HERA Memo #126
Received 29 Sept 2023

Cubic & Fourier Interpolation Over Frequency for the HERA Vivaldi Beam

Naomi Carl, Steven Murray

August 11, 2023

Abstract

This memo addresses the structural issue that arose during H4C end-to-end visibility simulations of the GSM of anomalously high power at high delay that obscured the cosmic signal. It examines different interpolation methods over frequency, spatial coordinates, and CPU versus GPU device usage. We find that in order to simulate delay-spectra with significantly improved dynamic range, the frequencies must be interpolated to at least the third order, (fifth order showed no improvement upon third order). Testing was also performed using a Fourier method for the frequency interpolation which proved to be far smoother than cubic interpolation. The issues arising from frequency interpolation are thus solved by cubic or Fourier interpolation, but new issues arising from spatial interpolation need to be explored in the future.

Introduction

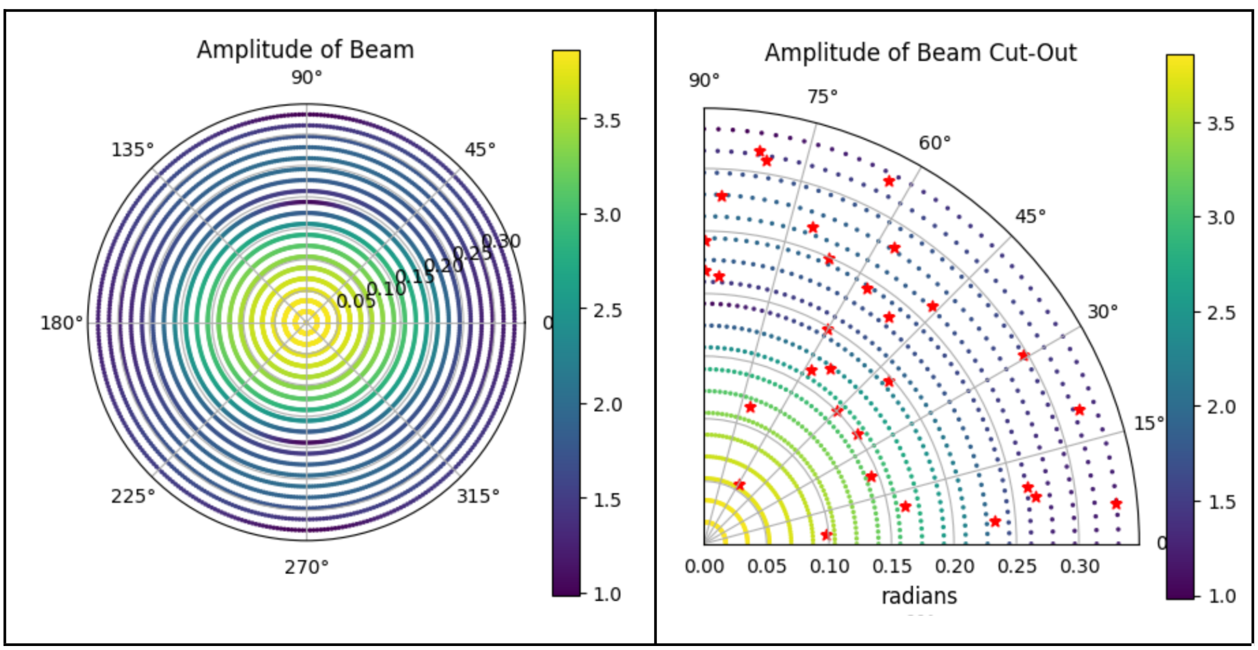

Figure 1: The left image demonstrates the beam visualization. The right image provides a magnified cut-out of a section of the beam with the point sources indicated as red stars. This illustrates how point sources often fall between the beam grid.

Interpolation is a necessary part of simulating HERA's beams onto locations in the sky. Recently, the HERA Validation team found that H4C end-to-end visibility simulations of the GSM had anomalously high power at high delay, obscuring the cosmic signal. While multiple potential causes were discussed, the most likely was due to beam interpolation -- either over frequency or angle. Simulation of the expected amplitude of sources on the sky depends on knowledge of the antenna beam. The best available models of the HERA beam have been calculated, using the CST E&M modeling program, on a regular angular grid but sources on the sky are at floating point locations. This problem is illustrated in Figure 1, where the point sources fall in between the grind points. Some kind of interpolation method is necessary to approximate the beam at the locations of known sources. In order to solve this problem, experiments were conducted running the simulation with various combinations of interpolation types. In this memo, the issue in the original simulation data will be discussed, methods of solving this issue will be explored, and conclusions will be made on how to correctly interpolate HERA's beams to reduce chromatic structure in simulated visibilities.

Implementation

The HERA Validation team has used the `vis_cpu` visibility simulator, which is driven by the `hera_sim` package, and depends on `pyuvdata`, to perform critical end-to-end simulations for H4C (and soon H6C). `hera_sim` must evaluate the beam attenuation at each sky-location and frequency at which a sky-model source appears at any given mock observation time. The original coordinates onto which the beam is simulated are at fixed, regular angles and thus requires an analytic beam model fit to the fixed simulation (De Los Santos 2022) or some form of interpolation. Recent simulations of the diffuse foregrounds for H4C validation were found to exhibit excess power at high delay, prompting an investigation into the cause of the spurious structure. There was reasonable suspicion that the issue lay with the interpolation methods for the beam, and our goal for this project was to find a way to correctly interpolate the beam so that it wouldn't create spectral structure on a cosmological scale that obscures the data. H4C uses the new Vivaldi feeds, so we use the Vivaldi CST simulations from Fagnoni+19 as our beam model for this study. There are several factors that play into the interpolation methods for the HERA Vivaldi beam, including CPU versus GPU devices, the frequency interpolation, and the spatial interpolation. All of these components are tested separately in this memo.

