

A TEST OF OPTIMAL MAPPING WITH SIMULATIONS OF A POINT SOURCE

Jacqueline N. Hewitt, Zhilei Xu, and Honggeun Kim

21 January 2021

ABSTRACT

We implement optimal mapping for simulated HERA H320 observations at 160 MHz. In this first test of optimal mapping we model only one point source and do not include the effects of noise or calibration errors. We assess the response of the instrument and the mapping algorithm in a large section of the HERA stripe spanning 186° in right ascension and 14° in declination. We recover the flux density of the point source, and the maps provide a quantitative assessment of the contribution of source grating lobes and near and distant sidelobes. We also compute the expected response of the instrument and algorithm to the point source by propagating the point source in the image plane through the instrumental point spread function matrix. The agreement between the simulated map and the model is very good, and we attribute the differences to the effects of pixelation of the sky model. The peak values of the maps and the models agree to better than 1%. In a pixel-by-pixel comparison, the maximum difference between the map and the model in image space is 10% for a HEALPix grid with NSIDE=512, 2% for a HEALPix grid with NSIDE=1024, and 4% for a HEALPix grid with NSIDE=512 convolved to $30'$ resolution.

1. Introduction

Many imaging algorithms are used to estimate sky maps from interferometric data. Historically, because of limitations in computation power, imaging algorithms have incorporated approximations to make the computation tractable. Some of these approximations are heuristic in nature, and involve subjective choices of algorithm parameters and of convergence criteria, obscuring any assessment of the statistical properties of the resulting map. As more computation power becomes available, it becomes feasible to take a fresh look at interferometric imaging and to develop more direct estimators of the sky brightness. These approaches can have the advantage that the statistical properties of the resulting sky image are well understood and can be rigorously quantified. One such approach is “Optimal Mapping” as discussed by Tegmark (1997), Morales and Matejek (2008), Dillon et al. (2015; henceforth D2015), Zheng et al. (2017) and others. Mapping performance was one of the design criteria for the HERA array configuration, and is extensively discussed within the

context of HERA by Dillon and Parsons (2016).

In this memo we present optimal maps of simulated HERA observations of a 1 Jy point source during a 12-hour period. These computations demonstrate some basic features of optimal mapping with the HERA-320 array, and provide a check of the mapping software. We explore the response to the point source over a 7.7° (radius) circular cap field of view centered on zenith at the time of observation. This includes maps with the source in the main lobe of the synthesized beam, maps with grating lobes due to aliasing that extend for many degrees across the sky, and maps dominated by the complicated sidelobe structure that is present even 90° away from the source. In these simulations and in the mapping we model the motion of the primary beam across the sky, and we demonstrate that, at least in this simple case, the flux of the source can be recovered by properly taking into account the weighting of visibilities due to the sampling of the array and the primary beam. For each map we also compute a model of the sky by propagating the point source position and flux through the point-spread-function matrix, and we quantify the differences between the sky map and model that we expect due to the numerical approximations that must be made to make the computation tractable. We expect that this work will form a foundation for studying more complicated and realistic models of the sky.

2. Review of Optimal Mapping

Following the notation and discussion of D2015, we write the radio interferometry equation as

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{n} \tag{1}$$

where \mathbf{y} is a column vector of visibility data, the “design” or “pointing” matrix \mathbf{A} is a two-dimensional matrix describing the instrument, \mathbf{x} is a column vector of sky pixels, and \mathbf{n} is a column vector of noise values associated with each visibility measurement. It can be shown that an optimal (in the sense of preserving information) estimate of the sky is

$$\hat{\mathbf{x}} = \mathbf{D}\mathbf{A}^\dagger\mathbf{N}^{-1}\mathbf{y} \tag{2}$$

where \mathbf{N} is the noise covariance matrix $\langle n_i n_j \rangle$ and \mathbf{D} is an invertible normalization matrix. It can be shown that the expected value of the estimator is

$$\langle \hat{\mathbf{x}} \rangle = \mathbf{D}\mathbf{A}^\dagger\mathbf{N}^{-1}\mathbf{A}\mathbf{x} \tag{3}$$

from which the point spread function matrix

$$\mathbf{P} = \mathbf{D}\mathbf{A}^\dagger\mathbf{N}^{-1}\mathbf{A} \tag{4}$$

follows. If we choose

$$\mathbf{D} = [\mathbf{A}^\dagger \mathbf{N}^{-1} \mathbf{A}]^{-1} \quad (5)$$

then $\langle \hat{\mathbf{x}} \rangle = \mathbf{x}$ and we have deconvolved the image; i.e., we have removed the effect of the instrument. In practice, the matrix to be inverted for deconvolution tends to be singular or nearly singular, and deconvolution is not possible. Certainly for the case we consider in this memo, snapshot images of HERA data, the matrix is ill conditioned (due to the aliasing), and we do not attempt to invert it.

