

Spectral redundancy for calibrating interferometers and suppressing the foreground wedge in 21 cm cosmology

Tyler A. Cox ,   Aaron R. Parsons,   Joshua S. Dillon ,  Aaron Ewall-Wice ,  and Robert Pascua 

¹Department of Astronomy, University of California, Berkeley, CA 94720, USA

²Radio Astronomy Laboratory, University of California, Berkeley, CA 94720, USA

³Department of Physics and Trottier Space Institute, McGill University, Montreal, QC H3A 2T8, Canada

Accepted 2024 June 20. Received 2024 June 18; in original form 2023 November 3

ABSTRACT

Observations of 21 cm line from neutral hydrogen promise to be an exciting new probe of astrophysics and cosmology during the Cosmic Dawn and through the Epoch of Reionization (EoR) to when dark energy accelerates the expansion of our Universe. At each of these epochs, separating bright foregrounds from the cosmological signal is a primary challenge that requires exquisite calibration. In this paper, we present a new calibration method called NUCAL that extends redundant-baseline calibration, allowing spectral variation in antenna responses to be solved for by using correlations between visibilities measuring the same angular Fourier modes at different frequencies. By modelling the chromaticity of the beam-weighted sky with a tunable set of discrete prolate spheroidal sequences, we develop a calibration loop that optimizes for spectrally smooth calibrated visibilities. Crucially, this technique does not require explicit models of the sky or the primary beam. With simulations that incorporate realistic source and beam chromaticity, we show that this method solves for unsmooth bandpass features, exposes narrow-band interference systematics, and suppresses smooth-spectrum foregrounds below the level of 21 cm reionization models, even within much of the so-called wedge region where current foreground mitigation techniques struggle. We show that this foreground subtraction can be performed with minimal cosmological signal loss for certain well-sampled angular Fourier modes, making spectral-redundant calibration a promising technique for current and next-generation 21 cm intensity mapping experiments.

Key words: instrumentation: interferometers – dark ages, reionization, first stars.

1 INTRODUCTION

Tomographic mapping of the redshifted hyperfine transition of neutral hydrogen (H I) holds great potential for studying the evolution of large-scale structure in the early Universe. Successfully observing the cosmological 21 cm signal will open a window for understanding the properties of the first stars and galaxies and constraining Λ CDM cosmology. For a review of 21 cm cosmology, see for example Furlanetto, Oh & Briggs (2006), Morales & Wyithe (2010), Pritchard & Loeb (2012), and Liu & Shaw (2020). Prior and ongoing experiments seeking to characterize patchy fluctuations in this 21 cm signal include the Precision Array for Probing the Epoch of Reionization (PAPER; Parsons et al. 2010), the Murchison Widefield Array (MWA; Tingay et al. 2013), the Low-Frequency Array (LOFAR; van Haarlem et al. 2013), the Hydrogen Epoch of Reionization Array (HERA; DeBoer et al. 2017), the Giant Metre Wave Radio Telescope (GMRT; Paciga et al. 2013), the Long Wavelength Array (LWA; Eastwood et al. 2019), the Canadian Hydrogen Intensity Mapping Experiment (CHIME; Bandura et al. 2014), and the Hydrogen Intensity and Real-time Analysis eXperiment (HIRAX; Newburgh et al. 2016).

The principal challenge faced by experiments aiming to measure the high-redshift 21 cm signal is to accurately separate the relatively weak cosmological signal from astrophysical foregrounds, particularly from continuum emission from our galaxy and other radio-bright galaxies, which are $\sim 10^{4-5}$ times brighter. To address this challenge, experiments such as HERA and CHIME rely on the distinct spectral characteristics of foregrounds and the 21 cm signal to separate the cosmological signal from the dominant foregrounds – an approach known as foreground avoidance.

For redshifted line emission, spectral frequency corresponds to a line-of-sight spatial distance, making the Fourier dual of frequency a probe of line-of-sight modes (i.e. k_{\parallel}) in the 3D Fourier space of spatial fluctuations. H I is expected to exhibit significant spatial variation on cosmological scales in density, ionization state, and spin temperature, giving the cosmological 21 cm signal structure at a variety of spectral scales (Morales & Hewitt 2004). Foregrounds, on the other hand, are spectrally very smooth, and so occupy a relatively small number of low-order spectral modes.

In the absence of instrumental systematics, these differences in spectral properties lead to a relatively clean separation between foreground contamination and higher order spatial modes of the cosmological signal (Morales & Wyithe 2010). However, this separation is degraded by the inherently chromatic nature of antenna beams and interferometric baselines, which modulate foreground

* E-mail: tyler.a.cox@berkeley.edu

power on spectral scales inversely proportional to the diameter or interferometric baseline length (Parsons et al. 2012b). The result is that foreground residuals tend to occupy a characteristic wedge shape in cylindrical k space (Datta, Bowman & Carilli 2010; Parsons et al. 2012b; Vedantham, Udaya Shankar & Subrahmanyan 2012; Pober et al. 2013; Thyagarajan et al. 2013; Liu, Parsons & Trott 2014a, b).

By avoiding modes within this wedge and focusing instead on the complementary Epoch of Reionization (EoR) ‘window’, detection of the cosmological 21 cm signal remains possible, albeit with a reduced number of modes – and thus, sensitivity – for estimating the 21 cm power spectrum. In practice, this picture is further complicated by spectral structure in signal chains, which convolutionally leak power from the foreground wedge into the EoR window, potentially overwhelming the cosmological signal. Preventing such spectral leakage requires precision calibration to invert any spectral structure that may have been imparted by the instrument. Given the high dynamic range between foregrounds and signal, chromatic effects larger than roughly one part in 10^5 must be mitigated.

Numerous approaches have been put forward to meet the challenge of accurately calibrating interferometers for 21 cm cosmology (e.g. Yatawatta et al. 2008; Liu et al. 2010; Sievers 2017; Dillon et al. 2018, 2020; Li et al. 2018; Byrne et al. 2019; Kern et al. 2020b). These can be roughly organized into the categories of sky-based calibration and redundancy-based calibration. Sky-based calibration relies on accurate models of the sky and antenna beams to solve for per-antenna gains in visibility measurements, but this approach is susceptible to model error and incompleteness (Barry et al. 2016; Ewall-Wice et al. 2017). Redundancy-based calibration attempts to circumvent model incompleteness by using the fact that visibility measurements are expected to be identical for interferometric baselines with the same separation vector, provided antennas have identical beam patterns – an assumption which can prove to be detrimental to the accuracy of derived calibration solutions if not realized in the field (Byrne et al. 2019; Orosz et al. 2019). In essence, redundancy-based calibration reduces the reliance on absolute knowledge of the sky and beam and focuses instead on symmetries internal to the antenna array.

In most implementations, both sky-based and redundant baseline calibration solve for antenna parameters independently versus frequency, relying on post-hoc methods such as averaging (Zheng et al. 2014), gain smoothing (Abdurashidova et al. 2022; HERA Collaboration 2023), or parametric fitting to limit the number of degrees of freedom in a calibration solution (Ali et al. 2015; Ewall-Wice et al. 2016a) in favour of calibrating all frequency channels in the data simultaneously. While this is not universally true (e.g. Yatawatta 2016), many calibration algorithms lack the capacity to calibrate all frequency channels simultaneously due to computing and memory requirements.

This is an unfortunate limitation, considering that the spectral axis provides the most powerful discriminant between astrophysical foregrounds and cosmological 21 cm signal. As experiments move from setting upper limits to making detections of the 21 cm power spectrum, there is a growing need for a rigorous accounting of the systematics that can arise from foregrounds interacting with low-level spectral structure in the instrument response. This motivates the need for calibration routines that can solve for the spectral structure while remaining robust to errors in sky and beam modelling.

In this paper, we introduce a new approach to calibrating radio interferometers for 21 cm cosmology, which we call *spectrally redundant calibration* or NUCAL, for short. This technique leverages the slow evolution of sky emission versus frequency to relate visibility measurements of the same angular Fourier modes at different frequencies. By extending the concept of redundant calibration along

the frequency axis, this approach allows the calibration of arrays with arbitrary bandpass structure, while minimizing potential calibration errors introduced by incomplete sky models.

This paper is structured as follows: in Section 2, we review the state of current calibration algorithms and motivate the need for new methods that optimize for the spectral smoothness of calibrated visibilities. In Section 3, we introduce the concept of spectral redundancy and how it can be leveraged to calibrate 21 cm arrays. We also introduce NUCAL – our approach to redundant spectral calibration that enforces spectral and spatial smoothness in calibrated visibilities. In Section 3.5, we apply NUCAL to simulated visibilities and show that we can faithfully solve for degenerate redundant-baseline calibration parameters and eliminate spurious structure into the calibration solutions. In Section 5, we demonstrate our ability to subtract foregrounds from simulated visibilities in a perfectly redundant array to recover the 21 cm signal within the wedge. In Section 6, we discuss the assumptions made by spectrally redundant calibration and identify areas for further refinement. We conclude with a summary of our results in Section 7.

2 REVIEW OF CALIBRATION TECHNIQUES FOR 21 CM COSMOLOGY

In this section, we review the current status of 21 cm calibration techniques and some of their limitations in order to motivate the need to improve direction-independent calibration methods by explicitly optimizing for spectral smoothness in calibrated visibilities. We begin with a brief review of the calibration problem.

2.1 Calibration problem

Per-antenna gain calibration is the process of solving for a single complex number per antenna as a function of time and frequency. These numbers model signal chain effects and enter into the observed visibility of a baseline between antennas i and j as the true visibility multiplied by the frequency- and time-dependent gain response of each antenna involved in the measurement,

$$V_{ij}^{\text{obs}}(\nu, t) = g_i(\nu, t) g_j^*(\nu, t) V_{ij}^{\text{true}}(\nu, t) + n_{ij}(\nu, t), \quad (1)$$

where V_{ij}^{obs} is the observed visibility between antennas i and j , g_i and g_j are the complex gain of antennas i and j , V_{ij}^{true} is the true visibility sampled by baseline, b_{ij} , and n_{ij} is Gaussian distributed thermal noise on that measurement (Hamaker, Bregman & Sault 1996; Smirnov 2011). Additionally, most 21 cm interferometers are dual-polarization instruments, which allow them to simultaneously measure electromagnetic signals from two orthogonal antenna polarizations. This requires that we also solve for frequency-dependent gains for each antenna polarization.

The process of correcting for these frequency, antenna, and polarization dependent gains is called *direction-independent calibration*, which we will refer to simply as *calibration* in this paper to distinguish it from *direction-dependent calibration*, which accounts for the spatial response of each array element. Direction-independent calibration methods used for solving for antenna gains traditionally fall within one of two categories: sky-based calibration and redundant calibration.

2.2 Sky-based calibration

One of the most common approaches to calibrating radio interferometers is through a method known as sky-based calibration. In sky-based calibration, detailed knowledge of the radio sky and antenna

beam pattern is employed to simulate a set of visibilities that ideally match the true, uncorrupted visibilities measured by the instrument. These simulated data products can then be used to set up an overconstrained system of equations to solve for antenna-dependent gain parameters. While this approach works quite well when the field of view is dominated by a single, bright, well-characterized source (Baars, Mezger & Wendker 1965), arrays optimized for 21 cm cosmology often have wide fields of view with elements that often have limited pointing ability, making it more challenging to build faithful sky and beam models for calibration. Additionally, diffuse Galactic emission can introduce calibration errors if not included in the model of the sky, particularly in short baselines which are more sensitive to diffuse foregrounds.

While sky maps have been slowly improving, existing experiments have not yet shown that the sky can be modelled to the precision necessary to calibrate visibilities to one part in $\sim 10^5$. In particular, delivering foreground maps which accurately account for the spatial distribution, amplitude, and spectral properties of diffuse continuum emission from the galaxy and extragalactic point sources to the precision required for 21 cm cosmology has proven very difficult. Barry et al. (2016) showed that unmodelled foreground point sources, even those below the confusion limit of current arrays, can introduce calibration errors which can limit the detection of the cosmological signal. These issues can be partially mitigated by down-weighting long baselines or only including relatively short, less-chromatic baselines when performing sky-based calibration (Ewall-Wice et al. 2017). However, this can make calibration less accurate because this approach relies on baselines dominated by diffuse emission, which is rarely accurate to better than the per cent level and is also highly polarized (Lenc et al. 2016). In addition to making a measurement within the EoR window challenging, these errors make the level of foreground subtraction necessary for measuring cosmological modes within the foreground wedge impossible.

Sky calibration is further complicated by the need to produce accurate models of the antenna primary beam, which typically have a complicated spatial response that is both polarization- and frequency-dependent. Significant work has gone into modelling instrument primary beam (Ewall-Wice et al. 2016b; Fagnoni et al. 2021a) and validating these models by attempting to measure beams *in situ* (Pober et al. 2012; Neben et al. 2016; Jacobs et al. 2017; de Lera Acedo et al. 2018; Eastwood et al. 2018; Line et al. 2018; Patra et al. 2018; Nunhokee et al. 2020). However, state-of-the-art beam models are still insufficient to acquire sky-based calibration solutions that are accurate enough, or to do foreground subtraction to level necessary for 21 cm cosmology. Even if the sky and beam are known to sufficient accuracy, ionospheric distortion can introduce frequency-dependent errors in the position and intensities of known sources that can corrupt calibration solutions, making this method difficult to execute properly for 21 cm experiments (Jordan et al. 2017; Gehlot et al. 2018; Yoshiura et al. 2021).

2.3 Redundant-baseline calibration

Another approach to calibrating an interferometer which largely skirts the issue of incomplete sky-models and imprecise knowledge of the beam is by calibrating an array assuming that antenna pairs which have the same baseline vector measure the same visibility value. Under this assumption, deviations from a shared measurement should be explainable by antenna-based gain terms that can be calibrated out of a raw visibility. This approach, known as *redundant-baseline calibration* (Wieringa 1992; Liu et al. 2010), can be used to calibrate antenna gains in the case when repeated measurements

of the same baseline vector are made without needing an a priori estimate of the true model visibilities.

Redundant-baseline calibration seeks to solve for antenna-gain and visibility parameters at each frequency by finding a solution to a system of equations of the form

$$V_{ij}(\nu, t) = g_i(\nu, t) g_j^*(\nu, t) V_{i-j}^{\text{sol}}(\nu, t) \quad (2)$$

which is done by minimizing χ^2 written as

$$\chi^2(\nu, t) = \sum_{i \neq j} \frac{|V_{ij}^{\text{obs}}(\nu, t) - g_i(\nu, t) g_j^*(\nu, t) V_{i-j}^{\text{sol}}(\nu, t)|^2}{\sigma_{ij}^2(\nu, t)}, \quad (3)$$

where V_{i-j}^{sol} is the visibility solution for redundant baselines with the same baseline vector separation as V_{ij}^{obs} and σ_{ij}^2 is the noise variance on baseline b_{ij} . In the case where the number of unique baseline separations made by a perfectly redundant array is significantly less than the total number of measurements made (such as with HERA; Dillon & Parsons 2016), a highly overconstrained system of equations can be set-up to solve for the gains and model visibilities. In practice, no array is perfectly redundant and non-redundancies in the array due to small deviations in antenna position, beam shape, and beam pointing errors can all introduce chromatic gain errors which can complicate a 21 cm measurement (Orosz et al. 2019). None the less, redundant calibration remains an attractive approach since it is mostly free of assumptions about emission from the sky, i.e. whether it is dominated by point sources, diffuse emission, or some combination of the two and does not initially require precision knowledge of the beam.