The Issues

The relevant anomaly in the simulated H4C GSM visibilities is clearly illustrated in Figure 2, in which two "lumps" are present at $\pm$1000 ns in the delay spectra, which is well outside the foreground wedge. These are evidence of artificial structures that are not seen in the real data. This means there is something wrong with the way the simulation is running. We suspected the issue was caused by the interpolation method of the beam. Artificial structure and non-smoothness in the structure could be caused by other things, but it would mainly be indicative of poor interpolation, since the goal of interpolation is to connect points smoothly. In linear interpolation, jagged lines connecting the points would lead to unsmooth structures, but the magnitude of non-smoothness that these structures could cause, and at which delays they would appear, was unknown. The lumps in the delay spectra in Figure 2 significantly interfere with the data on a cosmological scale.

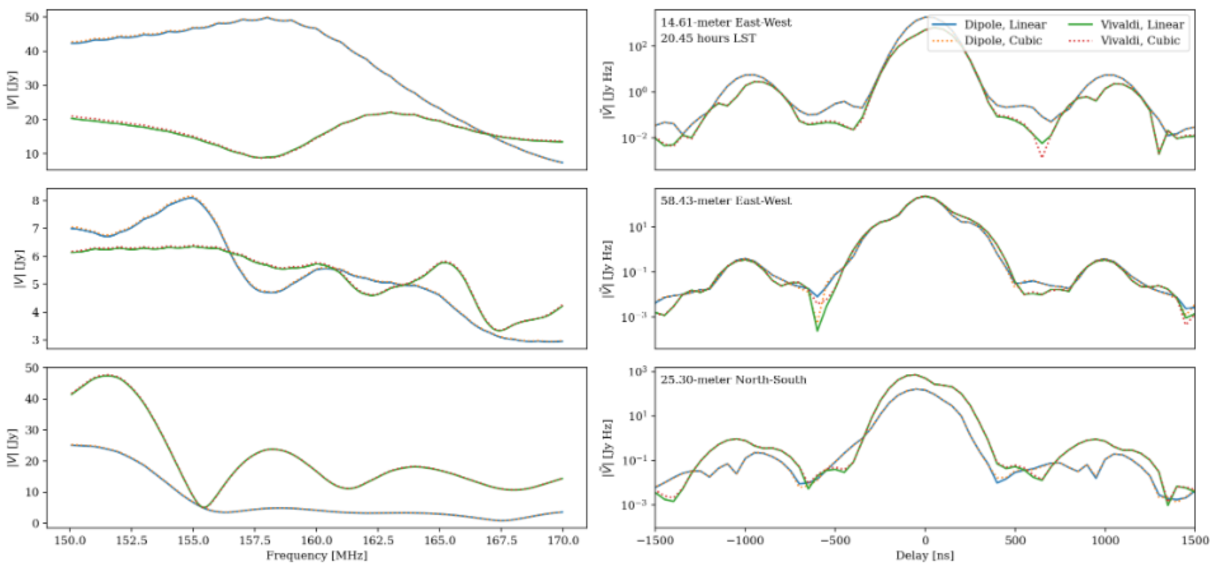

Figure 2: Visibility (left) and delay (right) spectra created previously for the HERA Validation project. This paper focuses on the Vivaldi beam, but the Dipole beam is included as well. In the delay spectra, there are “lumps” located on either side of the main lobe, in both linear and cubic interpolations. [Figure created by Bobby Pascua.]

Furthermore, the amplified high-delay power revealed in the right-hand panels of Figure 2 is, if you look closely, visually manifested as 1MHz-scale "waves" in the frequency-spectrum shown in the left-hand panels. The scale of these waves, at 1 MHz, is a clue that the issue might be the frequency-interpolation, as the input beam model was simulated at 1 MHz spacing. Thus, these waves are suspected to be an effect of linear interpolation. Additionally, a potentially separate problem with frequency discontinuities shows up in at least a few of the visibility plots; “blips” are visible in the Diffuse Foregrounds graph in Figure 3. Although interpolation is not an obvious causation or solution to these blips, the hope is that they will be corrected with a new interpolation (or perhaps just a re-run) of the simulations.

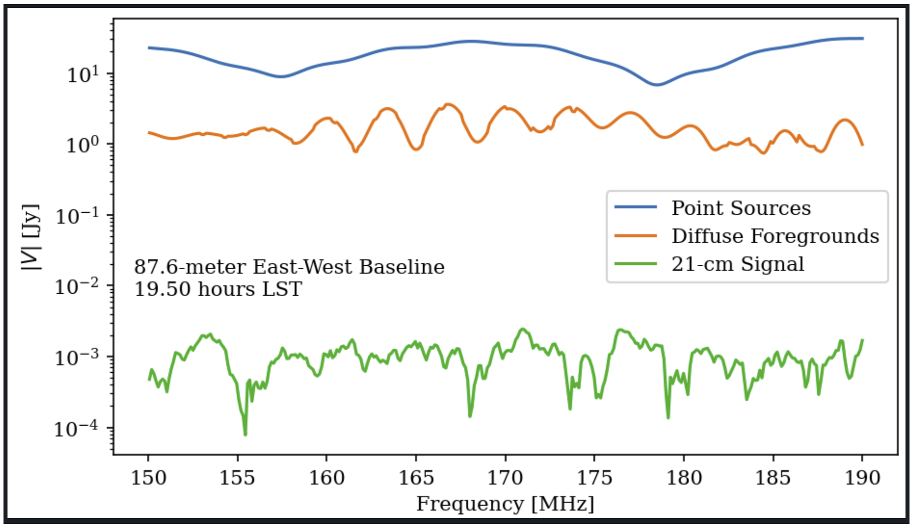

Figure 3: The visibility spectrum of the Diffuse Foregrounds has noticeable blips in the data that need to be investigated. Though the blips are rather small, it is important to understand where they are coming from and how they may affect the data. [Figure created by Bobby Pascua.]

Beam Interpolation Over Frequencies

We begin by simply testing methods of interpolation over frequency on the beam itself, selecting certain pixels to check the simulation, the results of which are shown in Figure 4. In the case of the delay spectra, cubic and quintic interpolation are significantly better than linear interpolation. Quintic interpolation appears to be smoother than cubic interpolation, though there is some evidence that the smoothness is saturating at the fifth order. The residuals compared to the quintic beam, for both real and imaginary components, are what we expect to see for both interpolation types over both feed and component, with linear having higher residuals than cubic. Although the higher-order frequency interpolation clearly reduces high-delay power, it is not yet clear how strongly this translates to full visibilities from a realistic sky model. Furthermore, we do not see the "1 MHz wave" structure that was observed in the H4C visibility simulations. To answer these questions, we turn to full visibility simulations.

