Draft version March 31, 2021
Preprint typeset using LaTeX style AASTeX6 v. 1.0

HERA MEMO #96: TIME CORRELATIONS OF EOR SIGNALS FOR HERA WITH VIVALDI FEEDS

Nicholas Kern[1], Scott Dynes[1], Bang Nhan[2], Honggeun Kim[1], Jacqueline Hewitt[1], Eleanor Rath[1], Nicolas Fagnoni[3], Eloy de Lera Acedo[3], David DeBoer[4]

[1]Dept. of Physics and Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, Cambridge, MA
[2]National Radio Astronomy Observatory, Charlottesville, VA
[3]Cavendish Laboratory, University of Cambridge, Cambridge, UK
[4]Radio Astronomy Laboratory, University of California, Berkeley, CA

ABSTRACT

We compute fringe-rate filter specifications for HERA visibilities outfitted with Vivaldi feeds for the suppression of cross-coupling systematics and for coherent time averaging. We perform ensemble mock EoR simulations with a Vivaldi beam response to compute the power spectral density (PSD) of an EoR-like sky in the fringe-rate domain. Using this, we can set specifications for how aggressive a fringe-rate filter can be before it begins to attenuate a certain fraction of the EoR signal. Said another way, we can determine the lossiness of a time filter based on which fringe-rate modes it keeps and which it rejects. We perform these tests for a fiducial beam model, and for the first time explore how primary beam perturbations impact these results. For fiducial feed displacement perturbations less than 3 cm, we find that the perturbed beam has a negligible impact on the EoR time correlations in the visibilities.

1. INTRODUCTION

Cross-coupling systematics in radio telescopes arise from the coupling of the voltage signals in the RF analog chain between antennas. They can arise from a variety of mechanisms, such as antenna-to-antenna reflections (i.e. mutual coupling) and capacitive crosstalk in analog-to-digital converter (ADC) units. These systematics can have non-trivial spectral structure, and are therefore necessary to control for 21 cm surveys that rely on spectral smoothness of the instrument response to isolate the weak cosmological signal. Parsons et al. (2016) presents an interferometric visibility filtering technique that uses the different time correlations of cross-coupling systematics and sky signals as a means to reject the former and isolate the latter. This relies on the assumption that cross-coupling systematics are slowly variable in time (e.g. see Kern et al. 2019), occupying mainly the common-mode (DC term) Fourier mode, whereas any sky signal of a drift-scan experiment will oscillate as the sky moves through the interferometric fringes, and therefore will occupy mainly non-zero Fourier modes.[1] Note that in this work, we define the Fourier domain as the Fourier transform across the time axis of the data (into the *frige-rate* domain), not the frequency axis. See also Shaw et al. (2014) for a description of how sky signals source time correlations into the visibilities of a drift-scan interferometer.

Parsons et al. (2016) show that under certain assumptions one can analytically compute the expected distribution of an EoR sky signal (independent of the EoR power spectrum model) in the fringe-rate Fourier basis (the Fourier dual of time) at each observed frequency channel. Given this distribution, they can tailor a high-pass time filter that filters off the slowly variable systematics, while largely preserving the EoR sky signal. The amount of EoR loss can be tuned to a nearly arbitrary amount, with increasing systematic rejection coming at the expense of increased EoR loss. In Ali et al. (2015), this was chosen to yield $\sim 28\%$ power loss, which was later corrected for at the end stage of the analysis.

However, the assumptions made in the analytic calculation of Parsons et al. (2016) begin to break down for experiments like the Hydrogen Epoch of Reionization Array (HERA), where for the shortest baselines the fringe separation is on the order of the primary beam full width half max (FWHM). In order to compute the power spectral density of an EoR signal in the fringe-rate domain, we therefore turn to Monte Carlo simulations of EoR sky signals, similar to Kern et al. (2019).

First let us define how we simulate the visibility response of an interferometer. The interferometric visibility between two antennas j and k at frequency ν and

[1] This is only true for baselines with non-zero projection along the East-West direction.

local sidereal time (t) is given as

$$V_{jk}(\nu, t) = \int d\Omega A(\nu, \hat{s}) I(\nu, \hat{s}, t) e^{2\pi i \boldsymbol{b}_{jk} \hat{s} \nu / c}, \quad (1)$$

where the integral is over the sphere, $\hat{s}$ is a unit pointing vector in local array coordinates (i.e. topocentric coordinates), A is the scalar primary beam response[2] of the antennas (assumed to be the same for each antenna), and I is the specific intensity of the sky that lies along $\hat{s}$ at time t. Note that for a fixed location on the sky in topocentric coordinates, the primary beam and fringe of the interferometer are fixed as a function of time, whereas it is the sky that is assumed to drift through the beam as a function of LST. We evaluate Equation 1 by representing the sky, beam, and fringe response as a NSDIE=128 HEALpix map (Górski et al. 2005) and directly sum their product, using the `healvis` visibility simulator (Lanman & Pober 2019; Lanman et al. 2020). For the frequencies and baselines explored in this work, representing the integral of Equation 1 with a NSIDE=128 HEALpix map has been shown to be adequate (Lanman & Pober 2019).

To model an EoR-like signal, we use an uncorrelated Gaussian random field for each pixel on the sky for each frequency channel (Figure 1). Because we are interested in computing the *time correlations* of an EoR-like, sky-locked signal, the frequency correlations of the model are not important. This means that our results will be largely insensitive to the power spectrum of the EoR model, and thus largely model independent.[3] While the mock EoR field is drawn with a variance of 25 mK^2, the amplitude of the field is irrelevant for determining the shape of the power spectral density in fringe-rate space. The frequency and sky angle response of the primary beam is simulated using CST Microwave Studio models of the HERA dish and feed, discussed more in the following section. We also compare against an analytic Airy disk model with an aperture size of 14-meters. We simulate visibilities for both XX and YY linear polarizations, which will each have a slightly different time correlations depending on the parts of the sky they upweight. However, due to the fact that we are mainly interested in the Stokes I content of the sky, the power spectral density and time averaging analyses presented here operate on the pseudo-Stokes I visibility, formed as $V_I = \frac{1}{2}(V_{XX} + V_{YY})$. The simulation parameters used in this work are tabulated in Table 1.

[2] The "power beam."

[3] One complication here is that the angular correlations of the model, at fixed frequency, will slightly change the measured correlations across time; however, because we are primarily interested in HERA's short baseline length, we expect these correlations to have a negligibly small impact on the results. Indeed, similar analyses have been done with different EoR power spectrum models that have yielded similar results.

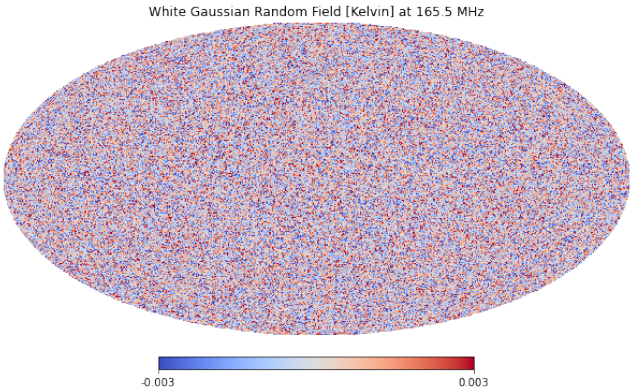

Figure 1. A HEALpix map of a single mock δT_b EoR realization at 165 MHz, which is drawn from a white Gaussian distribution.

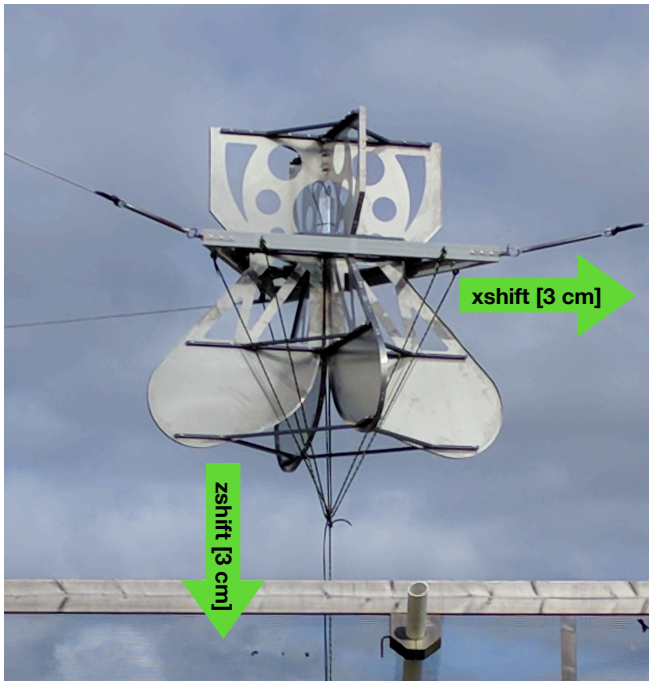

Figure 2. A Vivaldi feed installed in the field on a HERA dish. The feed is suspended by three kevlar cables attached to wooden poles surrounding the dish. Its vertical height is secured by a cable pulling it down from the center of the dish. In Section 2 and Section 3 we use a fiducial model where the feed is centered, whereas in Appendix A we explore how feed translation along the z and x direction (green arrows) impact our results.

