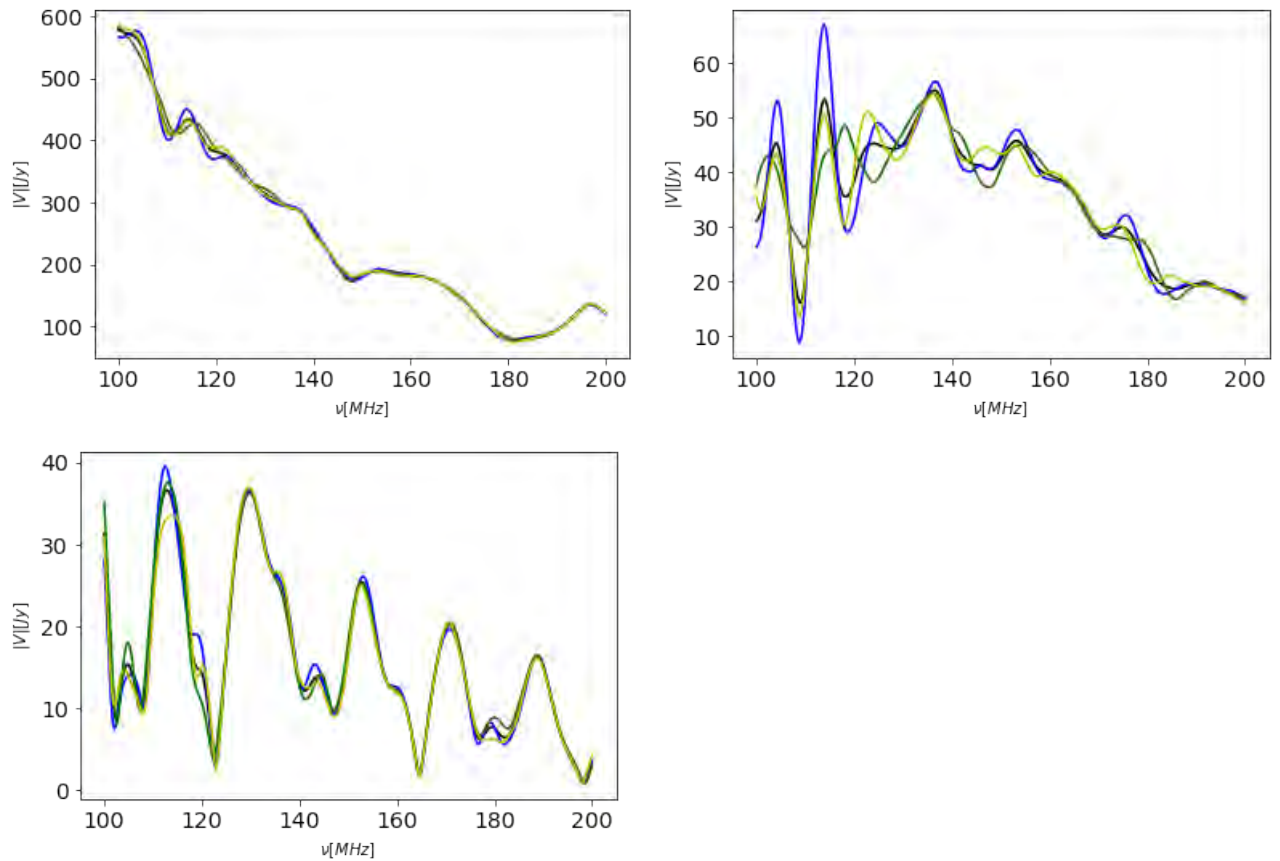


Figure 2.11: Simulated visibility spectra corresponding to field 1 (top panel), field 2 (middle panel) and field 3 (bottom panel). Colours indicate baselines with different primary beams: *EE* beams (blue), *EC* beams (green) and *CC* beams (yellow). Black shows the average visibility spectra.




Figure 2.12: Cross power spectra phases (see equation 2.15) of simulated triads with mutual coupling for field 2 for triad $(\nabla_{ECC}, \nabla_{EEC})$ (blue), $(\nabla_{EEC}, \nabla_{CCC})$ (green) and $(\nabla_{ECC}, \nabla_{CCC})$ (yellow), respectively.

are not tremendously different, visibility spectra of the diffuse emission component oscillates significantly across the 100 – 200 MHz range, with peak-to-peak variations occurring with a ~ 5 MHz period. The full foreground model spectrum, after attenuation by the mutual coupling beams, is far from being spectrally smooth.

Figure 2.9 shows the closure spectra and the corresponding power spectra of each sky model component for triad ∇_{CCC} . In field 1 Fornax A is the dominant source and appears point-like for the 29 m triad, with an approximately zero closure spectrum. Diffuse emission, however, shows closure spectra with pronounced frequency structure, likely due to emission from the Galactic plane in the beam sidelobes at $\alpha = 6 - 7^{\text{h}}$ (see top left and right panel of Figure 2.7).

This results in a $10^4 - 10^8 \text{ mK}^2 (h^{-1} \text{ Mpc})^3$ excess power at $k_{\parallel} > 0.5 h \text{ Mpc}^{-1}$ above the Fornax A model. As Fornax A is the brightest source, the power spectrum of the full foreground model closely resembles the Fornax A one: the foreground power is largely confined at small k_{\parallel} values, i.e. $|k_{\parallel}| < 0.5 h \text{ Mpc}^{-1}$, and remains flat at larger k modes.

Closure spectra and power spectra for the field 2 case, are similar to field 1. GLEAM sources are the dominant foreground component at all frequencies and their closure spectra is fairly smooth in frequency. This results in a power spectrum of the full foreground model similar to field 1, with power contained at $|k_{\parallel}| < 0.5 h \text{ Mpc}^{-1}$. Due the Galactic plane in the sidelobe region, the closure spectrum of diffuse emission shows pronounced frequency structure, compared to the closure spectrum of GLEAM sources, resulting in a $\sim 10^8$ times higher power than GLEAM sources at $|k_{\parallel}| > 0.5 h \text{ Mpc}^{-1}$.

Field 3 is a different case, where contributions from diffuse emission and GLEAM sources are at a comparable level. The closure spectrum of the full foreground model has frequency structure due to the coupling of diffuse emission and beam sidelobes. The power spectrum is different compared to the two other fields as there is foreground power up to $k_{\parallel} \sim 0.5 h \text{ Mpc}^{-1}$ and even beyond in the case of negative k modes, with the asymmetry due to the asymmetric brightness distribution of the Galactic plane. This inevitably results in power that is between $10^2 - 10^7$ times higher than power from GLEAM sources at $|k_{\parallel}| > 0.5 h \text{ Mpc}^{-1}$. With no bright source on the main lobe to “mitigate” the leakage from diffuse emission, the composite sky model shows an excess power that can be $10^4 - 10^6 \text{ mK}^2 (h^{-1} \text{ Mpc})^3$ at $|k_{\parallel}| \sim 0.5 h \text{ Mpc}^{-1}$

compared to the other two fields.

Figure 2.10 displays closure spectra and power spectra P_{∇} for the three simulated pointings for all triads. In the case of field 1, closure spectra that include mutual coupling beams, i.e. ∇_{CCC} , ∇_{ECC} and ∇_{CEE} , have a more pronounced frequency structure compared to the ∇_{HHH} triad. Power spectra of triads with mutual coupling beams show a slight broadening in the $0.1 < k_{\parallel} < 0.2 h \text{ Mpc}^{-1}$ compared to the triad with ideal beams, whereas all the triads have similar power spectra beyond $k_{\parallel} \sim 0.2 h \text{ Mpc}^{-1}$. Power spectra that include triads with mutual coupling beams have fairly similar power spectra across the whole k_{\parallel} range, independent of the beam type.

In the case of field 2, closure spectra from triads ∇_{CEE} , ∇_{ECC} and ∇_{CCC} show a more pronounced frequency structure compared to field 1, due to the presence of the Galactic plane in the sidelobe region, together with Fornax A. Their power spectra shows excess power up to $10^4 \text{ mK}^2 (h^{-1} \text{ Mpc})^3$ at $0.1 < |k_{\parallel}| < 0.2 h \text{ Mpc}^{-1}$ compared to the triad with ideal primary beams. Like field 1, power spectra that include triads with mutual coupling beams have fairly similar power spectra across the whole k_{\parallel} range, independent of the beam type.

Field 3 has bright emission in the sidelobe region and, therefore, closure spectra from triads with mutual coupling beams show even more frequency structure compared with the other fields. Power spectra are brighter compared two the previous two fields for $|k_{\parallel}| > 0.1 h \text{ Mpc}^{-1}$, with an excess power up to $10^6 \text{ mK}^2 (h^{-1} \text{ Mpc})^3$ compared to the ideal beam triad. They also tend to show some level of asymmetry between positive and negative k_{\parallel} modes.

We also present visibility spectra of simulated triads for completeness (Figure 2.11) and use them to provide an estimate of the deviation from redundancy, which may have an impact on calibration. We use the absolute difference between visibility spectra of antennas affected by mutual coupling and the average visibility spectra as a metric to quantify deviations from redundancy, averaged over the 160 – 190 MHz range. In the case of field 1, where most of the foreground emission is within the main beam, visibility spectra are fairly similar to each other and their deviation from redundancy is $\sim 2\%$. In field 2 and 3 where there are bright emissions on the sidelobe, the non-redundancy proves to be higher, with a non-redundancy value of $\sim 6\%$ and $\sim 7\%$ respectively, qualitatively previous works (e.g., Choudhuri et al., 2021) have also shown similar results. These values are also well within the 10% non- redundancy estimated by

previous studies (e.g., [Dillon et al., 2020](#)).

The impact that systematic effects induced by beam-to-beam variations may have on the detection of the 21 cm need a further investigation that we leave for the future. However, we looked at their effect when different triads are averaged together, like it is in actual observations ([Thyagarajan et al., 2020](#)). Rather than the power spectrum, we computed the cross-spectra P_{∇}^c between two 29 m triads with different primary beams ∇ and ∇' :

$$P_{\nabla}^c(k_{||}) = \tilde{\Psi}_{\nabla} \tilde{\Psi}_{\nabla'}^*, \left(\frac{\lambda^2}{2k_B} \right)^2 \left(\frac{D_c^2 \Delta D_c}{B_{\text{eff}}} \right) \left(\frac{1}{\Omega B_{\text{eff}}} \right). \quad (2.15)$$

We compute cross power spectra between triads affected by mutual coupling, namely ∇_{ECC} , ∇_{EEC} , and the ideal triad ∇_{CCC} , and show their phase in [Figure 2.12](#). The cross spectrum phase shows a certain degree of incoherency across triad pairs, with variations as large as π at the same k modes. This suggests that averaging cross power spectra together may lead a suppression of systematic effects induced by mutual coupling beams - in particular considering the large number of different beams in the final HERA configuration.

2.6 Discussion and conclusions

In this work we investigate the impact that primary beams affected by mutual coupling have on closure phase, used to detect the EoR signal. We use electromagnetic simulations of HERA dishes and both a simplified and a realistic foreground model in order to perform simulations of closure spectra and its power spectra. In the simulations we specifically include antenna pairs where primary beams are different from each other. We focus only on triad separated by 29 m baselines, already used in the early analysis of HERA closure spectra ([Carilli et al., 2020](#); [Thyagarajan et al., 2020](#)). As realistic foreground models, we include both point sources and diffuse emission. We simulated three different fields who range from a high to a low ratio between the foreground emission in the main beam lobe and in the sidelobe region. Our main conclusions may be summarized as following:

- in the presence of beams distorted by mutual coupling effects, closure spectra exhibit more pronounced frequency structure with respect to ideal beams, i.e. not affected by mutual

coupling. The effect on the power spectrum of the bispectrum phase is that foreground power bleeds from small k modes to intermediate modes, e.g., $0.1 < |k_{\parallel}| < 0.2 h \text{ Mpc}^{-1}$. Such excess power is $\sim 10^3 \text{ mK}^2 (h^{-1} \text{ Mpc})^3$ and does not vary significantly with different mutual coupling beam model or foreground model. Power spectra are not significantly different between models with or without mutual coupling at $|k_{\parallel}| > 0.2 h \text{ Mpc}^{-1}$. The presence of diffuse foreground emission that is brighter in the sidelobe region than in the main lobe exacerbates the leakage at all k modes, representing a worst case scenario amongst the foreground cases simulated in this work. This result is in agreement (at least at a qualitative level) with observed ripples in closure spectra that are present when the Galactic plane appears at low elevation (Carilli et al., 2020). Wide-field, diffuse foreground emission is known to be a relevant source of power leakage outside the wedge in standard power spectrum **measurements** too (e.g., Thyagarajan et al., 2015, 2016; Kern et al., 2020a);

- power spectra from triads that include mutual coupling beams do not significantly vary at any k_{\parallel} mode whether they include beams that are different for different baselines or not. In other words, the main source of foreground leakage at high k_{\parallel} modes - compared to the unperturbed beam case - is not the beam-to-beam variation for each baseline: power spectra that have essentially any combination of mutual coupling beams (including the same type for the triad) yield power spectra that are similar to each other, with a maximum difference of $\sim 10^2 \text{ mK}^2 (h^{-1} \text{ Mpc})^3$;
- the presence of strong foreground emission in the main lobe of the primary beam helps reducing the foreground leakage at $|k_{\parallel}| > 0.2 h \text{ Mpc}^{-1}$ in case of mutual coupling beams, although more complete simulations which include the 21 cm signal are needed to prove that this could be an effective observing strategy;

On the use of temporal filtering for mitigating galactic synchrotron calibration bias in 21 cm reionization observations

Published as: Ntsikelelo Charles, Nicholas Kern, Gianni Bernardi, Landman Bester, Oleg Smirnov, Nicolas Fagnoni, Eloy de Lera Acedo, *Monthly Notices of the Royal Astronomical Society*, Volume 522, Issue 1, June 2023, Pages 1009–1021

3.1 Abstract

Precision antenna calibration is required for mitigating the impact of foreground contamination in 21 cm cosmological radio surveys. One widely studied source of error is the effect of missing point sources in the calibration sky model; however, poorly understood diffuse galactic emission also creates a calibration bias that can complicate the clean separation of foregrounds from the 21 cm signal. In this work, we present a technique for suppressing this bias with temporal filtering of radio interferometric visibilities observed in a drift-scan mode. We demonstrate this

technique on mock simulations of the Hydrogen Epoch of Reionization Array (HERA) experiment. Inspecting the recovered calibration solutions, we find that our technique reduces spurious errors by over an order of magnitude. This improved accuracy approaches the required accuracy needed to make a fiducial detection of the 21 cm signal with HERA, but is dependent on a number of external factors that we discuss. We also explore different types of temporal filtering techniques and discuss their relative performance and tradeoffs.

3.2 Introduction

The Hydrogen Epoch of Reionization Array (HERA, [DeBoer et al., 2017](#)) is a newly built low-frequency radio interferometric array primarily designed to probe the Epoch of Reionization (EoR). The EoR period is one of the very difficult areas in cosmology to probe, and as of this writing, it is one of the least known areas experimentally. There are several other low-frequency arrays aimed at probing the EoR; these include the Low-Frequency Array (LOFAR, [van Haarlem et al. \(2013a\)](#)), Murchison Widefield Array (MWA, [Tingay et al. \(2013\)](#)), Giant Metrewave Radio Telescope (GMRT) EoR experiment ([Paciga et al., 2013](#)), and the Square Kilometer Array (SKA, [Koopmans et al., 2015](#)).

The 21 cm emission from the hyperfine transition of neutral hydrogen is one of the best observational probes of the EoR period. EoR experiments all face the common problems of bright foregrounds from Galaxy and extragalactic sources, which can be up to five orders of magnitude brighter than the 21 cm emission (e.g., [Jelić et al., 2008](#); [Bernardi et al., 2010](#); [Thyagarajan et al., 2016](#)). Fortunately, foreground emission is spectrally smooth compared to 21 cm emission, and therefore the 21 cm signal can be separated from foreground emission, for example, using the *delay spectrum* approach ([Parsons et al., 2012](#); [Thyagarajan et al., 2013](#); [Liu et al., 2014a](#)). The delay spectrum uses interferometric delays to isolate the power spectrum of the 21 cm signal from the foregrounds. Due to the spectral nature of the 21 cm signal, its power spectrum has emission at all Fourier k wavevector modes, whereas the foreground power spectra is limited to a wedge-like region in a two-dimensional $k_{\parallel} - k_{\perp}$ -space, referred to as the foreground wedge

(Datta et al., 2010; Parsons et al., 2012; Trott et al., 2012; Vedantham et al., 2012; Thyagarajan et al., 2013; Liu et al., 2014a,b; Morales et al., 2019; Mevius et al., 2022). This separation, in principle, leads to a region of data at high k_{\parallel} where the EoR 21 cm signal should cleanly separate from foreground contamination, what is referred to as the “EoR window.”

However, before power spectrum analysis is performed, interferometric data require calibration in order to correct for instrumental effects (see, for example, Smirnov, 2011, for a review on interferometric calibration). Calibration often includes the use of sky models, and typically, sky models are built from source catalogues. Source catalogues (e.g. Hurley-Walker et al., 2017; Shimwell et al., 2017; Hurley-Walker et al., 2022), however, do not fully characterise the sky emission due to the inevitable limited sensitivity and angular resolution (Grobler et al., 2014; Wijnholds et al., 2016; Trott & Wayth, 2016; Procopio et al., 2017; Barry et al., 2021; Gehlot et al., 2021). In addition, sky models often exclude diffuse Galactic synchrotron emission and are, therefore, not a full accounting of the low-frequency sky brightness distribution. This “incompleteness” in our understanding of the true sky leads to errors when performing calibration, which can compromise our ability to invert the telescope response and prohibit a clean measurement of the 21 cm cosmological signal (Barry et al., 2016; Patil et al., 2016; Ewall-Wice et al., 2017; Morales et al., 2019; Byrne et al., 2019; Kern et al., 2020c; Dillon et al., 2020; Byrne et al., 2021).

The instrumental response of the antenna can also add frequency structure on top of smooth foreground spectra. This is particularly apparent when bright foreground emissions appear in the sidelobes of the primary beam pattern (Choudhuri et al., 2021; Charles et al., 2022; Trott, 2022; Gan et al., 2022). Sidelobes of primary beams are typically non-smooth as a function of frequency and also not well-characterised by simulations or lab measurements, particularly in the presence of mutual coupling effects due to closely packed array configurations (e.g., Fagnoni et al., 2020; Borg et al., 2020; Bolli et al., 2022a,b). This inevitably leads to calibration errors, which, in turn, can also cause leakage of foreground power.