However, it is important to highlight that while redundant-baseline calibration reduces reliance on accurate models of the sky and beam, it ultimately still requires a sky-model to solve for a small number of remaining calibration degrees of freedom. Regardless of how overdetermined the system of equations used to minimize equation (3), the structure of χ^2 guarantees that there will always be a few parameters as a function of frequency that cannot be solved by redundant-baseline calibration. These frequency-dependent parameters are degeneracies of this system of equations, which arise as a result of a set of transformations that one can apply to gains and visibilities that leave the value of χ^2 unchanged (Zheng et al. 2014; Dillon et al. 2018, 2020). These transformations that can be applied to the estimated gain parameters are an overall amplitude degeneracy,

$$\begin{aligned} g_i g_j^* V_{ij}^{\text{model}} &\rightarrow (A g_i) (A g_j^*) (A^{-2} V_{ij}^{\text{model}}) \\ &= g_i g_j^* V_{ij}^{\text{model}}, \end{aligned} \quad (4)$$

and a phase gradient degeneracy,

$$\begin{aligned} g_i g_j^* V_{ij}^{\text{model}} &\rightarrow (g_i e^{i\Phi \cdot \mathbf{r}_i}) (g_j^* e^{-i\Phi \cdot \mathbf{r}_j}) (V_{ij}^{\text{model}} e^{-i\Phi \cdot \mathbf{b}_{ij}}) \\ &= g_i g_j^* e^{i\Phi \cdot \mathbf{b}_{ij}} V_{ij}^{\text{model}} e^{-i\Phi \cdot \mathbf{b}_{ij}} \\ &= g_i g_j^* V_{ij}^{\text{model}}. \end{aligned} \quad (5)$$

Here, both A and Φ are real-valued frequency- and time-dependent quantities and \mathbf{r}_i is the position of antenna i . When each polarization is calibrated independently, there are three degeneracies¹ per frequency and polarization: an average gain amplitude and a phase gradient in both the North–South and East–West directions.

¹An additional overall phase degeneracy, where $g_i \rightarrow g_i e^{i\phi}$ does not generally need calibration because it leaves calibrated single-polarization visibilities unaffected.

After performing redundant-baseline calibration, an additional *absolute calibration* must be performed to fill in this small set of degenerate calibration parameters per frequency, which amounts to setting the absolute flux scale and phase centre for each frequency observed (Liu et al. 2010; Zheng et al. 2014; Dillon et al. 2018; Byrne et al. 2019; Kern et al. 2020b). While this approach does significantly reduce the number of free parameters required to be solved for by a referencing a sky model, sky-model incompleteness errors can introduce calibration errors in redundant-calibration degeneracies which leak power into the EoR window that exceeds the 21 cm signal (Byrne et al. 2019) if further measures are not taken to mitigate the effect.

2.4 Hybrid approaches

It is worth noting that a number of hybrid calibration have been developed in recent years to address the issues detailed with both sky-based and redundant calibration. In Sievers (2017), a new approach CORCAL was presented which modelled covariance between baselines to bridge the gap between redundant-baseline calibration and sky-based calibration. Gogo et al. (2022) were able to show that this approach led to moderate improvement in the calibration of PAPER data over redundant calibration alone. Other approaches, such as BAYESCAL (Sims, Pober & Sievers 2022a, b) and Unified Calibration (Byrne et al. 2021), have developed Bayesian frameworks to incorporate model uncertainty into calibration. They have shown a reduction in the spurious spectral structure of gain solutions compared to traditional sky-based and redundant calibration when estimating gain solutions in simulated data. The newly proposed Delay-Weighted Calibration (Byrne 2023), a variant of sky-based calibration, has also showed significant promise in reducing spurious spectral structure in gain solutions by down-weighting Fourier modes thought to have significant sky-model error.

2.5 Mitigation of spectral structure in calibrated visibilities

One of the primary challenges of 21 cm cosmology is managing frequency-dependent calibration errors introduced when performing estimates of antenna-gain parameters. Assuming antenna gains are estimated accurately by some calibration method, one can simply divide the measured data by the estimated gains to recover the true sky signal. However, gain calibration is often performed inaccurately due to the issues mentioned above (i.e. inaccurate sky/beam models and non-redundancies). These errors often have fine-scale frequency structure which, due to the multiplicative nature of gain calibration (as in equation 1), are convolved with the true sky in Fourier space, leading to a leakage of foreground power into modes outside the wedge. Therefore, it is crucial that this calibration parameters be accurately estimated to prevent bright foreground from contaminating 21 cm signal.

Many of the frequency-dependent calibration errors introduced by both sky-based and redundant-baseline calibration have been mitigated by imposing the a priori assumption that the final calibrated gain parameters be smooth as a function of time and frequency. This assumption is generally imposed in a two-step process in which gains are estimated via sky-based or redundant calibration, then smoothed by fitting some set of slowly varying basis functions to the gain estimates (Barry et al. 2016; Dillon et al. 2018; Gehlot et al. 2018; Eastwood et al. 2019; Li et al. 2019). Because instruments are designed to be stable in time and have smooth bandpass responses, imposing this a priori assumption is fairly reasonable in theory. It is therefore assumed that fine-frequency structure in gain solutions is more likely the result of calibration errors than true spectral structure

in the instrumental response, and that it is therefore safer to ignore it than to introduce new chromaticity into calibrated visibilities. However, it remains to be seen whether these assumptions about smooth instrument responses hold in real-world antennas in the field. It is worth noting that there are exceptions to the approach of post-hoc smoothing calibration solutions, most notably Yatawatta (2016) which fits a parametrized version of antenna gains rather than smoothing post-calibration, but at the cost of greater computational expense and larger memory requirement.

These post-hoc approaches are suboptimal in the sense that they neglect a true accounting of what amount of true spectral structure may be found in the instrument bandpass or due to systematics such as cable reflections, which can be factored out of the data by modelling them as a gain-like term (Kern et al. 2019, 2020a). One can imagine that if the gains actually contain a significant amount of spectral structure, that post-hoc smoothing will remove those features from estimates of the gains, leading to structure which is not properly deconvolved from the calibrated visibilities potentially polluting the EoR window. For experiments to be successful in detecting the 21 cm signal, calibration routines must be designed to handle imperfect knowledge of the sky and primary beam while remaining flexible enough to correct for true fine-scale frequency structure in instrumental gains. In the next section, we introduce our approach to calibrating fine-scale frequency structure using spectral redundancy – the repeated sampling of modes in the uv -plane by baselines at different frequencies.

3 SPECTRAL REDUNDANCY

In this section, we explore the idea of spectral redundancy, an extension of traditional spatial redundancy to the spectral axis, which enforces consistency between measurements by comparing visibility measurements which sample the same uv -modes at different frequencies. We begin our discussion of spectral redundancy with a brief description of the visibility simulations used throughout the rest of this work (Section 3.1) and demonstrate the high degree of redundancy between baselines at frequencies which sample the same modes in the uv -plane (Section 3.2). We then discuss our approach to modelling spectrally redundant baselines by representing them as the sum of a flexible set of basis functions known as discrete prolate spheroidal sequences (DPSS). Finally, we revisit the antenna calibration problem and present our approach to calibrating data using spectral redundancy, NUCAL, in Section 3.5.

3.1 Visibility simulation review

Before we begin with a description of spectral redundancy, we first give a brief description of the visibility simulations used in this work. In order to explore the extent to which measurements which made at the same angular Fourier mode at different frequencies are correlated, we use the antenna layout of the HERA array for our simulations. HERA was designed to have a high degree of spatial redundancy, but due to its regular layout of baselines in the same orientation, also has a significant amount of spectral redundancy as well, making it a good candidate for explorations of model real world arrays with spectral redundancy. HERA’s layout and corresponding sampling of the uv -plane are shown in Fig. 1.

For our foreground model, we generate a set of 10^5 point sources randomly distributed across the sky. The flux density of each point source follows the power-law relation,

$$S(v) = f \left(\frac{v}{v_0} \right)^\alpha \quad (6)$$

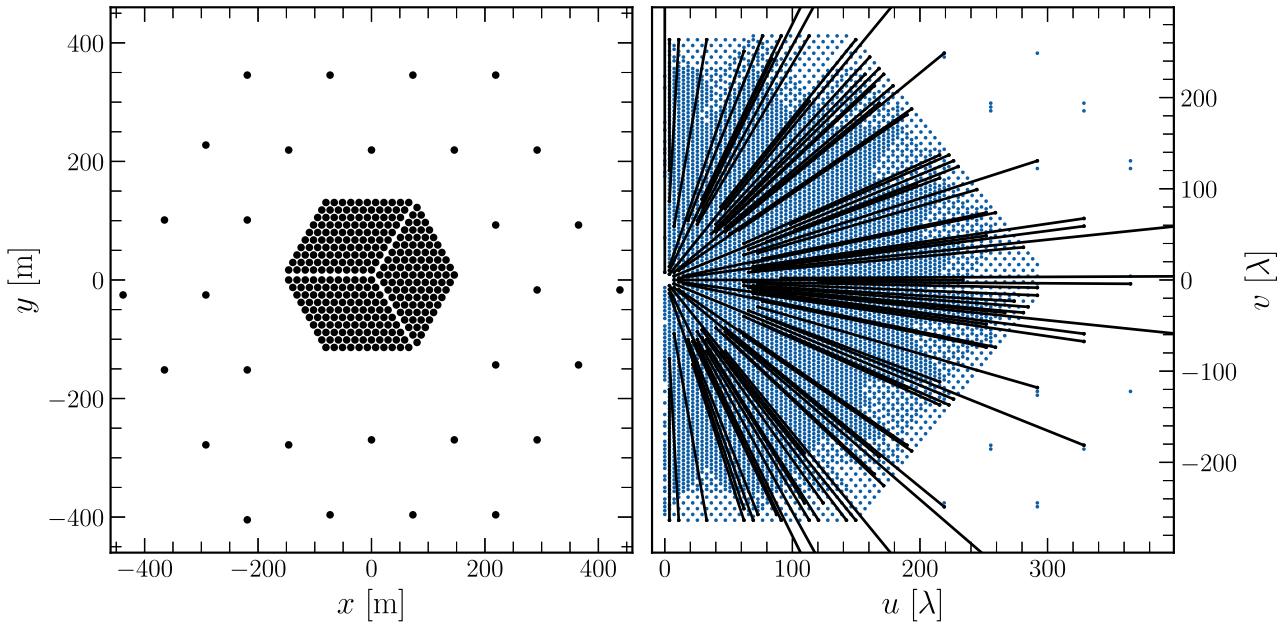


Figure 1. **Left Panel:** Array layout for the Hydrogen Epoch of Reionization Array (HERA). HERA was designed for redundant-baseline calibration, but its split-core configuration and 30 outrigger antennas effectively triple the number of unique baselines along any given heading in the uv plane compared to a simple hexagonally packed array (Dillon & Parsons 2016). **Right Panel:** Corresponding samples in the uv -plane for the antenna positions shown in the left panel. Black lines overlaid on the uv -plane designate radial headings which contain baselines with significant overlap of same uv -modes at different frequencies (at least 15 unique baseline types within the same orientation, though we will use all orientations with at least 10 unique baseline types for the NUCAL algorithm in Section 3).

where α is the spectral index, ν_0 is the arbitrary pivot frequency, which we define to be 150 MHz, and f is the flux density of the point source at frequency, ν_0 . Our foreground sky model is defined as the sum of each of these randomly generated point sources,

$$I(\hat{s}, \nu) = \sum_i S_i(\nu) \delta(\hat{s} - \hat{s}_i) \quad (7)$$

where δ is the Dirac-delta and \hat{s}_i is the unit vector describing the direction to that point. For each source in our simulation, we draw a random spectral index and flux density, where α is drawn from a uniform distribution ranging from -1.5 to -0.5 , and the flux density is drawn from an exponential distribution with a mean of 300 mJy. In addition to a point source foreground model, we also simulate a flat power spectrum EoR by generating uncorrelated, Gaussian random HEALpix maps (Górski et al. 2005) at each frequency. For these maps, we take the map size to be $NSIDE = 128$ ($N_{\text{pix}} = 1.96 \times 10^5$) and with a variance of 25 mK^2 , consistent with the expected fiducial brightness temperature of the signal at EoR redshifts (Mesinger, Furlanetto & Cen 2011).

For our beam model, we choose an azimuthally symmetric Airy function above the horizon, which roughly captures the spatial and spectral structure of a realistic instrument primary beam. This takes the form

$$A(\nu, \hat{s}) = \left(2 \frac{J_1(\pi d_{\text{ant}} \nu \sin \theta / c)}{\pi d_{\text{ant}} \nu \sin \theta / c} \right)^2, \quad (8)$$

where θ is the angle between zenith and the unit vector \hat{s} , J_1 is the Bessel function of the first order and first kind, d_{ant} is the diameter of the antenna, which for HERA we take to be $d_{\text{ant}} = 14$ m, and λ is the wavelength at which the beam is evaluated. While there are differences between the HERA beam and the Airy beam, they have similar spatial and spectral variability, making them a good candidate for realistic visibility simulations (Neben et al. 2016).

With the beam and sky-model set, we use the visibility equation,

$$V_{ij}(\nu) = \int d\Omega A(\nu, \hat{s}) I(\nu, \hat{s}) e^{-2\pi i \nu \mathbf{b}_{ij} \cdot \hat{s} / c}, \quad (9)$$

to perform a discrete sum over each source in our sky-model weighting by the value of the Airy beam at the position of the source for each unique baseline orientation in the HERA array. Each visibility is simulated over the frequency range 46.9 – 234.9 MHz with 1536, 122 kHz wide channels to match HERA’s extended frequency range and channel width (Fagnoni et al. 2021a).

In some figures later in this paper, we find it useful to generate thermal noise for visibilities simulated to test the dynamic range capabilities of our technique. We generate this noise by randomly sampling from a complex Gaussian distribution with zero mean and a standard deviation set by the radiometer equation,

$$\sigma_{ij} = \sqrt{\frac{V_{ii} V_{jj}}{\Delta t \Delta \nu}}, \quad (10)$$

where V_{ii} is the autocorrelation amplitude of antenna i , Δt is the integration time, which we set to HERA’s integration time, $\Delta t = 10$ s, and $\Delta \nu$ is the channel width.

3.2 Spectrally redundant sampling of the uv -plane

Traditional redundant-baseline calibration focuses on using similarities between baselines which share an orientation and physical length. However, within an array, there exist additional symmetries that can be leveraged to establish a concept of redundancy among measurements that are not conventionally considered redundant. Specifically, baselines that sample the same modes in the uv -plane at different frequencies – those sharing an orientation – may not yield identical visibilities values, but may still exhibit a form of redundancy in the sense that they are strongly correlated.

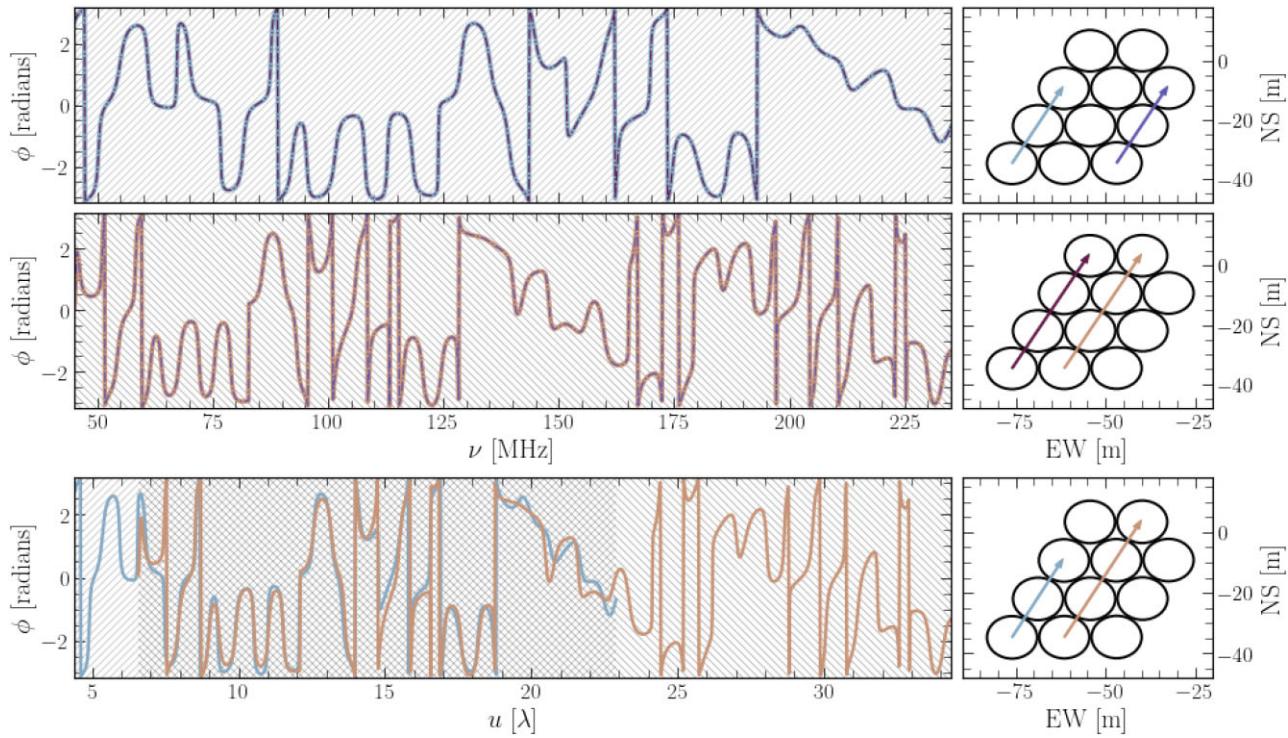


Figure 2. Traditional spatial redundancy involves comparing visibilities from antenna pairs separated by the same physical baseline vector, as we see in spectral phase plots of visibilities simulated with a chromatic beam and sky for 29 m (top) and 44 m (middle) baselines. Here, in these top two panels, the phase of each of the visibilities is plotted in the left column. In the right column, we show the corresponding antenna pairs and the baseline vectors connecting them. Because all antennas in the simulation share identical beams and the simulations are noiseless, the phases of the visibilities in the top rows align perfectly. In the lower panel, we overlay phase as a function of $u \equiv \nu |\mathbf{b}|/c$ for measured visibilities whose baselines have the same orientation but different lengths. Spectral redundancy is evidenced by the fact that both baselines see largely the same sky (at least in phase) at different frequencies. Their amplitudes will be quite different, since continuum foregrounds are might brighter at low frequency, but that spectral evolution is quite smooth at fixed u , as shown in Fig. 4. Note that though both baselines span the same range in ν , they span different but overlapping ranges in u (hashed regions).