Table 1. Healvis Simulation Parameters

Parameter	Value
Sky Resolution	HEALpix NSIDE 128
EoR Model	White Gaussian Field
Primary Beam Model #1	Fagnoni Vivaldi CST Simulation

Table 1 continued

Table 1 *(continued)*

Parameter	Value
Primary Beam Model #2	Airy Disk, 14 m aperture
Frequency Range	60 – 140 MHz
Baseline Lengths	0 – 90 meters
Polarizations	Linear XX & YY

NOTE—The power spectral density and time averaging analyses discussed below operate on the pseudo-Stokes I visibility.

2. POWER SPECTRAL DENSITIES IN THE FRINGE-RATE DOMAIN

To determine how much signal is lost when averaging EoR visibilities across LST (i.e. in applying a time-based filter), we need to compute the *power spectral density* (PSD) of the EoR visibilities in fringe-rate space (f), which is the square of the Fourier transformed visibilities. We first define the Fourier transform of the visibilities across LST,

$$\widetilde{V}(\nu, f) = \int dt V(\nu, t) e^{2\pi i t f}, \tag{2}$$

and then square the result and average across independent realizations to get an approximation of the ensemble-averaged PSD,

$$\langle P(\nu, f) \rangle \approx \frac{1}{N} \sum_{i}^{N} |\widetilde{V}_i(\nu, f)|^2 \tag{3}$$

where i indexes independent EoR simulations with the same beam model but with a different statistical realization of the sky.

We can gain some intuition for the expected fringe-rate behavior by inspecting the Fourier transform of the visibility (across time). Assume the observer is situated in local topocentric coordinates and let us simplify our sky model to a single point source. Over the course of an observing night, we will observe the source move across the sky and move through the primary beam response as well as the baseline fringes (for a drift-scan telescope). At any given moment, the movement of a source through the fringes creates a complex sinusoid in the visibility as a function of time, with a characteristic period given by the projection of the source's angular velocity onto the fringe pattern. Assuming the beam is smooth and constant, this is the only source of time correlations in the visibility. However, relaxing this assumption, we see that the beam (being multiplicative in the time domain) acts to convolve the delta-function fringe-rate response of a point source in the visibility with its own fringe-rate "beam kernel." The faster the beam evolves spatially at a given point on the sky the wider this beam kernel will be in fringe-rate space, which will spread the intrinsic fringe-rate response of the point source moving through the fringes. To first order, the total fringe-rate response of the visibility is thus the sum of the fringe-rate response of all point sources on the sky (diffuse sources can

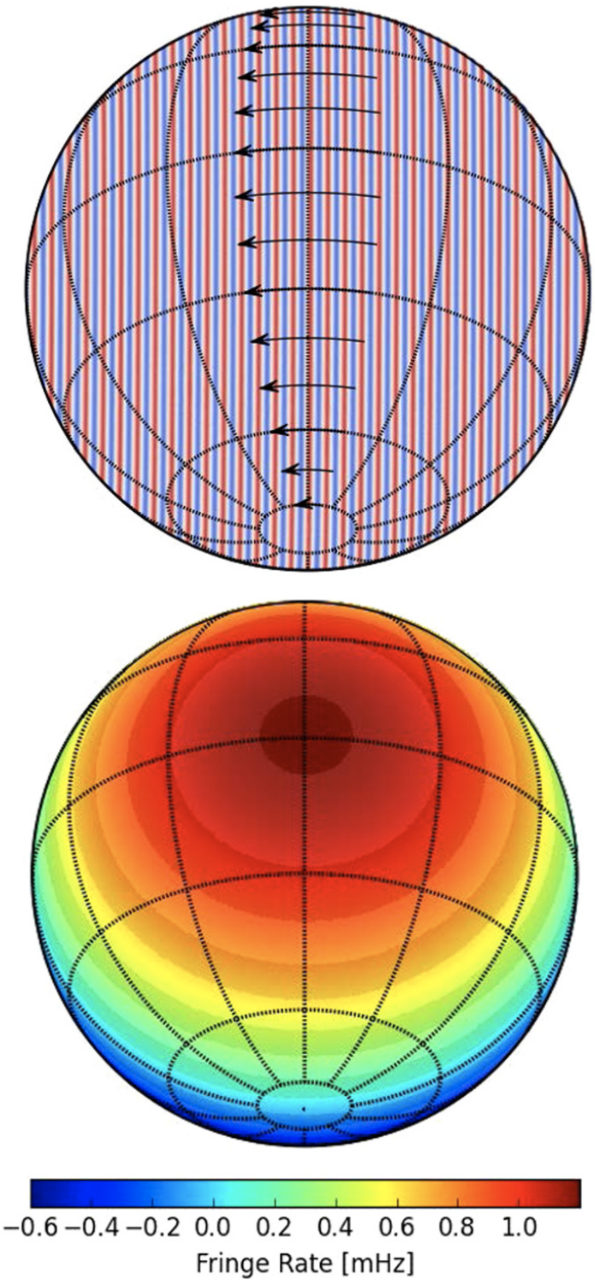

Figure 3. *Figures 1 & 2 from Parsons et al. (2016)*, showing the fringe response of a 30-m East-West baseline at 150 MHz on the sky (top) and the angular velocity of point sources at different sky coordinates (vectors). The instantaneous movement of the sky through the fringes therefore creates a "fringe-rate" response in the visibilities that is direction dependent based on HERA's latitude (bottom). To first-order, visibility fringe-rate filtering can be thought of as a sculpting of the field-of-view.