There are attempts to reduce the reliance on sky models in calibration, such as in redundant calibration (Liu et al., 2010; Dillon et al., 2020). However, redundant calibration still requires a sky model to constrain degenerate parameters that cannot be solved for using the redundant calibration equa-

tions (Zhang et al., 2016; Dillon et al., 2018; Kern et al., 2020c). Redundant calibration also assumes array redundancy, where multiple baselines that have the same orientation and length measure the same Fourier mode of sky signal. However, with actual observations, including the HERA array, baselines inevitably deviate from redundancy. This non-redundancy can be caused by multiple factors, such as antenna placement errors, antenna-to-antenna primary beam variations, and instrumental coupling (e.g. Orosz et al., 2019; Fagnoni et al., 2020; Kern et al., 2020c; Joseph et al., 2020; Choudhuri et al., 2021; Josaitis et al., 2022). Non-redundancy eventually leads to foreground power leaking into the EoR window (Joseph et al., 2018; Dillon et al., 2020).

Due to HERA’s compact array layout, its observations are more sensitive to large angular structures in the sky, such as the diffuse galactic synchrotron emission (Haslam et al., 1982). Although diffuse emission models have improved over the years (e.g. through the Global Sky Model - GSM); de Oliveira-Costa et al., 2008; Zheng et al., 2017), point source catalogues remain our most well-understood calibration models, particularly with the advent of low-frequency all-sky surveys (Hurley-Walker et al., 2017; Riseley et al., 2018; Shimwell et al., 2022). Thus, one of the limiting factors in HERA calibration is our relatively poor understanding of the diffuse galactic foreground emission (which we will refer to interchangeably as the GSM).

In this work, we aim to improve HERA calibration by filtering out diffuse emission in the data before running calibration in order to mitigate its impact on the recovered antenna gain solutions. This idea dates back to the application of fringe rate filters for calibration in Parsons & Backer (2009) and builds upon existing HERA calibration methodologies (Offringa et al., 2012; Kern et al., 2020c; Dillon et al., 2020). We look at two specific kinds of filters: a simple high-pass filter (referred to as a “baseline-independent notch filter”) and a baseline-dependent filter that upweights sky emission coming from within the main field-of-view of the primary beam (referred to as a “baseline-dependent main-lobe filter”). We assess the calibration improvement by computing the reduced chi-squared after calibration and by further examining the structure of the recovered gains in Fourier space to assess the amount of residual spectral structure. Lastly, we show how our technique helps mitigate the foreground leakage in the EoR window.

The paper is organized as follows: [section 3.3](#) gives an overview of the calibration problem, [section 3.4](#) describes our simulations, [section 3.5](#) presents an overview of fringe rate filtering and

an analysis of our results, and in [section 3.6](#) we present our conclusions.

3.3 Calibration overview

3.3.1 Sky-model based calibration

When observing the sky with a radio interferometer, the electromagnetic waves from celestial sources are measured by two antennas **and** are correlated, forming interferometric visibilities. The measured complex-valued visibilities are related back to the intrinsic sky emission and the instrumental response of the antennas via the *radio interferometer measurement equation* (RIME; [Hamaker et al., 1996](#); [Smirnov, 2011](#)). We can break the RIME into two parts, one that is dependent on the position of the source emission on the sky (direction-dependent) and another that is independent of sky position (direction-independent). The direction-dependent part of the RIME is written as

$$V_{ij}(\nu) = \iint |A(\mathbf{s}, \nu)|^2 I(\mathbf{s}, \nu) e^{-2\pi i \frac{\nu}{c} \mathbf{b}_{ij} \cdot \mathbf{s}} \frac{dldm}{n(\mathbf{s})}, \quad (3.1)$$

where V_{ij} is the visibility formed between antenna i and j , $A(\mathbf{s}, \nu)$ is the average antenna primary beam response, an example of the only direction-dependent corruption that will be considered in this paper. $I(\mathbf{s}, \nu)$ is the brightness distribution of the sky, ν is the frequency of the incoming electromagnetic wave, \mathbf{b}_{ij} is a baseline vector connecting antenna i and j , $\mathbf{s} = [l, m, n]^T$ is the position vector on the celestial sphere with an origin centred at the target field (phase centre), and (l, m, n) are the direction cosines, with $n = \sqrt{1 - l^2 - m^2}$. Any practical telescope also imparts direction-independent corruptions to the visibilities, which can be modelled using antenna-based gain terms, for example, the bandpass of the antenna front-end response. Thus, the corrupted visibilities are related to intrinsic visibilities as

$$V_{ij}^c = g_i V_{ij} g_j^* + n_{ij}, \quad (3.2)$$

where V_{ij}^c are the corrupted (or measured) visibilities, g_i and g_j are the complex gain terms of antenna i and j , and n_{ij} is any complex thermal noise generated by the telescope.

Antenna gain calibration is the process of estimating the direction-independent gain terms (g_i) from the data and then correcting the data to remove their effect. This is done by constructing a model of the true visibilities and then minimising a χ^2 statistic:

$$\chi^2 = \sum_{i,j} \frac{|V_{ij}^c - g_i g_j^* V_{ij}^m|^2}{\sigma_{ij}^2}, \quad (3.3)$$

where V_{ij}^m are the constructed model visibilities, and σ_{ij}^2 is the variance of the visibility, and the sum runs over all unique antenna pairs.

The chi-square minimisation depends on the accurate modelling of the intrinsic sky visibilities V_{ij} . It is crucial, therefore, that the sky model matches the intrinsic sky brightness as closely as possible. This is also referred to as ‘‘completeness’’ of the sky model: if our sky model does not contain certain components of the intrinsic sky brightness distribution, then it would be ‘‘incomplete’’. In practice, we never know the true sky distribution exactly, so all sky models are intrinsically incomplete, but the severity of this incompleteness can have different impacts on the calibration, which is the focus of this paper.

Note that self-calibration can, in principle, correct for some incompleteness effects, specifically due to phase variations in the gains and misestimated point source properties (Bhatnagar et al., 2008; Eastwood et al., 2018). However, self-calibration has shown success where we either have good priors on the nature of the underlying emission (e.g. they are point sources) or fairly accurate first guesses of the underlying sky brightness distribution. It is, therefore, not yet a reliable solution for the problems addressed in this work.

The problem of poorly modelled diffuse emission is not new, and some techniques for mitigating its impact exist. The most common and simple approach is to exclude baselines that are most heavily contaminated by this emission in the process of calibration. Short baselines are most sensitive to large angular scales and are, therefore, routinely thrown out prior to calibration against a sky model comprised mostly of point sources. However, in the case of HERA, it is not possible to completely mitigate this as even longer baselines are still sensitive to diffuse emission (e.g., Thyagarajan et al., 2015; Trott, 2022; Mevius et al., 2022).

Other attempts to mitigate spurious spectral structure in the calibration solutions rely on smoothing the gains across frequency (Barry et al., 2016; Kern et al., 2020c) or enforcing smooth-

ness in the fitting process (Yatawatta, 2015; Beardsley et al., 2016), but this does not fundamentally solve the calibration problem. Smoothing removes specific Fourier modes from the gains entirely, regardless of whether those terms held true calibration solutions or were the result of model incompleteness errors. In other words, some Fourier modes are fundamentally contaminated by model incompleteness errors, preventing us from applying calibration solutions at those scales. This is problematic if we need to invert the response of the telescope to high precision, which is generally needed for 21 cm measurements.

3.3.2 Redundant calibration

Redundant calibration does not need a sky model to solve for a subset of the antenna gains, and obtains them by exploiting the internal redundancy of an interferometric array. Redundant baselines have the same length and orientation and, therefore, measure the same Fourier mode of the sky brightness distribution. For example, assume we have a single baseline type, A , uniquely identified via its baseline vector \mathbf{b}_A . For all antenna pairs i, j that share this baseline vector, the calibration equation Equation 3.2 from before now becomes

$$V_{ij}^c = g_i V_A g_j^* + n_{ij}, \quad (3.4)$$

where V_A is now a parameter of the model, called the “redundant model visibility”. This is repeated for all unique baseline types, eventually building up an overconstrained system of equations, which can be solved via a χ^2 minimization (e.g., Liu et al., 2010):

$$\chi^2 = \sum_{i,j} \frac{|V_{ij}^c - g_i g_j^* V_A|^2}{\sigma_{ij}^2}, \quad (3.5)$$

where V_A is the corresponding redundant model visibility for the \mathbf{b}_{ij} baseline vector. Note that in this paper, the χ^2 is both a function of frequency and time. The gain solutions obtained from the system of linear equations are not unique; there are at least four degenerate parameters that need to be solved for after redundant calibration: the overall gain amplitude, the model visibility amplitude and the phase gradient across the array in the east-west and north-south orientations (Liu et al., 2010; Zheng et al., 2014; Dillon et al., 2018). These parameters are, in principle,

a function of polarization, frequency and time and require a sky model to be constrained. Thus redundant calibration itself is not entirely independent of a sky model. In this study, we will refer to this step in calibration as absolute calibration to differentiate this from sky calibration or the typical sky-based calibration discussed previously.

In the case of perfect calibration, where the sky model is complete, model visibilities only differ from the data by noise. The expectation value of chi-square $\langle \chi^2 \rangle$ is thus two times the degrees of freedom (DoF) for complex data (Dillon et al., 2020). The number of DoF, in general, is given by the difference between the number of data points and the number of fitted parameters. If we consider a single polarization and the case of sky-based calibration, the DoF are given by the difference between the number of visibilities and the number of antennas:

$$\text{DoF} = \frac{N(N-1)}{2} - N = N_{\text{bl}} - N, \quad (3.6)$$

where N is the number of antennas and N_{bl} is the number of baselines given N antennas. If we histogram the χ^2/DoF quantity, it is expected to follow a theoretical chi-square distribution (Dillon et al., 2020). In the case of redundant calibration, the DoF is given by:

$$\text{DoF} = N_{\text{bl}} - N_{\text{ubl}} - N + 2, \quad (3.7)$$

where N_{ubl} is the number of unique baselines that have a single sky model visibility V_A , i.e. the number of redundant baseline groups (Dillon et al., 2020).

In this work, we do not simulate any non-redundancies; thus, we expect redundant calibration to work perfectly, regardless of the sky model incompleteness (Dillon et al., 2018). The required absolute calibration step that follows redundant calibration will, conversely, introduce spurious gain errors due to missing components in the sky model. In the remainder of the paper, we will consider χ^2 statistics that include the full calibration solutions (both redundant and absolute calibration together), unlike Dillon et al. (2020), who only plot the reduced χ^2 for redundant calibration. Our calibration pipeline is a simplified version of the `hera_cal`¹ package used for HERA observations (HERA Collaboration et al., 2022).

¹github.com/HERA-Team/hera_cal

3.4 Foreground simulations

Here we describe our simulated HERA observations. We adopt a HERA-like array layout consisting of 91 antennas in a compact hexagonal configuration (Figure 3.1), with a dish-to-dish spacing of 14.6 meters (DeBoer et al., 2017). We use an electromagnetic simulation of the farfield HERA Phase I dish and dipole feed response (not including the effects of cross-coupling) to model the antenna primary beam (Fagnoni et al., 2020). In order to simulate a quasi-realistic case, we assume the following representation of the gain g for the j th antenna:

$$g_j(\nu) = A_j(\nu)e^{i\phi_j(\nu)}, \quad (3.8)$$

where the amplitude A follows a frequency power law

$$A_j(\nu) = A_j \left(\frac{\nu}{150 \text{ MHz}} \right)^{b_j}, \quad (3.9)$$

and A_j and b_j are drawn for each antenna from a Gaussian distribution $\mathcal{N} \sim (\bar{x}_A = 0.30, \sigma_A = 0.001)$ and $\mathcal{N} \sim (\bar{x}_b = -2.6, \sigma_b = 0.2)$, respectively. The phase $\phi_i(\nu)$ is modeled as:

$$\phi_j(\nu) = \sin(w_a \nu) + \cos(w_b \nu), \quad (3.10)$$

where w_a and w_b are drawn for each antenna from a Gaussian distribution $\mathcal{N} \sim (\bar{x}_w = 0.005, \sigma_w = 0.0005)$. The mean values of the parameters A and b , as well as their standard deviation σ_A and σ_b , are informed by the gain solutions from actual HERA observations (Kern et al., 2020c). The mean phase variation \bar{w} is based on single antenna phase dependency of actual HERA gains where the cable delay and geometric phase offset have been removed, and the variation of the mean phase between antenna stations, i.e. σ_w , is chosen to be within 10%. This exact choice does not strongly affect the results discussed below, but is a reasonable starting point for performing calibration after an initial delay calibration step.

We simulate HERA visibilities for three distinct sky models: i. shallow GLEAM + A-Team model; ii. deep GLEAM + A-Team model, and iii. GSM diffuse model. The GLEAM catalogue is a point source catalogue covering the HERA field of view considered in this work (Hurley-Walker et al., 2017). The shallow GLEAM catalogue includes GLEAM sources only down to a




Figure 3.1: Simulated HERA-like array layout with 91 antennas and 14.6 meter spacing between antennas. X and Y axes indicate distances aligned along East-West and North-South directions, respectively.




Figure 3.2: Simulated HERA visibilities of a 45 m baseline that includes the sky model ii., GLEAM sources and A-Team sources (top row) and only diffuse Galactic emission from the GSM (bottom row, see text for details). The left column shows the visibility amplitude as a function of LST and frequency, whereas the middle column amplitude of the fringe-rate visibilities as a function of fringe rate and time and the right column shows the fringe-rate, delay-transformed visibilities as a function of fringe rate and delay. Point source emission is more tightly confined to low delays and non-zero fringe rates. The diffuse emission has a significant build-up of power at near-zero fringe rates due to the pitchfork effect (Thyagarajan et al., 2016). The boundaries of the notch (white) and main lobe (orange) filters described in subsection 3.5.3 are shown as dashed lines.

flux density limit of 100 mJy (at 180 MHz), while the deep GLEAM catalogue goes down to a flux density limit of 25 mJy. In both cases, we also add models for bright, extended radio galaxies that we label as the A-Team (Fornax A, Hydra A, Pictor A, Hercules A, Virgo A, Cygnus A, Cassiopeia A; [McKinley et al., 2015](#); [Byrne et al., 2022](#)). Lastly, we also include the diffuse emission from the GSM ([Zheng et al., 2017](#)).

For each of the three sky models, we use `healvis` ([Lanman & Kern, 2019](#)) to simulate HERA visibilities using the array layout in [Figure 3.1](#) and numerically evaluate [Equation 3.1](#), assuming the entire observable hemisphere as the field of view. For the primary beam response, we use an electromagnetic simulation of the Phase I HERA dipole antenna response ([Fagnoni et al., 2020](#)). Our simulations span 150–180 MHz in frequency with a 97 kHz channel resolution, and are made for local sidereal times between $-2^{\text{h}} < \text{LST} < 7^{\text{h}}$ range at a 60 s cadence, corresponding to the time when the Galactic center falls partially in the far sidelobes.

We form two separate sets of visibilities, i.e. a mock calibration model that includes only the shallow GLEAM and A-Team components and a “true” mock dataset that includes all three sky models. Finally, we apply the simulated gains to the visibilities corresponding to the true dataset and added Gaussian thermal noise at a level comparable to HERA Phase I measurements, with a signal-to-noise ratio of ~ 100 per visibility ([HERA Collaboration et al., 2022](#)).

[Figure 3.2](#) shows an example of the deep GLEAM + ATeam (top row) and the GSM (bottom row) visibility data product for a 45 meter baseline. It highlights the different spectral and temporal structure of the two sky components in the native telescope measurement space (frequency vs time, left panels), as well as in a 1D and 2D Fourier space representation that we discuss in [section 3.5](#). The dashed lines show the extent of the two kinds of filters that we apply to the data, defined in detail in [section 3.5](#). This clearly illustrates how the two models manifest in the different spaces and hints at how we may be able to isolate them, at least partially, by applying suitable filters.

3.5 Fringe Rate Filters for Improving Calibration

Here we demonstrate how time-based filtering (what we call fringe-rate filtering) can improve the data-to-model alignment and thus partially mitigate some of the erroneous spectral features caused specifically by poorly modelled diffuse emission². First, we discuss the Fourier representations of the data and then discuss calibrating our mock HERA data with and without the two different types of fringe rate filters.

3.5.1 Fourier representations

The native measurement space of the interferometric visibilities is observing frequency (ν) and local sidereal time (LST). For telescopes that operate in drift-scan mode, the local sidereal time is effectively just the right ascension directly overhead at any given time. We refer to the Fourier dual of frequency as the delay domain, formed by taking a Fourier transform of the visibilities across the frequency axis:

$$\tilde{V}(\tau) = \int V(\nu) e^{-2\pi i \tau \nu} d\nu, \quad (3.11)$$

where τ is a delay (in seconds) and $\tilde{V}(\tau)$ is the Fourier pair of the frequency; similarly, we can also define delay-transformed antenna gains to be the Fourier transform of the antenna gains along the frequency axis:

$$\tilde{g}(\tau) = \int g(\nu) e^{-2\pi i \tau \nu} d\nu. \quad (3.12)$$

We refer to the Fourier dual of time (or LST) as the *fringe rate* and the fringe-rate visibility as the Fourier transform of the visibility along the time axis:

$$\tilde{V}(f) = \int V(t) e^{-2\pi i f t} dt, \quad (3.13)$$

where f is the fringe rate (in units of Hz)³.

²This technique can also be extended to suppress similar spurious features in the recovered gains created by non-redundancy in the data (e.g. from crosstalk or per-antenna beam errors), which we defer to future work.

³Referred to as delay-rate by [Parsons & Backer \(2009\)](#), but subsequently called fringe rate in works thereafter.




Figure 3.3: Recovered gain solutions after calibrating with an incomplete sky model. Left: The amplitude ratio g/g_c between the gain g obtained from the three simulated cases and the gain g_c obtained from calibration when a complete sky model is assumed. Both terms are the average gain across all the antennas at $LST = 1.2^h$. Red is the case of a complete sky model (which results in a ratio of one), blue is for the incomplete sky model, and green is for the incomplete sky model with a 40 m minimum baseline cut. Right: Here we show the amplitude of the gains in delay space after averaging over all antennas. The excess gain structure as a function of frequency due to the sky model incompleteness is evident even when a baseline cut is adopted (the bumps for $\tau > 200$ ns). This excess “shoulder” is particularly visible in delay space and its profile is fairly similar in both cases where the sky model is incomplete (with or without a baseline cut). For delays $\tau > 200$ ns this spurious structure is at its peak over an order of magnitude above the noise floor from the complete-model-derived gains (red).