To understand how both spatial and spectral redundancy arise in interferometric measurements, we can generalize equation (9) to express a visibility as a function of position in the uv plane, $\mathbf{u} \equiv \nu \mathbf{b}/c$, equivalent to the baseline vector in wavelength units. It enters into the equation as the Fourier dual of the projected angular brightness on the sky,

$$V(\mathbf{u}, \nu) = \int d\Omega A(\hat{\mathbf{s}}, \nu) I(\hat{\mathbf{s}}, \nu) \exp[-2\pi i \mathbf{u} \cdot \hat{\mathbf{s}}]. \quad (11)$$

From equation (11), it is clear that pairs of antennas that have the same separation vector \mathbf{b}_{ij} should measure the same visibility, assuming that beam responses are identical. In practice, the redundancy that arises from the spatial arrangement of these antennas is broken by the unique signal delay and amplifier gain versus frequency of each antenna. However, as is mentioned in Section 2.3, it is straightforward to set up a system of equations to solve for such per-antenna calibration terms in order to restore the redundancy of the underlying measurements.

Similarly, baselines which have dramatically different lengths may still be redundant in the sense that they measure very similar skies at different frequencies and thus be highly covariant. Consider the simple case of a linear array in which all the baseline vectors are oriented in the same direction. Let us also make the simplifying assumption that the sky and beam are achromatic, i.e. $I(\hat{\mathbf{s}}, \nu) = I(\hat{\mathbf{s}})$ and $A(\hat{\mathbf{s}}, \nu) = A(\hat{\mathbf{s}})$. This allows to use equation (11) to rewrite

equation (9) as

$$V_{ij}(\nu) = \int d\Omega A(\hat{\mathbf{s}}) I(\hat{\mathbf{s}}) \exp[-2\pi i \mathbf{u}_{ij} \cdot \hat{\mathbf{s}}] = V(\mathbf{u}_{ij}). \quad (12)$$

By ignoring the frequency-dependent of the sky and the beam in this toy example, we observe a given baseline's spectral characteristics are only dependent on the angular Fourier mode, u , that it measures. In this example, any two baselines in this linear array which sample the exact same angular Fourier modes ($\mathbf{u}_1 = \nu_1 \mathbf{b}_1/c = \nu_2 \mathbf{b}_2/c = \mathbf{u}_2$), would produce identical visibility values despite potentially having vastly different physical lengths and making measurements at different frequency channels. These measurements could be compared and leveraged to constrain antenna gains, similarly to what is done in traditional spatial redundant-baseline calibration.

While the sky and beam are not truly achromatic as described in this toy example, it is well known that the chromatic response of the foreground sky is quite smooth, as it is dominated by smooth spectrum processes such as galactic synchrotron emission. Likewise, antenna beams are designed to be as spectrally smooth as possible to prevent leakage of smooth spectrum foregrounds into the EoR window. Therefore, it is reasonable to assume that the sky-beam product $A(\hat{\mathbf{s}}, \nu) I(\hat{\mathbf{s}}, \nu)$ evolve smoothly with frequency, which would allow for a high degree of phase coherence between baselines which measure the same uv -modes. As shown in Fig. 2, baselines in the same orientation can in fact agree quite well in phase, provided the sky beam product evolve slowly as a function of frequency.

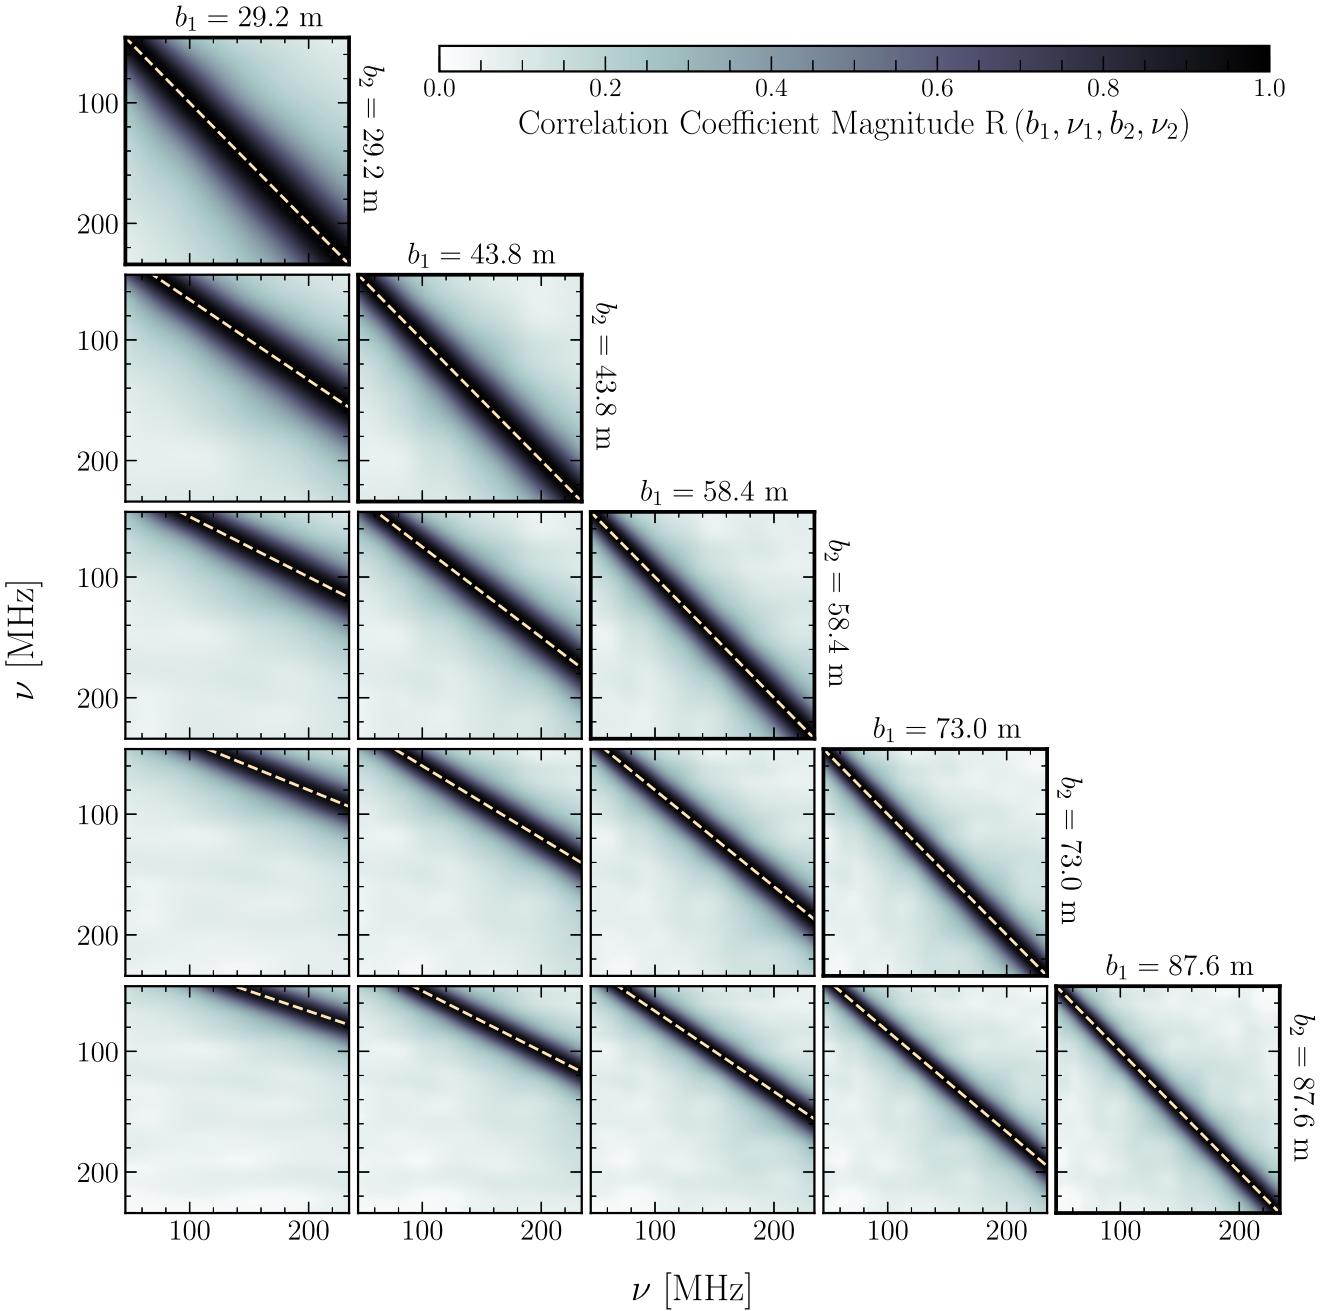


Figure 3. Cross-frequency correlation-coefficient matrices, $R_{12}(\nu_1, \nu_2) = C_{12}/\sqrt{C_{11}C_{22}}$, computed by averaging over 1000 randomly generated, noise-free point source foreground visibility simulations. Each correlation-coefficient matrix is computed between baselines in the same orientation, but with different lengths (except for the five panels on the diagonal, which show frequency-frequency correlations for the same baseline). The dashed line in each plot follows, $\nu_2 = (b_1/b_2)\nu_1$, and marks the location where baselines sample the same uv -mode, $\mathbf{u}_1 = \mathbf{u}_2$. Here, a clear correlation can be seen between baselines of different lengths, indicated by the near unity value of the cross-correlation coefficient, at frequencies where pairs of baselines redundantly sample the same angular Fourier modes. This strongly implies that different baselines at a fixed \mathbf{u} are sensitive to the same fundamental sky-beam structure and, in some sense, redundant.

We can further explore the degree to which sample the same modes in the uv -plane are redundant by examining the cross-frequency, cross-correlation coefficient between pairs of baselines. Here, we define cross-frequency, cross-correlation coefficient as

$$R_{ij}(\nu_1, \nu_2) = \frac{\mathbf{v}_{i,\nu_1} \cdot \mathbf{v}_{j,\nu_2}^*}{\sqrt{(\mathbf{v}_{i,\nu_1} \cdot \mathbf{v}_{i,\nu_1}^*) (\mathbf{v}_{j,\nu_2} \cdot \mathbf{v}_{j,\nu_2}^*)}} \quad (13)$$

where R is a matrix whose elements are the cross-correlation coefficients of pairs of frequencies sampled by baselines i and j and \mathbf{v}_{i,ν_1} a mean-subtracted visibility vector containing the visibility values from $N = 2000$ randomly generated, point source simulations (as described in Section 3.1) for baseline i at frequency ν_1 .

In Fig. 3, we show the cross-correlation coefficients for a sample of baselines within a group of baselines with the same azimuthal orientation in the HERA array. Along the main diagonal, we plot

the cross-correlation matrices of each baseline with itself, while the lower triangle features baselines crossed with another baseline of a different length. As one might expect, baselines which have been cross-correlated with themselves are highly correlated along their main diagonal. More interesting is the existence of a pronounced correlation between different baselines in the same orientation, particularly at frequencies where $\mathbf{u}_1 \approx \mathbf{u}_2$. This manifests in the near-unity cross-correlation coefficient in baselines of different lengths.

This correlation means that different measurements are sensitive to the same underlying information. A question naturally arises: why would this new sort of redundancy be *useful*? The answer comes from the fundamental difference between visibilities Fourier transformed along the frequency axis (or equivalently the line-of-sight axis) at fixed \mathbf{u} – the η transform – and visibilities Fourier transformed along the same axis for a fixed baseline \mathbf{b} – the τ or delay transform. The delay transform is, in some sense, more natural, since it is performed directly on the measured visibilities. However, as we illustrate in Fig. 4, the η transform is the space where foregrounds, even when multiplied by the beam, are the least chromatic.

In the centre two panels Fig. 4, we plot the amplitude and phase of a simulated slice of the uv -plane, modelled with smooth spectrum foregrounds and a chromatic Airy beam as described in Section 3.1, as a function of u and v . Overlaid on these panels are a series of large (small) white dotted lines, which represent every hundredth (tenth) frequency and the u mode sampled there by a group of nine baselines oriented in the East–West direction within the HERA array.

In the bottom panel, we construct the 2D power spectrum of the baselines depicted in the central two panels. We form this power spectrum by taking the Fourier transform over the frequency axis of each of the simulated visibilities individually (known as the delay spectrum) over the frequency range 130 – 170 MHz, as indicated by the magenta data points in the second and third panels. While the amplitude of the power spectrum is dominated by smooth spectrum sources, the process of forming the power spectrum through the delay transform effectively mixes rapidly varying uv -modes. Consequently, this mixing results in the contamination of foreground power at lower delays, or τ (the Fourier dual of frequency in the delay regime) extending to higher delays. The extent of the power leakage is contingent on the range of uv -modes sampled by a given baseline, which is dependent on the baseline’s physical length. This baseline-dependent mode-mixing effect is precisely what leads to the widely observed wedge effect that challenges the measurement of the 21 cm signal.

In the top panel, by contrast, we form the 2D power spectrum over the same frequency range, but this time applying a Fourier transform directly to the visibility simulation itself perpendicular to the uv -axis– η -transform. Unlike the delay transform, the η -transform does not introduce mixing between different uv -modes. Instead, it performs a Fourier transform along the frequency axis while keeping the u coordinate fixed. Here, we see that the η -transform isolates beam-weighted foregrounds to many fewer modes than the delay-transform due to keeping the more chromatic axis of this plane u fixed, leaving much of the Fourier space clean to measure the 21 cm signal. Put another way, the sky-beam product can accurately be described by many fewer parameters at fixed \mathbf{u} than at fixed \mathbf{b} for all but the shortest baselines.

While in practice we generally do not have access to all the required modes necessary to perform an η -transform and achieve this level of foreground isolation, this fundamental difference between the delay and η -transforms underscores the potential of redundant sampling of fixed uv -modes. If we can perform repeated sampling of fixed u -modes across frequency, as is done by comparing overlapping

measurements between baselines in the same orientation which redundantly sample the same modes in the uv -plane, one can effectively differentiate between the intrinsic chromatic structure of a baseline that samples a broad range of u -modes and the true chromaticity of the beam-weighted sky. Leveraging the power of spectrally redundant sampling of the uv -plane for precise estimates of the sky chromaticity holds great promise for a number of applications. In particular, it proves to be a powerful tool in calibration applications, solving for antenna gains which cause visibilities to deviate from the expected chromaticity at fixed u -modes, and for improved foreground suppression. We explore both applications later in this paper.