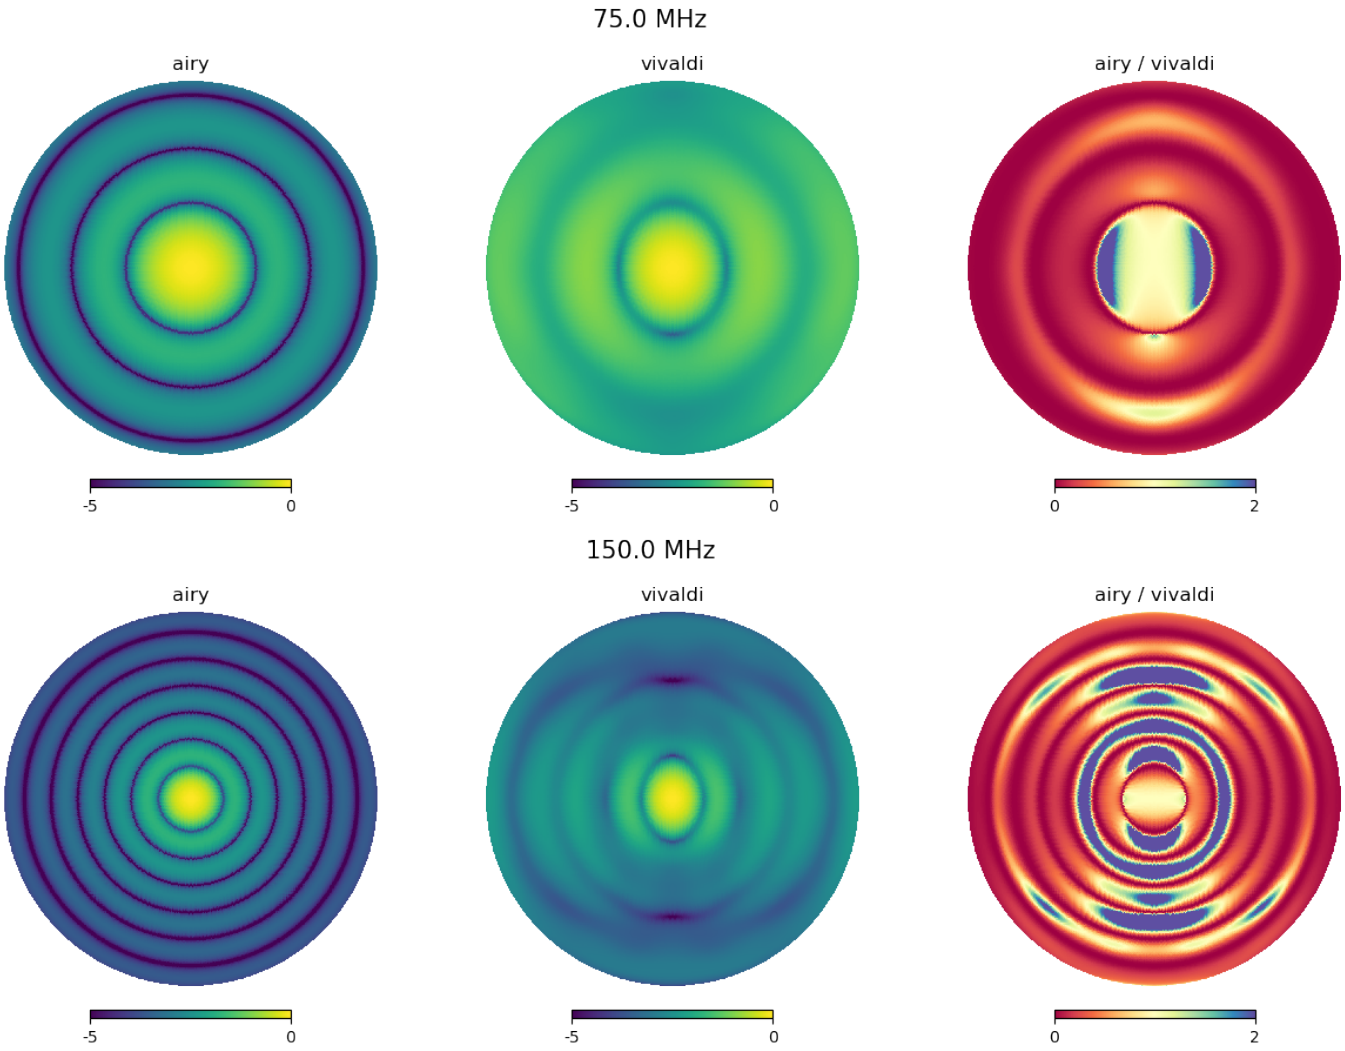

Figure 4. Comparison of the idealized Airy disk model (left column), the CST Vivaldi model of Fagnoni et al. (2020) (center column), and their ratio (right column). These are the power beams for one particular visibility polarization, where the left and center column are cast into log10 units, and the right column is a direct ratio. One thing that is very clear, supported by Figure 11, is the up-weighting of the horizon by the Vivaldi model at low frequencies compared to the Airy disk model.

be thought of roughly as the sum of many independent point sources). This is demonstrated in Figure 3, reproduced from Parsons et al. (2016), showing the fringe-response on the sky of a 30-m East-West baseline at 150 MHz (top) and the angular velocity of point sources at different parts of the sky (vectors). The projection of the two create the direction-dependent fringe-rate response of the visibility on the sky (bottom), highlighting how fringe-rate filtering acts to manipulate the field-of-view. As noted previously, the two beam models explored here are a CST simulation of the HERA dish and Vivaldi feed (Fagnoni et al. 2020), as well as an idealized Airy disk model with a 14-m aperture. A comparison of the power pattern of these two beams are shown in Figure 4 for a frequency bin in the low band and mid band.

The overarching goal of this exercise is to determine where in fringe-rate space the EoR signal occupies, which tells us the fringe-rate modes we do not want to filter out. If systematics and foregrounds occupy independent fringe-rate modes, then we can enact a fringe-rate filter to isolate the EoR from systematics. See Kern et al. (2019) for more details on why we expect cross-coupling systematics to be slowly time variable.

Here we report on the results of the EoR visibility PSD using a CST model of the Vivaldi beam from Fagnoni et al. (2020). Figure 5 shows the PSD of the EoR model for a handful of HERA baselines of various lengths and orientations (the array-layout is shown in the upper-left panel). The majority of the EoR power is sourced along a "cone" in fringe-rate & frequency space, whose midline increases both with increasing frequency and with increasing East-West baseline length. This can easily be understood in both case as the narrowing of the fringes on the sky, meaning a source moving with fixed angu-

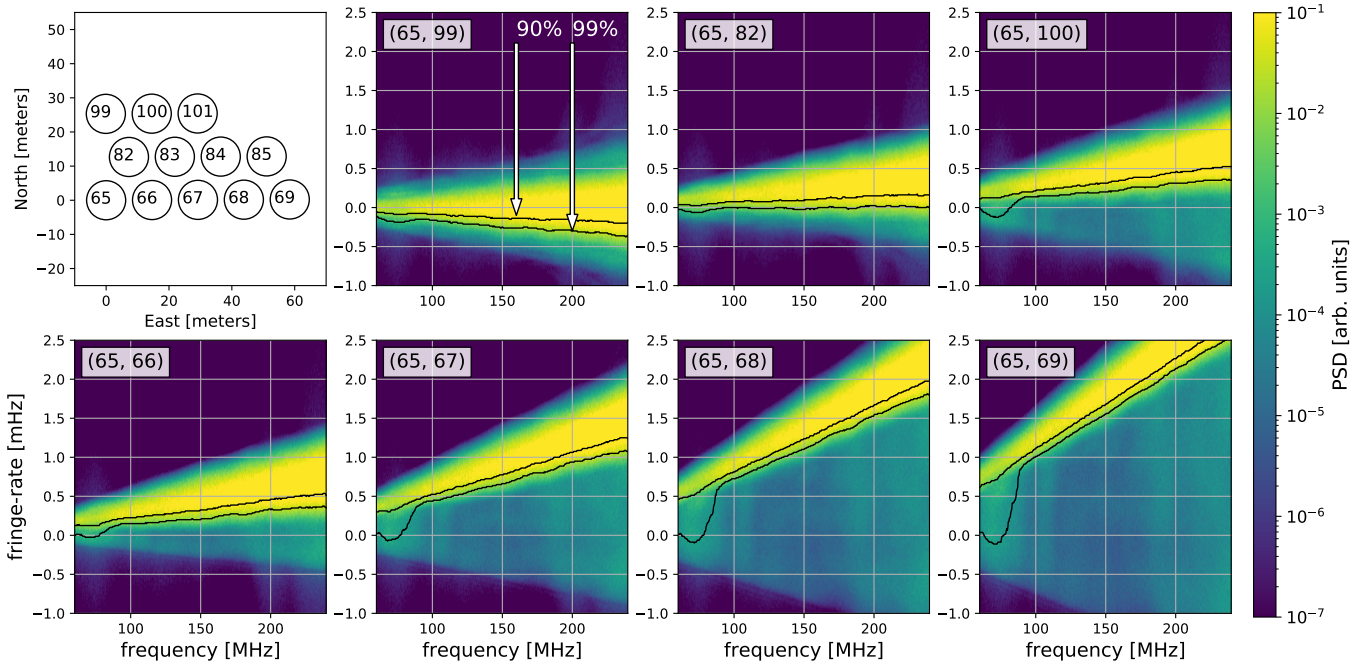

Figure 5. Simulated PSD of an EoR-like sky signal with the **Vivaldi beam model** in fringe-rate space for a handful of HERA baselines (the array layout is shown in the left panel) at two characteristic frequencies (center and right panels). Note that the peak power shifts to positive fringe-rates based on the projected East-West separation of the baseline: for baselines with no projected East-West length, the PSD is centered at $f = 0$ mHz.

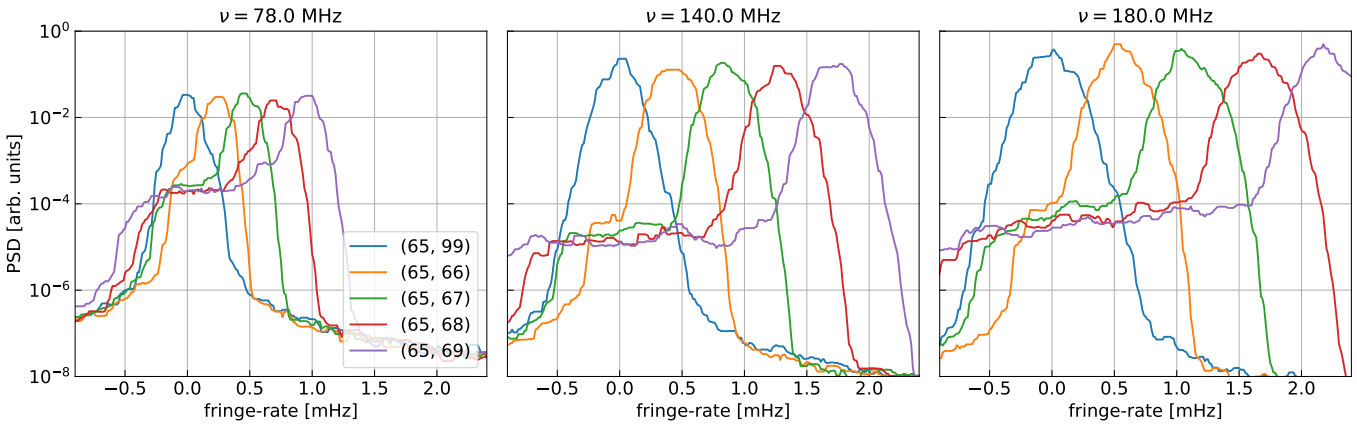

Figure 6. Vertical cuts through the PSD distributions in Figure 5 simulated with the **Vivaldi beam** at three frequency channels, showing more clearly the extent in fringe-rate space of the PSD.

lar velocity at a particular point on the sky and time of night creates a faster moving sinusoid in the visibilities over time. For baselines with little to no projected East-West length, the PSD of the EoR signal straddles a fringe-rate of $f = 0$ mHz. Also plotted are two lines, which represent a dividing line above which 90% (or 99%, as marked) of the EoR power resides in fringe-rate space. In other words, if one were to filter out all fringe-rates below the black line from the data, one would only attenuate 10% (1%) of the total EoR power in the visibilities.