The Fourier space of the visibilities (delay and fringe rate) is useful for separating different signals in the data. Delay space separates signals based on their incident angle from the pointing of the telescope, and is the basis for the delay spectrum foreground avoidance technique (Parsons et al., 2012; Liu et al., 2014a; HERA Collaboration et al., 2022). Similarly, the fringe rate basis (for drift-scan observations) also separates signals based on their relative motion through the fixed interferometric fringes, acting as another form of separation of signals on the sky (Parsons & Backer, 2009; Parsons et al., 2016).

A fringe rate filter can, therefore, be used to suppress sky signals that appear at specific fringe rate values, for instance, outside the primary beam field of view (Figure 3.2). Parsons et al. (2016) indeed describe how sky signals within the field of view appear at specific, baseline-dependent fringe rates. In contrast, sources outside the field of view appear to have a fringe rate that is closer to zero (i.e., a near-constant time response). In other words, fringe rate filtering is equivalent to “sculpting” the primary beam. We refer the reader to Parsons et al. (2016) for a more detailed description of fringe rate space for drift-scan observations.

Fourier representations have been used widely to separate signals and systematics in 21 cm observations (e.g. Kolopanis et al., 2019; Kern et al., 2020b; Josaitis et al., 2022); however, these have generally been applied to pre-calibrated data. Here we explore using these basis to separate signals before calibration.

3.5.2 Calibration Without Filtering

To gain an intuition for how calibration would work without applying any visibility filtering, we ran three different kinds of simulated observation cases. First, we used a complete sky model, i.e. the same components used for the simulations, the deep GLEAM catalogue, the A-team sources and the GSM to calibrate raw data. In the second case, we calibrated by using an incomplete sky model, i.e. including only the shallow GLEAM catalogue and the A-team sources. In the third case, we use the same sky model as the second case, but with the addition of a baseline cut at 40 m, i.e. all the baselines shorter than 40 m were excluded before calibration. This is a commonly used approach in order to mitigate the contribution of diffuse emission that appears

mostly on short baselines.

The recovered gains in each case are shown in [Figure 3.3](#). The gains recovered in the complete sky model case represent the true simulated gains, as they only differed from the true gains by noise. Notice this is not the case for the incomplete sky model case. There is an excess frequency structure that is caused by the unmodelled foreground emission, coupled with the sidelobe primary beam response. During calibration, gains absorb a portion of this unmodelled term, resulting in increased frequency structure. This indeed has been shown in literature, in both simulations and observations (e.g., [Wijnholds et al., 2016](#); [Barry et al., 2016](#); [Ewall-Wice et al., 2017](#); [Byrne et al., 2019](#); [Kern et al., 2020c](#)). Notice that there is no appreciable difference in the gains whether a baseline cut is adopted or not. The frequency structure caused by an incomplete sky model introduces excess power in the gains at delays in the range $100 < \tau < 400$ ns, right where the sensitivity to the EoR signal is the highest.

To demonstrate that our calibration pipeline works as intended in an ideal case, we plot the reduced χ^2 obtained from the visibilities calibrated with the recovered gains in the complete sky model case ([Figure 3.4](#)). The reduced χ^2 of the calibrated data shows good agreement with a theoretical unit- χ^2 distribution (dashed line), showing that our calibration pipeline is both unbiased and recovers the gains down to the thermal noise floor in the ideal circumstance of a complete sky model.

The bias introduced by the sky model incompleteness, with a baseline cut is evident in the reduced χ^2 distribution ([Figure 3.5](#)), which shows a strong deviation from the theoretical one. The mean of the reduced χ^2 distribution is biased toward high values, with a mean of $\chi^2/\text{DoF} \sim 19$, and a tail extending up to $\chi^2/\text{DoF} \sim 100$. [Figure 3.5](#) also shows the χ^2 distribution as a function of time and frequency spanning LSTs of 1.3 – 3.0 h (right panel). The reduced χ^2 varies noticeably as a function of frequency and LST; in particular, its bias is higher at low frequencies where the diffuse emission is significantly brighter. The reduced χ^2 is less biased around LST $\sim 2.8^{\text{h}}$, where bright, compact sources dominate the sky emission yielding to a more complete sky model.




Figure 3.4: Histogram of the reduced χ^2 (blue), i.e. χ^2/DoF , obtained when a complete sky model is used for calibration. The dashed line is the theoretical reduced χ^2 distribution with $k = 2 \times DoF$, i.e. the chi-square distribution degrees of freedom, and $x = \chi^2/DoF$. This is a simple demonstration that our calibration pipeline works as expected in the limit of a complete sky model. In other words, our gains are accurate down to the expected thermal noise.

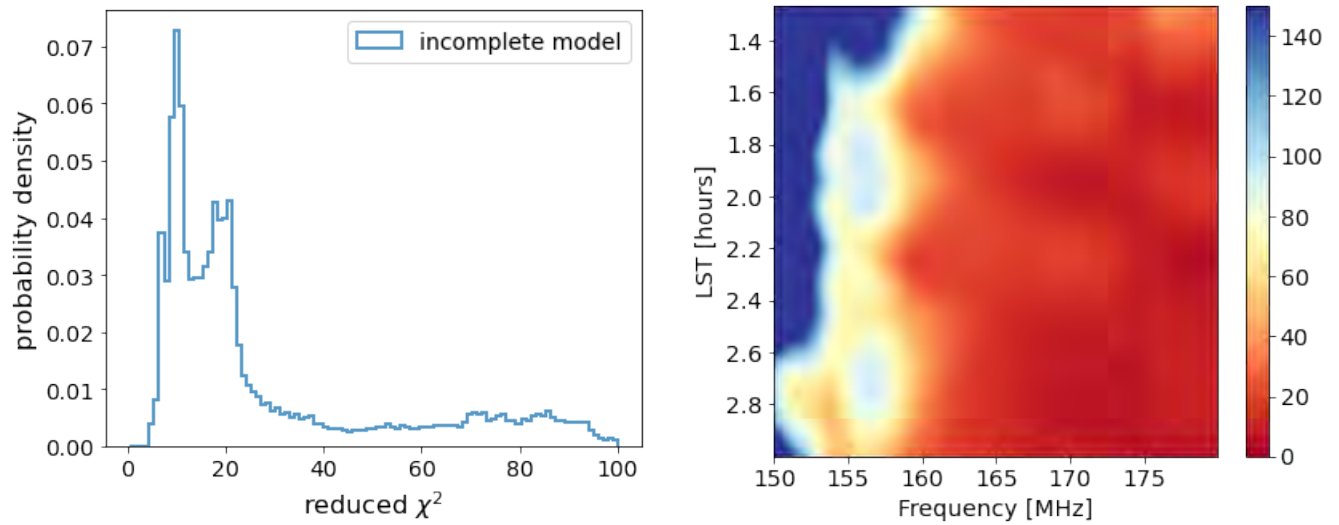


Figure 3.5: The resultant reduced chi-square after calibrating with an incomplete sky model and including a 40-meter minimum baseline cut. Left: Histogram of the reduced chi-square with baseline (blue). Right: The reduced chi-square as a function of time and frequency. The missing components in our sky model leads to a clear bias in the reduced chi-square after calibration. The observed chi-square is higher than what we find in actual data (e.g., [Dillon et al., 2020](#)), but this only because the simulation has a higher signal-to-noise ratio compared to actual data.

3.5.3 Fringe Rate Filtering Before Calibration

Here we propose to fringe rate filter the data before calibration as a way to improve the data-to-model match and thus mitigate some of the effects seen in [subsection 3.5.2](#). Note that fringe rate filtering is applied to *both* the raw data and the sky model visibilities in the same manner for each of the different filters described below. Also, note that the derived gains are applied to the unfiltered visibilities. In other words, our fringe rate filtering is only used en route to derive gains that are more robust to an incomplete sky model and are not used downstream in our analysis (thus mitigating concerns of a possible cosmological signal loss).

Note that, for the two filters described below, we apply them to the data using the DAYENU filtering formalism described in [Ewall-Wice et al. \(2020\)](#), which relies on the Discrete Prolate Spheroidal Sequences (DPSS; [Slepian, 1978](#)).

3.5.3.1 Notch Filters

We first consider a symmetric baseline-independent notch filter $F(f)$ centred at $f = 0$ mHz fringe rate, i.e. a high-pass filter, defined as

$$F(f) = \begin{cases} 10^{-8}, & |f| \leq f_{\max} \\ 1, & |f| > f_{\max} \end{cases}, \quad (3.14)$$

with $f_{\max} = [0.25, 0.40, 0.60]$ mHz, respectively. We refer to the three filters as f_{25} , f_{40} and f_{60} , respectively. The bounds f_{\max} of the f_{25} filter are the white dashed lines shown in [Figure 3.2](#) and indicate the breadth of the emission suppressed in fringe rate space. [Figure 3.6](#) shows the impact of baseline-independent notch filters on foreground emission in time and frequency space. As a result of filtering out some of the diffuse foreground component, the visibility spectra of the filtered data more closely match that of compact source sky model.

3.5.3.2 Main Lobe Filters

We also considered a second type of filter, which we refer to as a ‘‘main lobe’’ filter, because it aims to suppress emission from outside the primary beam field-of-view. In contrast to the

baseline-independent notch filter, which only filters out emission in a region near $f \sim 0$ mHz, the baseline-dependent main-lobe filter suppresses the signal everywhere *except* near the peak emission of the sky model in fringe rate space (Figure 3.2).

In principle, the baseline-dependent main-lobe filter should be a frequency-dependent filter, but because we operate over a relatively small bandwidth (30 MHz), we approximate it as frequency-independent, with little impact on our final results. Note that the response of the filter within the pre-defined fringe rate bounds is uniform, such that it can be thought of as a top-hat filter (similar to the baseline-independent notch filter). The filter presented here is similar to the filter described in Josaitis et al. (2022) for understanding mutual coupling.

The bounds of the baseline-dependent main-lobe filter are determined by its centre f_0 and its half-width f_w . These parameters are determined, for each baseline, by fitting a Gaussian profile $G(f)$ to the sky model visibilities in fringe rate space:

$$G(f) = A e^{-\frac{(f-f_0)^2}{2\sigma^2}}, \quad (3.15)$$

where A is the amplitude of the Gaussian, σ is its standard deviation, and f_0 is its mean. After the fit, we set the main lobe filter centre to be f_0 and its half-width to be $f_w = 2\sigma$, such that the full width of the main lobe filter is $4\times$ the fitted Gaussian's standard deviation. Figure 3.7 shows the foreground visibilities for a number of selected baselines after applying the main lobe filter. The filter not only removes emission near $f \sim 0$ mHz, but also removes all emission that is not centrally peaked at positive fringe rates. The impact of the filter can also be seen clearly in image space (Figure 3.8), where patchy structure coming from Galactic emission in the primary beam sidelobe is largely suppressed while compact source emission is retained.

3.5.3.3 Computing the Power Spectrum

In addition to the reduced χ^2 and the delay transform of the gains, we also used the power spectrum as a metric to assess the impact of fringe rate filters on calibration. We compute the per-baseline power spectrum of the visibilities, $P(\tau, b)$, following the delay approximation (Parsons et al., 2012). This simple formalism relates the Fourier transformed visibilities of each baseline,




Figure 3.6: Simulated visibility amplitude as a function of frequency and LST for a 45-meter HERA baseline. The left panel includes the deep GLEAM model and the GSM. Same is shown at the centre, but after applying a 0.25 MHz notch fringe-rate filter. The right panel shows the real part of the visibility at an LST = 2^h, for the unfiltered (dashed green line) and filtered case (dashed red line), alongside the shallow GLEAM model (black) used for calibration. After fringe-rate filtering, visibilities are in much better agreement with the model, with the large frequency ripples significantly suppressed.

$\tilde{V}_b(\tau)$, directly to the 21 cm power spectrum

$$P(\tau, b) = |\tilde{V}_b(\tau)|^2 \left(\frac{\lambda^2}{2k_B} \right)^2 \left(\frac{D_c^2 \Delta D_c}{B_{\text{eff}}} \right) \left(\frac{1}{\Omega B_{\text{eff}}} \right), \quad (3.16)$$

where λ is the center wavelength of observing bandwidth, k_B is the Boltzman constant, B_{eff} is the effective bandwidth, D_c is the comoving distance at the redshift of our measurement, ΔD_c is comoving distance parallel to line of sight, Ω is the field of view solid angle, \tilde{V}_b is the Fourier transformed visibility, and b is the visibility baseline length. We can map τ to the line-of-sight cosmological Fourier wavevector k_{\parallel} using the relation (Thyagarajan et al., 2013):

$$k_{\parallel} = \frac{2\pi\nu_{21}H_0E(z)}{c(1+z)^2} \tau, \quad (3.17)$$

where $\nu_{21} = 1420$ MHz, H_0 is the Hubble constant, $E(z) = [\Omega_m(1+z)^3 + \Omega_\Lambda]^{1/2}$, Ω_m is the normalized matter content and Ω_Λ is the normalized dark energy content. Note that while the baseline length maps to the perpendicular Fourier wavevector (k_{\perp}) and thus contributes to the overall $|\mathbf{k}|$ magnitude, for short baselines its contribution is negligible and we will drop it for simplicity and simply quote $P(k_{\parallel})$ hereafter.




Figure 3.7: Model visibility amplitude (GSM + GLEAM + A Team) as a function of fringe rate for a few different baselines in the HERA array with a projected East-West component greater than 30 meters (Figure 3.1). Solid (dashed) lines indicate visibilities before (after) applying the baseline-dependent main-lobe filter, showing that the filter retains only the emission at positive fringe-rate values associated with the primary beam. The legend marks the East and North extent of each baseline vector in meters.




Figure 3.8: Simulated dirty image that includes the GLEAM catalogue, the A-team sources and the GSM before (left) and after (right) applying the baseline-dependent main-lobe filter. The simulated observation is 10 minute long and centred at $\alpha = 0^\circ$. The apparent absence of sky emission outside the central $\sim 10^\circ$ area is due to the primary beam attenuation. The patchy structure in the left panel is the sidelobe structure from Galactic emission located outside the primary beam, which is largely removed after filtering.

3.5.3.4 Summary

The two filters described above both aim to limit the influence of the poorly modelled diffuse emission (GSM) component in the data. We can readily see that a significant amount of power from the GSM is found at and around $f \sim 0$ mHz, although there is still a non-zero amount of power at positive fringe rates. The baseline-independent notch filter acts to remove emission at zero and near-zero fringe rates, thus improving the prominence of the point sources in the visibilities. The baseline-dependent main-lobe filter, on the other hand, filters out everything outside the main lobe, i.e. a larger amount of signal and noise.

To gain more intuition as to why these fringe rate filters not only improve the data-to-model match but also *reduce spurious spectral structure in the recovered gains* we plot the pre-filtered and post-filtered true visibility in [Figure 3.6](#), and compare them to the incomplete model visibility used for calibration. We see that the unfiltered true dataset has large ripples (green dashed) that is not reflected in the sky model (black solid); however, after applying a fringe rate filter (red dashed), the true data is brought into closer alignment with the sky model. Specifically, the large ripples that are in the true data but absent in the sky model are greatly suppressed. We will see in the next section that this improvement in the data-to-model match translates to smaller errors in the gain solutions at the Fourier scales of the ripple.

3.5.4 Effects of fringe rate filters on calibration

Before we consider the application of fringe rate filters, we first discuss the side effects that fringe rate filters may have on visibilities. After filtering, the visibility noise is suppressed and also becomes correlated from time-to-time. Both of these will affect the computed reduced χ^2 : the noise correlation can be corrected by adjusting the degrees of freedom (DoF), whilst the signal loss by adjusting the weights σ_{ij} . In this study we will only consider adjusting the weights to correct for the reduced noise amplitude, as adjustments to the DoF do not strongly impact the results of the study.

White thermal noise is uncorrelated in the visibilities and thus occupies all Fourier modes in the data uniformly. The noise amplitude will in theory depend on the frequency and observing

time, but for our studies, we assume that the noise is both time and frequency independent. Thus we can use a single number to describe the noise in the visibilities, which is its total variance at each time and frequency pixel, σ^2 . A fringe rate filter suppresses noise depending on how many Fourier modes it removes from the data (Ali et al., 2015; Parsons et al., 2016). The resulting noise variance after filtering is related to the integral of the filter shape, $F(\cdot)$, in fringe rate space:

$$\sigma_f^2 = \sigma^2 \frac{\int F(f)\phi(f)df}{\int \phi(f)df}, \quad (3.18)$$

where $\phi(f)$ is a binary windowing function that isolates the fringe rates measured in the data (i.e. 1 if $f_{\min} < f < f_{\max}$ otherwise 0), and σ_f^2 is the pixel variance after filtering. When computing the χ^2 after filtering, we use this updated noise variance.

When we filter the raw data with a chosen fringe rate filter, we also apply the same filter to the model. This way, we ensure that the impacts of the filter are consistent between the data and the model. The chi-square for the filtered data is then given by:

$$\chi^2 = \sum_{i,j} \frac{|V_{ij}^{d,f} - g_i g_j^* V_{ij}^{m,f}|^2}{\sigma_{f,ij}^2}, \quad (3.19)$$

where $V_{ij}^{d,f}$ and $V_{ij}^{m,f}$ are filtered data and model visibilities respectively.

3.5.5 Calibration with the notch filter

We first investigate calibration after applying the baseline-independent notch filter, which suppresses emission centred at zero fringe rates. For baselines with small East-West projected length, however, emission from the main lobe appears near to zero fringe rate too (Parsons et al., 2016). For instance, the main lobe peak occurs at a fringe rate < 0.8 mHz for a baseline with an East-West projection length smaller than 14 m. After filtering, the SNR for short baselines is reduced as the filter also removes relevant foreground emission. We, therefore, excluded from calibration all baselines with E-W projection shorter than 30 m to circumvent this issue, resulting in a 46% baseline loss for our simulated HERA configuration. We will later see that the benefits of applying the filter outweigh the sensitivity loss.