The covariance matrix gives the covariance of the estimator $\hat{\mathbf{x}}$ and is

$$\mathbf{C} = \langle (\hat{\mathbf{x}} - \langle \hat{\mathbf{x}} \rangle)(\hat{\mathbf{x}} - \langle \hat{\mathbf{x}} \rangle)^\dagger \rangle = \mathbf{P} \mathbf{D}^\dagger \quad (6)$$

We choose \mathbf{D} to be the identity matrix, and leave more sophisticated choices of \mathbf{D} for future work. At times, for purposes of display only, we choose \mathbf{D} to be composed of the diagonal elements of \mathbf{P} , which has the effect of converting the units of the maps to Janskys (equivalent to the “primary beam correction” of radio astronomy).

The above formalism is elegant and mathematically well defined, but in practice not computationally feasible for maps that finely sample the entire sky above the horizon. As worked out in detail by D2015, for HERA it should be feasible to limit the computation to maps that cover “facets” of the sky. We match the field of a view of a facet to the width of the main lobe of the primary beam by introducing an additional matrix, $\mathbf{K}_{\text{facet}}$, that selects the sky pixels that are in the facet:

$$\hat{\mathbf{x}} = \mathbf{D} \mathbf{K}_{\text{facet}} \mathbf{A}^\dagger \mathbf{N}^{-1} \mathbf{y} \quad (7)$$

However, the flux from sources outside the facet will make significant contributions to the map through the sidelobes of the primary beam. To construct high fidelity models of the facet’s map, we must introduce another matrix \mathbf{K}_{PSF} that selects the pixels in the larger area over which sources contribute to the facet through convolution with the PSF:

$$\mathbf{P} = \mathbf{D} \mathbf{K}_{\text{facet}} \mathbf{A}^\dagger \mathbf{N}^{-1} \mathbf{A} \mathbf{K}_{\text{PSF}}^\text{T} \quad (8)$$

Finally, since we are only interested in the correlation between pixels in the facet, we take the noise covariance matrix to be

$$\mathbf{C}^\text{N} = \mathbf{P} \mathbf{K}_{\text{facet}}^\text{T} \mathbf{D}^\text{T} \quad (9)$$

3. Simulations of the Interferometric Visibilities

We simulate visibilities measured with the 320 antennas of the HERA core configuration. This configuration is described by Dillon & Parsons (2016), and the corresponding antenna

positions were retrieved from the HERA wiki¹. The antenna positions correspond to the design of the array, and not to the actual positions in the field. We choose the time span of the simulations to match a HERA observation date (JD2458792). We place a one-Jansky point source at the mid-point of the observation at RA=126.75° and Dec=30.72° (J2000), precessing the coordinates to the observation epoch J2019.8424. The declination is the latitude of the HERA array, so the source transits at zenith. The full source flux density of 1 Jy will be realized in a map that sums the maps made from both of the two polarizations; the simulation assumes 0.5 Jy in each polarization. The source is assumed to be unpolarized, and only the Stokes I component enters into the simulation.

The simulated visibilities are computed according to

$$V_{ij}^{\text{sim}}(\nu) = \sum_{n=1}^{N_{\text{sources}}} B_{ij}(\hat{n}, \nu) S_n(\nu) \exp\left(-\frac{2\pi i \nu}{c} \mathbf{b}_{ij} \cdot \hat{n}\right) \quad (10)$$