To take a closer look at these PSD profiles, we can inspect vertical slices of the PSD at specific frequency channels. Figure 6 shows this for the Vivaldi beam simulations, showing in more detail the relative extent of the PSD across various fringe-rates. The noise-like structure of the PSD is simply due to sample variance on our estimate of the mean PSD: in other words, if we were to simulate more draws of our EoR field and average their PSDs, we would see the noise-like structure decrease; however, with the amount of simulations performed here, we can get estimates of the integrals of the PSD that are accurate-enough to set the filtering specifications sought-out by this analysis. Near peak center,

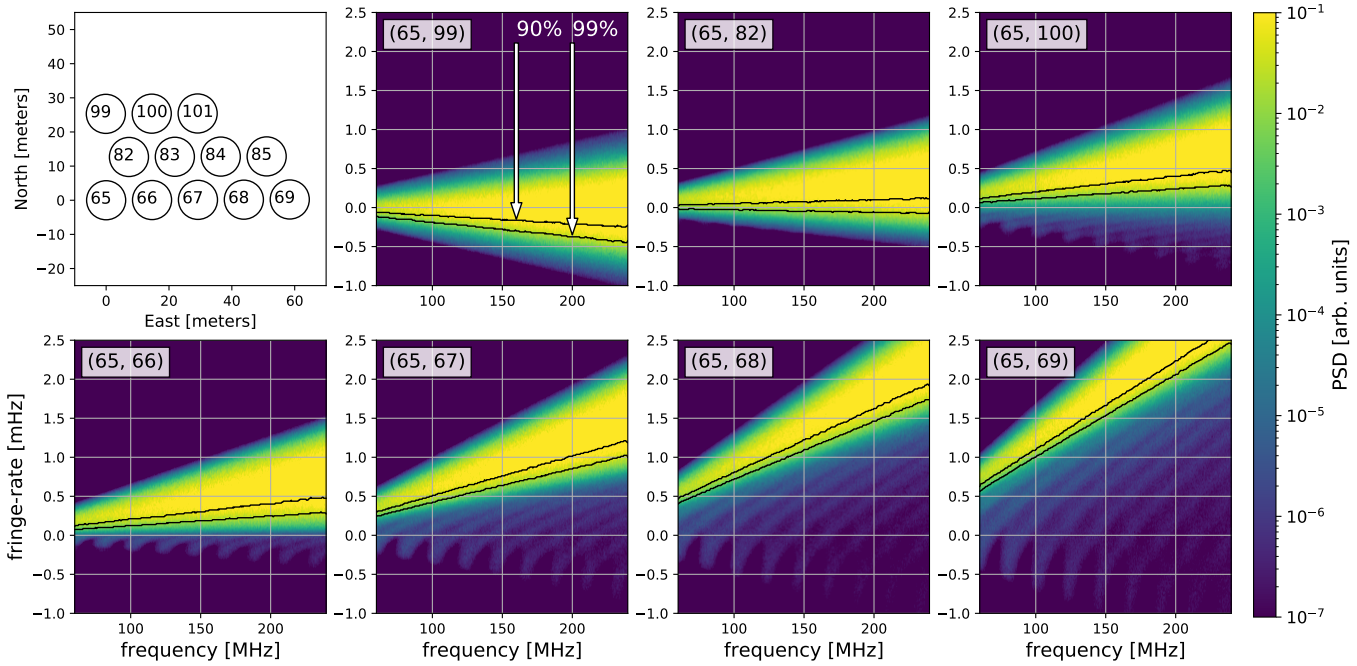

Figure 7. Simulated PSD of an EoR-like sky signal with the **Airy disk beam model** in fringe-rate space for a handful of HERA baselines (the array layout is shown in the left panel) at two characteristic frequencies (center and right panels). Note that compared to the Vivaldi model, these PSDs are much more compact in fringe-rate space and have less bleed to low and negative fringe-rates.

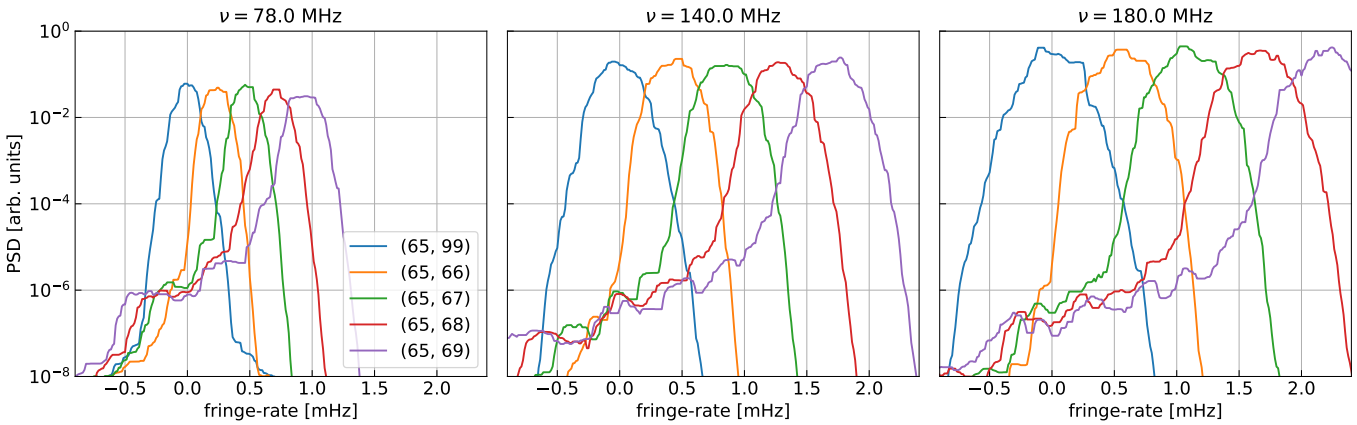

Figure 8. Vertical cuts through the PSD distributions in Figure 7 simulated with the **Airy beam** at three frequency channels, showing more clearly the extent in fringe-rate space of the PSD.

one can imagine that the PSD is well-fit by a Gaussian; however, the tails of the distributions show interesting non-Gaussian and asymmetric behavior, which is driven mainly by the relatively strong negative tails in the distributions. The strong negative tails of the PSD distributions are some of the interesting phenomenology observed specifically in the Vivaldi beam simulations. Although low-level tails were observed in similar simulations for the HERA Phase I dipole beam model–Figure 6 of Kern et al. (2019)–they were not observed at such high amplitudes. These tails are particularly large at very low frequencies ($\nu < 100$ MHz), which is also evidenced by the integral bounds given by the black lines in Figure 5, showing a clear dip in the 99% integral bound at low frequencies. **The consequence this will have for cross-talk filters is that they will be more lossy to the EoR signal with a fixed filter width Δf centered at $f = 0$ mHz.**

To drive home the comparison of the Vivaldi beam model, we repeat these simulations with an Airy disk model with a 14-meter aperture. These PSDs are shown in Figure 7, which are more compact in fringe-rate space than the Vivaldi model with less power appearing at low and negative fringe-rates. This is due to the fact that,

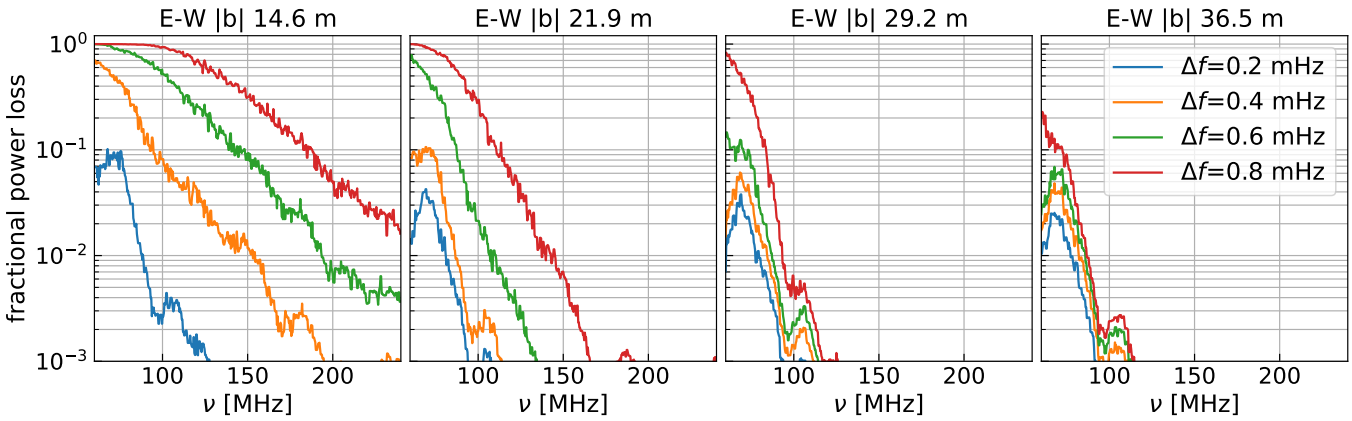

Figure 9. Signal loss specifications for high-pass cross-coupling filters of **Vivaldi beam** simulations with increasing filter width (Δf) for baselines of various East-West projection.