Not all the LST intervals are ideal for calibration; typically, fields where compact sources are prominent, are better suited for calibration. Here we focus on the LST = 2^{h} interval, the typical




Figure 3.9: Histogram of the reduced chi-square obtained from calibration with an incomplete sky model with and without our fringe rate baseline-independent notch filters for field centred at $LST = 2^h$. Colours indicate the different filters used, i.e. f_{25} (orange), f_{40} (green), f_{60} (red) and no filter (blue). The reduced χ^2 improves significantly after the use of a baseline-independent notch filter, both in terms of its peak value being closer to one, and a shallower tail at high values.




Figure 3.10: Reduced chi-square as a function of frequency and LST computed from calibrated visibilities after filter f_{40} is applied. A much better agreement between the data and model (i.e. lower reduced χ^2) is achieved compared to the case of no filtering (Figure 3.5).




Figure 3.11: The recovered gains after calibration (similar to [Figure 3.3](#) but for a field centred at $LST=2^h$). Here we plot the gains derived from a complete model (red), an incomplete model with no filtering (cyan), and an incomplete model having applied baseline-independent notch filters ([subsection 3.5.3.1](#)) of increasing width (green, blue, yellow). The significant amount of spectral structure seen when calibrating against an incomplete model (black) is heavily suppressed after applying the baseline-independent notch filters. All filters explored seem to have effectively the same performance.




Figure 3.12: The residual of the squared visibilities in delay space as a function of baseline length. We average all baselines of equal length (similar to the commonly used power spectrum “wedge” plots in $k_{\parallel} - k_{\perp}$ space). The residual is taken with respect to the noise-free complete visibility model. We show residuals in the case where calibration assumes a complete sky model and is thus perfect, leaving only noise in the residual (left panel), an incomplete sky model (middle panel) and a f_{25} notch filtered scenario (right panel). The visibilities are taken from a single time integration at LST = 1.2^h. The white line marks the horizon limit of the baselines, which bounds the natural extent of foreground emission in the data. Note that the increased noise amplitude with increasing baseline length comes simply because there are fewer longer baselines than short baselines in the data. The right panel shows significant suppression of foreground leakage due to gain errors, as expected.




Figure 3.13: The residual delay power spectrum of the calibrated raw visibilities with respect to the complete model (noise-free) visibilities, averaged over redundant baselines and 8 time integrations centred at $LST = 2^h$. Left: A 14 m redundant baseline group, calibrated using a complete (red) and incomplete (cyan) sky model, and having applied the f_{25} notch filter (green), the f_{40} notch filter (blue), f_{60} notch filter (yellow) and the baseline-dependent main-lobe filter (black) prior to calibration. Right: Same, but for the 28 m redundant baseline group. Note that for these short baselines, the k_{\perp} contribution is negligible to the total k magnitude, such that k_{\parallel} is effectively equivalent to the total magnitude $|\mathbf{k}|$.

LST range used for HERA calibration. In this LST range, bright, compact sources dominate the sky emission, and the bias due to an incomplete sky model is less pronounced. We considered a total of eight integration times around $LST = 2^h$. **Figure 3.9** shows the chi-square obtained from the calibration with an incomplete sky model and after applying the baseline-independent notch filter (equation 3.19; the same case as **Figure 3.5**). Compared with the case of unfiltered visibilities, the χ^2 is now less biased, with a mean value of 12, 9 and 6.5 for filters f_{25} , f_{40} and f_{60} , respectively, compared to 48 for the unfiltered visibilities. The wider the filter, the larger the suppressed emission, leading to a better match between the intrinsic and sky model visibilities.

As mentioned in **subsection 3.5.3**, the reduced chi-square changes with LST. **Figure 3.10** shows the reduced chi-square obtained after filtering visibility before calibration with an incomplete sky model. The chi-square has a maximum of 50, even in regions where the diffuse emission is dominant in the $1^h < LST < 1.8^h$ range, compared to 140 obtained without filter, the improvement is at most a factor of 3 in the reduced chi-square. Another way to look at the improvement is that the filter has effectively increased the LST range that can be used for calibration. Even better, the reduced chi-square is smaller than six in the LST range used for calibration, i.e. $[2.2^h - 2.8^h]$.

Figure 3.11 shows the gains recovered from calibration using an incomplete sky model and after applying a baseline-independent notch filter. Even at the LST range used for calibration, where compact sources dominate, the sky model is still significantly incomplete. Without applying the filter, gains absorb some of the unmodelled diffuse emission on the sidelobes of the primary beam. As a result, gains acquire a pronounced frequency structure, particularly in the $100 < \tau < 500$ ns delay range. The use of a fringe rate baseline-independent notch filter suppresses much of this frequency structure, therefore reducing the excess power seen at delays greater than ~ 100 ns. A noticeable result of the filter is that gains now closely match the actual gains up to the noise floor. The most noticeable difference is now a nearly delay-constant offset due to noisier, filtered visibilities. The gains obtained by using different filters have a very similar profile in delay space, indicating that the filter with a 0.25 MHz width is already sufficient to suppress unwanted emission.

The use of a baseline-independent notch filter has a distinct impact on the residual calibrated

visibilities in Fourier space. Let's consider pointing at LST = 1.25^h, where the diffuse emission is the dominant component of the sky emission. **Figure 3.12** shows the effect of the filter in a typical wedge plot in $k_{\parallel} - k_{\perp}$, or in the delay spectrum approach equivalently τ - baseline length space. We see that residuals are just noise-like when a complete sky model is used for calibration, as expected. When an incomplete sky model is used, however, the calibrated visibilities differ from the model visibilities due to biased gain calibration solutions. As a result, there is foreground emission in the residual confined to the $-500 < \tau < 500$ ns delay range. This contamination is significantly beyond the horizon limit for baselines shorter than ~ 70 m.

Residual foreground emission is reduced when calibration is preceded by fringe filtering. The bleed of foreground emission at delays of $|\tau| \sim 300 - 500$ ns is significantly suppressed by more than one order of magnitude. There is no foreground contamination beyond the horizon limit, leaving the EoR window effectively uncontaminated. Foreground residuals remain present at small delays, well within the horizon. Thus even in the worst field, i.e. where the sky model is most incomplete, filtering proves to be very effective at mitigating foreground contamination of the EoR window. This could, in theory, broaden the LST range that could be used for calibration.

Let's now consider the LST range used for calibration; field centred at LST = 2^h. **Figure 3.13** shows the residual power spectrum at $z = 7.6$. In this field, the foreground contamination can be as high as two orders of magnitude above the noise floor in the $0.15 < k_{\parallel} < 0.2$ h Mpc⁻¹ range when calibration is carried out with no fringe rate filter and an incomplete sky model. The behaviour is fairly similar for both the 14 and 28 m baselines considered. Such contamination is concerning, as it appears at k modes where the sensitivity to the EoR is the highest.

The contamination is greatly suppressed if the baseline-independent notch filters are used, leaving the smooth unmodelled emission at $k_{\parallel} < 0.15$ h Mpc⁻¹. Overall, the foreground leakage is suppressed at $k_{\parallel} \sim 0.2$ h Mpc⁻¹ by at least two orders of magnitude, down to the noise level in 14 m baselines. The filters help suppress unsmooth foreground even on longer baselines, i.e. 28 m baselines, but only up to a factor of ten. This is somewhat expected as the larger baselines are less sensitive to diffuse emission. Notably, the performance of the different filters in suppressing foreground spectral leakage is the same, although evidently, the filters filter varying amounts of foreground emission. Thus, the key to mitigating foreground leakage is to suppress




Figure 3.14: The recovered gains after calibration (similar to [Figure 3.11](#)), now including the application of the baseline-dependent main-lobe filter. Left: The antenna-averaged gain amplitude divided by the average gains estimated with a complete sky model. Right: The antenna-averaged gains in Fourier space. We now show the recovered gains having applied the baseline-dependent main-lobe filter, both to the complete sky model (purple) and an incomplete sky model (black). The baseline-dependent main-lobe filter not only suppresses the spurious structure in the incomplete model (cyan), but also reduces the overall noise floor of the recovered gains, as evidenced by the lower plateau at high delays.




Figure 3.15: Same as [Figure 3.14](#) but for a noise-free simulation. This now allows us to probe the performance of the technique down to higher dynamic ranges. The spurious structure in the gains due to model incompleteness (cyan) is suppressed by a over an order of magnitude for $\tau > 150$ ns by the baseline-dependent main-lobe filter, which also shows slightly better performance than the baseline-independent notch filters for this same delay range. In all, the gains recovered with the baseline-dependent main-lobe filter reach a dynamic range of $\sim 10^5$ for $\tau \gtrsim 400$ ns.

spectrally unsmooth foreground emission, i.e. foreground emission that is coupled with primary beam sidelobes and is also bright enough after primary beam attenuation to drive calibration solution.

3.5.6 Calibration with the main lobe filter

We now consider the more aggressive baseline-dependent main-lobe filter. [Figure 3.14](#) shows the gains obtained after calibrating with a filtered complete and incomplete sky model. As expected, the baseline-dependent main-lobe filter attenuates more of the missing foreground emission than the baseline-independent notch filters ([Figure 3.7](#)), resulting in gains that have less spectral structure. In addition, the baseline-dependent main-lobe filter also suppresses more of the thermal noise in the visibilities than the baseline-independent notch filters, resulting in gain solutions that themselves also have lower noise floors, as evidenced by the lower plateau in [Figure 3.14](#). However, as previously noted, the baseline-dependent main-lobe filter correlates noise between

different time bins, resulting in gain solutions that are now more correlated between different times. We defer exploration of these consequences on real HERA analyses to future work.

To understand the performance of these filters to deeper dynamic ranges, we repeated this full analysis on noise-free visibility simulations, shown in [Figure 3.15](#). Here we can see just how deeply the spurious spectral structure introduced by model incompleteness (cyan) is suppressed. We see that the baseline-independent notch filters suppress these structures by over an order of magnitude for $\tau > 150$ ns, with the baseline-dependent main-lobe filter gaining an additional factor of 2 in suppression, however, recall this comes at the expense of more time-to-time noise correlations. The choice of which filter is best likely depends on the specific analysis at hand. Nonetheless, we have demonstrated that both are highly effective at suppressing spurious spectral structure in HERA calibration solutions caused by poorly modeled diffuse galactic synchrotron emission.

3.5.7 What does this mean for detecting the EoR?

Having shown how fringe rate filtering the visibilities before calibration can suppress spectral structure in the recovered gains by over an order of magnitude, the question we are left with is, how much is enough? Is an order of magnitude improvement enough to actually make a 21 cm EoR detection? The answer to this question is highly dependent on a number of factors, including the kind of telescope used for observations (i.e. its primary beam response) and to which part of the sky it is pointed at (i.e. at a hot or cold spot of the foreground sky). Nonetheless, roughly speaking, it is generally assumed that we need to achieve calibration precision with a dynamic range in the recovered gains of $\sim 10^5$ or more to make a fiducial EoR detection for cosmological modes of $k \gtrsim 0.13 h \text{ Mpc}^{-1}$ ([DeBoer et al., 2017](#)). At $z = 8$, this roughly translates to $\tau = 400$ ns. From [Figure 3.15](#), we can see that, indeed, the recovered gain solutions after fringe rate filtering roughly hit this 10^5 dynamic range requirement for $\tau > 400$ ns, meaning that our technique can (in principle) produce gain solutions with the required precision for a fiducial 21 cm detection given the assumptions made in this study. Specifically, this is only a statement on our technique's ability to mitigate spectral structure due to unmodeled diffuse

foregrounds, not on its ability to mitigate other real-world factors, such as crosstalk and/or mutual coupling, which we defer to future work.

3.6 Conclusions

In this paper, we have presented a technique for mitigating the impact of poorly modelled Galactic diffuse foregrounds in the calibration of 21 cm drift-scan observations. The technique relies upon the use of temporal filters that isolate and suppress diffuse emission in the sidelobes of the primary beam, while retaining the well-understood point source emission in the main lobe of the primary beam. We explored two different types of filters: a baseline-independent “notch” filter, and a baseline-dependent “main lobe” filter. Notch filters suppress only foreground emission centred at the zero fringe rate mode in the visibilities, whereas the baseline-dependent main-lobe filters aim to suppress all emission that is inconsistent with point sources in the field of view of the primary beam.

To test our technique, we simulated realistic HERA observations where we included point sources, extended sources, and a diffuse galactic component in our sky model (Hurley-Walker et al., 2017; Zheng et al., 2017), and used an electromagnetic simulation of the HERA primary beam response (Fagnoni et al., 2020).

The simulated visibilities were corrupted with realistic mock gains and then pushed through HERA’s redundant and absolute calibration pipeline. We calibrate against a shallow point source catalogue in order to simulate a realistic observation where diffuse emission is not considered and the point source catalogue is incomplete. We test whether our visibility filters improve the quality of the subsequent recovered gain solutions, relative to the case where no filters are used. We found biased calibration solutions when no filters are used, even if we adopt the standard procedure of discarding short baselines that are most sensitive to diffuse emission. This bias takes the form of spurious frequency structure in the gains that appears at $\tau \gtrsim 150$ ns in the gains in Fourier space. This directly impacts the 21 cm power spectra, resulting in increased foreground leakage for $k > 0.15$ h Mpc⁻¹ when these calibration solutions are applied to the data.

The use of fringe rate filters before calibration substantially mitigates these biases, reducing the excess spectral structure in the recovered gain solutions by over an order of magnitude. This reduces the observed foreground leakage down to the simulated noise floor in the data at $k > 0.15$ h Mpc⁻¹. When repeating this analysis on noiseless simulations, we find that our technique yields gain solutions that are accurate at 1 part in 10⁵ in dynamic range for delays $\tau \gtrsim 400$ ns, which is roughly at the estimated calibration requirement for detecting the EoR with HERA at $k \gtrsim 0.2$ h Mpc⁻¹ (DeBoer et al., 2017).

This technique can be directly applied to improve the calibration of drift scan observation from radio interferometric arrays, such as the Canadian Hydrogen Intensity Mapping (CHIME, CHIME Collaboration et al., 2022) aimed at mapping neutral hydrogen over the redshift range $z = 0.8 - 2.5$. However, for other interferometric arrays probing the EoR, which are mainly observing in tracking mode, such as Square Kilometer Array (SKA, Koopmans et al., 2015) and Low-Frequency Array (LOFAR, van Haarlem et al., 2013a), our technique cannot be applied.

Future work will explore applying these fringe rate filters to the problem of mitigating non-redundancies, for example, due to primary beam variations across the array (Orosz et al., 2019), and mutual coupling (Fagnoni et al., 2020; Kern et al., 2019, 2020b; Josaitis et al., 2022), as well as looking at ways of implementing these filters in a way that is robust to radio frequency interference.

CHAPTER 4

Assessment of the absolute calibration in the H4C HERA data release

As described in the previous chapters, calibration is the process that removes the corruption of the celestial signal that occurs as it is collected by the telescope and then further propagates through the electronics. Standard calibration observations are carried out by pointing the array towards a point source with known properties. A redundant, non-tracking array like HERA, however, adopts a different approach, relying on a combination of redundant and absolute calibration (see Chapter 3 for a discussion about the HERA calibration strategy). HERA recently deployed new broadband feeds, covering the 46 – 234 MHz range and extending the observing range to the Cosmic Dawn epoch.

In this chapter, we assess the quality of wideband HERA observations, with a particular focus on the absolute flux density scale. During an observation, a calibrator source (a source with known properties) is observed and interspersed with observations of the target field. Throughout the observation of the calibrator, the source is tracked as it moves across the sky, antenna gain solutions are derived and subsequently transferred to calibrate the observations of the target field.

This is the case for arrays with steerable antennas. However, HERA, being a transit array and also only observing in drift-scan mode, means that it does not track any source. Thus, the calibration requirements are different. The observed sources used for calibration move across the primary beam as the Earth rotates relative to the celestial sphere. Thus, the primary beam response of HERA antennas needs to be modelled accurately to avoid any bias in the derivation of antenna gain solutions during calibration. However, existing primary models do not fully characterise the primary beam response (e.g., [Fagnoni et al., 2021b](#)), and this inevitably imparts errors to the calibrated data. These errors can potentially affect the spectrum of the calibrated visibilities, either its shape or its normalization. An error in the spectrum slope may, once again, contaminate the EoR window. An error in the spectrum normalization would not affect the EoR window but would still bias the comparison with theoretical models.

4.1 Imaging of wideband HERA observations

We used a subset of data (one night covering LST range from 0.28 till 6.50 hours) from H4C observations. H4C observations were conducted over nights between Julian Date (JD) 2459114 and 2459141. Table 4.1 shows the LST coverage, integration time, frequency resolution and observing frequency range. Visibilities were flagged following the standard HERA metrics¹. Although 102 antennas were correlated, only about 40 antennas were kept after flagging. Visibilities were redundantly calibrated ([Liu et al., 2010](#); [Zheng et al., 2014](#); [Li et al., 2018](#)), preceded by an initial estimate of antenna delays - *firstcal* (see [Dillon et al., 2020](#), for a description of the *firstcal* algorithm). Afterwards, an absolute sky calibration step was performed to pin down degenerate parameters from redundant calibration ([Kern et al., 2020a](#)). At this point, the visibility spectrum should be calibrated on an absolute flux density scale.

After calibration, we divided the bandwidth in seven sub-bands, each ~ 27 MHz wide. We Fourier transformed visibilities into an image for each sub-band, using the *wsclean* software package ([Offringa et al., 2014](#)). We used the multi-frequency synthesis algorithm ([Cornwell et al., 2012](#)) with a uniform imaging weighting scheme. The bandwidth choice was a com-

¹[HERA-Team/H4C_Notebooks/tree/master/data_inspect_all_ants](#)

Minimum Frequency	46.92 MHz
Maximum Frequency	234.3 MHz
Channel Width	122.07 kHz
Number of Channels	1536
LST Coverage	19 hours - 7.31 hours
Integration time	9.37 seconds

Table 4.1: Parameters of the HERA wideband observations.