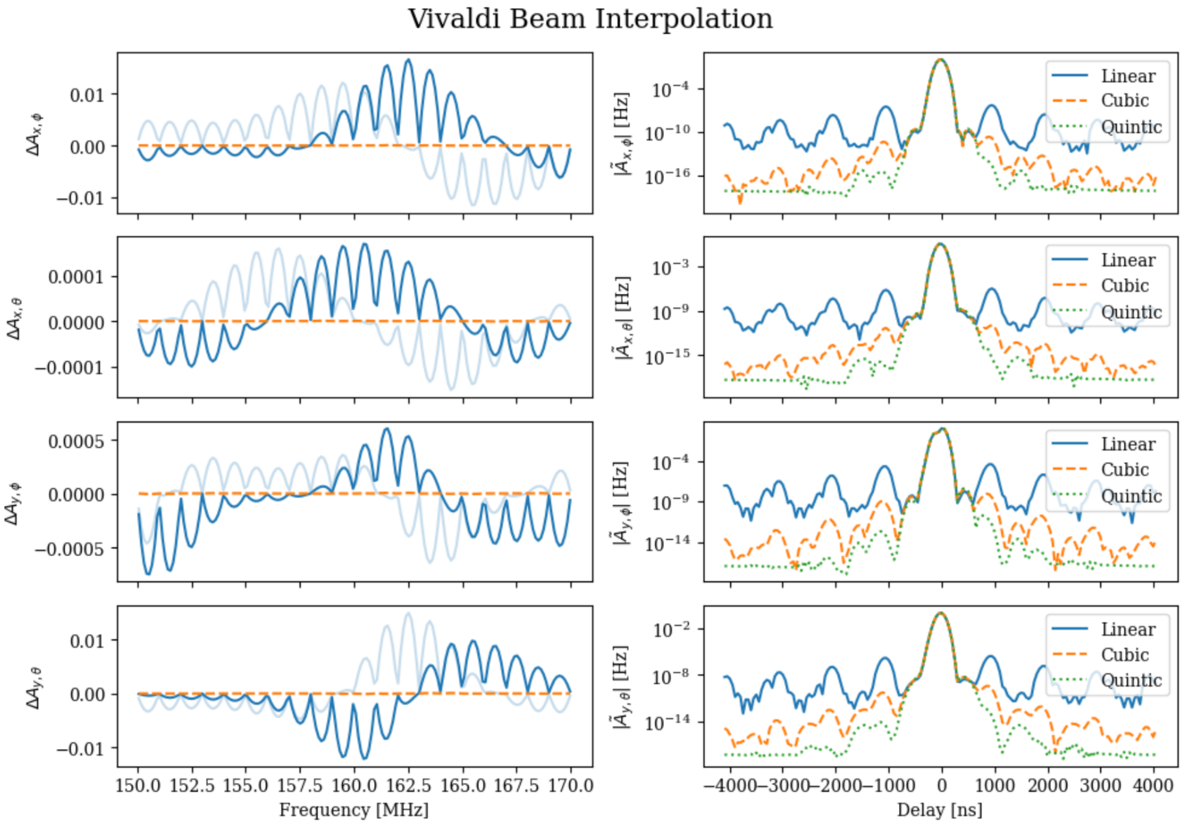

Figure 4: Various spline interpolations of the Vivaldi beam at a particular angular pixel, interpolated over frequency. Each row shows a different feed and vector component while different colors show different spline orders. The pixel here is at: theta=90°, phi=90°, but other pixels were tested and show qualitatively similar results. Left-hand panels are the interpolations in frequency space, and right-hand panels are the magnitude of the delay transforms. Linear interpolation has the highest magnitude of non-smoothness, while both cubic and quintic are significantly improved, with quintic being slightly smoother than cubic.

Fixing the Issues

To continue, we ran visibility simulations numerous times with different combinations of interpolation methods (all using `vis_cpu`), the results of which are displayed in Figure 5. The first thing we tested was if the simulation was running differently on GPU than CPU. The differences between these simulations were negligible, so only two examples of GPU-based simulations are shown in the figure. The next test was to perform linear and cubic spatial (‘xy’) interpolation. The differences in the high-delay power were negligible, not significant enough to be creating large lumps in the delay spectra, but it does create a significant disparity in the residuals, particularly in the center panel on the left; understanding this disparity is left for future work. However, when comparing frequency interpolation methods, we found that the high-delay power was significantly reduced for the cubic and quintic cases, in both cases yielding 5 orders of magnitude of dynamic range. The simulations suggest that increasing the spline order beyond 3 is not effective, with the quintic interpolation showing qualitatively similar delay spectra. After investigation, it was found that the configuration file used to run the H4C GSM simulations, (but not the EoR or point-source sky models), had erroneously set frequency interpolation to be linear, rather than the default (cubic), thus creating the noted issues.