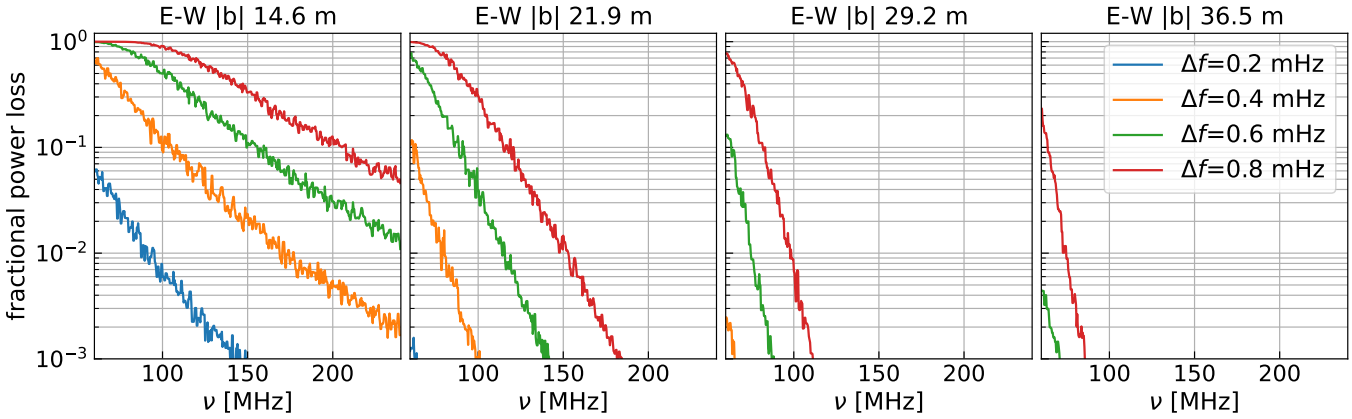

Figure 10. Signal loss specifications for high-pass cross-coupling filters of **Airy disk beam** simulations with increasing filter width (Δf) for baselines of various East-West projection.

compared to the Airy model, the Vivaldi beam is both wider and has more power at the horizon where low and negative fringe-rates in the visibility are sourced, e.g. Figure 2 of Parsons et al. (2016). Note that the dip in the 99% power bound at low frequencies is not observed in the Airy disk model. Inspecting the Vivaldi beam model compared to the Airy beam model, we see that indeed below 100 MHz the Vivaldi model has a particularly strong response at the horizon, which is likely responsible for the increase in the low and negative fringe-rate response at low frequency, and thus also this dip in the 99% bound.

While Figure 5 helps to give us a visual sense of where the power of EoR visibilities lies in fringe-rate space, it doesn't easily tell us how much signal loss we are inducing by actually applying a fringe-rate filter (the 90% and 99% lines do tell us this if we filter-out all fringe-rates below those lines). To better quantify how a fringe-rate filter attenuates the signal (specifically, a high-pass fringe-rate filter used for cross-coupling suppression), we can compare the integral of the PSDs with

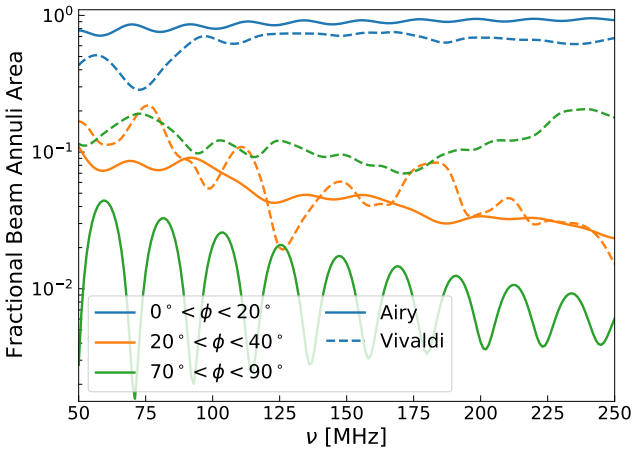

Figure 11. Fractional beam areas of the Airy disk (solid) and Vivaldi beam (dashed) for three distinct annuli centered at increasing zenith angle, ϕ. This shows that the Vivaldi model has more of its fractional total response by the horizon (green) than the Airy model, which explains the wider footprint in fringe-rate space of its PSD distributions.

and without filtering. Here, we assume that one applies

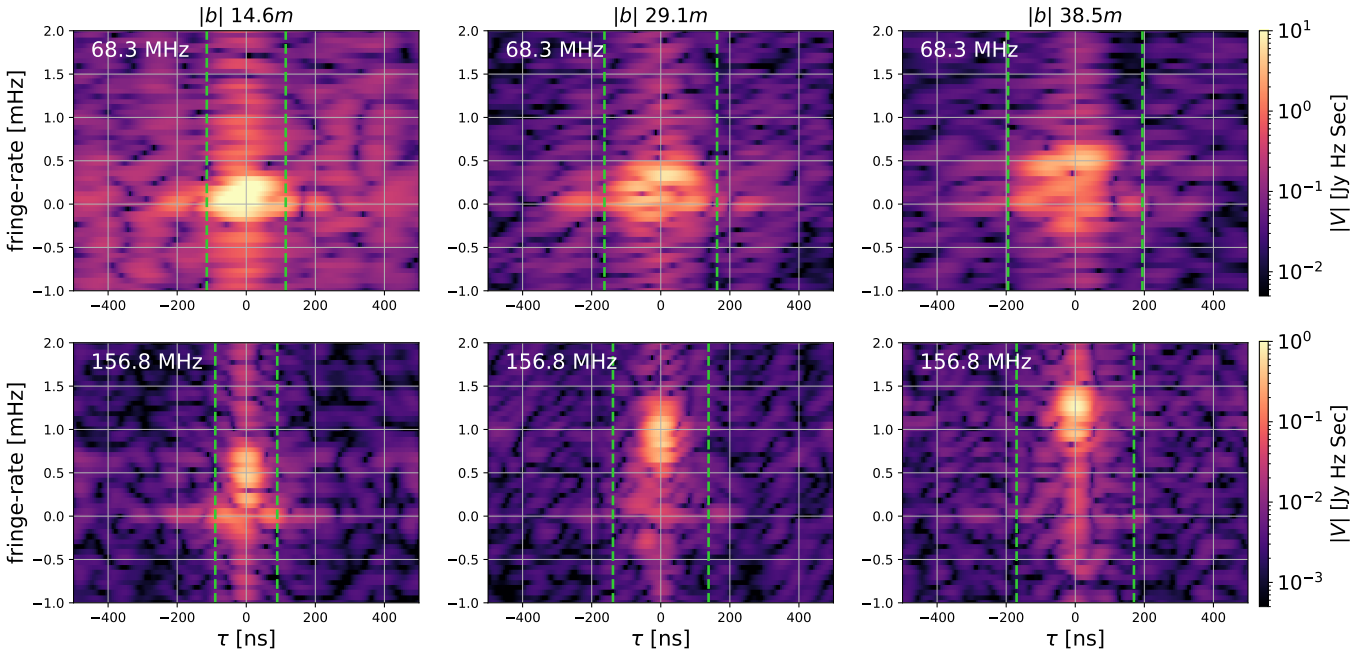

Figure 12. HERA visibilities (H4C observing season) in delay and fringe-rate space for a handful of East-West baseline groups in east-east visibility polarization, taken at two characteristic frequency bands (top and bottom panels). We show the amplitude of the visibilities after redundantly averaging of all physical baselines in each group (roughly a dozen baselines in each group). The dashed green lines show the baseline horizon with an additional buffer to account for the Blackman-Harris windowing before taking the Fourier transform. Both the low-band and mid-band show evidence for low-fringe rate emission beyond the foreground horizon that is particularly strong for shorter baselines. By eye, it seems that high-pass fringe-rate filters with a $\Delta f \sim 0.4$ mHz will be needed to adequately suppress the low-fringe rate systematics, however, this may be an overestimate as the windowing was also applied across the time domain as well, which will bleed structures to larger fringe-rates. This may make fringe-rate filters for cross-coupling systematics at low frequency particularly lossy to EoR signals.

a high-pass filter that suppresses all fringe-rate modes $|f| < \Delta f/2$. Note that this filter only suppresses low-fringe rate modes, is symmetric about $f = 0$ mHz, and leaves higher fringe-rate modes unaffected. If we integrate the PSDs along the fringe-rate axis with and without such a filter and take their ratio, we can derive the amount of signal power lost due to the filter. This is shown for the Vivaldi simulations in Figure 9, where each panel shows a different baseline of increasing East-West length, and each line represents a filter with an increasingly large width Δf. Note that the amount of loss is naturally the largest for the largest filter size, $\Delta f = 0.8$ mHz, and that this loss increases at lower frequencies for all filter sizes and baselines.