where i and j are the antenna numbers of the antenna forming the baseline, ν is the frequency, B_{ij} is the primary beam, S_n is the flux of a source, \mathbf{b}_{ij} is the baseline vector, and \hat{n} is a unit vector in the direction of a source. In this simulation, $N_{\text{sources}} = 1$ as we consider only one point source. The primary beam is calculated using a model of the antenna beam provided by Nicolas Fagnoni in 2019². The model includes one dish combined with a Vivaldi feed, and was computed with the commercially available CST antenna modeling software. In the electromagnetic simulation, one pair of blades is excited with a signal composed of frequencies in the bandpass of interest, and the far field components of the transmitted electric field are computed on a grid of altitude and azimuth. We model the orthogonal polarization by rotating the model by 90°. The vector components of the electric field are evaluated in the $(\hat{\theta}, \hat{\phi})$ basis of the spherical polar coordinate system. For this simulation of interferometric visibilities we assume all antenna beam patterns are the same, and we take the square magnitude of the total electric field ($E^2 = E_{\theta}^2 + E_{\phi}^2$) to represent the beam pattern, appropriate for observations of an unpolarized source. To compute B_{ij} of equation 10, we interpolate the complex electric fields of the θ and ϕ components to get the value of the beam at the position of the point source using RectBivariateSpline routine of the scipy³ package. We do not simulate thermal noise, gain errors, or radio frequency interference. For the mapping, we select just one visibility from each baseline-redundant group using the

¹<http://hera.pbworks.com/w/page/116555034/FrontPage>

²Vivaldi 1.8m - detailed mecha design - E-field - 100 ohm load - Pol X
from <https://drive.google.com/drive/folders/0B2jOp1twrZPrdGxIVXUzMnRXemM>

³<https://www.scipy.org>

“compress_by_redundancy” function of pyuvdata⁴. The simulation code was verified through comparison with a simulation of visibilities of twelve point sources computed with pyuvsim⁵. The two simulations agree at the 0.01% level.

4. Data Selection

In carrying out the simulations presented in this memo, we have two scientific goals in mind:

- Fitting models of the primary beam and of the sky, with the emphasis on measuring the primary beam pattern of the HERA antenna
- Fitting models of the 21cm emission associated with the Epoch of Reionization (EoR)

These two goals are quite different, with the first involving measurements of bright foregrounds, and the second involving measurements of faint structures. However, neither requires high resolution, so we limit the data to the 575 H320 redundant baseline groups that have length 150 meters or less; Figure 1 shows the sampling of the uv plane. Only one frequency channel (160.08 MHz) and one polarization (east-west) is mapped. We map the entire twelve-hour span of the data, which gives information on the structure of the sidelobes far from the position of the source in the primary beam along with the maps that have the source in or near the main lobe of the primary beam.

5. Map Geometry

We work in a HEALPix⁶ coordinate system with coordinates (θ, ϕ) , with θ representing the polar angle and ϕ the azimuthal angle. In this coordinate system, the declination on the celestial sphere is $\delta = \pi/2 - \theta$, and the right ascension is equal to ϕ . Fringes are “stopped” by casting the baseline coordinates to be a function of time, and the sky coordinates to be fixed. Estimates of the sky pixels on the celestial sphere are computed directly from the visibility data according to Equation 3. Ten minutes of data are used to compute each of 75 maps,

⁴<https://github.com/RadioAstronomySoftwareGroup/pyuvdata>

⁵<https://github.com/RadioAstronomySoftwareGroup/pyuvsim>

⁶<https://healpix.sourceforge.io>

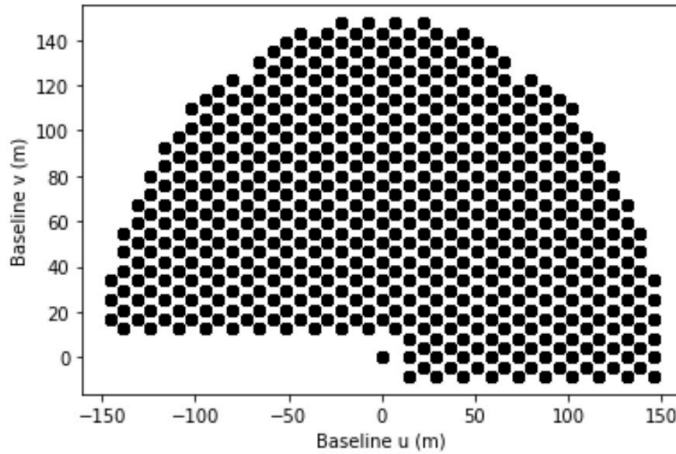


Fig. 1.— The uv coverage of the selected simulated data; $w = 0$ for all baselines. The autocorrelation point at $(0,0)$ was not used in the mapping.