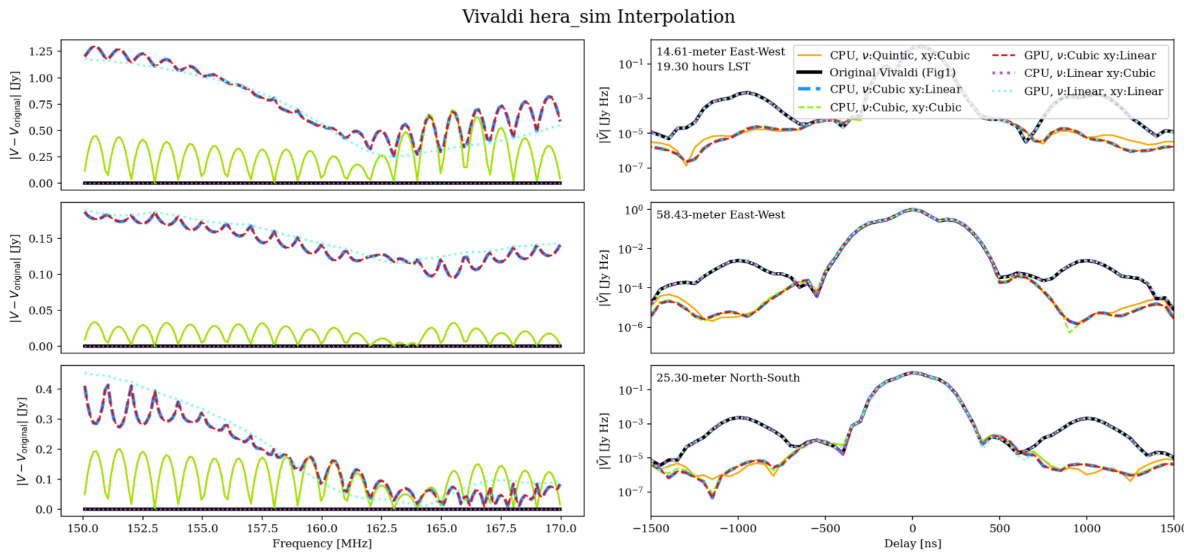

Figure 5: Graphs displaying the residuals (left) and delay (right) spectra of the HERA Vivaldi beam interpolating using various methods through `hera_sim`. Simulations that used linear frequency interpolation are indicated with dotted lines. Simulations that used cubic frequency interpolation are indicated with dashed lines. Red, blue, and teal show linear spatial interpolation simulations, with noticeably higher residuals. The simulation that was run using quintic frequency interpolation is marked with a solid orange line.

Furthermore, resolving this issue appeared to clear up the blips seen in the diffuse beam in Figure 3, as shown below in Figure 6. While the problem seems to have been resolved for the baseline, LST and frequency range in which the problem originally appeared, there is a possibility that the blips are a random issue in the simulation that may show up at other LSTs, frequencies, or baselines. If the blips show up again, the issue can be addressed on a case by case basis, but it seems the blips are fixed.

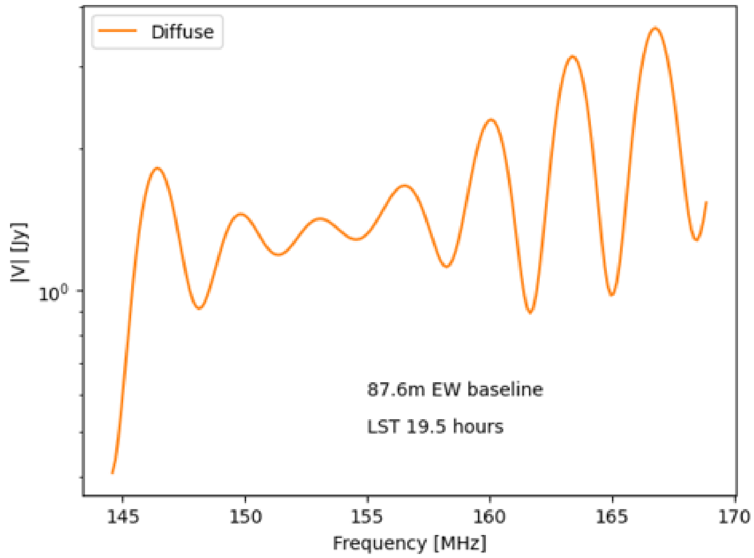

Figure 6: A recreation of the diffuse beam visibility graph in Figure 3, after the adjustments were made to correct the interpolation method of the H4C GSM. The blips are not visible in this graph.

As was mentioned before, the erroneous linear interpolation setting was not present in the EoR configuration file. While it is clear that using the correct frequency interpolation is vital for foreground simulations, we wanted to see the effects frequency interpolation has on the EoR. The EoR sky model was thus tested with linear, cubic, and quintic frequency interpolation methods. The differences between the simulations run using these different methods had negligible effects on the visibility/delay spectra, so much so that it is not noticeable in the results graph in Figure 7. This provides evidence that frequency interpolation methods can be evaluated merely based on smooth foregrounds.

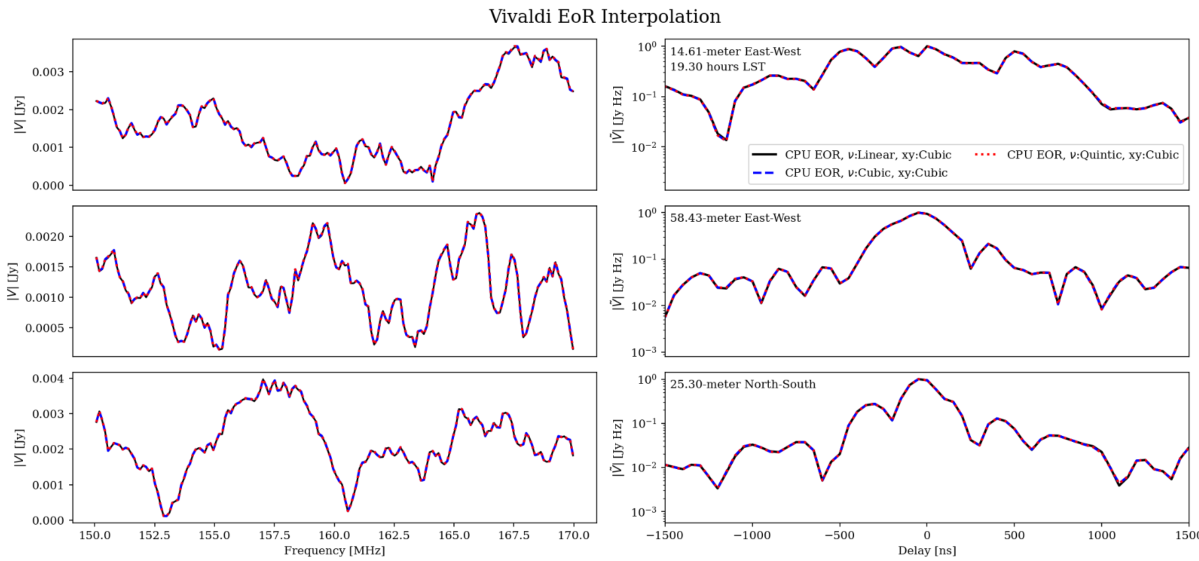

Figure 7: The visibility (left) and delay (right) spectra of the EoR sky simulations tested with linear, cubic, and quintic frequency interpolation. Though the lines appear to match up exactly, there are miniscule disparities between each graph, which we conclude are negligible.