If we wanted to set a specification for the most amount of loss we would be willing to tolerate for a filter (say 1% power loss), then Figure 9 tells us that at 100 MHz we can safely apply a $\Delta f = 0.2$ mHz cross-coupling filter. However, this same filter becomes more lossy below 100 MHz, so to not exceed our tolerance we should use a less aggressive filter (or not filter at all!). For large baseline lengths and at higher frequencies, we can boost the size of the filter Δf without exceeding our preset tolerance.

We also show these specifications for the Airy disk simulations in Figure 10. In general we see that for fixed frequency, baseline, and filter width, the Airy disk shows less signal loss than the Vivaldi feed. This again is related to the fact that the Vivaldi feed projects more EoR power to low and negative fringe-rates than the idealized Airy beam. However, this also demonstrates that these specifications are somewhat sensitive to the model of the primary beam. While it is clear that the difference between a Vivaldi and idealized Airy beam are substantial, it is not clear how realistic feed perturbations impact these metrics, which we return to in Appendix A.

To understand why the PSDs of the Vivaldi model differ from that of the Airy disk model, we said above that this is due to the fact that, compared to the Airy disk model, the Vivaldi feed up-weights parts of the sky near the horizon where low and negative fringe-rates are sourced in the visibility. First, Figure 2 of Parsons et al. (2016) shows that low and negative fringe-rates are sourced near the horizon (for a telescope at HERA's declination). Second, Figure 11 plots the fractional beam area in distinct annuli of increasing zenith angle, ϕ, normalized by the total integral of the beam. The solid lines show the Airy disk model, while the dashed lines show the Vivaldi model. Compared to the Airy disk model, the Vivaldi model has less of its total response

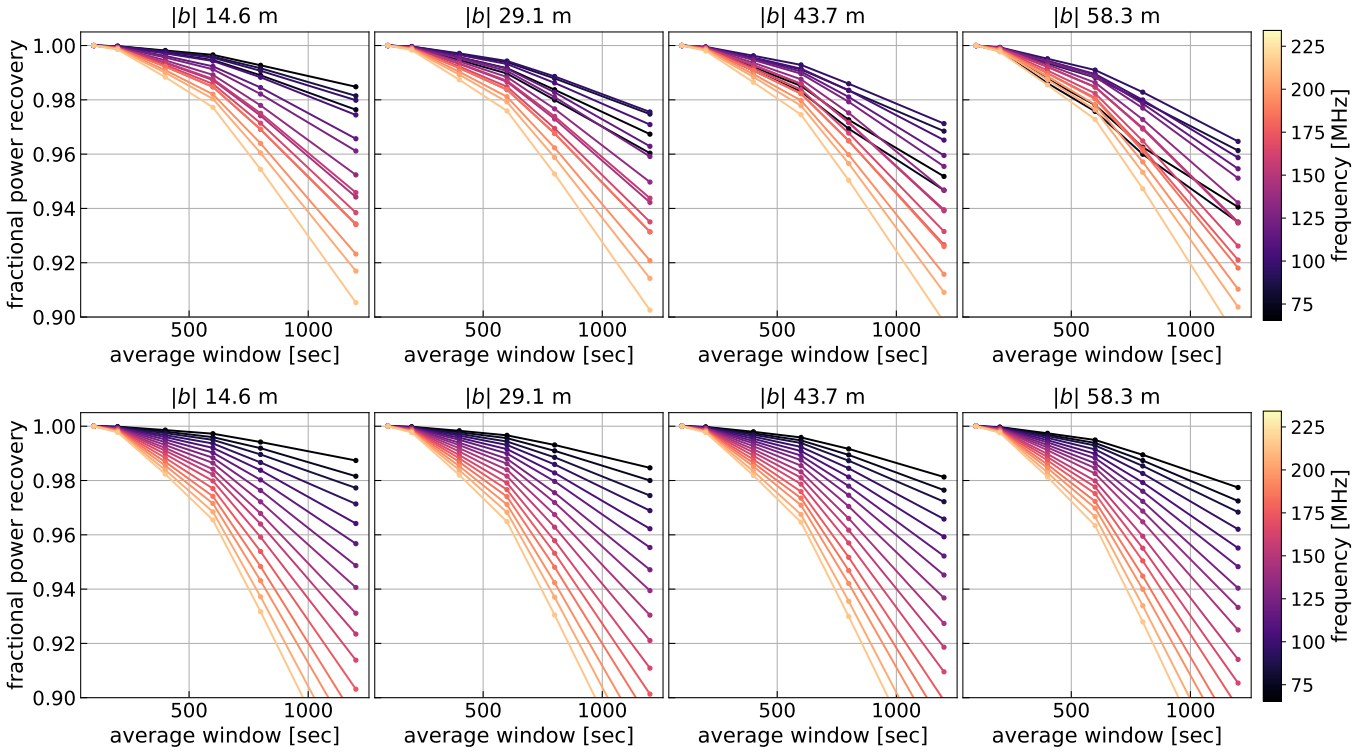

Figure 13. Fractional EoR power recovery after coherent LST averaging of the visibilities with increasingly large time windows for various East-West aligned baseline lengths. The top panels show the **Vivaldi beam** simulations, while the bottom panels show the **Airy disk beam** simulations. To limit signal loss to $\lesssim 2\%$ across the band, one should coherently average the visibilities no longer than ~ 400 seconds.

centered near zenith (blue), but this is compensated by more of its response sourced near the horizon (green), whereas midway through the Vivaldi and Airy model have roughly the same fractional response (orange).

We note that another way to think about the process of correcting for signal loss induced by a fringe-rate filter is to view it as a primary beam correction due to the "sculpting" of the primary beam from the filter. Indeed, this is how the formalism outlined in Parsons et al. (2016) would have us view it. In this case, the fringe-rate filter is not so much of a process by which EoR signal is lost, as it is a shaping of the primary beam response, which would naturally be corrected for in the process of power spectrum normalization (i.e. the Ω_{pp} term that is the integral of the squared primary beam response). This is a perfectly natural way of looking at this process, however, recall that the analytic tools used in Parsons et al. (2016) to compute the sculpted beam response are not suitable for HERA baselines due to a breakdown in the "constant beam" approximation. Therefore, we would be left to compute it in a similar manner as is done here, by imaging an ensemble set of fringe-rate filtered EoR visibilities and taking the average of their amplitude to converge to the "sculpted" primary beam response. Although this is a somewhat more complicated process to achieve what has already been derived here, it would be necessary, say, for proper normalization of fringe-rate filtered images.

To understand the kind of fringe-rate filtering necessary for systematic rejection of low-fringe rate systematics in HERA, we also show some representative commissioning data from the most recent observing season (H4C; observed October 16, 2020). In Figure 12, show the Fourier transformed visibility in both fringe-rate and delay space for a few East-West oriented baselines in east-east visibility polarization (which shows stronger systematics than north-north polarization). We also plot the baseline horizon delay (green dashed) with an additive delay buffer that accounts for the Blackman-Harris apodization (or windowing, or tapering) kernel in delay space (note that BH windowing is also applied along the time axis before taking the Fourier transform). We clearly see an excess of low-fringe rate structure beyond the foreground horizon, which could be indicative of cross-coupling systematics. These appear to decrease in amplitude for larger baseline lengths, which is the expected behavior of over-the-air reflections (or mutual coupling). By eye, it appears that high-pass fringe-rate filters with $\Delta f \sim 0.4$ may be enough to adequately suppress these low-fringe rate systematics, although we previous signal loss analysis suggests that this may induce large amounts of loss of the EoR signal at low frequen-

cies and short baselines.

3. COHERENT TIME AVERAGING SIGNAL LOSS

Here we discuss calculations for the amount of signal loss induced when coherently averaging different time (or LST) bins of drift-scan interferometric visibilities. The "coherent" in coherent LST averaging means we are averaging the complex visibilities, thus preserving the phase information. By LST averaging, we mean averaging different but nearby LST bins within the same observing night. For short time differences, these time bins will see largely the same sky, and thus after rephasing the visibilities we can average them together. However, because the sky has moved (even by a small amount for truly adjacent time bins), there will be some amount of decoherence (and thus power loss) when coherently averaging the visibilities. This is particularly the case when trying to measure the sky power across a wide field-of-view, where the fringe rephasing cannot be performed exactly for all points on the sky. Thus, over longer time differences there will be more decoherence and thus more signal loss. While we want to perform as much coherent LST averaging as possible to beat down thermal noise, there will be a point over long averaging windows where decoherence (and thus signal loss) will be too large to tolerate. Again, for averaging the visibility from a single point source this computation is relatively straightforward and can be done analytically; however, for quantifying the signal loss due to an isotropic field like an EoR signal this is more complicated, and we thus turn to direct numerical simulation as a test.