3.3 Constructing a model for spectrally redundant visibilities with DPSS

One important distinction between spatial redundancy and spectral redundancy is that while two or more spatially redundant baselines can be directly compared when calibrating, it is much less straightforward to compare measurements that sample the same modes in the uv -plane with spectral redundancy – a point we highlight in Figs 2 and 4. In traditional spatial redundancy, one expects that two baselines with the same separation vector should be redundant up to the per-antenna, per-frequency complex gains, assuming perfect redundancy in the array (i.e. identical beams and perfect antenna placement). Because each baseline makes the same measurements at the same frequency channels, comparing these measurements to enforce spatial-redundancy between baselines is trivial.

However, comparing measurements in an attempt to enforce the less-than-perfect spectral redundancy between two or more baselines that sample the same uv -mode is not as simple. Simply put, while the true visibility evolves smoothly with frequency at fixed u , it still does evolve and that evolution must be parametrized and modelled. Further, we must deal with the complexity introduced by the fact that baselines in the same orientation that cross a given mode in the uv -plane at different frequencies rarely sample the *exact* same modes. Both challenges must be tackled in order to compare measurements across baselines.

One prior approach to incorporating spectral redundancy into calibration tried to solve this problem by statistically modelling the correlations between baselines in order to solve for the models of the visibilities which were spectrally redundant. In Byrne et al. (2021), the authors introduced a technique for incorporating spectral redundancy into a unified Bayesian framework, citing private communication with some of us during the development of this work. Their technique forward-modelled the cross-frequency covariance between baselines, incorporating information about the sky and beam model into the estimate of the covariance matrices. These covariance matrices were then used in an optimization loop which penalized non-redundant features between pairs of visibilities when solving for a set of spectral redundant visibility models. This technique has the benefit of modelling the redundancy between baselines which may not strictly sample the same modes in the uv -plane, but nearby modes. This approach makes the technique very flexible, but extremely computationally challenging for arrays with wide-bandwidths and many frequency channels to take advantage of spectral redundancy. Additionally, this approach requires accurate knowledge of the sky and primary beam to produce estimates of the cross-baseline frequency–frequency covariance matrices used in the optimization loop.

Ideally, we would like our technique to make assumptions which are informed by realistic properties of the sky and the beam, but not *require* precise knowledge of either to perform accurate

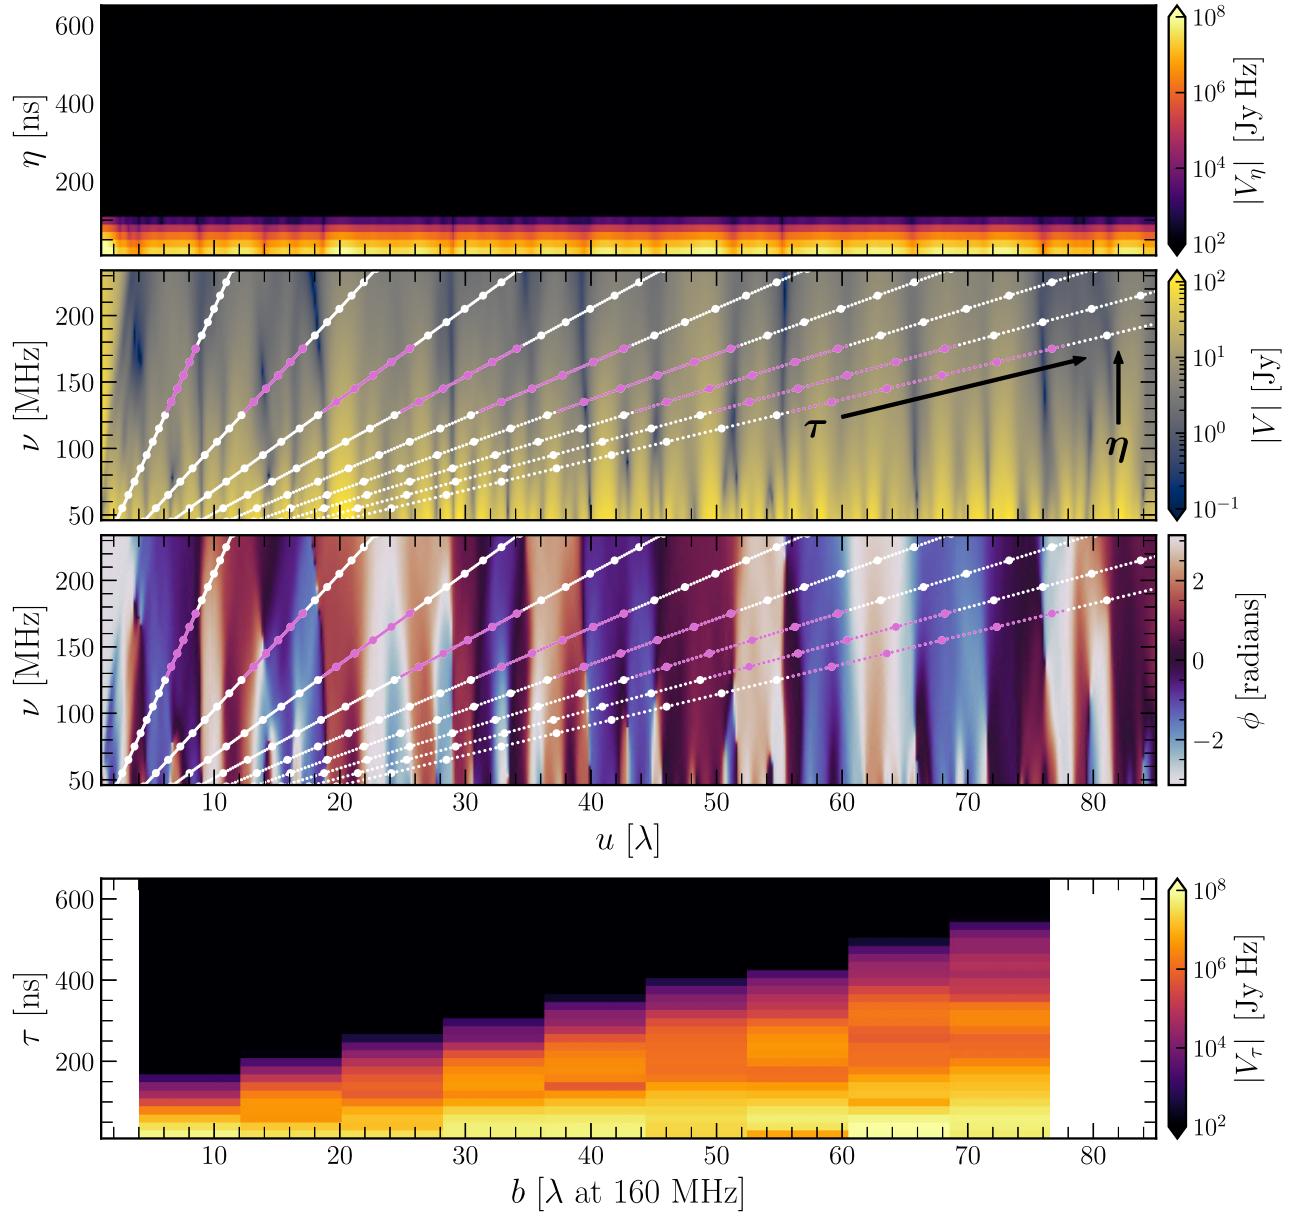


Figure 4. A demonstration of the advantages of spectral redundant sampling of the uv -plane. In the centre two panels, we plot the $u - v$ footprint of a group of nine baselines oriented in the East–West direction within the HERA core in white over a simulation of visibility amplitudes and phases modelled with smooth spectrum foregrounds and an Airy beam, as a function of u and v . Here, the large dots represent visibility measurements for every 100th channel and small dots represent every 10th channel (the true channel width of 122 kHz). All visibility measurements fall on these lines. In the bottom panel, we form the 2D power spectrum of the visibilities from the baselines plotted in the centre two panels by taking the delay transform of each of the simulated visibilities individually over 130 – 170 MHz (as shown by the magenta points between frequencies 130–170 MHz). While the amplitude of the power spectrum is dominated by smooth spectrum sources, forming the power spectrum by taking the delay transform of the visibilities, V_τ , mixes quickly varying uv -modes, which bleeds power, from the intrinsically foreground-dominated lower-order Fourier modes to the higher order Fourier modes, leading to the wedge effect. In the top panel, we form the 2D power spectrum over the same range of frequencies, but instead take the Fourier transform of the visibility simulation perpendicular to the uv -axis, also known as the η -transform, V_η . Here, the η -transform highlights the localization of foreground power to low- η modes which results from the Fourier transform being taken over a fixed angular scale. By comparing the measurements of baselines which redundantly sample the same modes in the uv -plane, one can infer the true frequency-structure of the foreground-beam product, despite the fact that the baselines themselves are highly chromatic as evidenced by the delay transform.

calibration and modelling of the spectrally redundant visibilities. Additionally, we would like to strictly enforce that our modelling procedure explicitly limits the modelled foreground visibilities to the level of spectral structure we expect. To achieve these requirements, we choose to model our spectrally redundant visibilities as a linear combination of a finite number of highly efficient basis

functions known as DPSS or the Slepian basis. These basis functions have the ideal property of having their power being maximally concentrated within a contiguous region of Fourier space with a half-bandwidth of η_w while also being orthonormal, making them a powerful set of basis functions for modelling band-limited signals.

In recent years, the DPSS basis has gained popularity in the 21 cm community as an efficient basis for per-baseline modelling of foregrounds (Ewall-Wice et al. 2021), in-painting missing data in channels flagged for radio frequency interference (RFI; Pagano et al. 2023), and for performing gain calibration (Ewall-Wice et al. 2022). As a mathematical basis, DPSS modes span the space of variations subtended by interferometric measurements of foregrounds while providing a rigorous formulation for limiting the scale of modelled structure – particularly in applications with irregular sampling such as in-painting or modelling flagged data – using pseudo-inverses of an analytically calculable covariance matrix.

The DPSS basis functions are defined with respect to some finite band within which we want to maximally concentrate the power of that function. The spectral concentration ratio for some function φ is the Fourier domain of v, η , is defined as follows:

$$\mu = \int_{-\eta_w}^{\eta_w} |\varphi(\eta)|^2 d\eta / \int_{-\eta_{\max}}^{\eta_{\max}} |\varphi(\eta)|^2 d\eta, \quad (14)$$

where η_w is the half-width in η that we wish to maximally concentrate the power of φ and η_{\max} is the Nyquist limit for the discrete grid of points over which the discrete band in v is defined. By maximizing this quantity, one obtains the first DPSS mode for a given η_w . The full set of DPSS basis functions is formed for this η_{\max} by repeatedly solving for the function that maximizes μ and is also orthogonal to all prior solutions to equation (14). This procedure provides a basis which is complete within the bounds $-\eta_w < \eta < \eta_w$ and also maximally concentrates its power within these bounds. Slepian (1978) showed that this optimization problem could be summarized by solving the eigenvalue problem,

$$\sum_{n=0}^{N_d-1} \varphi_n(N_d, \mathcal{W}) \frac{\sin[2\pi\mathcal{W}(m-n)]}{\pi(m-n)} = \varphi_n(N_d, \mathcal{W}) \lambda_n(N_d, \mathcal{W}), \quad (15)$$

where $\mathcal{W} = \eta_w \Delta v$ is a parameter which sets the filter half-width of the filter in Fourier space, $\varphi_n(v)$ is the n th DPSS function, N_d is the number of discrete points where φ is evaluated, and λ_n is the corresponding eigenvalue of the n th DPSS function.

While the number of functions in a complete set of DPSS basis functions is equal to the number of samples in the basis, N_d , typically only a few basis vectors are actually needed to represent a given band-limited signal to high precision if it is spectrally smooth. Which basis functions are used for modelling a given band-limited signal is typically determined by the eigenvalue spectrum of the DPSS eigenfunctions. The eigenvalue spectrum of DPSS modes has a characteristic step shape where significant eigenvalues have a value near 1 and whose power is primarily concentrated within η_w and insignificant eigenvalues have a value near 0 and have significant power leakage beyond $|\eta| > \eta_w$, with a few eigenvalues with intermediate values ($1 > \lambda > 0$). The approximate number of eigenmodes with near unity eigenvalue values asymptotes to $\sim 2B\mathcal{W}$ when $N \rightarrow \infty$ (Slepian 1978; Karnik, Romberg & Davenport 2020). The shape of this eigenvalues spectrum allows us to select a finite number of basis functions to fit or filter data by taking only those with eigenvalues greater than some eigenvalue cut-off ϵ . We perform this fitting and/or filtering of data presumed to be band-limited using the linear least-squares approach laid out in the DAYENUREST formalism in Ewall-Wice et al. (2021). In this paper, we take $\epsilon = 10^{-12}$.

For this paper, we utilize the DPSS basis to modelling groups of baselines within the same orientation as the outer product of two sets of DPSS vectors in order to model the spatial and the spectral axes.

Namely,

$$V_{i-j}^{\text{model}}(v) = \sum_m \sum_n a_{mn} \varphi_m(\|\mathbf{u}_{ij}\|, \ell_{\max}) \varphi_n(v, \eta_{\max}) \quad (16)$$

where φ_m and φ_n are the sets of basis functions that model the uv -axis and frequency axis respectively and ℓ_w and η_w are parameters which set the Fourier half-width of the uv -axis and frequency axis respectively. By taking the outer product of these set of basis functions, we effectively form a new set of basis functions whose power is concentrated in a rectangular region of 2D Fourier space within the Fourier domains of u and v (ℓ and η respectively). Representing baselines in this way allows us to precisely limit the model to the level of structure that we expect to be present in each axis of a set of spectrally redundant visibilities. For the spatial axis, we restrict the DPSS basis to evolve no faster than $-1 < \ell < 1$ as limited by the horizon (i.e. half-wavelength scales in u). For the spectral axis, we inform of choice of DPSS based on stringent limits set by antenna chromaticity (Ewall-Wice et al. 2016b; Neben et al. 2016; Thyagarajan et al. 2016; Patra et al. 2018; Nunhokee et al. 2020; Fagnoni et al. 2021a, b) and foreground smoothness (Tegmark et al. 2000; Wang et al. 2006; de Oliveira-Costa et al. 2008; Liu & Tegmark 2012; Zheng et al. 2017). For this paper, we take the Fourier half-width of the spectral axis to be $\eta_w = 25$ ns.

We compute the DPSS basis functions along the frequency axis by using the SCIPY implementation (Virtanen et al. 2020), which employs the Lanczos algorithm to solve the eigenvalue problem (equation 15). This approach necessitates that the samples along the axis where the basis functions are generated be uniformly spaced. Given that each baseline within the HERA array sample the same frequencies and the frequency channels are sampled uniformly, this condition is satisfied.

In contrast, the spatial axis lacks a uniformly spaced grid that is consistent across all baselines in the group that would allow us to use the SCIPY implementation. This discrepancy arises from the fact that the spacing of discrete samples along the uv -axis is determined by the length of the baseline (i.e. $\Delta u = \Delta v \mathbf{b}/c$). Therefore, each baseline in the spectrally redundant group samples the uv -axis non-uniformly, which requires that we generate values from the Slepian basis that exactly correspond with the location each baseline samples in the uv -plane. To satisfy this requirement, we instead generate samples from the continuous extension of the DPSS – the prolate spheroidal wave functions.²

We generate these functions by utilizing an implementation that solves the continuous eigenvalue problem described by

$$\int \frac{\sin[2\pi\eta_w(v - v')]}{\pi(v - v')} \varphi_n(v', \eta_w) dv' = \lambda_n \varphi_n(v, \eta_w). \quad (17)$$

This method uses a normalized Legendre polynomial expansion to sample from the prolate spheroidal wave functions at the exact values of u each baseline measures. Further details regarding the process for generating these samples can be found in Moore & Cada (2004).

To demonstrate that this basis and parametrization is sufficiently flexible to model astrophysical foregrounds as observed by a realistic beam, we explore how well we can model the entire u - v plane (similar to Fig. 4). The fit, which we show in Fig. 5 for one of the best-sampled spectrally redundant groups in the array, is constrained

²Note that discrete prolate spheroidal wave functions are discrete sequences sampled from the continuous prolate spheroidal wave functions. If the two sets of functions had the same Fourier half-width support and sampled the same points, the values of the functions would be identical.