Fourier Interpolation

Though cubic interpolation satisfied the main concerns arisen from this issue, there was interest in whether or not more advanced methods of interpolation, such as Fourier, would create even smoother spectra. This involved taking the fast Fourier transform over the frequency of the original beam, padding the Fourier-space by a factor of ten with zeros, and, finally, performing an inverse Fourier transform to return the original signal with resolution power of ten times the original. The results of this method on a particular pixel of the beam are displayed in Figure 8 (cf. Fig. 3). Despite the improved smoothness that resulted from this, there is reasonable concern that the non-periodicity of the beam itself may break the Fourier interpolation, particularly at the ends of the spectra; the shoulders are evidence of this. Several tapers were considered to eliminate these shoulders, but we decided upon a Gaussian taper, as a Gaussian taper can be corrected for exactly when we go to higher resolution. This may have been possible with other tapers, but Gaussian was simple, and, unlike other tapers, it is not dependent on the precise number of channels, rather it is dependent on the position of those channels within the range of frequencies, making it easier to correct for. So, a Gaussian taper was added to the Fourier interpolation, as it was a simple first choice, and it eliminated the issue of the shoulders in the Fourier interpolation.

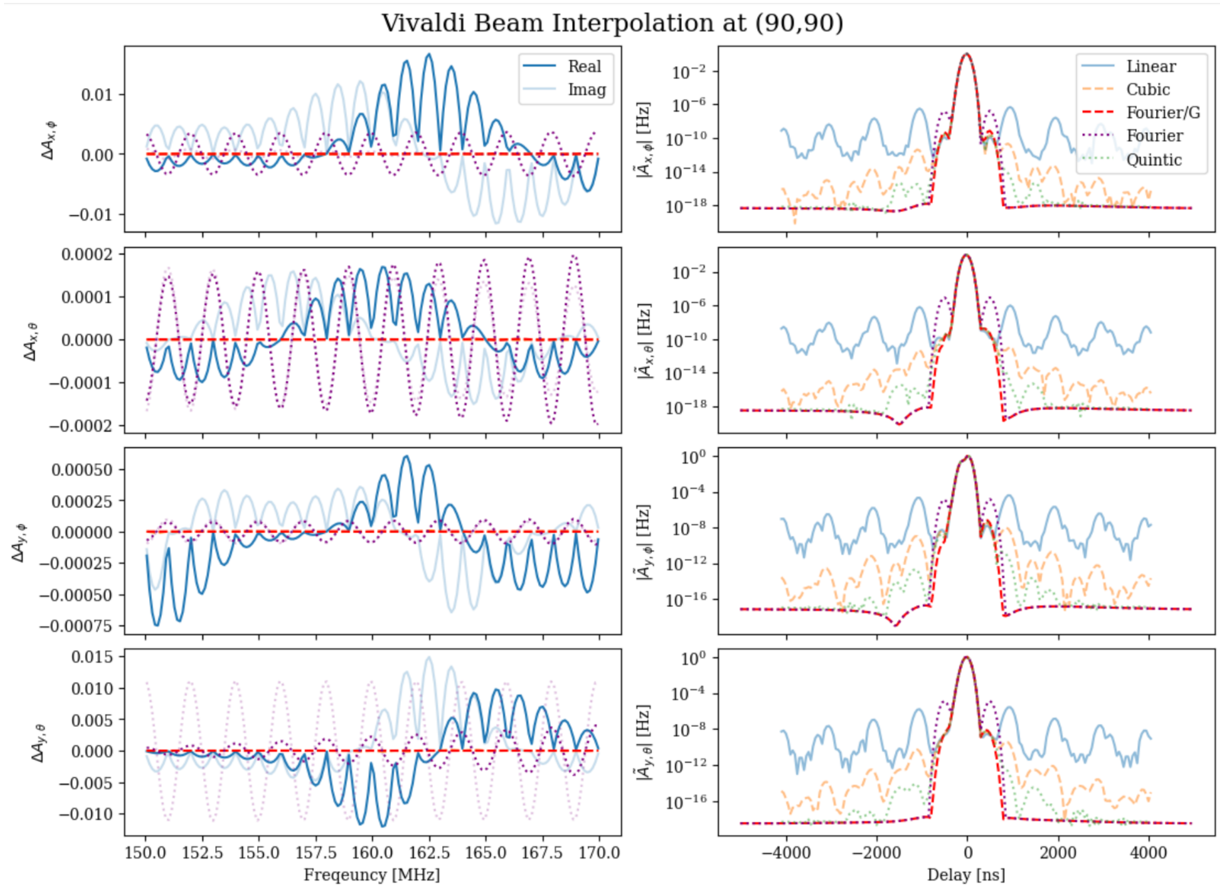

Figure 8: This is the same as Figure 4, graphing a pixel at (90,90), with the addition of two methods of Fourier interpolation; Fourier/G, in red, represents Fourier interpolation with a Gaussian taper, while the other Fourier, in purple, has no taper. The issue with the untapered Fourier interpolation can be seen in the significant shoulders off of the main node in the spectra, which are effectively reduced when we apply the Gaussian taper. Both appear much smoother than the other methods of interpolation, though there is concern the ends of the spectrum may not be accurate due to the nonperiodic nature of the beam.

This testing of the Fourier interpolation method was extended to the simulation, which is shown below in Figure 9. Fourier interpolation with a Gaussian taper yields far smoother results, nearly four orders of magnitude lower than cubic interpolation. This shows promising potential for this method, so we will be adding the option for Fourier interpolation (with the Gaussian taper as default) to `hera_sim`. However, there are some caveats to the Fourier method and other things that should be tested, which we will discuss next.