The LST averaging scheme we will look at is a simple scheme, where we scroll down in LST and take all drift-scan time bins within a specified time window, rephase them to a common pointing center, take their direct, unweighted average, and then move on to the next set of time bins over the same window. This is not exactly the same as the optimal fringe-rate filtering discussed in Parsons et al. (2016), where one applies a fringe-rate filter tailored to minimize the error bars on the power spectrum. However, one can think of the averaging scheme we will present as a sinusoidal fringe-rate filter (with a width equal to one over the time window) with post-filtering decimation.

To measure the amount of signal loss induced by the averaging, we take the coherent average of the 24-hour visibilities independently for each frequency bin and for four representative East-West baselines over increasingly large time windows. For each set of averaging windows, we square the averaged visibilities and take the time average of the remaining time bins. We then normalize the result of each of these averages by the power of the shortest time average window, assuming it is small enough to induce negligible amounts of loss, which is a fair assumption for the timescales studied here. This yields the fractional power recovered by the time averaging, which we show in Figure 13, showing the frequency and baseline dependence of the signal loss. The top panels show this exercise on the **Vivaldi beam** simulations, while the bottom panels show the **Airy disk beam** simulations, which interestingly exhibit slightly more loss for fixed frequency and baseline length. Overall, Figure 13 tells us that if we wish to limit signal loss to, say $\lesssim 2\%$ in power, then the longest window over which we can coherently average the visibilities is ~ 500 seconds.

4. SUMMARY

We have performed numerical simulations of mock EoR fields with a series of primary beam models for HERA, including an idealized Airy disk with a 14-meter aperture, a CST simulation of the Vivaldi feed and HERA dish, and CST simulations of the translationally-perturbed Vivaldi feed and HERA dish. We use these to quantify the amount of signal loss induced when applying Fourier filters along the time axis of the drift-scan visibilities, in other words, when enacting various forms of fringe-rate filters. We quantify the following:

- We first compute the power spectral density (PSD) distributions of the EoR visibilities in the fringe-rate domain, showing that, compared to an idealized Airy disk, the Vivaldi beam model up-weights more low and negative fringe-rate modes, which expands the fringe-rate footprint of the PSD. We also show that 3 cm feed perturbations do not appreciably change these PSDs.

- We then look at cross-coupling high-pass filters, which reject near-zero fringe-rate modes to suppress cross-coupling systematics. We show that the extended footprint of the Vivaldi beam simulations make such filters more lossy, setting specifications for the lossiness of cross-coupling filters as a function of frequency and baseline length.

- We then look at signal loss introduced when coherently time (or LST) averaging the drift-scan visibilities, showing that signal loss can be kept to $\lesssim 2\%$ across the band when averaging the visibilities over a time window no larger than ~ 500 seconds.

- Lastly, in the appendix, we show how sensitive these metrics are to slight translation perturbations of the feed position, showing that 3 cm translational feed perturbations have a negligible impact on our results.

APPENDIX

Table 1. Perturbed Beam Models

Beam Name	X, Y, Z Offset [cm]
base	0, 0, 0
zshift	0, 0, -3
xshift	3, 0, 0

A. THE IMPACT OF PRIMARY BEAM PERTURBATIONS

Here we look at the impact that slight translational perturbations to the feed position have on the measured PSDs of Section 2. This is in theory important because changes to the primary beam weighting pattern on the sky will impact the fringe-rate footprint of the EoR power spectral densities, and thus impact how a fringe-rate filters interacts with the total EoR power in the visibilities. A example of this already demonstrated in this analysis are the large differences seen in the PSD between the idealized Airy disk model and that of the fiducial Vivaldi feed model. Here, we seek to understand how robust the simulated PSD distributions under minor translational feed perturbations. We look at two particular perturbations, summarized in Table 1 with a visual in Figure 2, which are consistent with observed feed movement in the field (Rath et al. 2021). Figure A.1 shows the directional response of the base beam in the low and midband, and the difference of the base model with the perturbed model, showing the kinds of directional structure sourced by the feed perturbation. Note that these simulations are performed in three distinct subbands (65–85, 160–180, and 200–220 MHz), because the perturbed beams are only available in these bands.

First we look to see how the feed perturbations impact the estimated PSDs. We show slices of the PSDs at fixed frequencies in Figure A.2, showing the fractional difference between the xshift model (top panels) and the zshift model (bottom panels) with the unperturbed base model for a few East-West baselines. We see that the fractional difference is generally $\leq 10\%$ between the base model and the perturbed model. We also compute the implied signal loss due to a crosstalk high-pass filter (analogous to Figure 9), and find that the fractional difference of this metric between the perturbed and unperturbed model is always $\leq 1\%$ for all baselines and frequencies. This suggests that the 3-cm feed perturbations have an overall negligible impact on the PSDs and the implied signal loss due to crosstalk filtering when compared to the base model.

Lastly, in Figure A.3, we compare the specifications derived for the amount of signal loss induced due to coherent time averaging (analogous to Figure 13). For the three representative frequencies in the low, mid, and highband, we show that the 3-cm perturbed beam models have a negligible impact on the coherent LST average signal loss specifications. showing that the fractional difference between the metric for the base model and that of the xshift model (top panels) and zshift model (bottom panels) are exceedingly small ($\leq 0.1\%$ for all baselines, frequencies, and integration windows).

Thus, we see that 3-cm translational feed perturbations have an overall negligible impact on all of the metrics looked at in this work, although we do not rule out the possibility that different kinds of perturbations (e.g. tip and tilt) or larger translational perturbations may have a larger impact.

REFERENCES

Ali, Z., et al. 2015, ApJ, 809, 61

Fagnoni, N., de Lera Acedo, E., Drought, N., DeBoer, D. R., Riley, D., Razavi-Ghods, N., Carey, S., & Parsons, A. R. 2020, arXiv e-prints, arXiv:2009.07939

Górski, K. M., Hivon, E., Banday, A. J., Wandelt, B. D., Hansen, F. K., Reinecke, M., & Bartelmann, M. 2005, ApJ, 622, 759

Kern, N. S., Parsons, A. R., Dillon, J. S., Lanman, A. E., Fagnoni, N., & de Lera Acedo, E. 2019, ApJ, 884, 105

Lanman, A. E., & Pober, J. C. 2019, MNRAS, 487, 5840

Lanman, A. E., Pober, J. C., Kern, N. S., de Lera Acedo, E., DeBoer, D. R., & Fagnoni, N. 2020, MNRAS, 494, 3712

Parsons, A. R., Liu, A., Ali, Z. S., & Cheng, C. 2016, ApJ, 820, 51

Rath, E., Dynes, S., Hewitt, J., & Molewa, M. 2021, Motion of HERA Antenna Feeds, Tech. rep., Massachusetts Institute of Technology, HERA Memo #98

Shaw, J. R., Sigurdson, K., Pen, U.-L., Stebbins, A., & Sitwell, M. 2014, ApJ, 781, 57