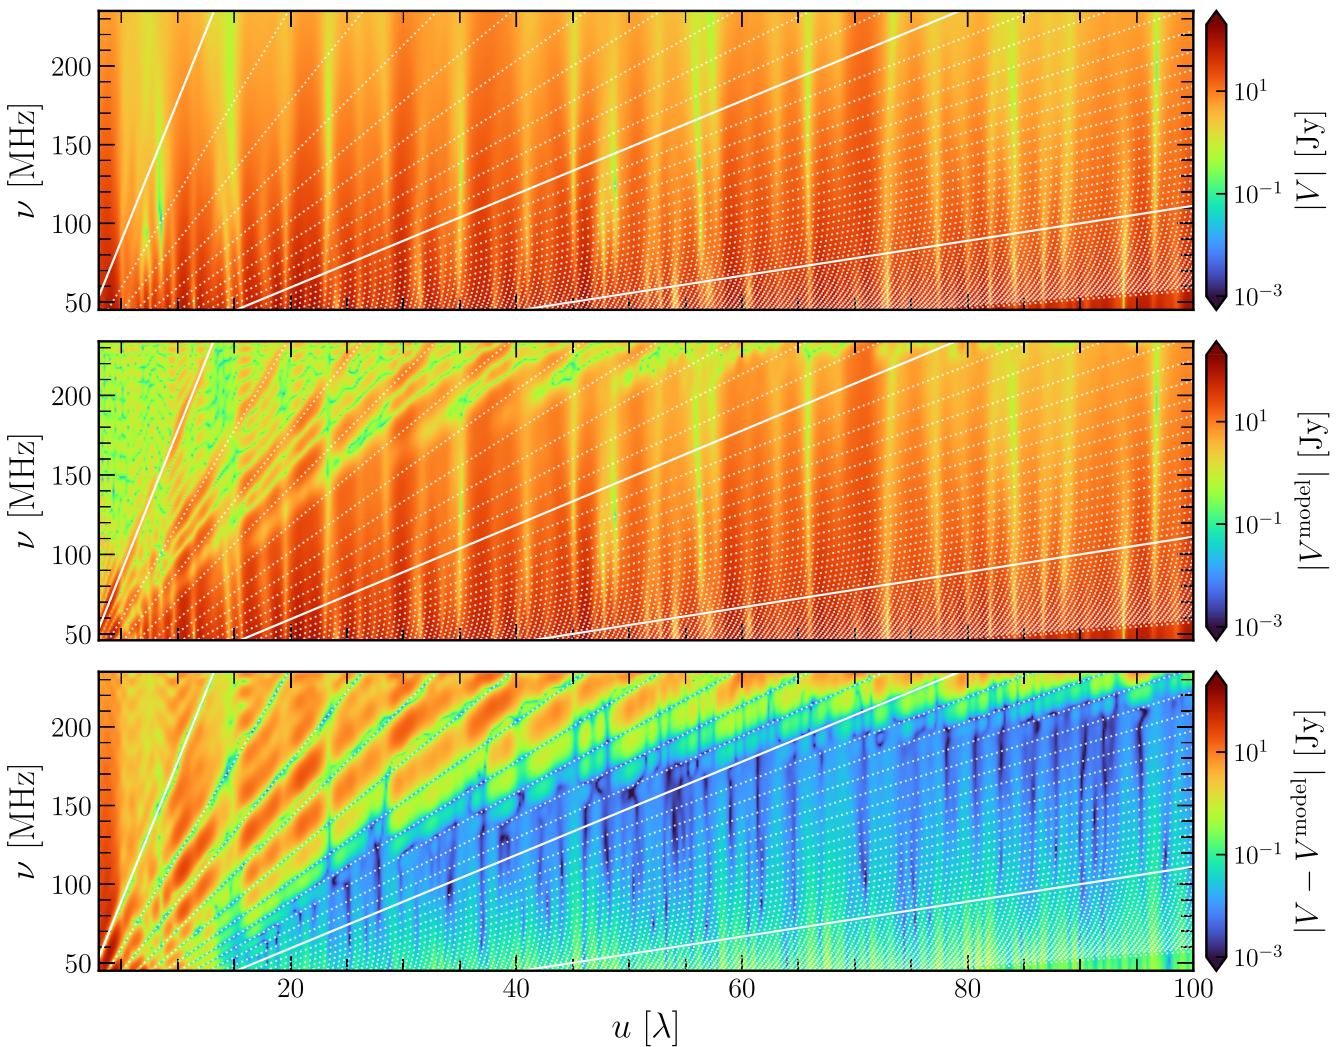


Figure 5. Here, we demonstrate the DPSS model’s ability to accurately fit the u - v plane with a relatively restricted set of basis vectors. In the top panel, we show true visibility without noise, evaluated using equation (9), for one of the best-sampled spectrally redundant groups in the array. This is done not just at the points actually measured by the array (shown in dotted and solid lines), but across the entire plane. We then show in the middle panel our model (equation 16), which is fit using only the discrete set of visibility samples. Finally, we show the magnitude of the difference in the bottom panel. The bold lines indicate the baselines we show the relative error of the model in Fig. 6. When the frequency density of visibility samples is high enough – i.e. where the spectral redundancy is large enough – the model is a good fit and can even accurately interpolate between measured points (blue regions). Where the spectral redundancy is lower, such as in the upper left, the fit is underdetermined and only the actual sampled visibilities are well-modelled. This figure can give us a sense of where we expect foreground modelling and subtraction with spectral redundancy to be effective and where we expect it to be lossy and to remove cosmological signal in the wedge – a point we return to in Section 5.

only by the discretely sampled visibilities in the spectrally redundant group. Despite that, not only are the observable visibilities fit well, but much of the unobserved parts of the plane are fit well too. This strongly implies that with sufficient spectral redundancy, the DPSS model is overdetermined by the visibilities and that extra information could be used to subtract foregrounds within the wedge. Of course, this raises a question of signal loss, which we return to in Section 5.

Regardless, by examining the quality of the fit to the actual measurable visibilities, which we show in Fig. 6, we can see that the model and its limited range of basis vectors is sufficient to model foregrounds to better than 1 part in 10^6 . At least to the level of realism reflected by our simulations, this is more than enough dynamic range to precisely model and remove foregrounds to below the expected EoR level. This is very promising, though there are certainly real-world complexities that any application of this technique would have to address, as we discuss in Section 6.1.

3.4 The robustness of the DPSS model to spectrally localized features

One downside to incorporating cross-frequency, cross-baseline information into a model fit, as is described in the previous subsections, is the potential for that fit to be corrupted by features in the visibilities which are compact in frequency and large in amplitude. These features, which violate the assumptions of spectral redundancy, can lead to model errors which can ripple across the frequency band of one or more baselines. One such spectrally localized feature commonly found in visibility data is RFI.

In Fig. 7, we demonstrate this effect by fitting the DPSS basis described in the previous section to the most spectrally redundant group of visibilities within the HERA array – the same group used in Figs 5 and 6. Prior to performing the fit, we simulate unflagged

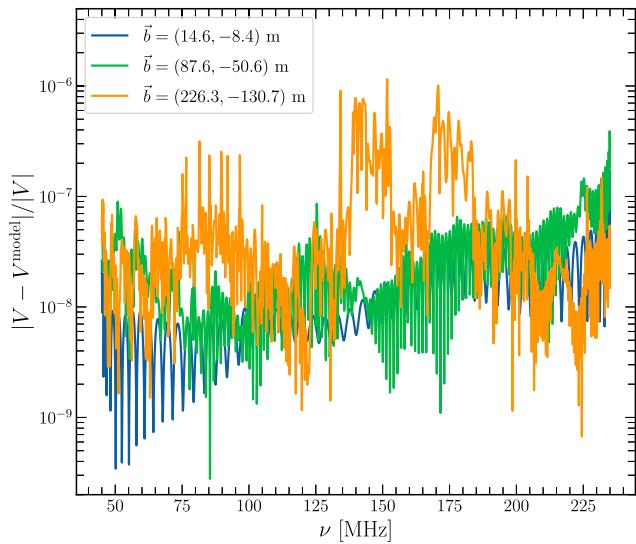


Figure 6. To more quantitatively assess the fidelity of the NUCAL model fit in reproducing the simulated visibilities, we plot the relative error of the model shown in Fig. 5 evaluated for the three baselines highlighted as solid lines. Regardless of the level of spectral redundancy at the u modes probed by these baselines, the model is able to fit a foreground-only simulation to a very high dynamic range – better than 1 part in 10^6 . This is despite the fact that the NUCAL is restricted to only modelling slow frequency evolution at fixed u . This demonstrates the flexibility of the model, though we expect that the level of spectral redundancy will still matter when assessing cosmological signal loss (see Section 5).

RFI in the measurement set by adding a 50 Jy RFI blip to all the simulated visibilities at a single frequency channel near 100 MHz. We then perform a linear-least-squares fit of the DPSS basis to the visibilities with the RFI added and plot the simulated visibilities, the DPSS model, and residuals. From the figure, we observe that fitting the DPSS basis to a set of baselines corrupted by RFI does in-fact lead to errors which impact nearby channels, as can be seen in the residuals. However, thanks to the stringent constraints imposed by spectral redundancy and the DPSS basis, our set of basis function are still able to model and remove foreground emission down to the thermal noise floor over the majority of the band.

While these types of errors are concerning if we wish to produce high-fidelity models of foreground emission over a wide-range of frequencies, there are techniques which can minimize the impact of the spectrally localized features, such as detecting and flagging the worst offending sources of RFI before fitting and implementing procedures which identify RFI-impacted channels during the fitting process. This potential pitfall may also to prove a strength of spectral-redundancy – if certain channels exhibit a particularly strong discrepancy between data and model, that maybe be evidence for narrowband RFI that should be flagged before re-running the fit.

3.5 Spectrally redundant calibration (NUCAL)

For highly spectrally redundant baseline groups (see Fig. 1), the parameters describing the smooth foregrounds in equation (16) are overdetermined. This means that the data can simultaneously constrain the foreground model parameters and nuisance parameters like calibration gains. Just as in redundant-baseline calibration, our goal when we perform that fit is to minimize χ^2 , which represents the noise-weighted sum of the difference between our measured and visibilities modelled with DPSS modes in equation (16). Here, we

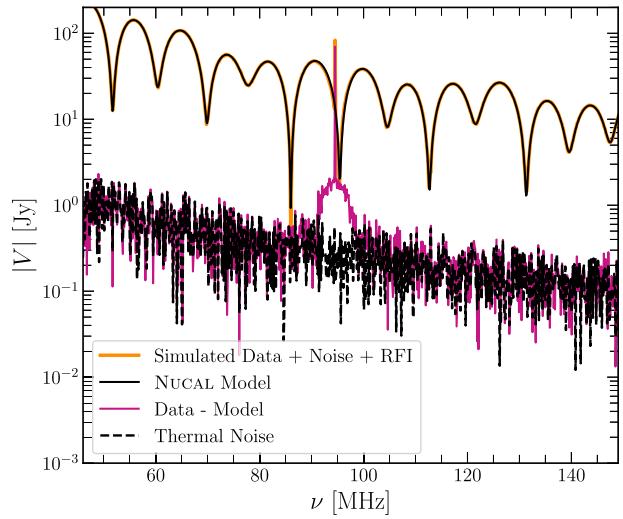


Figure 7. Demonstration of the robustness of the DPSS basis being used to model a group of spectrally redundant baselines to real-world effects. This modelling basis is restricted to modelling structure which is smooth as a function of frequency at a fixed u , and rejects shapes in visibilities that are spectrally compact. To demonstrate this ability, we simulate RFI in one frequency channel across all baselines in the spectrally redundant group (58 baselines with the same orientation), and fit our model to the simulations with noise and RFI added. Leveraging spectral redundancy, we find that the DPSS basis functions are able to produce a model (solid black line) of smooth spectrum foreground emission of one of those baselines (in this case 9.3 m east, 12.0 m south) to the level of the thermal noise (dashed black line) while mostly rejecting the RFI spike, which can be seen in the data-model residuals. The residuals show remarkable agreement with the thermal noise simulated across most of the band, demonstrating the ability of the DPSS basis to model and remove foregrounds. In particular, this technique may be adept at identifying low-amplitude narrowband RFI, which is difficult to locate using traditional methods and has the potential to introduce high-delay structure which can swamp the cosmological signal. A more thorough investigation of that potential is left for future work.

generalize equation (3),

$$\chi^2 = \sum_{ijv} \frac{|V_{ij}^{\text{obs}}(\nu) - G_{ij}(\nu) V_{i-j}^{\text{model}}(\mathbf{u}_{ij}, \nu)|^2}{\sigma_{ij}^2(\nu)}, \quad (18)$$

where we seek to solve for both for our visibility model and antenna gains. Here, σ_{ij}^2 is a weighting function that incorporates thermal noise variance and any flags, which are treated as having infinite variance, and G_{ij} are a set of gain parameters to be solved for in calibration. While in redundant-baseline calibration, V_{i-j}^{sol} is allowed to have arbitrary spectral structure considering that each frequency is solved for independently, NUCAL restricts $V^{\text{model}}(\mathbf{u}_{ij}, \nu)$ to be the sum of a limited number modes at fixed u that vary smoothly in frequency.

By using multiple orientations within the uv -plane which have baselines with significant overlap and modelling each with the DPSS basis, equation (18), the simultaneous solution of $G_{ij}(\nu)$ and $V^{\text{model}}(\mathbf{u}_{ij}, \nu)$ becomes even more overdetermined – depending on the degrees of freedom we assign to our model for $G_{ij}(\nu)$. Restricting the visibility model to a set of relatively smooth DPSS vectors in u and v allows this calibration to minimize spectral structure in calibrated visibilities and attempts to put as much of the observed spectral structure in uncalibrated (or partially calibrated) visibilities into the gains. Where traditional spatially redundant calibration uses internal symmetries in an array to calibrate antenna gains to one

another at each frequency independently, NUCAL uses analogous symmetries in the spectral response to explicitly calibrate the array along the frequency axis. This allows spectrally redundant baselines to be brought into alignment with a physically motivated model for the angular arrangement and chromatic variation of the bright foregrounds that dominate their measurements.

This is one of the primary benefits of spectral redundancy: it provides us a new-axis for enforcing self-consistency between visibilities. Although spectral redundancy does not require that two different baselines precisely measure identical visibility values at the same angular Fourier modes across various frequencies, it does impose stringent constraints on the extent of spectral variations at a constant value of u across baselines of vastly different lengths. Because we anticipate these measurements to exhibit gradual changes with frequency at fixed angular scales, deviations from the anticipated level of spectral smoothness observed in the raw visibilities can often be ascribed to non-smooth signal chain effects. These deviations are subsequently accounted for as direction-independent antenna gain terms during the fitting process.

While NUCAL can in principle be applied to solve for per-antenna gains, for this paper we restrict NUCAL to only work within the degenerate space of spatially redundant baseline calibration. NUCAL is a somewhat natural solution to fixing the degeneracies of redundant-baseline calibration without introducing spectral structure into the degenerate parameters, as many arrays which have leaned heavily on spatial redundancy, such as HERA and PAPER, are also capable of being calibrated using spectral redundancy. The procedure for solving for these degeneracies is quite straightforward within the framework of NUCAL. We restrict $G_{ij}(\nu)$ to take the form

$$G_{ij}(\nu) = A(\nu) \exp [i\Phi(\nu) \cdot \mathbf{b}_{ij}], \quad (19)$$

where A and Φ are the frequency-dependent amplitude and phase tip-tilts degeneracies described in Section 2.3. In this way, NUCAL replaces (or at least refines) the absolute calibration step that must follow after redundant-baseline calibration. More importantly, it does so in a way that optimizes for the spectral smoothness of the calibrated visibilities and thus reduces the impact of gain errors due to sky-model incompleteness.

Because NUCAL relies on fitting many spectrally redundant groups simultaneously to constrain these array-wide degeneracy parameters, and because adding gains makes the fit no longer linear, the computational challenge is dramatically increased compared to fitting a DPSS model to a single spectrally redundant group. Therefore, we implement the calibration routine described above as a first-order gradient descent optimization in JAX and OPTAX (Bradbury et al. 2018) to take advantage of the built-in autodifferentiation, just-in-time compilation, XLA acceleration, and library of popular optimizers. As written, the gradient descent algorithm makes the problem of calibrating every frequency channel simultaneously tractable, but only converges to the local minima closest to our initial guess of the model parameters. Therefore, NUCAL is best applied after an initial sky-based absolute calibration has already been performed.