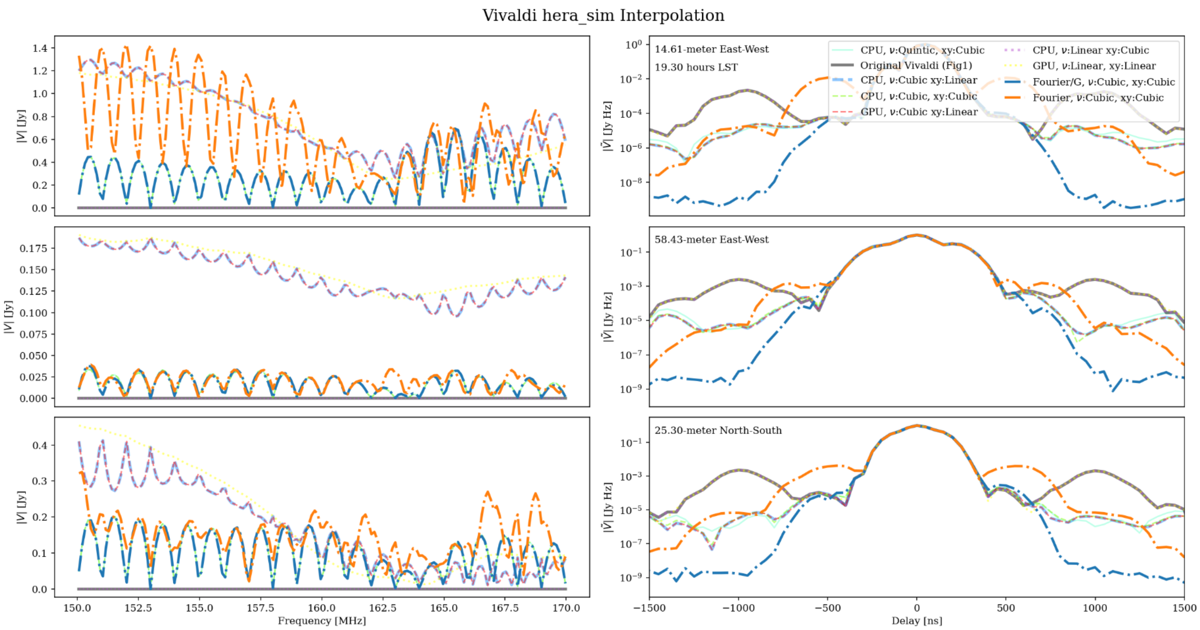

Figure 9: This is the same as Figure 5, displaying the residuals and visibility/delay spectra, with the addition of two methods of Fourier interpolation; Fourier/G, in blue, represents Fourier interpolation with a Gaussian taper, while the other Fourier, in orange, has no taper. This shows that the patterns displayed in the beam in Figure 7 apply to the simulation data. The delay spectra graphs show that Fourier interpolation with the Gaussian taper reduces the data by about four orders of magnitude more than the cubic interpolation method.

There is concern that this smoothness is not necessarily accurate, so further testing comparing these results and a downsampled version of these results were compared to the original cubic interpolated beam, this is shown in Figure 10. The downsampled Fourier interpolated simulation only has every other point, compared to the regular Fourier interpolated simulation that uses all the given points. Both of these simulations are then compared to the cubic interpolated simulation, which we know goes through the simulated data points and is therefore accurate in the vicinity of the simulated frequencies. The results of Figure 10 imply that the smoothness created by the Fourier interpolation method (with the Gaussian taper) is accurate, as the difference in visibility matches up every 1 and 2 MHz, where the original given points are. When using the taper, there is near exact correspondence with the cubic interpolated beam every 1 MHz, meaning there is a periodic disagreement that is less than 1%. Given that the Fourier interpolation is by construction smooth, deviations of the Fourier interpolation from the cubic interpolation imply a commensurate level of structure in the cubic interpolation. The fact that the Gaussian tapered interpolations have very little difference from each other indicates that the intrinsic beam model doesn't have structure below the ~2 MHz scale, which further justifies the use of the Fourier Gaussian model. This is a good sign for the legitimacy of the Fourier interpolation method and further displays a taper is necessary for use of the Fourier method.

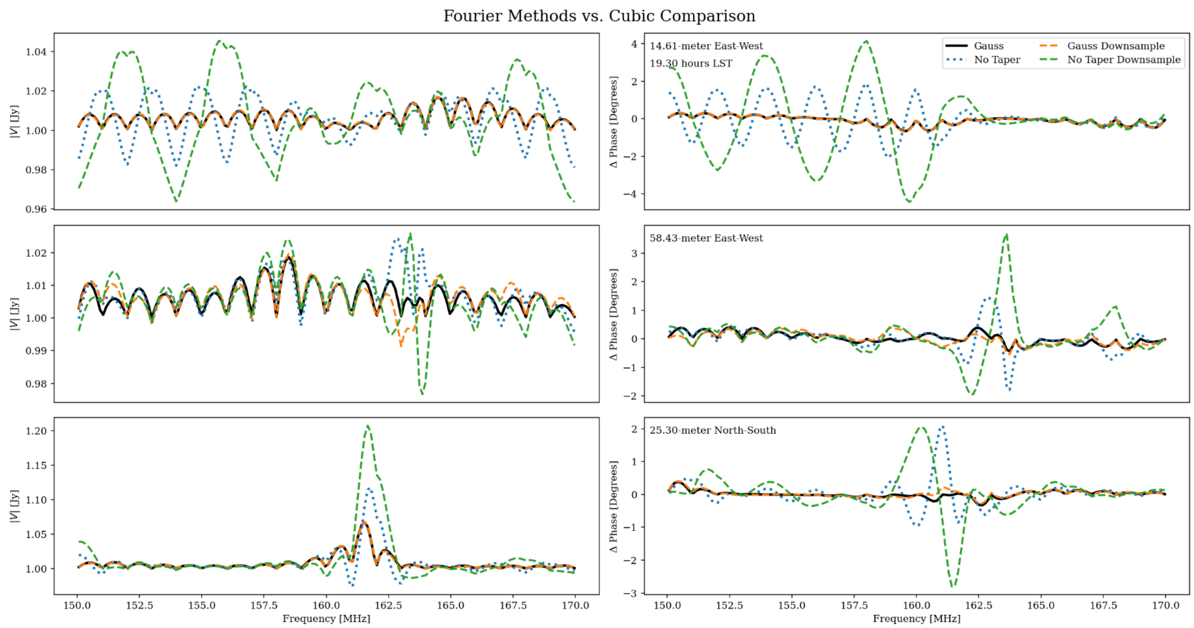

Figure 10: For the graphs on the left, the ratio of magnitudes of each method, listed in the legend, was compared to the original cubic interpolation method. The Fourier interpolations without tapers, whether downsampled or not, performed worse than the Gaussian tapered samples. The right side displays the difference in phases between the four methods and the cubic method. There are a few anomalies present, particularly in the bottom visibility graph around 160 MHz, which may need further investigation.