Central frequency of each sub-band image (ν) (MHz)	Synthesized beam size (arcmin)
60.2	58.78×39.97
86.6	49.38×34.18
113.8	35.43×24.59
140.6	28.88×20.26
167.3	24.81×17.40
194.1	22.27×15.42
220.9	20.31×13.93

Table 4.2: Details of the sub-band images.

promise between a value that minimizes synthesized beam sidelobes and still offers a sufficient number of data points across the bandwidth. Table 4.2 shows the reference frequency of each sub-band image, together with their synthesized beam size. Each image was deconvolved using $N = 2000$ iterations, achieving noise rms values in the $[276 - 3600]$ mJy beam⁻¹ range. In the following analysis, we selected only the snapshot observation at LST = 2^h , where two relatively bright sources are within a couple of degrees from zenith, where the primary beam attenuation is small. The same field was also used to derive the flux scale correction of the first HERA data release (Kern Nicholas, 2016). Figure 4.1 and 4.2 shows the images at the various sub-bands, with an 9.6 minute integration time. The two brightest sources in the images are

GLEAM J020012-305324 and GLEAM J015035-293158, with coordinates (RA = $2^{\text{h}} 0^{\text{m}} 12.7^{\text{s}}$; DEC = $-30^{\circ} 53' 27''$) and (RA = $1^{\text{h}} 50^{\text{m}} 35.7^{\text{s}}$; DEC = $-29^{\circ} 2' 7.2''$), respectively. The sources have flux densities of 17.95 Jy and 16.59 Jy at 151 MHz (Hurley-Walker et al., 2022). We used the Python Blob Detector and Source Finder *pybdsf* (see, Mohan & Rafferty, 2015) to extract the flux density of both sources in each image. The *pybdsf* software package is a general-purpose source extractor that identifies and best fits sources in a given image, returning their properties. It decomposes an image into a set of Gaussians, shapelets, or wavelets and fits for the source position and flux density.

We created a catalogue of all the sources in the image at the various frequencies by running *pybdsf* on each image. We set the detection threshold to be 5σ in each image. We eventually retained only the two sources of interest from the catalogue and corrected their apparent flux density by the HERA primary beam model (Fagnoni et al., 2020). Figure 4.3 shows the primary beam response at the position of the GLEAM J020012-305324 source. The source is within the main lobe in all pointings, with a maximum attenuation of 0.3 dB. Although actual HERA beams will likely differ from the simplified beam model used in this analysis, the difference remains small within the main lobe (see Fagnoni et al., 2020; Fagnoni et al., 2021b).

Figure 4.3 shows the primary-beam corrected spectrum of the GLEAM J020012-305324. We extracted the source flux densities using *pybdsf* at each frequency. The extracted source flux density at each frequency is then corrected for the HERA primary beam in order to compare its spectrum with the catalogue value. The source is isolated, thus limiting errors that may be due to blending effects. However, the background of diffuse emission may, somewhat, affect the best fit. We expect this effect to be smaller at high frequencies and larger at low frequencies due to the increasing brightness of the diffuse emission.

The right panel of Figure 4.3 shows the measured and model spectra, respectively, for the GLEAM J020012-305324 source. The model spectrum follows the standard power law:

$$S(\nu) = S_0 \left(\frac{\nu}{151} \right)^{\alpha}, \quad (4.1)$$

where $S_0 = 18.0$ (16.6) Jy and $\alpha = -0.86$ (-0.8) for sources GLEAM J020012-305324 (GLEAM J015035-293158). We note that the model and measured spectra are overall similar, albeit with




Figure 4.1: Images of the field centred at $\sim 2^{\text{h}}$ at 60 MHz (top left panel), 87 MHz (top right panel), 114 MHz (bottom left panel) and 140 MHz (bottom right panel), respectively. Two bright sources can be identified near the centre: they are used to set the flux scale of the H4C calibrated data. Images are $17^{\circ} \times 17^{\circ}$ wide.




Figure 4.2: Same as Figure 4.1 but at 167 MHz (top left panel), 194 MHz (top right panel) and 220 MHz (bottom left panel), respectively.




Figure 4.3: Left: Primary beam value as a function frequency at the location of source GLEAM J020012-305324. Colours denote the different positions of the source on the primary beam as the source drifts across it. All source positions are within two degrees from the zenith pointing. Right: observed spectrum of the source GLEAM J020012-305324. At each frequency, asterisk are flux density measurements for different transit times. The green line shows the source spectrum from the GLEAM catalogue.




Figure 4.4: Left: Flux scale ratio r_d between calibrated and model (black stars) flux density for the GLEAM J020012-305324 source. Each asterisk denotes the recovered spectra of source from the different LST/snapshot observations. The green light is the best fit polynomial model. Right: Same, but for source GLEAM J015035-293158.

some discrepancies, mostly at low and high frequencies. We quantify the discrepancy between the model and measured spectra using the relative difference defined as:

$$D_r(\nu) = \frac{|S_m(\nu) - S_d(\nu)|}{S_m(\nu)} \times 100\% \quad (4.2)$$

We found that the average relative difference between the measured and model spectra is 9.4% across the entire band. This can be assumed as an estimate of the accuracy of the spectral calibration. Calibration can be improved by improving the sky and beam models used, although here we attempt to provide a correction that is antenna independent and that can be applied to fix the overall visibility spectrum scale in a fashion similar to [Jacobs et al. \(2013\)](#) and [Ali et al. \(2015\)](#). In order to do so, we define the ratio r_d as:

$$r_d(\nu) = \frac{S_d(\nu)}{S_m(\nu)}, \quad (4.3)$$

where $S_d(\nu)$ and $S_m(\nu)$ are the data and model spectra respectively. We fit the ratio with a second order polynomial in frequency r :

$$r(\nu) = a + b\nu + c\nu^2, \quad (4.4)$$

so that a calibrated/corrected S_c spectrum can be obtained as:

$$S_c(\nu) = \frac{S_d(\nu)}{r} \quad (4.5)$$

The polynomial coefficients are determined by minimising the chi-square χ^2 between the data and the model:

$$\chi^2(a, b) = \sum_{\nu} \frac{|r(\nu, t)_d - r_m(a, b, \nu, t)|^2}{w(\nu, t)^2}, \quad (4.6)$$

where r_d is the ratio computed from data (i.e., using equation 4.4), r_m is the fitted line, t is time of the observation, and w are the weights assigned to each data point, the inverse of the primary beam value, $w = 1/E(l, m, \nu, t)$ (e.g., [Jacobs et al., 2013](#)). The weights are chosen so that the flux extracted from the image when the source is near the zenith is up-weighted. This significantly reduces the contribution of primary beam errors to the fitted model, thus reducing the error in the derived flux scale. The left panel of Figure 4.4 shows r_d and r for the GLEAM

J020012-305324 source. The fitted flux scale ratio, i.e., polynomial fit r , ranges between 0.89 and 1.17, i.e. the measured spectrum differs from the model by up to 17%.

We repeated the same procedure for the GLEAM J015035-293158 source, which transits at approximately one degree from zenith (Figure 4.5). We notice that the primary beam response at the source location changes fairly rapidly at frequencies higher than ~ 180 MHz, and that the difference with respect to the model spectrum is more marked below 100 MHz, which may be due to the larger impact of mutual coupling, unaccounted for in the H4C calibration. The relative average difference between the measured and model spectra is 14.5% across the entire band (right panel of Figure 4.5). The derived flux scale ratio r (right panel of Figure 4.4) for this source ranges between 0.026 and 0.32. Thus, the maximum deviation between the model and measured spectra for this source is 32%.

In an ideal calibration scenario, i.e., perfect calibration and source extraction, we expect the flux scale ratio to be the same for both sources, $r = 1$, because the model and measured spectra would be the same. Since the flux scale should be ideally the same for both sources, we combine the data, and we perform a joint polynomial fit to both r_d values (Figure 4.6). This procedure also reduces bias in the derived flux scale ratio r as it is no longer derived from a single source. Notably, the flux ratio has a significantly less frequency structure compared to the flux scale ratio derived from individual sources. Table 4.3 shows the best-fit values for a , b and c for the flux scale r as well as the standard error. The derived flux scale ratio r , ranges between a minimum of 0.94 ± 0.29 at $\nu = 60$ MHz and 1.17 ± 0.36 at $\nu = 220$ MHz.

Figure 4.7 shows the average flux of both sources, corrected using the derived flux scale. Errors bars in the corrected flux scale are derived from standard error recovered from the source extractor. The corrected spectra of source GLEAM J020012-305324 closely match the model spectra of the source from the GLEAM catalogue. The relative average difference between the measured and model spectra is 3.0% across the entire band, which represents an improvement of 6.4%. Also, the corrected flux of source GLEAM J015035-293158 closely matches the spectra of the source from the GLEAM catalogue in mid frequencies $100 < \nu < 200$ MHz. However, the derived flux scale ratio does not correct well the calibrated flux near the edge frequencies, 60 MHz and 220 MHz, due to increased inaccuracies in our primary beam model at these fre-

quencies. After correction, the relative average difference between the measured and model spectra is 4.4% across the entire band, and this is an improvement of 10.1%. The maximum bandpass error after flux scale correction for sources GLEAM J020012-305324 and GLEAM J015035-293158 is 9.4% and 19.1%, respectively.

In summary, the average bandpass error before flux scale correction in the entire band is 11.9%, and there is a variation with frequency, with a maximum bandpass error of 32%. After flux scale correction, the bandpass error reduces significantly, with an average bandpass error of 3.7% and a maximum of 19%. The bandpass error in the flux of H1C data after flux scale correction reported by [Kern Nicholas \(2016\)](#) is below 10% level. However, we found that the maximum bandpass error is 19% after correcting for the flux using the derived flux scale correction. This increase in bandpass error is partially due to the wider frequency band used in H4C data [46-234] MHz, as our flux scale correction technique does not work well at edge frequencies 60 MHz and 220 MHz, and this is due to increased primary beam errors at these frequencies. The derived flux scale offset is antenna-independent, and it can be applied to fix the overall visibility spectrum scale of H4C data post-calibration in a fashion similar to [Jacobs et al. \(2013\)](#).

Fitted coefficient	Value	Standard error
a	0.858	0.02
b	0.0015	4.00×10^{-04}
c	-3.6×10^{-07}	1×10^{-06}

Table 4.3: Flux scale ratio fitted parameters and the standard error in each fitted parameter.




Figure 4.5: Left: Primary beam response of source GLEAM J015035-293158. The colours denote the primary beam values of the source as a function of frequency at different zenith pointings of the array. All array pointings are at a minimum of 1 degree to a maximum of 2 degrees from the location of source GLEAM J015035-293158. Right: Extracted source flux from imaged calibrated data; the source flux is primary beam corrected (black). The green line shows the source spectra based on the GLEAM catalogue.




Figure 4.6: Black: Ratio r_d of calibrated data flux to the model flux. Green: shows the fitted quadratic model to the r_d data, where r_d data is taken from both sources GLEAM J020012-305324 and GLEAM J015035-293158.




Figure 4.7: Left: Corrected average flux of source GLEAM J020012-305324 (black), using the derived flux scale ratio (r) (Figure 4.6) and the associated errors bars. Source spectra (GLEAM J020012-305324) according to the GLEAM catalogue (green). Right, same, but for source GLEAM J015035-293158.

Mitigating the impact of mutual coupling in the calibration of HERA observations with the use of fringe rate filters

In this chapter, we simulate the impact that mutual coupling has on the calibration of HERA observations. Mutual coupling occurs when an incident electromagnetic signal from the sky is partially absorbed by an antenna and reflected to a nearby antenna, where it is then also partially absorbed and reflected to another nearby antenna. This phenomenon occurs predominately in closely packed antenna arrays like HERA. As discussed in Chapter 1, mutual coupling complicates the calibration process in HERA by introducing spectral non-smooth errors, which then impart spectral structure to smooth foreground spectra. This overall leaks foreground power into the EoR window. In this chapter, we propose the use of fringe rate filters to mitigate the impact of mutual coupling. We investigate the use of a similar filtering and calibration technique explored in Chapter 3 as a way to minimise errors in calibration resulting from the effects of mutual coupling.




Figure 5.1: Simulated array layout with 174 antennas and 14.6 m spacing between antennas. X and Y axes indicate distances aligned along east–west and north–south directions, respectively.

5.1 Simulation framework

We make use of simulated visibilities from [Josaitis et al. \(2022\)](#). Their simulations included modelling the first-order antenna-antenna coupling in radio interferometers. Specifically, they derive a semi-analytic model of the interferometric visibilities (for details see, [Josaitis et al., 2022](#)), in which first-order coupling effects are considered. In their simulation, an incident radiation enters all interferometric elements in the array, and then each antenna in the array simulation absorbs most of the incident radiation, but some of the power is then reflected due to impedance mismatch at the terminals of the antenna feeds. Thus, each element in the array not only absorbs the incident array from the sky but re-radiates it across the array, which is subsequently absorbed




Figure 5.2: Simulated visibility amplitude from 14.6 m redundant baselines. The black line denotes the visibility spectra without mutual coupling. All the other colours denote the visibility spectra from baselines with mutual coupling.




Figure 5.3: Histogram of the reduced χ^2 for visibilities without mutual coupling (blue), i.e. χ^2/DoF . The dashed line is the theoretical reduced χ^2 -distribution with $k = 2 \times \text{DoF}$, i.e. the chi-square distribution DoF , and $x = \chi^2/\text{DoF}$. This is a simple demonstration that our calibration pipeline works as expected when the effects of mutual coupling are excluded.

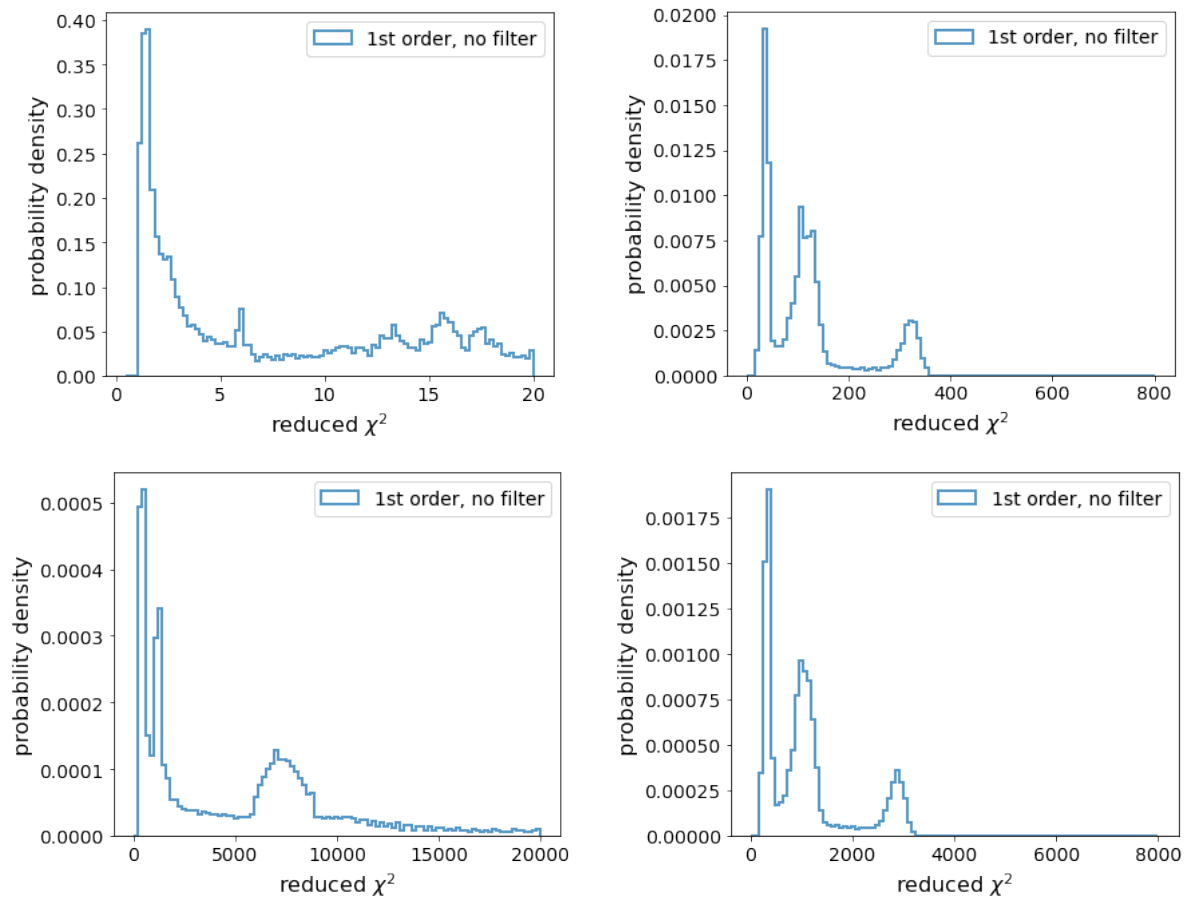


Figure 5.4: Top row: Histogram of the reduced χ^2 for visibilities set with mutual coupling obtained from redundant calibration for both sub-bands, low (left) and high (right). Bottom row: same but for chi-square obtained from absolute calibration step.




Figure 5.5: Reduced chi-square from redundant calibration for data set with mutual coupling as a function of both frequency and LST.

by other elements. This simulation only considers the first-order perturbations, and thus, it only considers the effects of re-radiated radiation being absorbed only once by all the array elements.

The simulated visibilities span a frequency range of [140 – 190] MHz, with a frequency resolution of 122.07 kHz. We consider observations in the LST range [1.9^h – 2.6^h], with an integration time $dt = 53.4$ s. The simulation is done for a HERA array consisting of 174 antennas shown in Figure 5.1. Simulated visibilities include both compact and diffuse components of the sky emission, as in Chapter 3.

Figure 5.2 shows the visibility spectra from 14.6 m redundant baselines with and without mutual coupling. The visibility spectra from redundant baselines have a higher amplitude at low frequency, as expected, due to the increasing brightness of diffuse emission at low frequencies. The visibility spectra from the 14.6 m baselines without mutual coupling differ significantly from the visibility spectra of the same baseline with mutual coupling. Also, notably, the spectra from redundant baselines with mutual coupling deviate from each other, and these deviations can be as high as 10% at ~ 140 MHz.

We corrupted the simulated visibilities using simulated gains following the same procedure as in Chapter 3. The gain g for the j th antenna is given by:

$$g_j(\nu) = A_j(\nu) e^{i\phi_j(\nu)}, \quad (5.1)$$

where the amplitude A follows a frequency power law

$$A_j(\nu) = A_j \left(\frac{\nu}{150 \text{ MHz}} \right)^{b_j}, \quad (5.2)$$

and A_j and b_j are drawn for each antenna from a Gaussian distribution $\mathcal{N} \sim (\bar{x}_A = 0.30, \sigma_A = 0.001)$ and $\mathcal{N} \sim (\bar{x}_b = -2.6, \sigma_b = 0.2)$, respectively. The phase $\phi_i(\nu)$ is modeled as:

$$\phi_j(\nu) = \sin(w_a \nu) + \cos(w_b \nu), \quad (5.3)$$

where w_a and w_b are drawn for each antenna from a Gaussian distribution $\mathcal{N} \sim (\bar{x}_w = 0.005, \sigma_w = 0.0005)$. The mean values of the parameters A and b , as well as their standard deviation σ_A and σ_b , are informed by the gain solutions from actual HERA observations (Kern et al., 2020a). The mean phase variation \bar{w} is based on single antenna phase dependency of actual HERA gains,

where the cable delay and geometric phase offset have been removed, and the variation of the mean phase between antenna stations, i.e. σ_w , is chosen to be within 10%. We then add noise to the data, such that SNR for simulated bandwidth [140-190] MHz $\sim 10^4$.