It is worth noting that this technique bears several similarities to the recently proposed calibration technique, CALAMITY (Ewall-Wice et al. 2022), which also utilizes the DPSS basis to model foregrounds and solves for antenna gains by minimizing a loss function using first-order gradient, but does so in a way that does not incorporate spectrally redundant information into the fit. For a review of the similarities and differences between these two approaches, see Appendix A.

3.5.1 NUCAL degenerate parameters

Like spatially redundant calibration, spectrally redundant calibration is not a full solution of all the calibration degrees of freedom. Given that NUCAL calibrates on the basis of internal consistency between baselines that redundantly sample the same modes in the uv -plane at different frequencies, there are a set of transformations that can be applied to either the gains or visibility models that ultimately leave the value of χ^2 in equation (18) unchanged. For example, one can take $A \rightarrow 2A$ from equation (19) and $a_{mn} \rightarrow \frac{1}{2}a_{mn}$ from equation (16) without modifying the product $G_{ij}V_{i-j}^{\text{model}}$ in equation (18). The precise number and shape of these degeneracies depends on one's array layout and DPSS parametrization.

However, because we enforce that the visibility model is spectrally smooth at a fixed uv -mode across baselines in a spectrally redundant group, the number of NUCAL degeneracies must be dramatically reduced from the number of degeneracies that must be fixed by traditional absolute calibration after spatially redundant calibration. There were at least three physically important numbers per frequency per polarization. Here, we expect only a handful of such numbers over the whole band. This largely eliminates the need for sky-based absolute calibration to faithfully calibrate fine-scale spectral structure, and thus reduces the impact sky or beam modelling errors can have on the final calibration solution.

In practice, one can use NUCAL by starting with a best-attempt at sky-based absolute calibration and letting the degeneracies simply be left unmodified by the gradient descent solver, which only takes steps in non-degenerate directions. Because the visibility model is restricted to one consistent with smooth spectrum foregrounds, the dynamic range of foreground subtraction (see Section 5) is not affected by absolute calibration errors in the degenerate subspace of NUCAL. This of course assumes that perfectly calibrated foregrounds are well-captured by the parametrized NUCAL model. While one can still get errors in the degenerate subspace that affect the magnitude of the recovered EoR signal as a function of redshift, these sorts of errors do not risk mixing the foreground signal into uncontaminated modes, as is the case in the standard direction-independent calibration problem.

4 DEMONSTRATION OF SPECTRALLY REDUNDANT CALIBRATION (NUCAL) ON SIMULATED DATA

In this section, we test the performance of NUCAL on a set of simulated visibilities which contain realistic spectral variation and calibration errors. We begin with a description of the gains applied to the simulated data to produce a set of ‘uncalibrated’ data.

4.1 Setting up the problem

For the simulations used in this section, we assume that the data have already been redundantly calibrated and that the only remaining calibration required is the removal of the degeneracies of spatially redundant calibration. This is done to reduce the computational cost of performing a per-antenna gain calibration, but is still a scenario of practical relevance given that the arrays designed to perform spatial redundant calibration also generally have some degree of spectral redundancy (e.g. PAPER and HERA). Here, we simulate the need for such a calibration by moving our model visibilities within the degenerate parameter space of redundant-baseline calibration. To highlight NUCAL’s ability to accurately account for arbitrary amounts of spectral structure in an instrument bandpass, we simulate the

Table 1. Visibility simulation and calibration details.

| Parameter | Value |
|---|--|
| Instrument layout | HERA-350 (Hexagonal Core and Outrigger Antennas; Fig. 1) |
| Frequency range | 46.9–234.4 MHz |
| Frequency resolution ($\Delta\nu$) | 122 kHz |
| Sky model | 10^5 Point Sources (equation 6) + Flat-Spectrum Power Spectrum |
| Beam model | Airy Beam (equation 8) |
| Spatial axis DPSS half-width parameter | 0.5λ |
| Frequency axis DPSS half-width parameter (η_w) | 25 ns |
| Minimum number of unique baselines per group | 10 |
| Model constraints (data points/parameters) | 11.94 |

degenerate amplitude (equation 4) and tip-tilt parameters (equation 5) to vary rapidly from channel-to-channel by drawing the amplitude from a normal distribution with a mean of 1 and standard deviation of 0.01 and the tip-tilt degeneracy from a normal distribution with a mean of 0 radians per meter and standard deviation of 0.01 radians per metre for both the north–south and east–west parameters.

We then use NUCAL to calibrate this corrupted data set by selecting baselines within orientations which have greater than 10 unique baseline types to participate in the estimation of the gains. For each unique orientation, we construct a foreground model, allowing each set of foreground components to vary independently of all other orientations included in the calibration and restrict the set of basis functions modelling the spatial axis to only model out $u = 100\lambda$. We then execute the first-order gradient descent to refine our initial estimates of the foreground model components and the redundant-baseline degeneracies. The final product of the NUCAL implementation produces a DPSS-based model of the sky, V_{i-j}^{model} , a set of calibrated visibilities, V_{ij} , and an estimate of the spatially redundant degeneracies, $A(\nu)$ and $\Phi(\nu)$. Since our primary goal is to test the dynamic range of this calibration procedure, our simulations are noise-free. See Section 3.1 for details on visibility simulations and Table 1 for a summary of simulation and calibration parameters.

4.2 Demonstration of NUCAL’s dynamic range in Fourier space

After NUCAL, we perform a delay transform, in which we substitute a Fourier transform along the frequency axis of a visibility for a line-of-sight Fourier transform (Parsons et al. 2012b; Liu et al. 2014a),

$$\tilde{V}(\tau) = \int d\nu V(\nu) W(\nu) e^{-2\pi i \tau \nu} \quad (20)$$

where $W(\nu)$ is a frequency taper function applied to the data prior to Fourier transforming, which we take to be a Blackman–Harris function (Harris 1978). We perform this delay-transform on both the NUCAL-calibrated visibilities, and the uncalibrated visibilities and the visibilities of the foregrounds and EoR separately.

In Fig. 8, we inspect the delay-transformed visibilities associated with three baselines that belong to the same spectrally redundant group. The primary goal of this analysis is to assess the effectiveness of NUCAL in mitigating spurious spectral structure imparted by a highly chromatic bandpass. For each of the baselines plotted, we find that NUCAL is effectively able to solve for and thus calibrate out the redundant-calibration degenerate parameters. More importantly, NUCAL is able to converge to a solution of the degeneracies at a precision necessary to recover the 21 cm signal outside the wedge for each of the baselines plotted. This demonstrates NUCAL’s utility for dramatically reducing the complexity of absolute calibration after redundant-baseline calibration and for isolating the result from

sky and beam modelling errors that would impart spurious spectral structure to contaminate the EoR window.

5 SPECTRAL REDUNDANCY AS A METHOD OF FOREGROUND MODELLING AND SUBTRACTION

Potentially the most interesting application of spectral redundancy is the production of incredibly accurate sky models. The model of the sky that we fit during spectrally redundant calibration, can also be used as a spatirospectral filter of beam-weighted smooth-spectrum sky emission. When the sampling of uv modes is sufficiently dense along a given orientation, this model constrains foreground emission at a given frequency using information from a wide range of other frequencies – much wider than the coherence scale of cosmological 21 cm emission. This means that, for certain orientations and frequency ranges, this technique has the ability to remove foregrounds while retaining 21 cm emission within the wedge, even with a realistic level of sky and beam chromaticity. Access to these modes would greatly increase the sensitivity of 21 cm experiments as a significant number of SNR modes lost to foreground contamination, particularly interferometers with many long baselines which must sacrifice more sensitivity than compact arrays to estimate the delay spectrum using only modes outside the wedge (the so-called foreground avoidance strategy).

In this section, we attempt to use spectral redundancy to excise foregrounds from a set of simulated visibilities while still leaving modes within the wedge. For ease of analysis, we assume that our simulated visibilities are perfectly calibrated so that the effectiveness of this basis for removing foregrounds while leaving 21 cm signal recoverable can be evaluated without the complication of instrumental systematics. In order to examine our ability to recover cosmological signal within the wedge for existing instruments, use the single most spectrally redundant group of visibilities within the HERA array. With this set of uncorrupted simulated visibilities, we use the set of DPSS modelling vectors to perform a least-squares fit to the simulated data allowing the set of filters modelling the u -axis the freedom to model structure that varies at scales of half a wavelength and the set of filters which model the spectral axis to model variation with a Fourier half-width in η of $\eta_{\text{max}} = 25$ ns as was done in Sections 3 and 3.5.

After computing this model, we form an estimate of the power spectrum by forming delay-spectra for each baseline within the spectrally redundant group for each of the simulated data products (simulated 21 cm signal, foregrounds, 21 cm signal and foregrounds, and the visibilities filtered by subtracting the NUCAL) using the delay-transform where the visibilities are multiplied by a Blackman–Harris window function (equation 20). We take the delay-spectrum over the

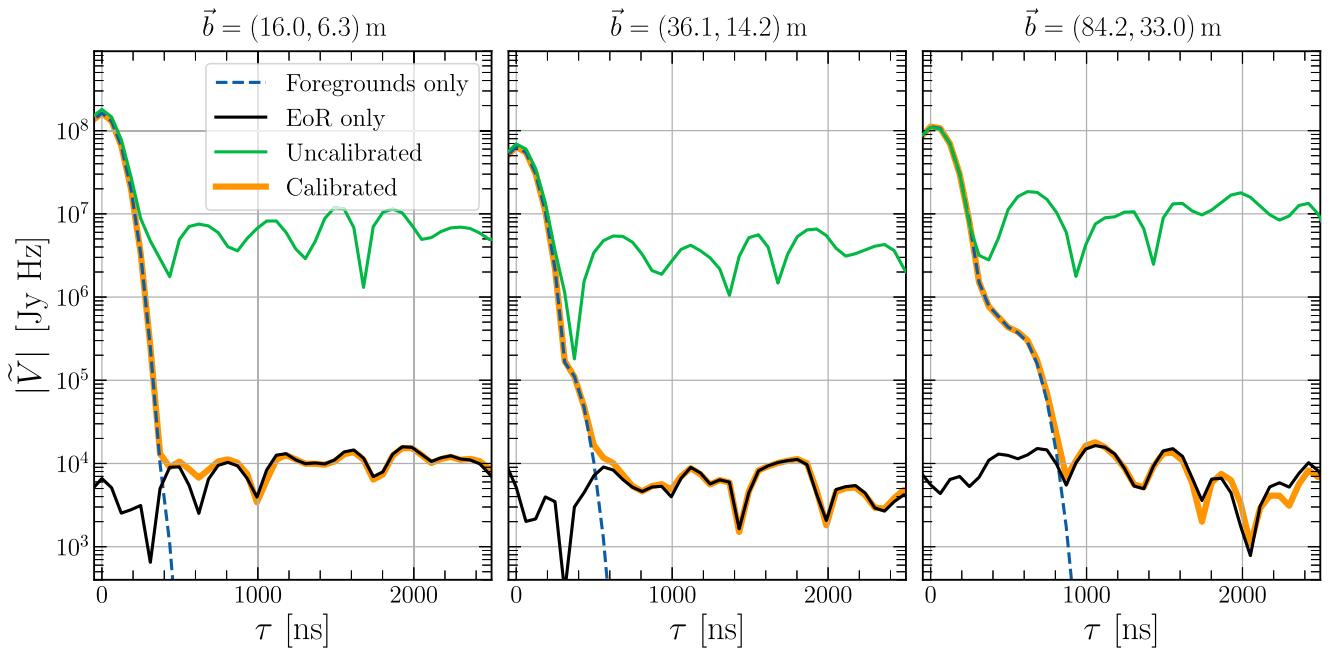


Figure 8. Demonstration of NUCAL’s ability to calibrate out arbitrary spectral structure introduced by an instrument for three different baselines within the same spectrally redundant group. Plotted in green is a noise-free but intentionally uncalibrated simulated visibility, with gains only in the degenerate subspace of redundant-baseline calibration. We also plot the power spectrum of noise-free simulated foreground visibilities (blue dashed), the EoR power spectrum (black), and the visibility after spectrally redundant calibration (orange). We find that NUCAL, by restricting the level of spectral structure allowed in the calibrated visibility, correctly removes the spatially redundant degenerate parameters in our simulated visibilities, leaving calibrated visibilities consistent with the EoR power spectrum outside the wedge.

frequency band, 60–70 MHz, where more baselines in the group have more significant spectrally redundancy. It should be noted that while the power spectrum is estimated over a small sub-band, the NUCAL model is fit over the full frequency range of 46.9–234.4 MHz. We then estimate the 2D power spectrum by squaring the delay spectra and scaling to cosmological power spectrum units using,

$$\hat{P}_{ij}(\mathbf{k}_\perp, k_\parallel) = \frac{X^2 Y}{\Omega_{pp} B} \left| \tilde{V}_{ij}(\tau) \right|^2, \quad (21)$$

(Parsons et al. 2014) where \hat{P}_{ij} is our estimate of the power spectrum as measured by baseline, b_{ij} , Ω_{pp} is the solid angle integral of the primary beam squared, X and Y are linear scaling factors which map comoving distance to angular separation and radial comoving distances to the frequency interval $\Delta\nu$,

$$X(z) = \frac{c(1+z)^2}{\nu_{21} H(z)} \quad (22)$$

$$Y(z) = \frac{c}{H_0 \nu_{21}} \frac{(1+z)^2}{E(z)} \quad (23)$$

and B is the bandwidth over which the delay transform was taken (Furlanetto et al. 2006; Parsons et al. 2012b). We also transform our baseline and delay coordinates into cosmological spatial modes perpendicular, $\mathbf{k}_\perp = 2\pi \mathbf{b}_{ij}/X$, and parallel to the line of sight, $k_\parallel = 2\pi \tau/Y$ (Parsons et al. 2012a). For all results shown, we adopt a Λ CDM cosmology derived from the Planck 2018 analysis when computing power spectra and cosmological coordinates (Planck Collaboration VI 2020).

In Fig. 9, we display the results of our foreground filtered visibilities as 2D power spectra for 21 cm signal, foregrounds, foregrounds and 21 cm and the filtered foregrounds and 21 cm signal. We find that NUCAL is able to substantially suppress foreground modes within

the wedge, while revealing relatively clean EoR signal. We find substantial signal loss – i.e. incorporation of EoR fluctuations into the NUCAL model, which is then subtracted off – at low k_\parallel and k_\perp . This is clearest for short baselines, which have the least spectral redundancy (see Fig. 4). This lack of spectral redundancy leads to degeneracies between the foreground filters and 21 cm signal that leads to an oversubtraction of cosmological signal within the wedge. Note that this result is for the single best-sampled spectrally redundant group with 58 unique baselines, all oriented in the same direction (see Fig. 1); we should expect less effective subtraction for other groups in HERA. Many orientations within HERA have too few baselines to fit the smooth foreground degrees of freedom in equation (16). Of the 61 075 total single-polarization baselines in HERA-350, which sample 6610 unique baseline vectors, 26.8 per cent belong to baseline groups with 10 or more unique baselines – enough to be useful for NUCAL (see Section 3.5). 13.5 per cent belong to highly spectrally redundant baseline groups with 25 or more unique baselines.

In Fig. 10, we quantify the level of signal loss from overfitting during foreground subtraction by comparing the filtered power spectrum to the input EoR power spectrum. We find that some amount of signal is lost throughout the wedge and slightly beyond the horizon, due to the small amount of chromaticity given the DPSS basis modelling the frequency axis. Across the range of k_\perp modes which are well-sampled, we find a signal loss value of ~ 20 per cent. This is a non-negligible amount of signal, to be sure, but it is certainly preferable to losing 100 per cent of the EoR signal inside the wedge that one has to accept with a strategy of foreground avoidance.