The non-periodic nature of the beam is in combat with the periodicity on which Fourier interpolation relies; this causes issues at the end of the spectra, which is exemplified in the left hand plots in Figure 11. The Fourier method fares well until the ends of the spectrum, where it deviates several orders of magnitude from the cubic spline method. Though this is not ideal, the effectiveness of the Fourier method at the peak of the spectra may cause it to be preferred over cubic interpolation, despite the disparity at the edges. We can see in the delay spectra in Figure 11 that the Fourier interpolation is smoother and has a noticeable higher difference in magnitude, particularly on the highest baseline, 58.43-meter. There are pros and cons to both the cubic and Fourier interpolation methods to be considered.

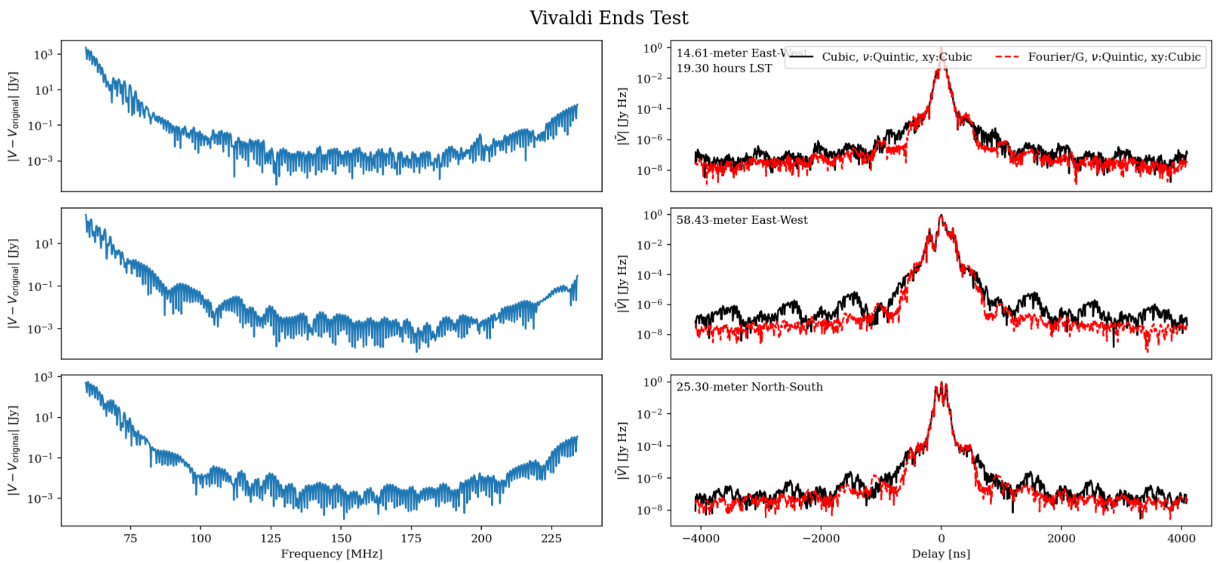

Figure 11: The right plots display the delay spectra for the cubic spline interpolation method and the Fourier interpolation method across the full range of frequencies, rather than just the peak area of the spectra as was done in earlier tests. The left plots show the magnitude of difference between the Fourier and cubic methods (log-scaled). At 75 MHz and below, there is a spike in the Fourier residuals that we would expect from the non-periodicity, the magnitude of difference jumping up by several orders of magnitude. However, for the delay spectra, the Fourier interpolation is at least a magnitude lower than the peak in the spectrum, showing its benefits.

Remaining Spatial Interpolation Issue

Due to the disparity that seemed to be caused by the spatial interpolation method in Figures 5 and 8, it was decided to run additional testing on spatial interpolation. Just testing linear vs. cubic spatial interpolation, with one frequency (156.7 MHz) over 200 five-second intervals, graphs of the visibility, difference in phase, and fringe rates were created and shown below in Figure 12. The primary concern arising from this additional testing is the “rumble” present in the center visibility graph. It is possible that this is naturally occurring from sources popping in and out of the horizon, but it could also be due to the interpolation method or the fringe and resolution of the simulation. Additionally, there is a shoulder that appears in the linear interpolation for the center fringe rate graph, which is somewhat concerning.