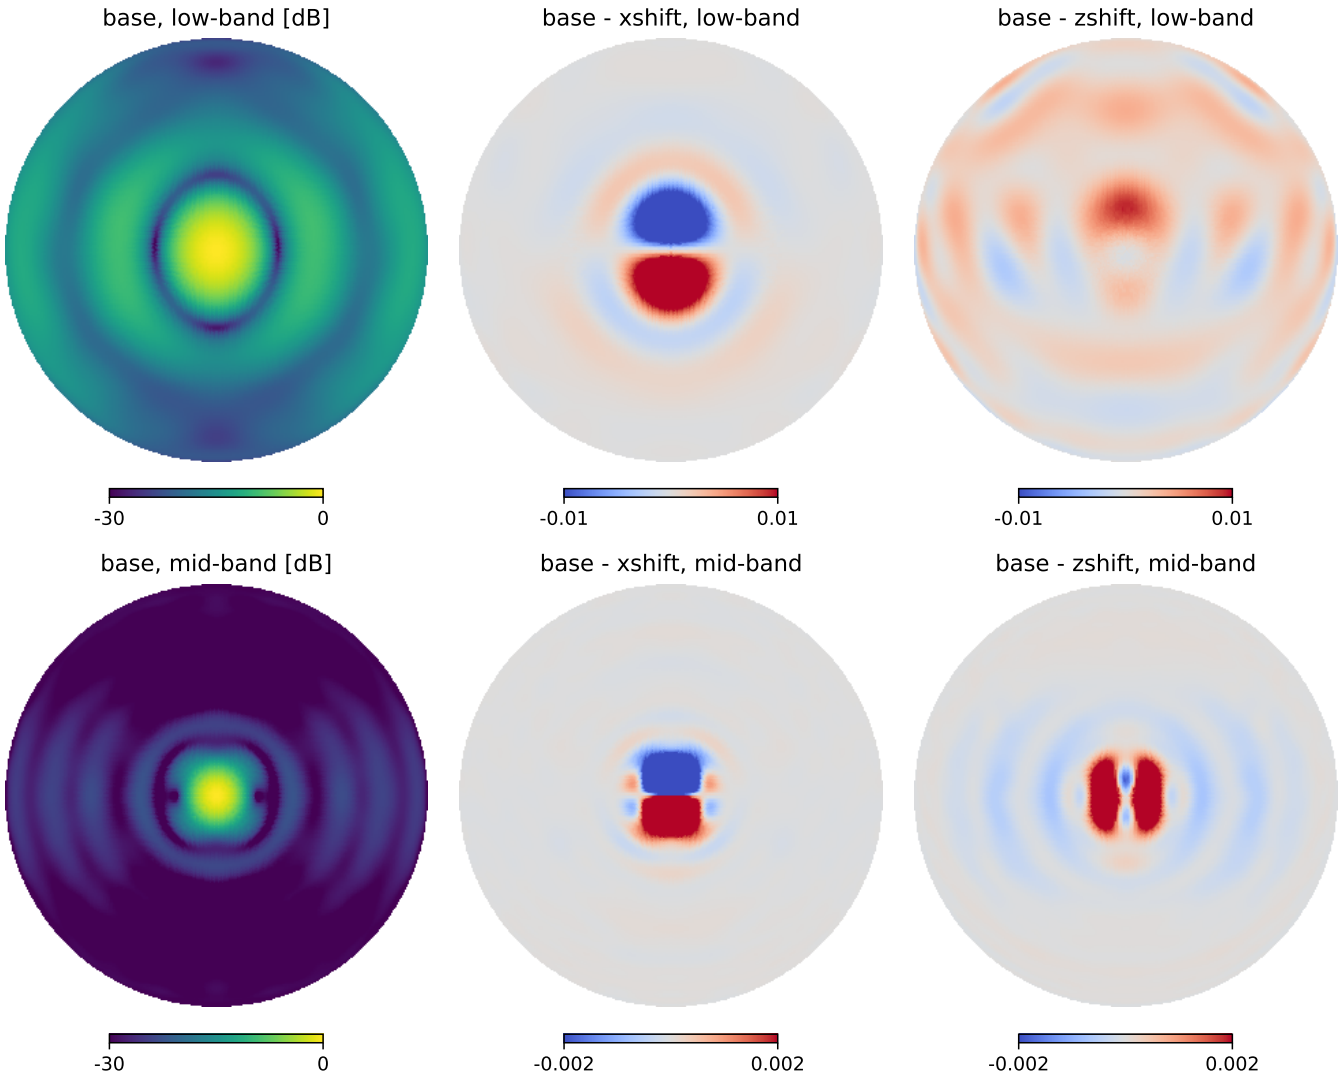

Figure A.1. This shows the base (unperturbed) Vivaldi beam model power response (left) in logarithmic units, and then shows the difference of the base beam with a x-displaced feed (center) and a z-displaced feed (right), showing the small but non-zero impact of the primary beam power pattern with such displacements.

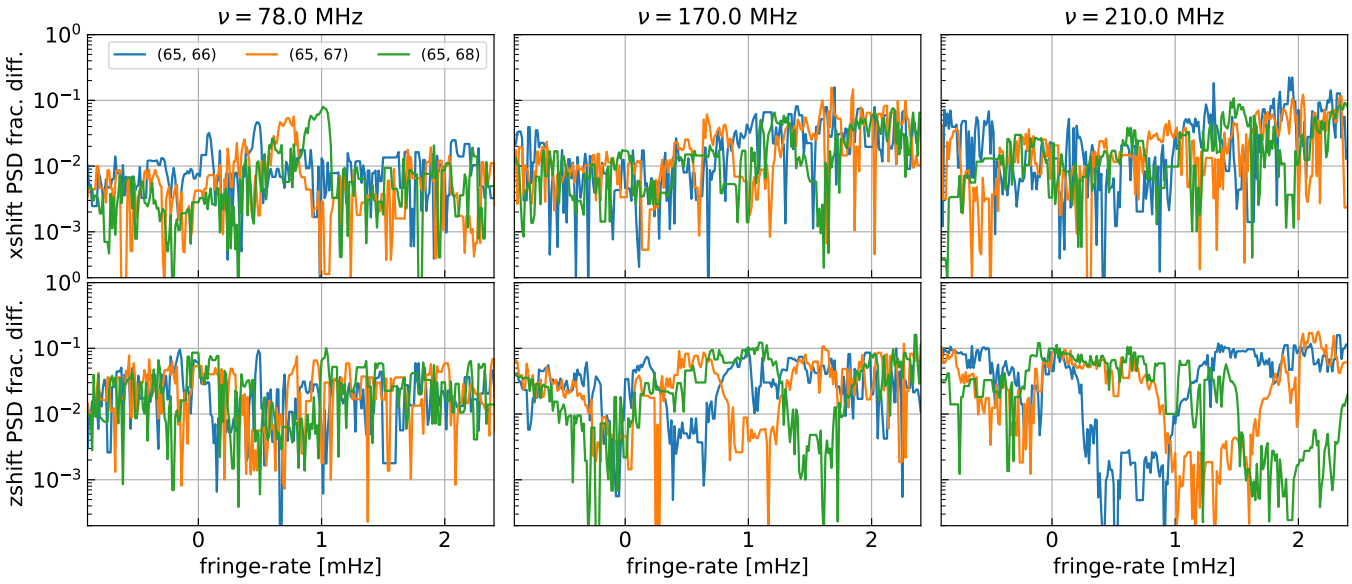

Figure A.2. The fractional difference between the base model PSD and the xshift (top-panels) and zshift (bottom-panels) PSD at three different frequencies, showing that the fractional difference is generally $\leq 10\%$. When comparing the expected signal loss due to high-pass fringe-rate filters from the perturbed PSDs (e.g. Figure 9), we find that they agree with the base model to within $\leq 1\%$ for all baselines and all frequencies.

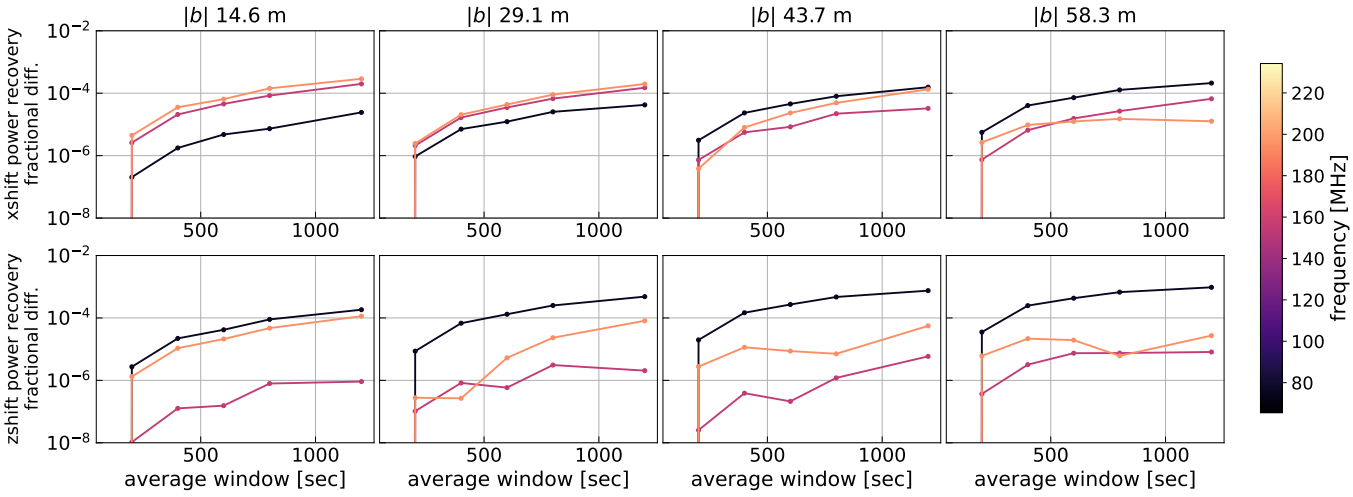

Figure A.3. The same LST averaging signal loss test as Figure 13 but now taking the fractional difference between the base model and the xshift perturbed Vivaldi beam model (top panel) and the zshift beam model (bottom panels). We see that the perturbed beam models have a negligible impact on the signal loss metric ($\leq 0.1\%$) for all frequencies, baselines, and integration windows.