for a total of 12 hours of simulated observations. As time progresses the mapped pixels shift toward the west along the HERA stripe. The circular field of view of a map facet has radius 7.67° (λ/D for the frequency mapped) and is centered on zenith. While only HEALPix pixels falling within the field of view are calculated, the position-dependent point spread function (PSF) is calculated over a larger field of view so that the model image, computed by multiplying the model sky pixels with the PSF at each point, includes sidelobes of sources outside the field of view of the facet. The radius of the PSF field is twice that of the facet, or 15.3° . The radii of the two fields are determined by the angle from zenith, and therefore there is a $\cos\delta$ distortion in the shape of the map if it is plotted in right ascension and declination coordinates. The angular size of the HEALPix pixels is set with the HEALPix parameter NSIDE. We choose NSIDE=512, giving 3,145,728 pixels on the 4π steradians of the celestial sphere, or an approximate pixel resolution of 0.11° . Sky estimates are calculated only for the pixels that fall within the facet, resulting in 14,070 pixels calculated for each facet. Figure 2 shows a representative single facet plotted on the celestial sphere. The PSF is calculated for 57,120 pixels in the PSF field of view (the number of pixels for the facet and the PSF region change by of order 5 pixels from map to map because of the effects of rounding).




Fig. 2.— An example of the area mapped in a facet, shown on the celestial sphere.

6. Primary Beam Weighting

For the beam-weighting of the mapping, we use the same CST beam model as in the data simulation. As in the simulations, we interpolate the amplitudes of the θ and ϕ components of the complex electric field vector (E_θ, E_ϕ) for the frequency under study (to the nearest MHz) with the `scipy` “`RectBivariateSpline`” routine, and the beam response computed as $|E_\theta|^2 + |E_\phi|^2$ normalized so the peak is unity. The normalized beam is evaluated at each point on a HEALPix grid (NSIDE=512) of azimuth and elevation. When the \mathbf{A} matrix is calculated (see below), the azimuth and elevation of each map point is calculated for the observation time, and the beam value is computed using the HEALPix `get_interp_value` function. Using a gridded beam (rather than recalculating the beam for each element of the \mathbf{A} matrix) is a possible source of error in the mapping, and we plan to interpolate directly to azimuth and elevation in a future implementation of the mapping algorithm. Figure 3 shows a region centered on the bore sight of the primary beam pattern used for the map weighting. This models the response of one antenna to an unpolarized source.

7. The Design Matrix \mathbf{A}

Consistent with the choice of map geometry, the components of \mathbf{A} are calculated according to

$$A_{ij} = B_j \exp \frac{2\pi i}{\lambda} \vec{b}_i(t_i) \cdot \hat{s}_j \quad (11)$$

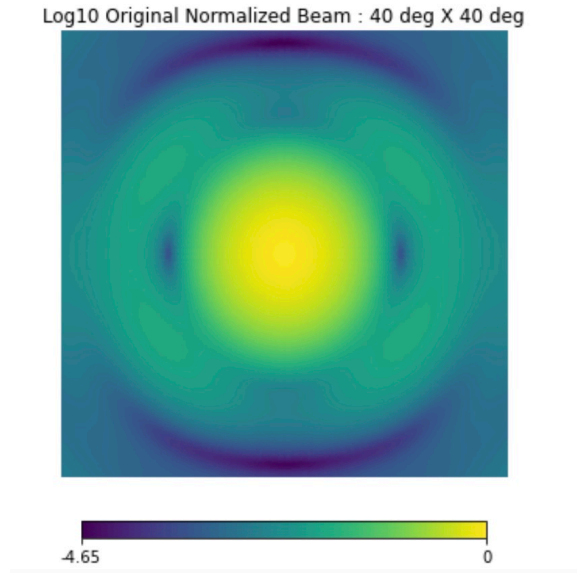


Fig. 3.— The primary beam used in the simulations and the map beam-weighting.

where the index i runs over the baseline-time labels of the visibilities, and the index j runs over the sky pixels. B_j is the response of the normalized primary beam in the direction \hat{s} of source pixel j . For the calculation of the interferometric phase, the sky is “fixed” during the computation of the map, and not updated for each time stamp. The coordinates of the baseline (b_x, b_y, b_z) are updated for every ten-second visibility, stopping the fringes at the phase center of the map. The beam weighting, $B_j = B(\ell(t), m(t))$ is updated for each pixel for every ten-second visibility by calculating the azimuth and elevation of each pixel at each time stamp, and selecting the beam weight from the HEALPix beam model using the healpy bilinear interpolation `get_interp_val`. In future implementations of this algorithm we plan to fix the baseline coordinates (so they do not vary with time), and use the time-varying $(\ell(t), m(t))$ coordinates of each pixel in the calculations of both the interferometric phase and the beam weighting, which will be computationally more efficient.