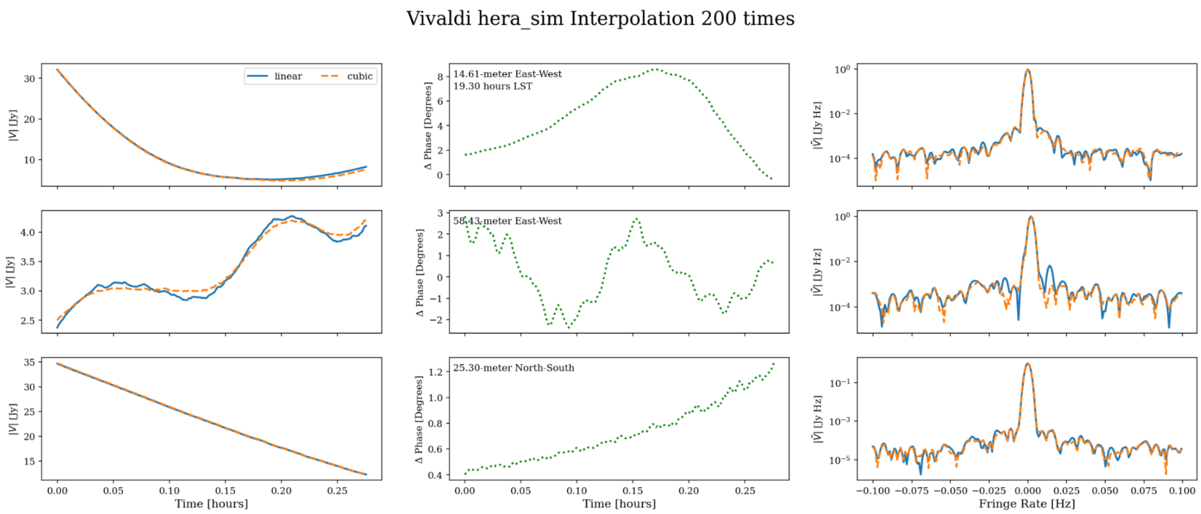

Figure 12: This simulation ran both cubic and linear interpolated beams at one frequency (156.7 MHz) over the length of 200 time intervals of five seconds. The leftmost graph displays the beam; the center graph shows the difference in phase between the linear and cubic interpolated beams; the right graph shows the fringe rates of the sample. The main concern from these results is the "rumble" in the center visibility graph, for the 58.43 meter baseline.

It has been decided that, since GPU is only capable of linear spatial interpolation, that H4C will use linear spatial interpolation. However, this is definitely an issue that should be further explored going forward. Different beam interpolation methods, including the Fourier method used for the frequency interpolation, should be tested to see if it resolves the issue. As explored by De Los Santos in 2022, a spherical harmonic upsampled model for the spatial interpolation may prove useful. Furthermore, to determine if the rumble is naturally occurring, we should use a beam that is artificially smoothed down to zero at the edges; there shouldn't be any bumps in this case, but if there are, we will know the rumble is not naturally occurring. Different resolutions for the simulation should also be taken into account, as well as the fringe rate. For example, a test could be to take the sky model at half of its current resolution and see the effects it has on the rumble.

Conclusions

This memo investigated and resolved the delay-structure issue, tested Fourier interpolation methods as opposed to spline methods, and began to investigate discrepancies within the spatial interpolation. The results are summarized below.

Solving the delay-structure problem with H4C:

- When looking at simple delay-spectra of individual beam pixels, third (or higher) order splines appear to be required to eliminate artificial structure in the data.

- Cubic interpolation over the frequency was effective in eliminating the artificial structure in the visibility delay spectra.
- Visibilities simulated with cubic and quintic frequency-interpolation show no significant difference in either frequency- or delay-space.
- CPU versus GPU run simulations were also tested and their differences were negligible.
- The issue in the original visibility and delay spectra was due to the configuration file used to run the H4C GSM simulations being set to linear interpolation over the frequency, as opposed to the default setting, cubic.
- Frequency interpolation has negligible effects on the EoR (since the dominant contribution to non-smoothness for the EoR comes from the sky model itself), so the frequency interpolation methods can be tested based on the smoothness of the foreground.
- Since cubic interpolation yields 5 orders of magnitude of dynamic range in the magnitude of delay-transformed visibilities, it has been deemed provisionally sufficient for H4C simulations.

Fourier-based interpolation

- An option for Fourier interpolation should be added to `hera_sim`, since it has proven useful in several of our tests, though cubic interpolation may be better in some aspects.
- For the Fourier method to be effective, a taper must be used, we used a Gaussian taper.
- The beam model doesn't seem to have structure lower than 2 MHz, which further justifies the use and accuracy of the Fourier method.
- There is a spike in the Fourier interpolation residuals under 75 MHz, but the benefits of its behavior around the peak are valuable despite the issues caused by non-periodicity

Spatial interpolation:

- Both types of spline interpolation yielded nearly identical visibilities, except for a slight (spectrally smooth) offset, noticeable in frequency-space.
- Particularly at the highest baseline tested (58.43m), there is a rumble and strange structure present in the spectra, the extent of which is not yet known.
- Although it will not be investigated in H4C, spatial interpolation appears to be an issue that is in need of more exploration going forward.

Overall, the main thing to take away from this study is that HERA simulations should be run with third-order or higher interpolation over frequency. However, there are some caveats to these results:

- It is possible that interpolation methods are altogether unsuitable for smooth foreground simulation, and that more sophisticated modeling methods are required (eg. models that are tuned to be compact in Fourier space). This might be required to achieve greater than 5 orders of magnitude in dynamic range.

- Not every individual frequency, antenna pair, or pixel was tested for each experiment; though unlikely, it is possible that there are disparities in the simulations.
- The blips seem to have disappeared, but since their origin is uncertain, we can't be sure that they won't pop up again at different LSTs or baselines.
- For Fourier interpolation, there may be better tapers or alternative methods that are more effective or accurate.
- The rumble indicates that there may be an issue in the method for spatial interpolation, the fringe rates, or the resolution of the simulation, or it may be nothing.

Though using cubic frequency interpolation fixes the main problem that motivated this work, and Fourier interpolation seems useful as well, there are many potential avenues for improvement in the future:

- Complex interpolation methods, other than Fourier, or different variations of Fourier, could be tested to see the effects it may have on the simulation.
- Going beyond H4C, methods other than linear interpolation should likely be used for the spatial interpolation; spherical-harmonic basis functions may be a beneficial option.
- The spatial testing done shows there is need to see if the rumble in the results is naturally occurring or arising from an error in `hera_sim`. Ideas discussed in this memo include: using a beam artificially smoothed down to zero or using half the resolution of the current simulation to see the effects on the rumble.

Acknowledgements

Thank you to Dr. Danny Jacobs and Dr. Jonathan Pober with the RENTU program at Brown University and Arizona State University and the HERA Validation team for providing support for this research, and thank you to Bobby Pascua and Piyanat Kittiwisit for their collaboration.

References

De Los Santos, C. (2022). Spherical Harmonic Decomposition of HERA's Electromagnetic Primary Beam Simulations. HERA Memo #114.