Such signal loss is not unexpected; the NUCAL model fits any power at low η regardless of its astrophysical origin. In a delay spectrum, that power gets spread across τ modes, which explains why we see signal loss throughout and even a bit beyond the wedge. However, the

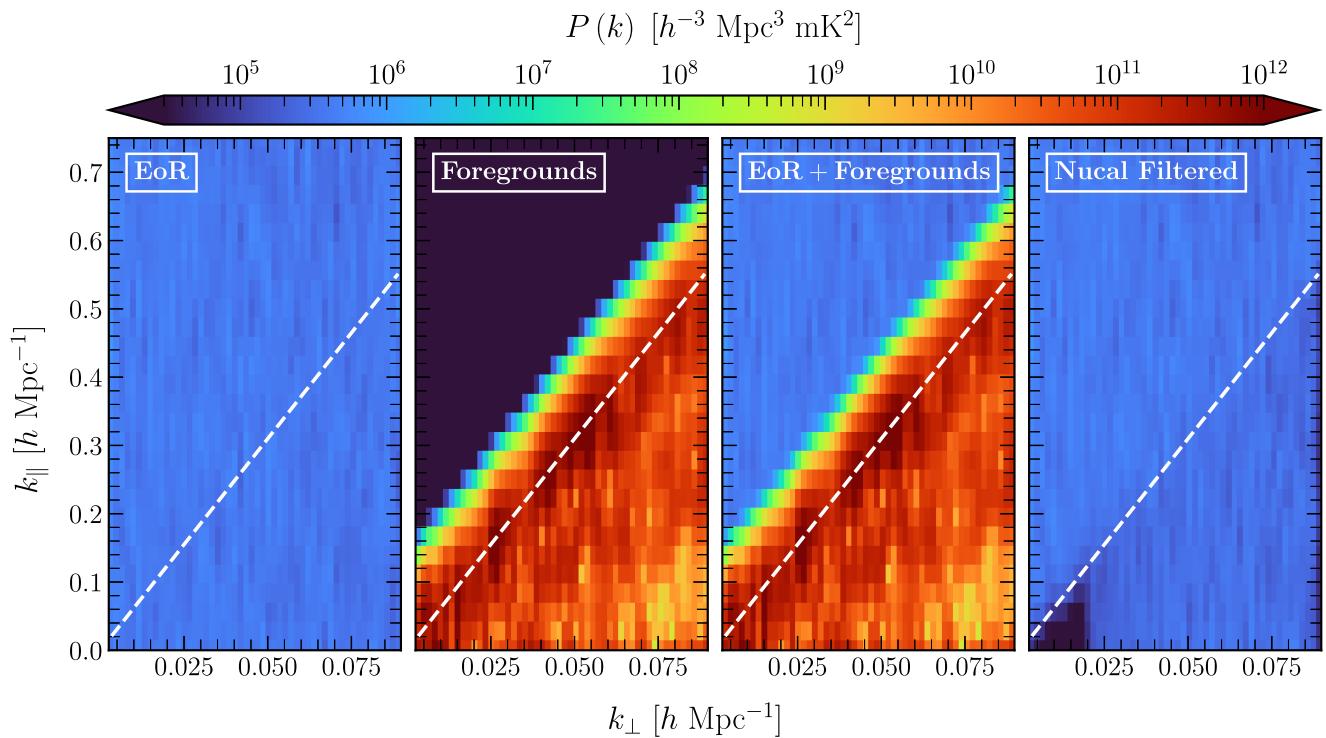


Figure 9. Demonstration of the foreground filtering capabilities of spectral redundancy on a set of simulated visibilities from the most spectrally redundant group in the HERA array. Fitting a NUCAL model suppresses foregrounds within the wedge, but is able to retain 21 cm signal in simulated, noise-free simulations. By fitting versus u and v , NUCAL is able to remove foreground emission from a point-source model with varying spectral indices and a realistic level of beam chromaticity. In the first three subplots, we show the 2D power spectra of EoR, foregrounds, and the sum of both components. In the right-most panel, we plot the 2D power spectrum of the full simulated data (third panel) but with the NUCAL model subtracted. By subtracting off our best-fitting model, we are able to substantially suppress foregrounds across a wide range of k_{\perp} modes while retaining 21 cm signal. Signal loss from overfitting appears in the last panel at low k_{\parallel} and k_{\perp} where spectral redundancy is minimal and 21 cm modes are degenerate with NUCAL foreground filters. We examine that loss more directly in Fig. 10.

part of the power of this technique is the ability to use information from many other frequencies to better estimate (and thus remove) low η structure, which is overwhelmingly foreground-dominated. While this technique suppresses the 21 cm signal to a larger extent than other techniques (Mertens, Ghosh & Koopmans 2018; Mevius et al. 2022), the benefit of this approach is that it makes very few assumptions about the foregrounds and primary beam while still modelling them to good precision.

Of course, to use this technique on real data, one would have to precisely and systematically model and account for the signal loss in observed spectrally redundant groups – a task beyond the scope of this work. Regardless, this level of foreground subtraction is a promising step towards mitigating the impact of foregrounds in certain array-types while retaining 21 cm signal within the wedge.

6 DISCUSSION

In the previous sections, we explored the promise of utilizing spectral redundancy as both a method of performing accurate direction-independent gain calibration and foreground subtraction. Here, we discuss both the real-world limitations of and potential future extensions to the technique.

6.1 Real-world challenges of leveraging spectral redundancy

This work remains an incomplete study of what needs to be done to evaluate the ability of NUCAL to calibrate real data and to

remove foregrounds and recover 21 cm signal within the wedge more generally. We have not fully explored important effects such as proper fixing of NUCAL’s degenerate parameters, which requires reference to a sky-model, which are always incomplete. In the event that the sky-model significantly deviates from the observed visibility, improper degeneracy removal may impact the efficacy of foreground subtraction and complicate the recovery of cosmological modes within the wedge. Future work is needed to explore our ability to subtract foregrounds in the presence of sky-model incompleteness.

As with redundant calibration, spectrally redundant calibration requires that the frequency-dependent uv -plane be self-consistent between measurements made by different baselines. Several known effects can lead to deviations between baselines, which lead to non-redundancies in spectrally redundant calibration. These non-redundancies arise as a result of beam non-redundancies and antenna position errors, mutual coupling between elements in the array, and polarized emission.

The first is the most straightforward, as they are also non-redundancies associated with spatially redundant-baseline calibration: all antennas in the array must have the same primary beam response and be placed accurately. Spectral redundancy requires that beams are uniform across the array, otherwise baselines assumed to be spectrally redundant will measure different uv -planes, considering that the measured uv -plane is a convolution of the true uv -plane and the antenna beam. Attempting to jointly model baselines with NUCAL that include antennas with non-redundant beams would likely introduce calibration errors into the antenna gains, complicating the

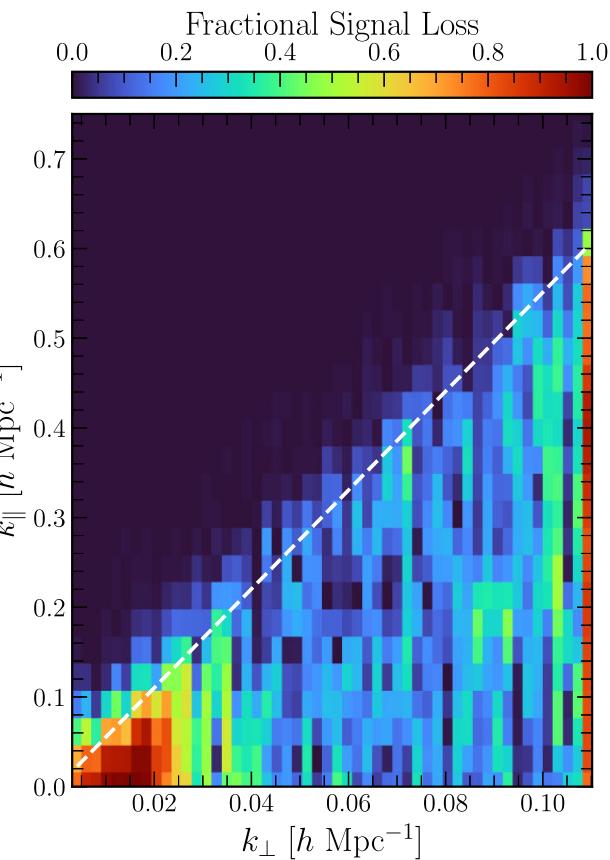


Figure 10. Fractional signal loss computed following foreground filtering, shown in Fig. 9. At low k_{\perp} and low k_{\parallel} baselines which contribute to these measurements of the power spectra have poor spectral-redundancy, leading to oversubtraction of the 21 cm signal. At intermediate k_{\perp} scales, more baselines overlap a given uv -mode, leading to better differentiation between foregrounds and 21 cm signal and reduced (~ 20 per cent) signal loss within the foreground wedge. Signal loss is largely negligible in the EoR window, given that the DPSS modelling basis is restricted to producing wedge-like shapes in delay space. There is some signal beyond the horizon given by the dotted white line due to either the Blackman–Harris taper used or due to the small amount of chromaticity allowed by the DPSS basis functions used to model the frequency axis.

recovery of 21 cm signal in the EoR window. Deviations in both the primary beam size and beam pointing centre have been shown to introduce spectral errors in redundant calibration (Orosz et al. 2019), and will likely be a point of concern when applying NUCAL to data. Kim et al. (2023) show that beam errors can be mitigated in traditional redundant calibration by performing fringe rate filtering the data before calibrating, and could potentially help mitigate similar errors in spectrally redundant calibration.

Similarly, antenna placement errors such as East–West/North–South positional errors or height offsets between elements may prove to introduce spurious structure in calibrated visibilities if one assumes that baselines sample the same modes in the uv -plane actually sample slightly different modes. Orosz et al. (2019) showed that gain errors introduced by antenna placement errors can be mitigated by down-weighting long baselines during calibration. However, the models produced by NUCAL are crucially informed by long baselines, and therefore down-weighting them may prove to be less effective for spectral redundancy than spatial redundancy. More work is needed

to understand how to best handle antenna placement errors within the NUCAL framework.

In addition to beam non-redundancies and antenna placement errors, leakage of long baselines’ visibilities into less chromatic short baselines due to mutual coupling between array elements is likely to pose a challenge to the convergence of NUCAL in its current form. Given that these systematics extend significantly beyond the wedge, they are unable to be modelled with the set of DPSS basis functions used to model beam-sky product, and are unable to be removed via per-antenna gain calibration (Kern et al. 2019; Josaitis et al. 2022; Charles et al. 2023; HERA Collaboration 2023). The inability to incorporate these effects into models of the foregrounds or gains will lead to chromatic gain errors if unaccounted for. However, one could imagine that NUCAL could be modified to solve for first-order coupling coefficients, such as those presented in the framework developed in Josaitis et al. (2022), in addition to the per-antenna gains, potentially reducing the impact of mutual coupling effects on the estimates of gain parameters.

The assumption that the foreground-beam product is spectrally compact may be violated by Faraday-rotated diffuse emission, particularly near the Galactic Centre, where the effect is most prominent (Taylor, Stil & Sunstrum 2009). Spectral structure introduced by this polarized emission will cause variations at a fixed u that we will be unable to model using our smooth set of basis functions, likely leading to chromatic errors in the gain estimates if that are not properly accounted for. One approach to calibrating visibilities which we expect to be highly polarized, is to require that the foregrounds in pseudo-Stokes I be smooth as a function of frequency rather than assuming that either of the linear polarizations are spectrally redundant. This formulation should help to reduce the impact of Faraday-rotated polarized emission on the estimates of the gains.

In addition to the non-idealities of real-world instruments and calibration, we did not provide a full study of the factors affecting signal loss when subtracting NUCAL modelled foregrounds from simulated visibilities. Quantifying the precise level of signal loss is crucial when performing this kind of analysis, as the inclusion of wedge modes with substantial signal loss in an estimate of the power spectrum will artificially lower the amplitude of the estimated power spectrum (Aguirre et al. 2022). The exact amount of signal loss in a given wedge mode will depend on the precise set of basis vectors fit with equation (16). It will also depend on the spectral smoothness of the instrument’s primary beam response, the true spectral smoothness of the foregrounds, and the degree to which a given mode contributing to an estimate of power spectrum is sampled redundantly in frequency. These effects are fundamental to the array on which NUCAL is being performed, and entirely independent of sky-model incompleteness and non-redundancy. The exploration of how variations in dish size, beam chromaticity, and antenna layout affect the level of signal loss within the wedge in 21 cm measurements is also left for future work.

6.2 Future applications of NUCAL

While we have shown the benefits of spectral redundancy and spatirospectral filtering for calibration and foregrounds subtraction, there are many more improvements of this technique and future applications which could make it more widely useful to the 21 cm community.

Perhaps the most conceptually straightforward improvement to this technique is extending NUCAL to solve for per-antenna gains. There are two potential ways in which this could be done. The first is by simply extending equation (18) to incorporate per-antenna

gains, instead of restricting the gain terms to the degenerate space of redundant-baseline calibration. While this approach is mathematically straightforward, extension to full per-antenna gains is extremely computationally challenging due to the massive increase in data volume and number of parameters required to solve for. Utilization of the JAX environment to efficiently compute gradients and its ability to use GPUs to perform would likely help, but it is unclear if the approach would be computationally tractable for such as HERA without major improvements to the algorithm itself. Another complication of extending NUCAL to per-antenna gains is the potential risk of introducing signal loss beyond the wedge. Because NUCAL attempts to produce the smoothest possible calibrated visibilities, extending calibration to incorporate more degrees of freedom may lead to per-antenna gains overfitting 21 cm signal, especially beyond the foreground wedge. More work is needed to understand how extending NUCAL to incorporate more calibration degrees of freedom introduces the signal loss outside the wedge.

The second approach is slightly more limited, but more computationally efficient. Instead of solving for per-antennas gains and NUCAL model parameters simultaneously, we could instead fit a NUCAL-based model to data which have already been calibrated and then use that model as a sky-model to recompute per-antenna gains. Because the NUCAL is highly constrained by the redundancy of an array and physically motivated to pick out foreground emission from sets of spectrally redundant baselines, it could be an excellent foreground model for sky-based calibration. One potential downside of this approach is that if the NUCAL model is derived from visibilities with significant errors due to prior miscalibration and there is a lack of highly spectrally redundant baselines within the array, this approach may be unlikely to fix those errors. However, both approaches are interesting avenues for further exploration and warrant future work.

As presented, NUCAL is currently limited to modelling foregrounds in baselines which are part of groups that redundantly sample the same spatial Fourier modes at different frequencies. While some arrays do take advantage of redundant baseline measurements and therefore often redundantly sample the same uv -modes along radial spokes of the uv plane, arrays which are not radially redundant will be unable to take advantage of the benefits of NUCAL as currently formulated. However, given the flexibility of our modelling basis functions, NUCAL's modelling equation (16) could in principle be modified such that the entire uv plane be modelled jointly in a higher dimensional representation of these DPSS eigenfunctions.

Doing so would significantly increase the number of free parameters incorporated into the foreground model, but will improve the generality of the approach to allow for the calibration of array types which do not explicitly sample the same u modes. Additionally, this approach may improve the ability for non-radially redundant arrays to better model and subtract foregrounds by empirically estimating the cross-frequency covariance between baselines which samples nearby modes in the uv -plane without explicitly simulating foregrounds with a model of the sky and instrument primary beam. This approach to using modelling the entire frequency-dependent uv -plane may also offer a reduction in signal loss incurred by subtracting foregrounds with a NUCAL model. Such an exploration and its implementation is outside the scope of this paper and left for future work.

Given the extreme dynamic range challenge 21 cm cosmology presents, 21 cm arrays must be co-designed to allow for an analysis which best allows for the separation of foregrounds from 21 cm signal. Next-generation 21 cm experiments will need to restore access to modes within the wedge in order to directly image cosmological hydrogen. Such three-dimensional images of neutral hydrogen support cross-correlation studies with other probes of large-scale

structure, providing verification of a detection of the 21 cm and unlocking stronger constraints on the physics of reionization (Beardsley *et al.* 2015; Kovetz *et al.* 2017; La Plante *et al.* 2023). Just as HERA's highly regular array configuration is designed to prioritize spatial redundancy (Dillon & Parsons 2016; DeBoer *et al.* 2017), future array designs may be optimized for both spatial and spectral redundancy, allowing for calibration and foreground modelling with NUCAL and potentially enabling for the recovery of cosmological modes within the foreground wedge that could be used for imaging. Future work is needed to investigate how new array configurations and element design can be leveraged to maximize the effectiveness of this technique.

7 SUMMARY

In this paper, we introduced a novel technique for precisely calibrating interferometers for 21 cm cosmology known as *spectrally redundant calibration*, or NUCAL for short. Inspired by traditional spatially redundant calibration, our approach solves for antenna gains by exploiting the correlations between baselines which redundantly sample the same angular Fourier modes in the uv -plane at different frequencies. Leveraging the fact that baselines which sample the same angular scales see highly correlated beam-sky products at different frequencies, we model this redundantly sampled $u - v$ plane as a linear combination of a highly efficient set of smooth basis vectors known as DPSS.