8. The Noise Matrix \mathbf{N}

These simulations are noiseless, so the noise matrix \mathbf{N} , before considerations of redundant baseline averaging, is the identity matrix. However, as discussed below we apply a factor of $(1/n_{\text{baselines}})$ to the diagonal elements of \mathbf{N} to take into account the number of baselines that enter into the average for each redundant baseline group. We leave simula-

tions of antenna-by-antenna noise and noise weighting of the mapping algorithm to a future extension of the analysis.

9. Computation of Ten-Minute Maps and Models

The following are the steps that enter into the computation of each ten-minute facet map and its model:

- Time selection: 75 images are computed, each from 9.5 minutes of simulated data. The initial LST for the first image is 36.8° , and subsequent maps each begin 2.4° after the beginning of the previous one.
- Frequency selection: only one channel; the frequency is 160.08 MHz.
- Polarization selection: only one polarization, east-west (numpy array index = 1)
- Baseline averaging: since the simulation is noiseless, the baselines are decimated by choosing one visibility from each redundant baseline group.
- Baseline selection: only baselines with length less than or equal to 150 meters.
- Facet field of view selection: a circular cap of radius $\lambda/D = 7.67^\circ$ (where $D = 14\text{m}$) centered on zenith, and placed on the HEALPix celestial sphere after computation. The radius is the arc length on the surface of the unit sphere.
- PSF field of view selection: a circular cap of radius 15.3° centered on zenith. The radius is the arc length on the surface of the unit sphere.
- Pixel lists: for each map, lists of the coordinates of the HEALPix pixels (NSIDE=512) that fall within the facet and within the PSF regions. These lists drive the computation and determine the $\mathbf{K}_{\text{facet}}$ and \mathbf{K}_{PSF} matrices.
- Calculation of A-Matrix: the elements of \mathbf{A} according to Equation 11. \mathbf{A} is a $N_{bt} \times N_{\text{PSF}}$ matrix, where N_{bt} is the number of baseline-time visibilities contained in the ten-minute dataset and N_{PSF} is the number of PSF pixels being mapped (these include all the facet pixels).
- Computation of map: the $\hat{\mathbf{x}}$ map pixels according to Equation 7 with $\mathbf{D} = \mathbf{I}$.
- Computation of PSF matrix: the \mathbf{P} matrix according to Equation 8 with $\mathbf{D} = \mathbf{I}$.

- Computation of model map: the point source is placed on the HEALPix grid according to its right ascension and declination (corrected for precession, consistent with the simulation), then the model map is computed by multiplication of \mathbf{P} and the model sky vector.

Notes:

The imaging is done directly on the surface of the celestial sphere. Therefore, there are no “w-term” corrections to be made; the w-terms are taken into account exactly with no need for a flat-sky approximation.

All three coordinates of the baselines are used in the computation of the interferometric phase, so the array is not assumed to be coplanar (though in this simulation it is in fact coplanar since the design antenna positions are used).

10. Results: The Sky Maps

The average of the sky maps along the HERA stripe is shown in Figure 4. The maps have been corrected for the factor of 2 for one polarization, and for the visibility and beam weighting, giving units of Janskys. Because the maps are noiseless, the flux density varies by several orders of magnitude from the point source to the far-off sidelobes. The recovered flux density of the point source is 0.99 Jy, close to the simulated value of 1.00 Jy. The maps




Fig. 4.— The average (in units of Janskys) of the sky maps along the HERA stripe. In these noiseless images the dynamic range from the far-off sidelobes to the 1 Jy flux density of the point source is very large, so the scale is logarithmic. The right ascension range plotted is 186 degrees and the declination range plotted is 14 degrees, centered on the source position.