We simulate two sets of raw data products: one that includes the effects of mutual coupling and one that does not. The latter visibility product (without mutual coupling) is used to verify our pipeline as well as to serve as a benchmark to quantify the effects of mutual coupling, i.e., we compare our results with mutual coupling to the case without.

We subdivide the band into two sub-bands, which we name low-band and high-band, with a frequency range of [140 – 165] MHz and [165 – 190] MHz, respectively. We make use of a sky model with zeroth visibilities. Zeroth visibilities are simulated sky model visibilities that do not include any effects of mutual coupling.

We calibrate the data by first running redundant calibration (see Chapter 3 for a discussion of redundant calibration), followed by absolute calibration. Here, the data is further calibrated using the sky model (zeroth order visibilities). We then apply the full set of gains from redundant calibration and absolute sky calibration to obtain calibrated data. We also compute the reduced chi-square from the redundant and absolute calibration steps.

Figure 5.3 and 5.4 shows the chi-square obtained from redundant calibration for both an ideal calibration, where the visibilities do not contain effects of mutual coupling and calibration, where the data contains the effects of mutual coupling, respectively, for both the low and the high band. The chi-square histogram obtained from the data set without mutual coupling (Figure 5.3), as expected, follows the theoretical χ^2 distribution, indicating that the recovered gains are perfect, also verifying the calibration pipeline works as expected. The recovered chi-squared from the data set with mutual coupling (Figure 5.4), however, is biased to higher values. The mean value of the χ^2/DoF is 11.34. This bias is due to differences/non-redundancy errors in visibility products caused by mutual coupling, as the redundant calibration makes no reference to the sky model. Also, the redundant calibration errors are larger at lower frequencies (low band) than at higher frequencies (high band). The mean value of the χ^2/DoF for the high sub-band is only 1.06. There are two reasons for this; the first is the increased brightness of radio emission at lower frequencies. Figure 5.5 shows that the chi-square is systematically worse at low frequencies,

consistent with the effect of a more pronounced mutual coupling and increased sky brightness. The enhanced difference would lead to larger errors in redundant calibration. The enhanced difference would lead to larger errors in redundant calibration. Secondly, mutual coupling effects, hence non-redundancy errors from the mutual coupling simulation, decrease with increasing frequency (Josaitis et al., 2022).

The bottom panels of Figure 5.4 show the chi-square obtained after the absolute calibration step. Notably, the recovered chi-square again is higher for the low-frequency sub-band compared to the high-frequency band. Also, the value of the recovered chi-square from the absolute calibration step is about ~ 10 times greater than that obtained from the redundant calibration step. The increase in chi-square is a result of the "incomplete" sky model that is used for calibration, the sky model does not account for effects of mutual coupling, and thus, the model and raw visibility products differ significantly (see Figure 5.2). In addition, errors in the redundant calibration step are carried through to the absolute calibration step, thus calibration errors accumulate, leading to the obtained chi-square that is significantly larger in the absolute calibration step.

Figure 5.6 shows the gains recovered from redundant calibration for both the low and high sub-band. Gains recovered from calibration in the case of mutual coupling have a pronounced frequency structure. This is the result of deviations from redundancies due to mutual coupling, leading to non-smooth calibration errors that result in the observed frequency structure in the gains. Thus, when taking the Fourier transform of the gains along the frequency axis, i.e., delay-transformed gains, they have excess power at delays $200 < \tau < 800$ ns compared to gains recovered from the raw data set without mutual coupling (red). The dynamic range of the gains obtained in the case without mutual coupling is $\sim 10^4$ at $\tau \sim 200$ ns. When these non-smooth gains are used to calibrate visibilities, the low delay modes of the EoR window will be contaminated with foreground emission. Notably also, the mutual coupling effects lead to excess power in the gains at a broader range of delays ($200 < \tau < 800$ ns) compared to calibration errors from incomplete sky modelling ($200 < \tau < 600$ ns) shown in Chapter 3. Indeed Josaitis et al. (2022) also found that mutual coupling tends to contaminate a wider area of the EoR window.

The recovered gains in the case of mutual coupling at the high frequency sub-band, however, have comparatively less excess power at $\tau = 400$ ns, with a dynamic range $\sim 8 \times 10^3$, compared

to $\sim 4 \times 10^3$ on the case of the low frequency sub-band. Notably, also the excess power in the high sub-band gains spans a narrower delay range, only extending up to 500 ns (compared 600 ns in low sub-band).

Figure 5.7 shows the gains recovered after the absolute calibration step for both data sets with and without mutual coupling. The gains obtained in the case of mutual coupling show a pronounced frequency structure. As a result, there is excess power in gains at $\tau > 200$ ns for both low and high sub-bands. The dynamic range of the high and low sub-band gains at $\tau = 400$ ns is $\sim 10^3$, whereas the gains from the data set without mutual coupling have a dynamic range $\sim 10^4$. The excess power in the high-band gains spans a narrower delay range, only extending up to a delay of 1000 ns, compared to excess power in low-band gains that extends up to 4000 ns. Notably, the excess power in the gains recovered from the absolute calibration step extends to a wider delay range $200 < \tau < 1000$ ns in high band, compared to $200 < \tau < 500$ ns obtained in the redundant calibration step, representing an increase of 500 ns. The low band gains from absolute calibration have excess power at all delays $\tau > 200$ ns, whereas the gains from redundant calibration step have limited excess power, only extending up to a delay of 800 ns.

5.2 Application of fringe rate filters

We investigate the use of a fringe rate filter as a way to reduce the impact of mutual coupling in calibration. The idea is to filter out some mutual coupling effects, thus bringing the raw visibilities closer to the sky model. This is a similar approach to that used in Chapter 3, i.e., we apply fringe rate filtering to *both* the raw data and the sky model visibilities. We then use the derived gains to calibrate the unfiltered raw visibilities.

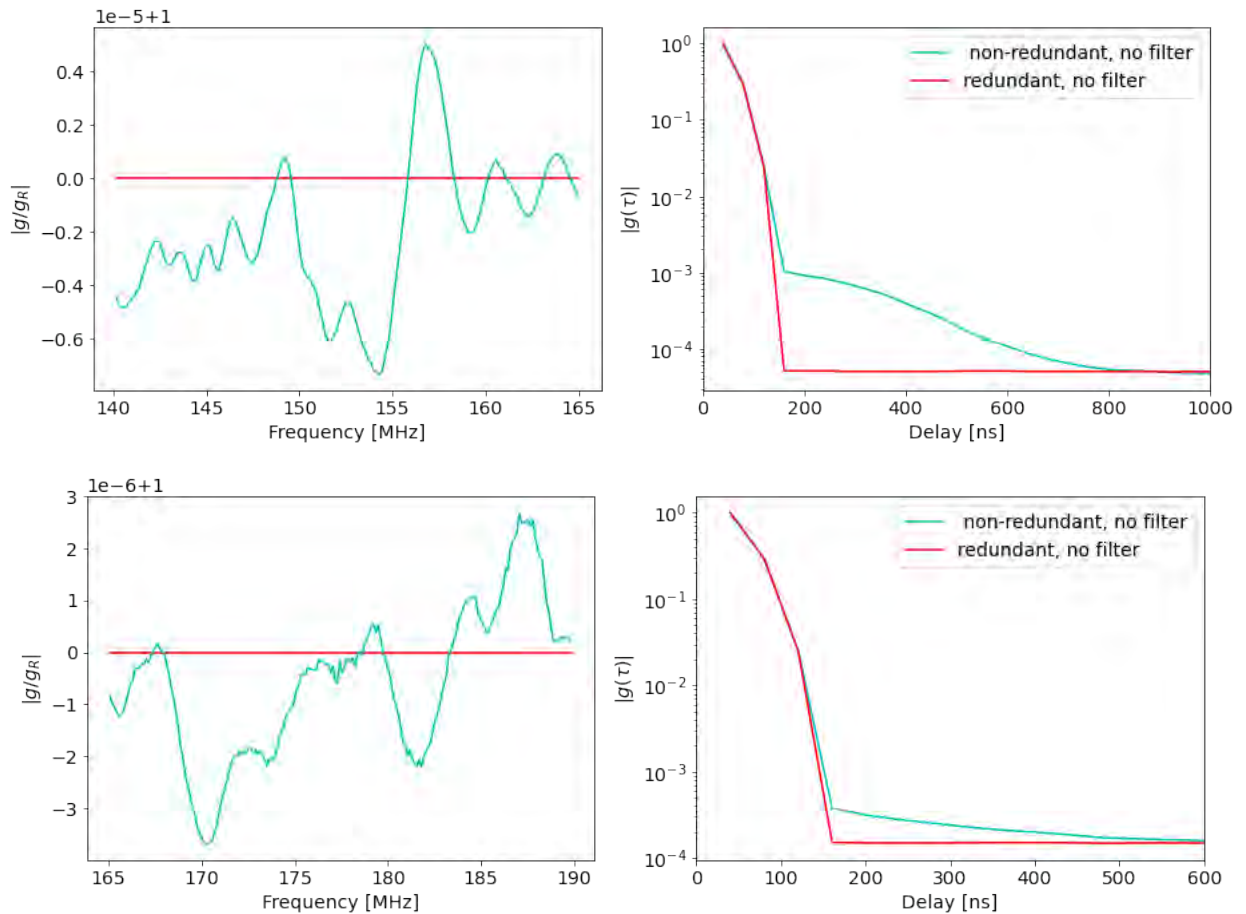


Figure 5.6: Amplitude ratio g/g_r between the gain g obtained using the data set with mutual coupling and the gain g_r obtained from calibrating the data set without mutual coupling. Gains recovered from redundant calibration for both data sets with mutual coupling (cyan) and without mutual coupling (red). Top row: Shows gains as a function of frequency from the lower band (top left) and delay transform of the gains (top right). Bottom: same as top row but for the high sub-band.

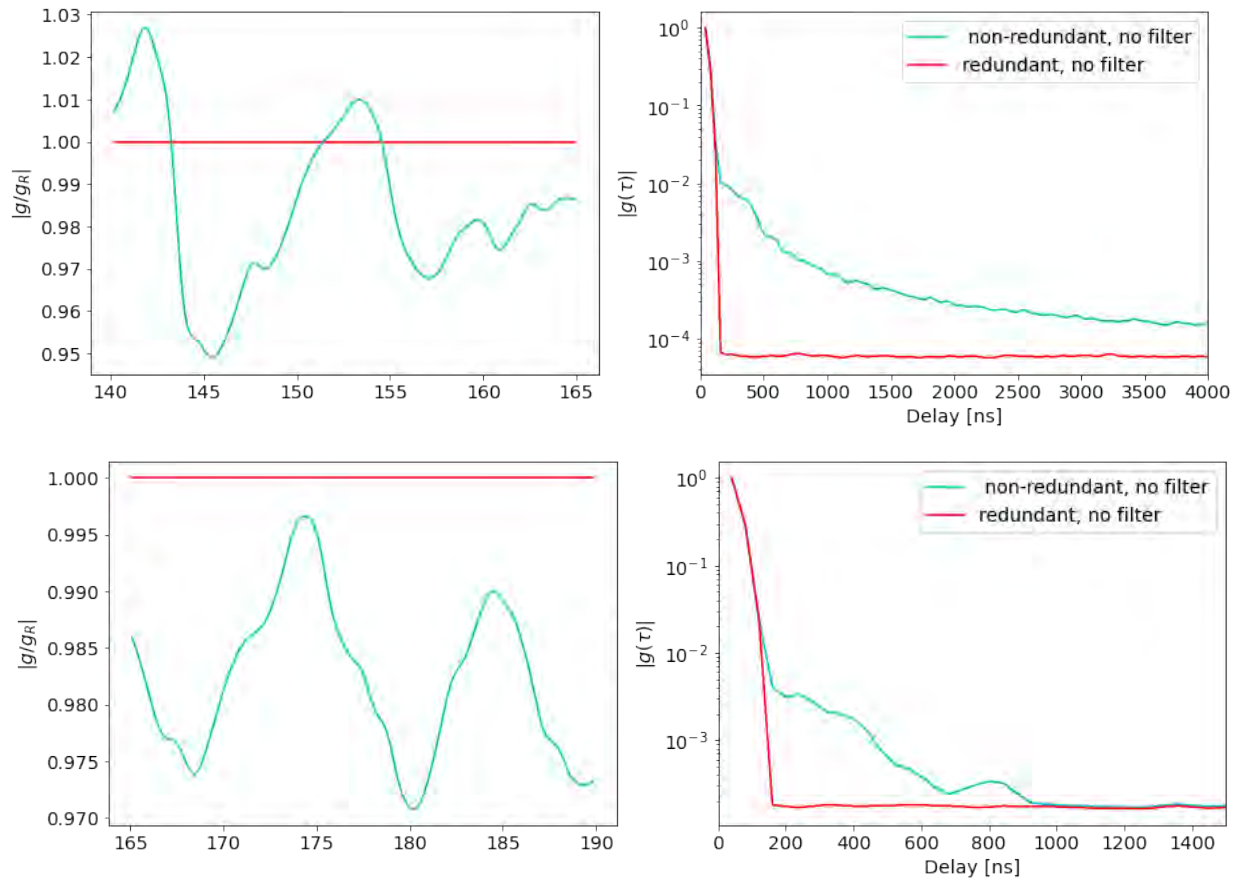


Figure 5.7: Amplitude ratio g/g_r between the gain g obtained using the data set with mutual coupling and the gain g_r obtained from calibrating the data set without mutual coupling. Gains recovered from the absolute calibration step for both data sets with mutual coupling (cyan) and without mutual coupling (red). Top row: Shows gains as a function of frequency from the lower band (top left) and delay transform of the gains (top right). Bottom: same as top row but for the high sub-band.

5.2.0.1 Notch Filters

We first consider a symmetric, baseline-independent notch filter $F(f)$ centred at $f = 0$ mHz fringe rate, i.e. a high-pass filter, defined as:

$$F(f) = \begin{cases} 10^{-8}, & |f| \leq f_{\max} \\ 1, & |f| > f_{\max} \end{cases}, \quad (5.4)$$

with $f_{\max} = 0.25$ mHz, we refer to the filter as f_{25} .

5.2.0.2 Main Lobe Filters

We also considered a second type of filter, which we refer to as a ‘‘main lobe’’ filter, because it aims to suppress emission from outside the primary beam field-of-view. In contrast to the baseline-independent notch filter, which only filters out emission in a region near $f \sim 0$ mHz, the baseline-dependent main-lobe filter suppresses the signal everywhere *except* near the peak emission of the sky model in fringe rate space.

The baseline dependent main lobe filter is a frequency-dependent filter, but because we operate over a relatively small bandwidth (25 MHz within each sub-band), we approximate it as frequency-independent. The response of the filter within the pre-defined fringe rate bounds is uniform, such that it can be thought of as a top-hat filter.

The bounds of the baseline-dependent main-lobe filter are determined by its centre f_0 and its half-width f_w . These parameters are determined for each baseline by fitting a Gaussian profile $G(f)$ to the sky model visibilities in fringe rate space:

$$G(f) = A e^{-\frac{(f-f_0)^2}{2\sigma^2}}, \quad (5.5)$$

where A is the amplitude of the Gaussian, σ is its standard deviation, and f_0 is its mean. After the fit, we set the main lobe filter centre to be f_0 and its half-width to be $f_w = 2\sigma$, such that the full width of the main lobe filter is four times the fitted Gaussian standard deviation. We denote the baseline dependent main lobe filter as M_{de} .

5.2.1 Use of a notch filter to improve calibration

We first examine the χ^2 obtained after absolute sky calibration. Figure 5.8 shows the chi-square obtained when notch fringe filters are applied before calibration. The chi-square obtained from the absolute calibration after applying filter f_{25} has reduced significantly, with a mean value of 4.01 and 1.04 for both high and low bands, respectively. This represents a reduction in χ^2 of factor 4 and 1.5 in the low and high band, respectively. This reduction occurs because the filters suppress some of the mutual coupling features in the raw data, thus bringing the data closer to the model. A smaller χ^2 value is indicative of robust calibration. Thus, we expect the recovered gains to improve as well. Figure 5.9 shows the recovered gains after applying fringe rate filters. The recovered gains have a less pronounced frequency structure, and as a result, the delay transform of the gains has significantly less excess power at delays $\tau > 200$ ns. The dynamic range is improved by factor ~ 10 at delays $200 < \tau < 500$ ns for both high and low band gains, particularly when the notch filter f_{25} is applied. However, the gains recovered after applying the main lobe filter do not show any improvement in the dynamic at delays $200 < \tau < 500$ ns, whereas the notch filter improves the dynamic range by a factor of ~ 5 in the low frequency sub-band. In the high sub-band, the main lobe filter improves the dynamic range by a factor of 2, and the notch filter improves it by a factor of ~ 10 in the $200 < \tau < 400$ ns range. Notably, the gains recovered from the filtered data have a higher noise floor, this is due to lower SNR as some of the sky signal is suppressed by filters (similar to what we obtained in Chapter 3).

We can see that the notch filter performs better than the main-lobe filter, given the relative dynamic range of the recovered gains. We can explain this behaviour by examining both the filtering and calibration process closely. The main key to getting robust gains after calibration is to minimise the differences between the sky model and raw data visibility products. The main lobe filter retains the main lobe emission. Although the effects of mutual coupling are smaller in the main lobe compared to the side lobes, there are significant deviations in the antenna primary beam main lobe due to mutual coupling effects in the main lobe, and these deviations are enhanced by the bright emission falling on the main lobe of the beam. Unlike the sky emission falling on the sidelobes, the sky emission falling on the main lobe is not heavily attenuated,




Figure 5.8: Histogram of the reduced chi-square obtained from the calibration of the data with mutual coupling and no filter (cyan), where we apply the notch fringe rate filter f_{25} before calibration (green) for both low (left panel) and high (right panel) sub-band. The reduced χ^2 improves significantly after using a notch filter, both in terms of its peak value being closer to one and a shallower tail at high values.

therefore, it dominates the calibration solutions. Differences in the main lobe of the primary beam caused non-redundancies between mutual coupling beams, and the differences between primary beams with and without mutual coupling have a much larger impact on the calibration compared to differences that come from the sidelobe of the primary beam. The key goal of the filters is to suppress a large portion of emission where mutual coupling effects are pronounced whilst retaining a good portion of the sky signal that can be used for calibration. On baselines with short a E-W projection, the baseline independent notch filter filters a portion of the emission in the primary beam main lobe, whereas the main lobe filter would retain this emission. An example of such a baseline emission is shown in Figure 5.10. The shoulder emission in the main lobe is filtered out by the notch filter (green), thus minimising the impact of non-redundancy in the main lobe and bringing the raw data closer to the sky model visibilities, thus reducing errors in calibration. However, the baseline-dependent main lobe filter (black) retains this emission.