We demonstrate in Section 3.5 our ability to calibrate out arbitrary spectral structure in redundant calibration degeneracies in a set of simulated HERA visibilities and show that NUCAL explicitly preserves the EoR window by preventing fine-scale calibration errors from coupling with foreground power. Because this approach relies on self-consistency between baselines at different frequencies, we do not need detailed models of the sky and beam to perform calibration while ensuring that our calibrated visibilities remain spectrally smooth.

While this technique does require reference to a sky model to solve for a set of additional parameters which are degenerate with our empirically motivated foreground model, the number of degenerate parameters that are required to be solved for are substantially reduced compared to that of traditional redundant calibration and are limited to shapes residing within the foreground wedge. Assuming that the algorithm reaches convergence to the global minima, this significantly reduces the amount of spurious spectral structure that can be introduced via calibration error due to sky model incompleteness.

One of the most exciting aspects of this work is NUCAL's ability to precisely model and subtract foregrounds from baselines, which redundantly sample the same modes in the uv -plane. As shown in Section 5, combining spectral redundancy with spatiосpectral filtering is able to accurately model the beam-weighted sky even in case in which realistic chromatic beams and foregrounds are used. Assuming that the uv -sampling is dense enough to break the degeneracy between variation in the spatial and spectral axes, modelling foregrounds baselines in the same radial heading using the DPSS basis shows promise for recovering 21 cm signal within the wedge.

One must be cautious when applying this technique to real data, as overfitting can lead to cosmological signal loss, especially when a given uv -mode is not sufficiently redundantly sampled by the baselines in its unique orientation. As this effect is dependent on the assumed chromaticity of the sky and beam, the antenna layout of the array, and observing frequencies. Future work will attempt to quantify this signal loss for different array types. Despite those

caveats and limitations, our technique is the first proposal to enable the estimation of 21 cm signal within the wedge without detailed knowledge of foregrounds or the primary beam and is thus perhaps the first step towards recovering ~ 80 per cent of the sensitivity that interferometers like HERA must sacrifice by avoiding the foreground wedge.

SOFTWARE

This work was enabled by a number of software packages including MATPLOTLIB³ (Hunter 2007), NUMPY⁴ (Oliphant 2006), SCIPY⁵ (Virtanen et al. 2020), JAX⁶ (Bradbury et al. 2018) for data analysis and modelling, and OPTAX⁷ (Babuschkin et al. 2020) for the gradient descent framework and library of optimizers.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge helpful discussions with Zachary Martinot and Mike Wilensky. This material is based upon work supported by the National Science Foundation under grants #1636646 and #1836019 and institutional support from the HERA collaboration partners. This research is funded in part by the Gordon and Betty Moore Foundation. HERA is hosted by the South African Radio Astronomy Observatory, which is a facility of the National Research Foundation, an agency of the Department of Science and Innovation. JSD gratefully acknowledges the support of the NSF AAPF award #1701536.

DATA AVAILABILITY

The data underlying this article (e.g. the visibility simulations) will be shared on reasonable request to the corresponding author.

REFERENCES

Abdurashidova Z. et al., 2022, *ApJ*, 924, 51
 Aguirre J. E. et al., 2022, *ApJ*, 924, 85
 Ali Z. S. et al., 2015, *ApJ*, 809, 61
 Baars J. W. M., Mezger P. G., Wendker H., 1965, *ApJ*, 142, 122
 Babuschkin I. et al., 2020, The DeepMind JAX Ecosystem. Available at: <http://github.com/deepmind>
 Bandura K. et al., 2014, in Stepp L. M., Gilmozzi R., Hall H. J.eds, Proc. SPIE Conf. Ser. Vol. 9145, Ground-Based and Airborne Telescopes V. SPIE, Bellingham, p. 914522
 Barry N., Hazelton B., Sullivan I., Morales M. F., Pober J. C., 2016, *MNRAS*, 461, 3135
 Beardsley A. P., Morales M. F., Lidz A., Malloy M., Sutter P. M., 2015, *ApJ*, 800, 128
 Bradbury J. et al., 2018, JAX: composable transformations of Python + NumPy programs, <http://github.com/google/jax>
 Byrne R., 2023, *ApJ*, 943, 117
 Byrne R. et al., 2019, *ApJ*, 875, 70
 Byrne R., Morales M. F., Hazelton B. J., Wilensky M., 2021, *MNRAS*, 503, 2457
 Charles N., Kern N., Bernardi G., Bester L., Smirnov O., Fagnoni N., Acedo E. d. L., 2023, *MNRAS*, 522, 1009
 Datta A., Bowman J. D., Carilli C. L., 2010, *ApJ*, 724, 526

DeBoer D. R. et al., 2017, *PASP*, 129, 045001
 Dillon J. S., Parsons A. R., 2016, *ApJ*, 826, 181
 Dillon J. S. et al., 2018, *MNRAS*, 477, 5670
 Dillon J. S. et al., 2020, *MNRAS*, 499, 5840
 de Lera Acedo E. et al., 2018, *Exp. Astron.*, 45, 1
 de Oliveira-Costa A., Tegmark M., Gaensler B. M., Jonas J., Landecker T. L., Reich P., 2008, *MNRAS*, 388, 247
 Eastwood M. W. et al., 2018, *AJ*, 156, 32
 Eastwood M. W. et al., 2019, *AJ*, 158, 84
 Ewall-Wice A. et al., 2016a, *MNRAS*, 460, 4320
 Ewall-Wice A. et al., 2016b, *ApJ*, 831, 196
 Ewall-Wice A., Dillon J. S., Liu A., Hewitt J., 2017, *MNRAS*, 470, 1849
 Ewall-Wice A. et al., 2021, *MNRAS*, 500, 5195
 Ewall-Wice A., Dillon J. S., Gehlot B., Parsons A., Cox T., Jacobs D. C., 2022, *ApJ*, 938, 151
 Fagnoni N., de Lera Acedo E., Drought N., DeBoer D. R., Riley D., Razavi-Ghods N., Carey S., Parsons A. R., 2021a, *IEEE Trans. Antennas Propag.*, 69, 8143
 Fagnoni N. et al., 2021b, *MNRAS*, 500, 1232
 Furlanetto S. R., Oh S. P., Briggs F. H., 2006, *Phys. Rep.*, 433, 181
 Gehlot B. K. et al., 2018, *MNRAS*, 478, 1484
 Gogo T. G. et al., 2022, *MNRAS*, 510, 1680
 Górski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelmann M., 2005, *ApJ*, 622, 759
 HERA Collaboration, 2023, *ApJ*, 945, 124
 Hamaker J. P., Bregman J. D., Sault R. J., 1996, *A&AS*, 117, 137
 Harris F. J., 1978, *Proc. IEEE*, 66, 51
 Hunter J. D., 2007, *Comput. Sci. Eng.*, 9, 90
 Jacobs D. C. et al., 2017, *PASP*, 129, 035002
 Jordan C. H. et al., 2017, *MNRAS*, 471, 3974
 Josaitis A. T., Ewall-Wice A., Fagnoni N., de Lera Acedo E., 2022, *MNRAS*, 514, 1804
 Karnik S., Romberg J., Davenport M. A., 2021, *ACHA*, 55, 97
 Kern N. S., Parsons A. R., Dillon J. S., Lanman A. E., Fagnoni N., de Lera Acedo E., 2019, *ApJ*, 884, 105
 Kern N. S. et al., 2020a, *ApJ*, 888, 70
 Kern N. S. et al., 2020b, *ApJ*, 890, 122
 Kim H. et al., 2023, *ApJ*, 953, 136
 Kovetz E. D. et al., 2017, preprint ([arXiv:1709.09066](https://arxiv.org/abs/1709.09066))
 La Plante P., Mirocha J., Gorce A., Lidz A., Parsons A., 2023, *ApJ*, 944, 59
 Lenc E. et al., 2016, *ApJ*, 830, 38
 Li W. et al., 2018, *ApJ*, 863, 170
 Li W. et al., 2019, *ApJ*, 887, 141
 Line J. L. B. et al., 2018, *Publ. Astron. Soc. Aust.*, 35, e045
 Liu A., Shaw J. R., 2020, *PASP*, 132, 062001
 Liu A., Tegmark M., 2012, *MNRAS*, 419, 3491
 Liu A., Tegmark M., Morrison S., Lutomirski A., Zaldarriaga M., 2010, *MNRAS*, 408, 1029
 Liu A., Parsons A. R., Trott C. M., 2014a, *Phys. Rev. D*, 90, 023018
 Liu A., Parsons A. R., Trott C. M., 2014b, *Phys. Rev. D*, 90, 023019
 Mertens F. G., Ghosh A., Koopmans L. V. E., 2018, *MNRAS*, 478, 3640
 Mesinger A., Furlanetto S., Cen R., 2011, *MNRAS*, 411, 955
 Mevius M. et al., 2022, *MNRAS*, 509, 3693
 Moore I. C., Cada M., 2004, *Appl. Comput. Harmon. Anal.*, 16, 208
 Morales M. F., Hewitt J., 2004, *ApJ*, 615, 7
 Morales M. F., Wyithe J. S. B., 2010, *ARA&A*, 48, 127
 Neben A. R. et al., 2016, *ApJ*, 826, 199
 Newburgh L. B. et al., 2016, in Hall H. J., Gilmozzi R., Marshall H. K., eds, Proc. SPIE Conf. Ser. Vol. 9906, Ground-based and Airborne Telescopes VI. SPIE, Bellingham, p. 99065X
 Nunhokee C. D. et al., 2020, *ApJ*, 897, 5
 Oliphant T. E., 2006, A Guide to NumPy. Vol. 1, Trelgol Publishing, Spanish Fork, UT, USA
 Orosz N., Dillon J. S., Ewall-Wice A., Parsons A. R., Thyagarajan N., 2019, *MNRAS*, 487, 537
 Paciga G. et al., 2013, *MNRAS*, 433, 639
 Pagano M. et al., 2023, *MNRAS*, 520, 5552

³<https://matplotlib.org/>

⁴<https://numpy.org/>

⁵<https://www.scipy.org/>

⁶<https://github.com/google/jax>

⁷<https://github.com/google-deepmind/optax>

Parsons A. R. et al., 2010, *AJ*, 139, 1468
 Parsons A., Pober J., McQuinn M., Jacobs D., Aguirre J., 2012a, *ApJ*, 753, 81
 Parsons A. R., Pober J. C., Aguirre J. E., Carilli C. L., Jacobs D. C., Moore D. F., 2012b, *ApJ*, 756, 165
 Parsons A. R. et al., 2014, *ApJ*, 788, 106
 Patra N. et al., 2018, *Exp. Astron.*, 45, 177
 Planck Collaboration VI, 2020, *A&A*, 641, A6
 Pober J. C. et al., 2012, *AJ*, 143, 53
 Pober J. C. et al., 2013, *ApJ*, 768, L36
 Pritchard J. R., Loeb A., 2012, *Rep. Prog. Phys.*, 75, 086901
 Sievers J. L., 2017, preprint ([arXiv:1701.01860](https://arxiv.org/abs/1701.01860))
 Sims P. H., Pober J. C., Sievers J. L., 2022a, *MNRAS*, 517, 910
 Sims P. H., Pober J. C., Sievers J. L., 2022b, *MNRAS*, 517, 935
 Slepian D., 1978, AT&T Tech. J., 57, 1371
 Smirnov O. M., 2011, *A&A*, 527, A106
 Taylor A. R., Stil J. M., Sunstrum C., 2009, *ApJ*, 702, 1230
 Tegmark M., Eisenstein D. J., Hu W., de Oliveira-Costa A., 2000, *ApJ*, 530, 133
 Thyagarajan N. et al., 2013, *ApJ*, 776, 6
 Thyagarajan N., Parsons A. R., DeBoer D. R., Bowman J. D., Ewall-Wice A. M., Neben A. R., Patra N., 2016, *ApJ*, 825, 9
 Tingay S. J. et al., 2013, *Publ. Astron. Soc. Aust.*, 30, e007
 van Haarlem M. P. et al., 2013, *A&A*, 556, A2
 Védantham H., Udaya Shankar N., Subrahmanyam R., 2012, *ApJ*, 745, 176
 Virtanen P. et al., 2020, *Nat. Methods*, 17, 261
 Wang X., Tegmark M., Santos M. G., Knox L., 2006, *ApJ*, 650, 529
 Wieringa M. H., 1992, *Exp. Astron.*, 2, 203
 Yatawatta S., 2016, preprint ([arXiv:1605.09219](https://arxiv.org/abs/1605.09219))
 Yatawatta S., Zaroubi S., de Bruyn G., Koopmans L., Noordam J., 2008, preprint ([arXiv:0810.5751](https://arxiv.org/abs/0810.5751))
 Yoshiura S. et al., 2021, *MNRAS*, 505, 4775
 Zheng H. et al., 2014, *MNRAS*, 445, 1084
 Zheng H. et al., 2017, *MNRAS*, 464, 3486

APPENDIX A: NUCAL VERSUS CALAMITY

NUCAL bears several similarities to the recently proposed CALAMITY technique (Ewall-Wice et al. 2022) due to its use of first-order optimization to determine gain solutions and the use of DPSS modelling. In this appendix, we highlight the specific similarities and differences between the two techniques by casting them both under the formalism adopted in the rest of the paper. CALAMITY and NUCAL differ primarily in two ways.

A1 Different foreground modelling functions

Let the set of all frequencies sampled by every baseline be the same, which we denote with the vector \mathbf{F} with length N_F . From

equation (16), we can describe any per-baseline modelling approach like the one described in Ewall-Wice et al. (2022) with a set of ϕ_k basis functions where the vectors are partitioned into b different sets where the m^{th} set has length n_m different functions, $\{\varphi_{0,0}, \dots, \varphi_{0,n_0}\}, \{\varphi_{1,0}, \dots, \varphi_{1,n_1}\}, \dots, \{\varphi_{b,0}, \dots, \varphi_{b,n_b}\}$ and each m^{th} baseline is centred at the set of \mathbf{u} values in the N_F length set \mathbf{U}_m . The first index of each function labels the baseline, and the second index labels the basis functions necessary to model that baseline.

$$\varphi_{m,n}(\mathbf{u}, v_i) = \begin{cases} \phi_{m,n}(v) & \mathbf{u} \in \mathbf{U}_m \text{ and } v \in \mathbf{F} \\ 0 & \text{otherwise,} \end{cases} \quad (\text{A1})$$

In the per-baseline modelling case, we work with a set of basis functions that only having support over a single baseline. In practice, we fit discrete vectors to our data the basis vectors in the per-baseline modelling case are non-zero over a single baseline. For the DPSS modelling used in CALAMITY, we choose $\phi_{m,n}$ equal to

$$\phi_{m,n} = u_n(v, \mathcal{W}_m) \quad (\text{A2})$$

where u_n is the n^{th} Slepian sequence (Slepian 1978) with normalized bandwidth of $\mathcal{W}_m = 2B\tau_m$ where τ_m is the delay width that is chosen to describe the foregrounds modulated by the primary beam.

A2 Different gain modelling assumptions

In NUCAL, we assume that the redundant degrees of freedom in the gains have already been removed through some redundant calibration strategy. All that is left to solve for in the gains is a common per-frequency absolute gain factor $A(v)$ and the tip-tilt term $\Phi(v)$. In CALAMITY, we retain and fit per-frequency complex gain for every antenna in the array. This limits the utility of NUCAL as specifically laid out in this paper to redundant arrays with identical beams while CALAMITY can, in principle, be applied to arbitrary array layouts with non-redundant beams. Reducing the gain degrees of freedom may reduce the amount of signal loss suffered by our measurement for a fixed array layout, though the larger number of independent uv modes sampled by non-redundant arrays could counteract this. We leave a detailed exploration of the tradeoffs between gain modelling assumptions and array layout to future work.

In summary, NUCAL reduces to CALAMITY in the limit that the antenna gains are described by a tip-tilt parameter and an amplitude degeneracy, the true visibility modelling functions have support only over individual baselines, and are described by discrete Slepian sequences instead of sampled prolate spheroidal wave functions.

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