are qualitatively different when (1) the point source is in the primary beam of the HERA antennas, (2) the maps are dominated by grating lobes of the point source, and (3) the

maps are dominated by distant sidelobes of the point source. We show representative maps of each of these three cases in Figure 5. The units of these maps are the flux density multiplied by the sum of the beam-weighted visibilities in each pixel. Peak values of these maps are very different, being approximately 10^6 when the source is in the primary beam, approximately 10,000 when a grating lobe dominates, and approximately 200 when distant sidelobes dominate. In the individual snapshot maps in which the point source is in the facet field of view, after correction for the weights used in the mapping the recovered flux density of the point source all fall within 3% of the simulated flux density of 1 Jy. The agreement is best (0.8%) for the map with the point source closest to zenith.




Fig. 5.— Representative $20^\circ \times 20^\circ$ sky maps. The units are flux density times the number of beam-weighted visibilities at each pixel, and the color scale has been chosen so that it saturates at 10% of the peak intensity of the map. The offsets in right ascension (map center relative to the source position) for the three maps are 65.8° , 17.3° , and 0.37° , respectively. The map on the left is representative of maps that are dominated by far-off sidelobes; the map in the center is representative of maps dominated by a grating lobe; and the map on the right is representative of maps in which the point source is within the field of view of the facet.

11. Results: The Point Spread Function Matrix

Each row of the point spread function (PSF) matrix encodes the PSF for one pixel in the sky within the area in which the PSF was calculated. The point source position is closest to the center of HEALpix pixel 2377426. In Figure 6 we show the PSF for that pixel for the map that has the point source closest to its phase center. The sidelobes surrounding the main lobe are at a level of a few percent.

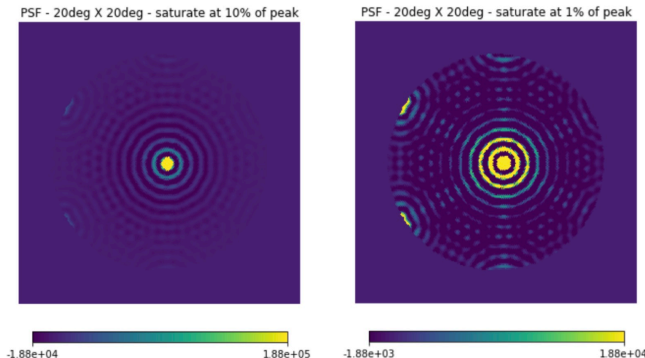


Fig. 6.— The point spread function for the pixel that contains the point source. We show two color scales: on the left the color scale saturates at 10% of the peak; on the right the color scale saturates at 1% of the peak.

12. Results: The Model Maps

In Figure 7 we show models for the representative maps that we display in Figure 5. The units are flux density times the number of beam-weighted visibilities at each pixel, and the color scale has been chosen so that it saturates at 10% of the peak intensity of the map (when the peak is nonzero). For the map on the left for which the point source is outside the PSF field of view, the model is zero. For the map in the center, the point source is outside the facet field of view (the apparent source is actually a grating lobe) but it is within the PSF field of view, so a model is produced. For the map on the right, the point source is in both the facet and the PSF fields of view, and a model is produced.

13. Comparison of Sky and Model Maps, and Algorithm “Noise”

For maps in which the point source falls within the PSF field of view, the sky and model maps appear to agree well when viewed by eye (see Figure 8). The agreement of the peak values of the map and the model are very good; their difference is only 0.8%.

To further assess the agreement quantitatively, we take the difference of the sky and model maps and plot these for our representative maps in Figure 9. For all our maps for which the point source falls within the PSF sky area, the difference maps consistently have a peak brightness of about 10% of the peak brightness of the sky map. We attribute this level of disagreement to the difference in the positions of the point source in the actual (simulated) data and its position on the pixelated HEALPix sky model. The difference