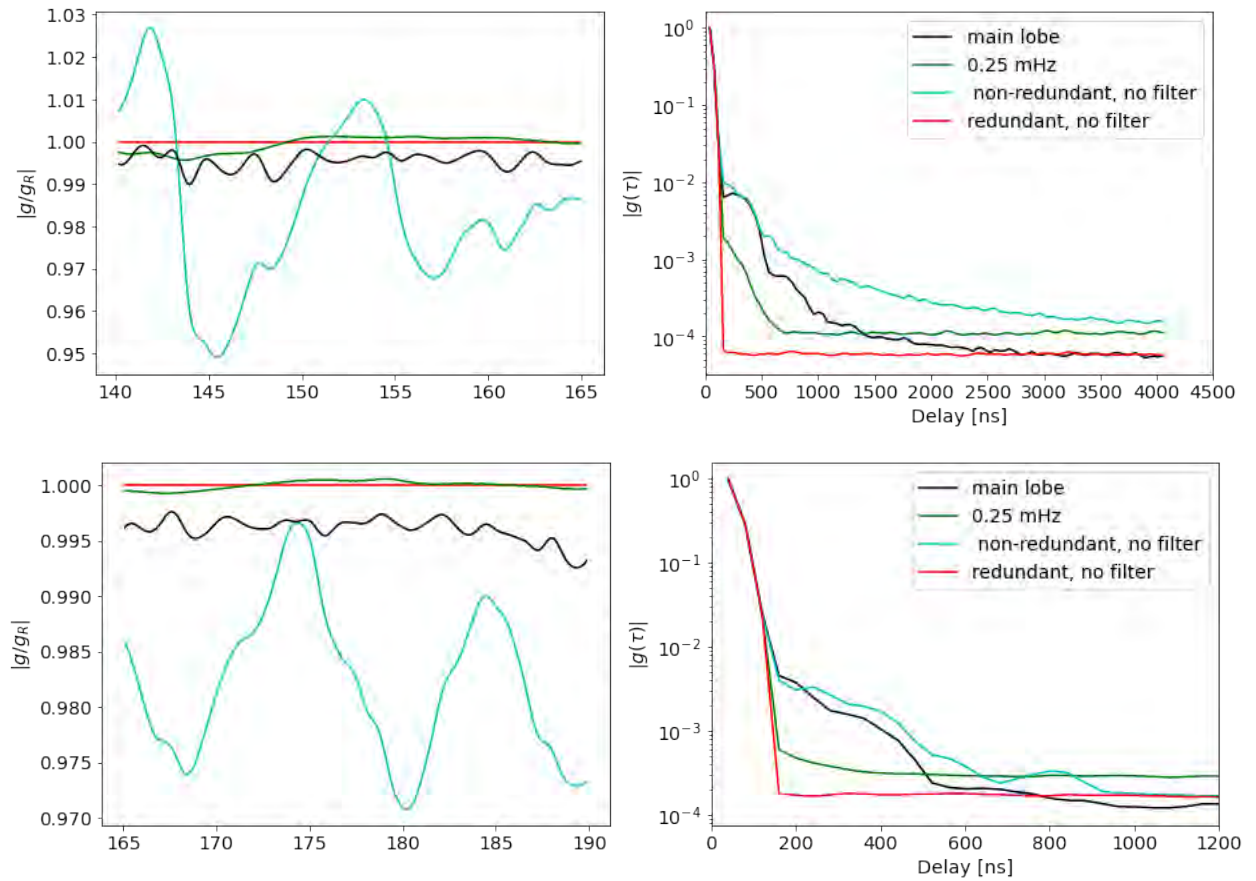


Figure 5.9: The recovered gains after mutual coupling visibilities are calibrated. The quantity on the y -axis is the same as Figure 5.7. Here, we plot the gains recovered from calibrating the data set without mutual coupling (red) and with mutual coupling (cyan) without applying filters. We then show gains after applying a notch filter of increasing width f_{25} (green) and the main lobe filter (black). The significant amount of spectral structure seen when calibrating the data set with mutual coupling is heavily suppressed after applying a filter, even more so when the baseline-independent notch filter is applied (f_{25}).

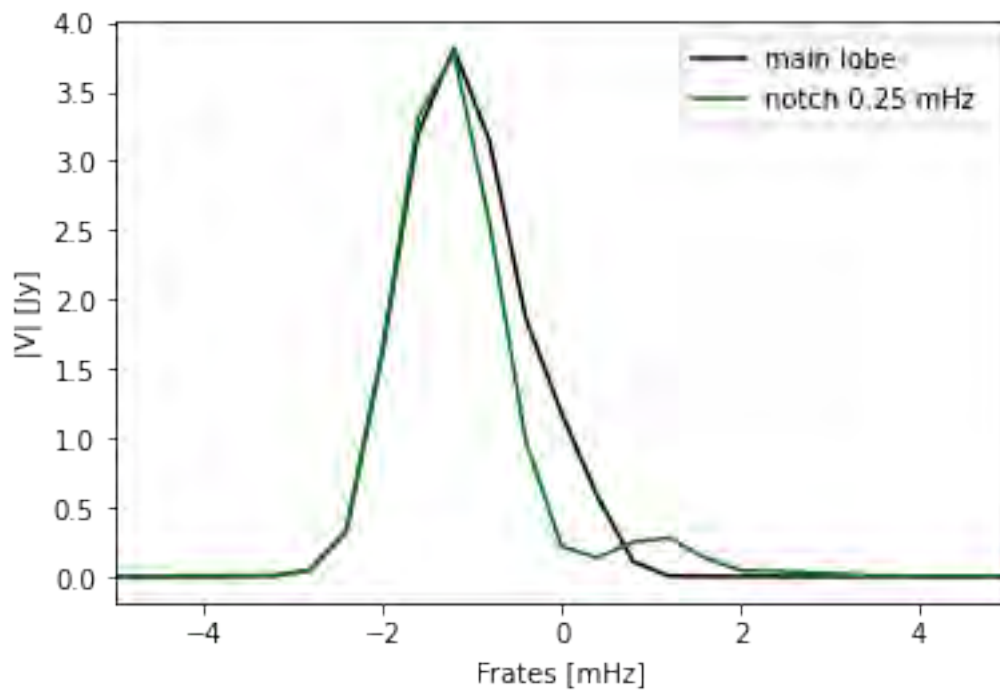


Figure 5.10: Simulated model visibility products of baseline (238, 239) (Figure 5.1) with mutual coupling in fringe rate space. The filtered visibility products using notch filter f_{25} (green) and baseline dependent main lobe filter M_{be} (black). The notch filter f_{25} filters a portion of the main lobe emission, whilst the main lobe filter (M_{be}) retains this emission.

5.3 Foreground Power Spectrum

We also compute the per-baseline power spectrum $P(\tau, b)$ (see Chapter 1 for detailed description):

$$P(k_{\parallel}, b) = |\tilde{V}_b(\tau)|^2 \left(\frac{\lambda^2}{2k_B} \right)^2 \left(\frac{D_c^2 \Delta D_c}{B_{\text{eff}}} \right) \left(\frac{1}{\Omega B_{\text{eff}}} \right), \quad (5.6)$$

where λ is the centre wavelength of observing bandwidth, k_B is the Boltzmann constant, B_{eff} is the effective bandwidth, D_c is the comoving distance at the redshift of our measurement, ΔD_c is comoving distance parallel to the line of sight, Ω is the field of view solid angle, \tilde{V}_b is the Fourier transformed visibility, and b is the visibility baseline length. The wavevector k_{\parallel} is given by (Thyagarajan et al., 2013):

$$k_{\parallel} = \frac{2\pi\nu_{21}H_0E(z)}{c(1+z)^2} \tau, \quad (5.7)$$

where τ is the delay, $\nu_{21} = 1420$ MHz, H_0 is the Hubble constant, $E(z) = [\Omega_m(1+z)^3 + \Omega_\Lambda]^{1/2}$, Ω_m is the normalized matter content and Ω_Λ is the normalized dark energy content.

In Figure 5.11, we show the residual visibilities obtained after calibrating raw data with and without mutual coupling. The residual visibilities are obtained by taking the difference between calibrated data and the sky model visibilities without mutual coupling effects. As expected, when visibilities do not include mutual coupling, we get a perfect calibration, and thus, the calibrated visibilities differ from the sky model visibilities by noise. However, when our simulated visibilities include the effects of mutual coupling, our calibration is not perfect, as demonstrated by the recovered chi-square (Figure 5.8) and gains (Figure 5.7). When these recovered gains are used to calibrate the raw data they impart on non-smooth spectral structure to the foreground emission, thus some of the foreground emission is pushed to higher delays beyond the horizon limit (white line). Hence, the calibrated visibilities differ from the sky model visibilities. This is most prominent on the short baselines, where the residual foreground emission extends up to delays of $\tau \sim 500$ ns at baselines length $b < 75$ m. The residual foreground power is 10^3 above the noise (note that residuals are squared). When the notch filter f_{25} is applied before calibration, more robust gains are used to calibrate the unfiltered raw data. We find that the calibrated visibilities closely match the sky model visibilities. As a result, the foreground residual emission beyond the horizon limit is significantly reduced by a factor of $\sim 10^2$. The excess residual foreground

within the horizon line poses no issue as this region in delay space is expected to be occupied by foreground emission.

Figure 5.12 further shows the residual power spectrum for 14 m redundant baselines for both low and high frequency sub-bands. The foreground residual power is significantly less in the high sub-band than in the low sub-band as the foreground residual emission is reduced to the noise level at $k_{\parallel} > 0.4 \text{ h Mpc}^{-1}$, but at the low sub-band, this only occurs at $k_{\parallel} > 0.6 \text{ h Mpc}^{-1}$. This is again due to the increased mutual coupling effects at the low band. We also note the main-lobe filter also suppresses foreground residual emission. However, the suppression is not as significant as the notch filter in both the low and high sub-band. In the low band, the notch filter suppresses residual foreground by a factor 10^2 to levels below the noise floor at k-modes in the range $0.4 < k_{\parallel} < 0.6 \text{ h Mpc}^{-1}$, whereas there is no reduction when main lobe filter is used. In the high sub-band, the notch filter suppresses foreground residual emission by a factor 10^2 better than the main lobe filter in the range $0.2 < k_{\parallel} < 0.6 \text{ h Mpc}^{-1}$. Thus, the use of the f_{25} notch filter robustly gains that significantly reduce the foreground power leakage due to the effects of mutual coupling. We have also shown in Chapter 3 that the use of this filter can also reduce foreground power leakage caused by the use of an incomplete sky model during calibration. In addition, since the filter is narrow, its impact on the SNR is comparatively small compared to the main lobe filter.

5.4 Conclusions

In this chapter, we investigated the impact of mutual coupling in the calibration of HERA observations. We showed that mutual coupling leads to erroneous recovered gains that are spectrally unsmooth. And when these gains are used to calibrate the data, they impart unsmooth spectral structure in the foreground emission, leading to foreground power leakage in the EoR window. We showed that the use of fringe rate filters suppresses foreground power leakage by providing more robust gains that can help calibration. Importantly, we showed that using a notch filter centred at zero fringe rates provides a significant foreground power leakage suppression compared to the main lobe filter, which suppresses all emission outside the main lobe of the primary beam.




Figure 5.11: Residual squared delay-transformed visibilities in delay space as a function of baseline length at low band. The residuals are average over all baselines of equal length. The residual is taken with respect to the noise-free visibility model without mutual coupling. We show residuals in the case where calibration assumes raw data visibility products without mutual coupling and is thus perfect, leaving only noise in the residual (left panel), and raw data visibility with mutual coupling (middle panel), and a f_{25} notch filtered scenario (right panel). The white line marks the horizon limit of the baselines, which bounds the natural extent of foreground emission in the data.




Figure 5.12: The residual delay power spectrum of the calibrated raw visibilities with respect to the sky model (noise-free) visibilities, averaged over 14 m redundant baselines and, over time, integrations. Left: Low band at redshift $z = 8.3$, where raw data is without mutual coupling (red) and with mutual coupling (cyan), and then having applied the f_{25} notch filter (green). Right: Same, but for the high sub-band at redshift $z = 6.9$.

We found that the notch filter suppresses foreground power leakage by a factor 10^2 at $z = 8.3$ in the range $0.4 < k_{||} < 0.6 \text{ h Mpc}^{-1}$. At $z = 6.9$, the leakage foreground power suppression is 10^2 at $0.2 < k_{||} < 0.6 \text{ h Mpc}^{-1}$.

CHAPTER 6

Conclusions

The high redshifted 21 cm signal is the most promising probe of the epoch of reionization (Furlanetto et al., 2006), however, observations of the 21 cm are challenged by the presence of bright foreground emission from our Galaxy and extra-galactic sources (Jelić et al., 2008; Bernardi et al., 2010; Thyagarajan et al., 2016). Fortunately, foreground emission is spectrally smooth and can be separated from the foreground emission using the delay spectrum approach (Parsons et al., 2012; Thyagarajan et al., 2013; Liu et al., 2014a).

Observations of the 21 m signal from the Hydrogen Epoch of Reionization Array (HERA), however, are corrupted by the non-smooth instrumental response of the array. In particular, effects of mutual coupling lead to non-smooth antenna primary beam response (Fagnoni et al., 2020; Josaitis et al., 2022). Often, these non-smooth responses are too complex to model and calibrate out. As a result, this leads to non-smooth calibration errors, which overall impart non-smooth spectral structure in foreground emission. In addition, sky models used in calibration are often incomplete due to incomplete and inaccurate existing models of the Galactic diffuse emission. The exclusion of the diffuse component of the sky in calibration also leads to non-smooth calibration errors, which impart spectral structure on smooth foreground spectra (Kern

et al., 2020a; Dillon et al., 2020).

In this thesis, we investigated the effect that antennas with different primary beams have on foreground closure spectra and, ultimately, on the power spectrum of the bispectrum phase - the metric used to detect the 21 cm signal (Thyagarajan et al., 2018, 2020). In particular, to quantify how much foreground avoidance is affected by the presence of mutual coupling beams, i.e. how much foreground power is spread into the EoR window. We found that primary beams affected by mutual coupling lead to a leakage of foreground power into the EoR window, which can be up to $\sim 10^4$ times, and is mainly caused by the unsmooth spectral structure primary sidelobes of mutual coupling beams. This power leakage is confined to $k < 0.3$ pseudo h Mpc $^{-1}$.

We investigated the impact that the use of an incomplete sky model, where the diffuse component of the sky is excluded, has on the calibration of HERA observations. We found, similar to other studies (e.g., Kern et al., 2020a), that this leads to unsmooth recovered gains, which, when used in calibration, lead to foreground power leakage in the EoR window. We proposed the use of fringe rate filters (Parsons & Backer, 2009; Parsons et al., 2016) to improve calibration, by suppressing diffuse emission in the sidelobes of the primary beam, while retaining the well-understood point source emission in the main lobe. We explored two different types of filters: a ‘notch’ filter, and a ‘main lobe’ filter. Notch filters suppress only foreground emission centred at the zero fringe rate mode in the visibilities, whereas the baseline-dependent main-lobe filters aim to suppress all emission that is outside the field of view of the primary beam. We found that both fringe rate filters reduced excess spectral structure in the recovered gain solutions by over an order of magnitude. And this provided more robust gains, that could be used to calibrate the data. This overall reduced the observed foreground leakage down to the simulated noise floor in the data at $k > 0.15$ h Mpc $^{-1}$.

We also showed that fringe rate filters can also be used to mitigate the impact of mutual coupling in calibration. We found that mutual coupling effects to erroneous that have spectral structure. These gains have excess spectral power at larger delay ranges $200 < \tau < 1000$ ns, significantly larger delay leakage range compared to excess power range in gains recovered from the use of an incomplete sky model during calibration. We found that the use of a fringe rate filter provided more robust gains with reduced spectral structure, and this overall reduced the

foreground power leakage into the EoR window by two orders of magnitude at $0.4 < k < 0.6 \text{ h Mpc}^{-1}$. We also found that the notch filter provided better suppression by a factor of $\sim 10^2$ of foreground leakage compared to the main lobe filter.

Lastly, we also performed an assessment of the calibration HERA H4C data release, and we found that the average bandpass error in the data for the entire band is 12%. We proposed a technique that can be used to derive a flux scale correction. We found that after flux scale correction, the bandpass error reduces significantly, with an average bandpass error of 3.7%. The derived flux scale offset was antenna-independent, and it can be applied to fix the overall visibility spectrum scale of H4C data post-calibration in a fashion similar to [Jacobs et al. \(2013\)](#).

In this thesis, we have, therefore, demonstrated and quantified the impact of systematic errors caused by the effects of mutual coupling and the exclusion of diffuse emission in calibration. We have demonstrated how fringe rate filters can be used to suppress foreground power leakage due to systematic errors. However, more investigation is still needed because flagged data, when included in actual observations, can potentially affect the filtering of data. When data containing flags is Fourier transformed along the time axis, we expect there to be “spectral leakage” resulting in the mixing of fringe rate modes of the sky due to the sidelobe structure introduced by the null/flags in the data. Thus, flagging on the time axis, can reduce the effectiveness of the filter. To reduce the impact of missing data points, the data can be “cleaned” prior to applying the filters, using a similar approach to that of [Kern et al. \(2020a\)](#), where the impact of flags is effectively reduced by applying a “clean” algorithm to remove sidelobe structure, thus preventing the mixing of fringe rate modes of sky prior to filtering the data. Also, we can explore new in-painting techniques that can be used to mitigate the impact of RFI in data (e.g., [Pagano et al., 2023](#)). Future work will be dedicated to exploring a combination of the main lobe and notch fringe rate filter, particularly in dealing with bias errors introduced by the use of an incomplete sky model and the effects of mutual coupling. Our work has shown that the main lobe filter is effective at reducing biases in calibration introduced by sky model incompleteness, and the notch filter is effective at reducing calibration from mutual coupling. We have also demonstrated a technique that can be used to correct spectral inaccuracies in HERA observations that are still left in after calibration. Our simulations provide a useful framework to interpret and improve EoR observations.

Bibliography

Ali, Z. S., Parsons, A. R., Zheng, H., et al. 2015, *apj*, 809, 61

Barry, N., Beardsley, A. P., Byrne, R., et al. 2019, *pasa*, 36, e026

Barry, N., Bernardi, G., Greig, B., Kern, N., & Mertens, F. 2021, arXiv e-prints, arXiv:2110.06173

Barry, N., Hazelton, B., Sullivan, I., Morales, M. F., & Pober, J. C. 2016, *mnras*, 461, 3135

Beardsley, A., Hazelton, B., Sullivan, I., et al. 2016, *apj*, 833, 102

Bennett, J. O. 2007, *The essential cosmic perspective*, 4th edn. (San Francisco ; London : Addison Wesley), "The kit contains instructions and a code for you to access this dynamic website"—Cover, p. 3

Bernardi, G., McQuinn, M., & Greenhill, L. J. 2015, *apj*, 799, 90

Bernardi, G., de Bruyn, A. G., Brentjens, M. A., et al. 2009, *aap*, 500, 965

Bernardi, G., de Bruyn, A. G., Harker, G., et al. 2010, *aap*, 522, A67

Bhatnagar, S., Cornwell, T. J., Golap, K., & Uson, J. M. 2008, *aap*, 487, 419

- Bolli, P., Bercigli, M., Di Ninni, P., Mezzadrelli, L., & Virone, G. 2022a, *jatis*, 8, 011023
- Bolli, P., Davidson, D. B., Bercigli, M., et al. 2022b, *jatis*, 8, 011017
- Borg, J., Magro, A., Zarb Adami, K., et al. 2020, *mnras*, 496, 933
- Bowman, J. D., Rogers, A. E. E., Monsalve, R. A., Mozdzen, T. J., & Mahesh, N. 2018, *nat*, 555, 67
- Byrne, R., Morales, M. F., Hazelton, B., et al. 2022, *mnras*, 510, 2011
- Byrne, R., Morales, M. F., Hazelton, B. J., & Wilensky, M. 2021, *mnras*, 503, 2457
- Byrne, R., Morales, M. F., Hazelton, B., et al. 2019, *apj*, 875, 70
- Callingham, J. R., Ekers, R. D., Gaensler, B. M., et al. 2017, *apj*, 836, 174
- Carilli, C. L., Nikolic, B., Thyagarayan, N., & Gale-Sides, K. 2018, *Radio Science*, 53, 845
- Carilli, C. L., Thyagarajan, N., Kent, J., et al. 2020, *apjs*, 247, 67
- Chapman, E., Bonaldi, A., Harker, G., et al. 2015, *PoS, AASKA14*, 005
- Charles, N., Bernardi, G., Bester, H. L., et al. 2022, *mnras*, 512, 2716
- Cheng, C., Parsons, A. R., Kolopanis, M., et al. 2018, *apj*, 868, 26
- CHIME Collaboration, Amiri, M., Bandura, K., et al. 2022, *apjs*, 261, 29
- Choudhuri, S., Bull, P., & Garsden, H. 2021, *mnras*, 506, 2066
- Cohen, A., Fialkov, A., Barkana, R., & Lotem, M. 2017, *mnras*, 472, 1915
- Cornwell, T. J., Voronkov, M. A., & Humphreys, B. 2012, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 8500, *Image Reconstruction from Incomplete Data VII*, ed. P. J. Bones, M. A. Fiddy, & R. P. Millane, 85000L
- Datta, A., Bowman, J., & Carilli, C. 2010, *apj*, 724, doi:10.1088/0004-637X/724/1/526

- Datta, A., Choudhury, M., & Chakraborty, A. 2017, in American Astronomical Society Meeting Abstracts, Vol. 230, American Astronomical Society Meeting Abstracts #230, 316.14
- De Breuck, C., Röttgering, H., Miley, G., van Breugel, W., & Best, P. 2000, *aap*, 362, 519
- de Oliveira-Costa, A., Tegmark, M., Gaensler, B. M., et al. 2008, *mnras*, 388, 247
- DeBoer, D. R., Parsons, A. R., Aguirre, J. E., et al. 2017, *pasp*, 129, 045001
- Dillon, J., Liu, A., Williams, C., et al. 2014, *prd*, 89, 23002
- Dillon, J. S., Tegmark, M., Liu, A., et al. 2015, *prd*, 91, 023002
- Dillon, J. S., Kohn, S. A., Parsons, A. R., et al. 2018, *mnras*, 477, 5670
- Dillon, J. S., Lee, M., Ali, Z. S., et al. 2020, *mnras*, 499, 5840
- Eastwood, M. W., Anderson, M. M., Monroe, R. M., et al. 2018, *aj*, 156, 32
- Ewall-Wice, A., Dillon, J. S., Liu, A., & Hewitt, J. 2017, *mnras*, 470, 1849
- Ewall-Wice, A., Bradley, R., Deboer, D., et al. 2016, *apj*, 831, 196
- Ewall-Wice, A., Kern, N., Dillon, J. S., et al. 2020, *mnras*, 500, 5195
- Fagnoni, N., de Lera Acedo, E., Drought, N., et al. 2021a, *IEEE Transactions on Antennas and Propagation*, 69, 8143
- Fagnoni, N., de Lera Acedo, E., DeBoer, D. R., et al. 2020, *mnras*, 500, 1232
- Fagnoni, N., de Lera Acedo, E., DeBoer, D. R., et al. 2021b, *mnras*, 500, 1232
- Fialkov, A., & Barkana, R. 2019, *mnras*, 486, 1763
- Fuller, G. M., Mathews, G. J., & Alcock, C. R. 1988, *Phys. Rev. D*, 37, 1380
- Furlanetto, S. R., Oh, S. P., & Pierpaoli, E. 2006, *prd*, 74, 103502
- Gan, H., Mertens, F. G., Koopmans, L. V. E., et al. 2022, arXiv e-prints, arXiv:2209.07854

- Gehlot, B. K., Mertens, F. G., Koopmans, L. V. E., et al. 2019, *Monthly Notices of the Royal Astronomical Society*, 488, 4271
- Gehlot, B. K., Jacobs, D. C., Bowman, J. D., et al. 2021, *mnras*, 506, 4578
- Ghosh, A., Mertens, F., Bernardi, G., et al. 2020, *mnras*, 495, 2813
- Grobler, T. L., Nunhokee, C. D., Smirnov, O. M., van Zyl, A. J., & de Bruyn, A. G. 2014, *mnras*, 439, 4030
- Hamaker, J. P., Bregman, J. D., & Sault, R. J. 1996, *aaps*, 117, 137
- Haslam, C. G. T., Salter, C. J., Stoffel, H., & Wilson, W. E. 1982, *aaps*, 47, 1
- Hazelton, B. J., Morales, M. F., & Sullivan, I. S. 2013, *apj*, 770, 156
- HERA Collaboration, Abdurashidova, Z., Aguirre, J. E., et al. 2022, *apj*, 925, 221
- HERA collaboration, Abdurashidova, Z., Aguirre, J. E., Alexander, P., et al. 2022, *apj*, 924, 51
- Hills, R., Kulkarni, G., Meerburg, P. D., & Puchwein, E. 2018, *nat*, 564, E32
- Hurley-Walker, N., Callingham, J. R., Hancock, P. J., et al. 2017, *mnras*, 464, 1146
- Hurley-Walker, N., Galvin, T. J., Duchesne, S. W., et al. 2022, *pasa*, 39, e035
- Jacobs, D. C., Parsons, A. R., Aguirre, J. E., et al. 2013, *apj*, 776, 108
- Jelić, V., Zaroubi, S., Labropoulos, P., et al. 2008, *mnras*, 389, 1319
- Josaitis, A. T., Ewall-Wice, A., Fagnoni, N., & de Lera Acedo, E. 2022, *mnras*, 514, 1804
- Joseph, R. C., Trott, C. M., & Wayth, R. B. 2018, *aj*, 156, 285
- Joseph, R. C., Trott, C. M., Wayth, R. B., & Nasirudin, A. 2020, *mnras*, 492, 2017
- Kamionkowski, M. 2007, *Dark Matter and Dark Energy*, arXiv:0706.2986
- Keller, P. M., Nikolic, B., Thyagarajan, N., et al. 2023, *mnras*, 524, 583

Kellermann, K. I., & Pauliny-Toth, I. I. K. 1969, *apjl*, 155, L71

Kern, N. S., & Liu, A. 2021, *mnras*, 501, 1463

Kern, N. S., Parsons, A. R., Dillon, J. S., et al. 2019, *apj*, 884, 105

Kern, N. S., Dillon, J. S., Parsons, A. R., et al. 2020a, *apj*, 890, 122

—. 2020b, *apj*, 890, 122

Kern, N. S., Parsons, A. R., Dillon, J. S., et al. 2020c, *apj*, 888, 70

Kern Nicholas, Gianni Bernardi, C. C. 2016, ABSOLUTE CALIBRATION OF HERA-47 IDR1,
Tech. rep., University of California Berkeley

Kohn, S. A., Aguirre, J. E., La Plante, P., et al. 2019, *apj*, 882, 58

Kolopanis, M., Jacobs, D. C., Cheng, C., et al. 2019, *apj*, 883, 133

Koopmans, L., Pritchard, J., Mellema, G., et al. 2015, *Adv. Astrophys. with Sq. Km. Array*, 1

Lanman, A. E., & Kern, N. 2019, healvis: Radio interferometric visibility simulator based on
HEALpix maps, Astrophysics Source Code Library, ascl:1907.002

Li, W., Pober, J. C., Hazelton, B. J., et al. 2018, *apj*, 863, 170

Liu, A., Parsons, A., & Trott, C. 2014a, *prd*, 90, 23018

—. 2014b, *prd*, 90, 23019

Liu, A., & Shaw, J. R. 2020, *pasp*, 132, 062001

Liu, A., Tegmark, M., Morrison, S., Lutomirski, A., & Zaldarriaga, M. 2010, *mnras*, 408, 1029

Liu, A., Zhang, Y., & Parsons, A. R. 2016, *apj*, 833, 242

Martinot, Z. E., Aguirre, J. E., Kohn, S. A., & Washington, I. Q. 2018, *apj*, 869, 79

McKinley, B., Yang, R., López-Caniego, M., et al. 2015, *mnras*, 446, 3478

- 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
- Mesinger, A., Furlanetto, S., & Cen, R. 2011, *mnras*, 411, 955
- Mesinger, A., Greig, B., & Sobacchi, E. 2016, *mnras*, 459, 2342
- Mevius, M., Mertens, F., Koopmans, L. V. E., et al. 2022, *mnras*, 509, 3693
- Mohan, N., & Rafferty, D. 2015, PyBDSF: Python Blob Detection and Source Finder, *Astrophysics Source Code Library*, record ascl:1502.007, ascl:1502.007
- Morales, M. F., Beardsley, A., Pober, J., et al. 2019, *mnras*, 483, 2207
- Natwariya, P. K., & Bhatt, J. R. 2020, *mnras: Letters*, 497, L35–L39
- Neben, A. R., Bradley, R. F., Hewitt, J. N., et al. 2016, *apj*, 826, 199
- Offringa, A. R., de Bruyn, A. G., & Zaroubi, S. 2012, *mnras*, 422, 563
- Offringa, A. R., McKinley, B., Hurley-Walker, N., et al. 2014, *mnras*, 444, 606
- Orosz, N., Dillon, J. S., Ewall-Wice, A., Parsons, A. R., & Thyagarajan, N. 2019, *mnras*, 487, 537
- Pacholczyk, A. G. 1970, *Radio astrophysics. Nonthermal processes in galactic and extragalactic sources*
- Paciga, G., Chang, T.-C., Gupta, Y., et al. 2011, *mnras*, 413, 1174
- Paciga, G., Albert, J. G., Bandura, K., et al. 2013, *mnras*, 433, 639
- Pagano, M., Liu, J., Liu, A., et al. 2023, *mnras*, 520, 5552
- Park, J., Mesinger, A., Greig, B., & Gillet, N. 2019, *mnras*, 484, 933–949
- Park, J., Mesinger, A., Greig, B., & Gillet, N. 2019, *mnras*, 484, 933

- Parsons, A., Pober, J., McQuinn, M., Jacobs, D., & Aguirre, J. 2012, *apj*, 753, 81
- Parsons, A., Liu, A., Aguirre, J., et al. 2014, *apj*, 788, 106
- Parsons, A. R., & Backer, D. C. 2009, *aj*, 138, 219
- Parsons, A. R., Liu, A., Ali, Z. S., & Cheng, C. 2016, *apj*, 820, 51
- Parsons, A. R., Pober, J. C., Aguirre, J. E., et al. 2012, *apj*, 756, 165
- Parsons, A. R., Backer, D. C., Foster, G. S., et al. 2010, *aj*, 139, 1468
- Patil, A., Yatawatta, S., Zaroubi, S., et al. 2016, *mnras*, 463, 4317
- Patil, A., Yatawatta, S., Koopmans, L., et al. 2017, *apj*, 838, 65
- Peebles, P. J. E. 1968, *apj*, 153, 1
- Perkins, S. 2021, *ska-sa* *dask-ms*
- Perkins, S. J., et al. 2021, in *ASP Conf. Ser.*, Vol. TBD, *ADASS XXX*, ed. J.-E. Ruiz & F. Pierfederici (San Francisco: ASP), 999 TBD
- Philip, L., Abdurashidova, Z., Chiang, H. C., et al. 2019, *Journal of Astronomical Instrumentation*, 08, 1950004
- Pober, J. C., Parsons, A. R., Aguirre, J. E., et al. 2013, *apjl*, 768, L36
- Pober, J. C., Hazelton, B. J., Beardsley, A. P., et al. 2016, *apj*, 819, 8
- Price, D. C., Greenhill, L. J., Fialkov, A., et al. 2018, *mnras*, 478, 4193
- Procopio, P., Wayth, R. B., Line, J., et al. 2017, *pasa*, 34, e033
- Remazeilles, M., Dickinson, C., Banday, A. J., Bigot-Sazy, M. A., & Ghosh, T. 2015, *mnras*, 451, 4311
- Riseley, C. J., Lenc, E., Van Eck, C. L., et al. 2018, *pasa*, 35, e043

Rybicki, G. B., & Lightman, A. P. 1979, Radiative processes in astrophysics

Santos, M. G., Cooray, A., & Knox, L. 2005, *apj*, 625, 575

Shimwell, T. W., Rottgering, H. J. A., Best, P. N., et al. 2017, *VizieR Online Data Catalog*, J/A+A/598/A104

Shimwell, T. W., Hardcastle, M. J., Tasse, C., et al. 2022, *aap*, 659, A1

Sims, P. H., Lentati, L., Alexander, P., & Carilli, C. L. 2016, *mnras*, 462, 3069

Singh, S., & Subrahmanyan, R. 2019, *apj*, 880, 26

Singh, S., Subrahmanyan, R., Shankar, N. U., et al. 2018, *Experimental Astronomy*, 45, 269

Singh, S., Subrahmanyan, R., Udaya Shankar, N., et al. 2017, *apjl*, 845, L12

Slepian, D. 1978, *AT T Technical Journal*, 57, 1371

Smirnov, O. M. 2011, *aap*, 527, A106

Storey, W. V., J., C. B. Ashley, M., Naray, M., & Lloyd, J. 1994, *American Journal of Physics - AMER J PHYS*, 62, 1077

The HERA Collaboration, Abdurashidova, Z., Aguirre, J. E., et al. 2021, *arXiv e-prints*, arXiv:2108.07282

The HERA Collaboration, Abdurashidova, Z., Adams, T., et al. 2022, *arXiv e-prints*, arXiv:2210.04912

Thyagarajan, N., & Carilli, C. L. 2020, *prd*, 102, 022001

Thyagarajan, N., Carilli, C. L., & Nikolic, B. 2018, *prl*, 120, 251301

Thyagarajan, N., Parsons, A. R., DeBoer, D. R., et al. 2016, *apj*, 825, 9

Thyagarajan, N., Udaya Shankar, N., Subrahmanyan, R., et al. 2013, *apj*, 776, 6

- Thyagarajan, N., Jacobs, D. C., Bowman, J. D., et al. 2015, *apj*, 804, 14
- Thyagarajan, N., Jacobs, D., Bowman, J., et al. 2015, *apj*, 804, 14
- Thyagarajan, N., Carilli, C. L., Nikolic, B., et al. 2020, *prd*, 102, 022002
- Tingay, S. J., Oberoi, D., Cairns, I., et al. 2013, *Journal of Physics: Conference Series*, 440, 012033
- Trott, C. M. 2022, *jatis*, 8, 011011
- Trott, C. M., & Wayth, R. B. 2016, *pasa*, 33, e019
- Trott, C. M., Wayth, R. B., & Tingay, S. J. 2012, *apj*, 757, 101
- Trott, C. M., Jordan, C. H., Midgley, S., et al. 2020, *mnras*, 493, 4711
- van Haarlem, M. P., Wise, M. W., Gunst, A. W., et al. 2013a, *aap*, 556, A2
- . 2013b, *aap*, 556, A2
- Vedantham, H., Udaya Shankar, N., & Subrahmanyam, R. 2012, *apj*, 745, 176
- Wang, J., Xu, H., An, T., et al. 2013, *apj*, 763, 90
- Wijnholds, S. J., Grobler, T. L., & Smirnov, O. M. 2016, *mnras*, 457, 2331
- Wouthuysen, S. 1952, *aj*, 57, 31
- Yatawatta, S. 2015, *mnras*, 449, 4506
- Yatawatta, S., de Bruyn, A. G., Brentjens, M. A., et al. 2013, *aap*, 550, A136
- Zeldovich, Y. B. 1972, *mnras*, 160, 1P
- Zhang, L., Bunn, E. F., Karakci, A., et al. 2016, *apjs*, 222, 3
- Zhao, J.-H., Morris, M. R., & Goss, W. M. 2019, *apj*, 875, 134

Zheng, H., Tegmark, M., Buza, V., et al. 2014, *mnras*, 445, 1084

Zheng, H., Tegmark, M., Dillon, J. S., et al. 2017, *mnras*, 464, 3486