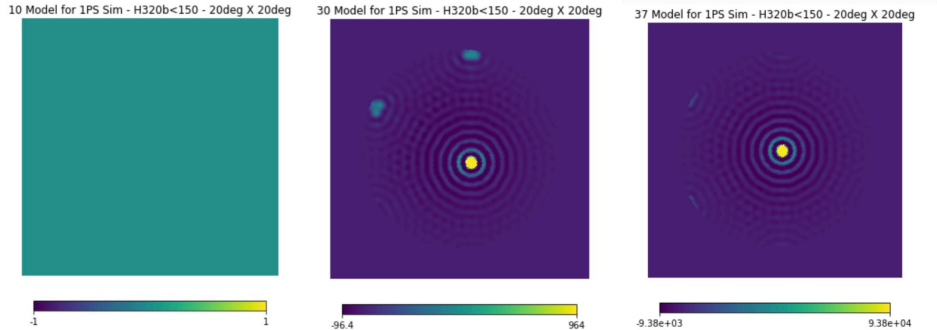


Fig. 7.— Representative $20^\circ \times 20^\circ$ model maps. As in Figure 5, the units are flux density times the number of beam-weighted visibilities at each pixel, and the color scale has been chosen so that it saturates at 10% of the peak intensity of the map (when the peak is nonzero). The offsets in right ascension (map center relative to the source position) for the three maps are 65.8° , 17.3° , and 0.37° , respectively. For the map on the left for which the point source is outside the PSF field of view, the model is zero. For the map in the center, the point source is outside the facet field of view (the apparent source is actually a grating lobe) but it is within the PSF field of view, so a model is produced. For the map on the right, the point source is in both the facet and the PSF fields of view, and a model is produced.

occurs because the model source is moved to the position of the center of the closest HEALPix pixel. Since the linear size of one pixel is 0.11° , and the FWHM of the synthesized beam is approximately $\lambda/b_{\max} = 0.7^\circ$, one would expect a difference at about this level. We tested this explanation for the magnitude of the difference maps by computing images with NSIDE=1024. A comparison of maps at these two different levels of pixelization are shown in the rightmost two maps in Figure 9. Increasing the pixel resolution reduces the difference between the map and model from 10% to 2%.

We anticipate that maps and models such as these will be used to assess the goodness of fit of primary beam models and sky models. In this noiseless simulation, the differences between the maps and the model are due to the approximations made to make the computation tractable. These differences are a source of “algorithm noise” which must be considered when assessing agreement between data and model. For the imaging parameters chosen here, it appears that the pixelization of the sky model is the main source of algorithm noise. This error might be reduced by convolving the maps to a lower resolution. We tested this by convolving with a gaussian kernel so that the convolved image had a resolution of $30'$; this reduced the difference between NSIDE=512 maps and models from 10% to 4%.

A straightforward way in which to assess the goodness of fit is to calculate a χ^2 statistic

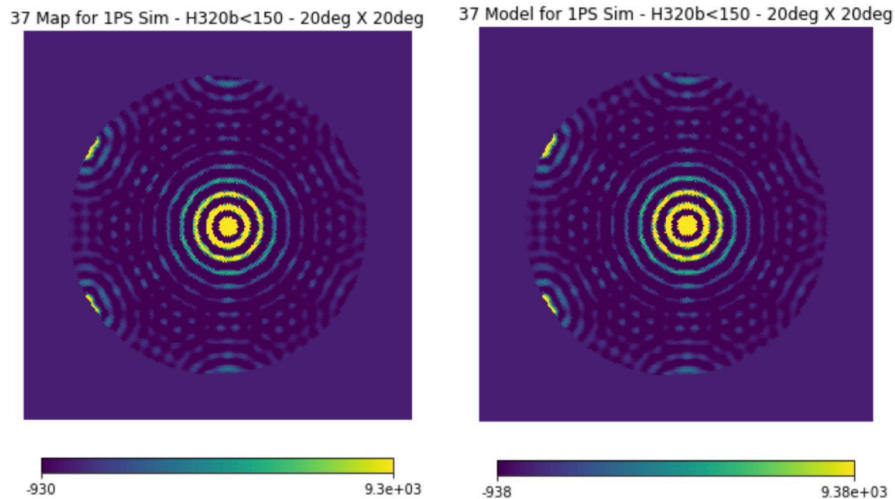


Fig. 8.— Side-by-side comparison of the map and the model for the map in which the point source is closest to the phase center. As in Figure 5, the units are flux density times the number of beam-weighted visibilities at each pixel. The color scale is chosen so that it saturates at 1% of the peak, facilitating a comparison of the sidelobes.

that is the sum of the squared deviations with appropriate weighting by the covariance matrix. For our choice of the normalization matrix \mathbf{D} , the covariance matrix is simply the point spread function matrix \mathbf{P} restricted to the facet pixels (see Equation 9). For the reasons discussed above, we do not have \mathbf{P} for maps in which the point source does not fall within the PSF field of view. Therefore, as an initial examination of the contribution of the 75 maps to a χ^2 statistic, we simply plot (Figure 10) the sum of the squared deviations for all the maps. The differences are due to algorithm noise, which scales with the peak beam-weighted flux density of the maps. Therefore, the magnitude of the algorithm noise follows the response of the primary beam. A proper assessment of goodness-of-fit requires the calculation of a true χ^2 statistic; we postpone that for future work.

14. Conclusions

We have simulated optimal mapping of HERA data as first proposed by D2015 and explored in more detail by Dillon & Parsons (2016). The simulations reproduce the 320 antennas of the HERA core, and one polarization of the Vivaldi beam. This first simulation explores the simplest case of twelve hours of observation of a 1 Jy point source located in the




Fig. 9.— Difference maps for our representative maps. The color scales are chosen so that each saturates at 1% of the peak of the original map in both the positive and negative extremes. For the map on the left, the main contribution to the difference between the map and the model is the unmodeled far-off sidelobes. For the remaining maps, the main contribution is the sub-pixel misalignment between the mapped point source and the model point source. Increasing the number of pixels in the HEALPix grid (NSIDE=512 to NSIDE=1024) reduces the difference between the map and model from 10% to 2%.

HERA stripe. We map “facets” of the sky in 9.5-minute snapshots centered on zenith. The area of the facet is matched in size to the primary beam, and we compute the point-spread-function matrix for an area a linear factor of two larger in size. We recover 99% of the flux density of the point source in the average map of the HERA stripe, and the flux densities of the point source within the individual snapshot facets are within 3% of the simulated value. We map the grating lobes and the sidelobes for a range of $\pm 90^\circ$ of right ascension in the HERA stripe.

We compute model maps by convolving the point source with the point spread function at its position. The peak flux densities of the maps and models agree to better than 1%. In a pixel-to-pixel comparison of the maps and model maps, we find agreement of 10% or better. The agreement improves to 2% when the pixel size is made a linear factor of two smaller. This indicates that the primary source of the “algorithm noise” is the pixelization of the sky, though there could be contributions due to inaccuracies in interpolating the primary beam model and in visibility simulation. The agreement improves to 4% if the maps and models are convolved to a resolution of $30'$ before subtraction. Clearly, there is a tradeoff between precision and computational cost of the calculation. The deviations (map minus model) for all the pixel form a statistical measure of goodness-of-fit that we plan to explore in future work. In principle, this could provide a metric for assessing models of the primary beam and of the sky.

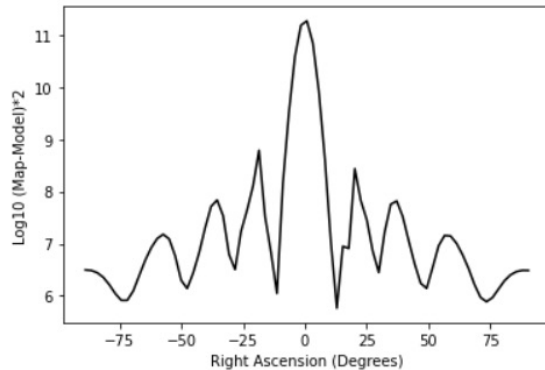


Fig. 10.— Sum of the squared deviations between the map and model for each map, as a function of the difference in right ascension between the source and the map center.

15. Future Work

- Measurement noise simulated for the data and included in the \mathbf{N} weighting
- Interpolate beam directly to model source position
- Quantitative calculation of goodness-of-fit from difference maps and covariance matrix for use in beam measurement and sky model estimation
- Alternate forms of \mathbf{D} for weighting.

REFERENCES

- Dillon, J.S., et al., “Mapmaking for Precision 21cm Cosmology.” 2015, Phys. Rev. D.
- Dillon, J.S., & Parsons, A. R., “Redundant Array Configurations for 21cm Cosmology.” 2016, ApJ.
- Morales, M. F., & Matejek, M., “Software Holography: Interferometric Data Analysis for the Challenges of Next Generation Observatories.” 2008, MNRAS.
- Tegmark, M., “How to Make Maps from Cosmic Microwave Background Data without Losing Information.” 1997, ApJ.
- Zheng, H., et al., “Brute-Force Mapmaking with Compact Interferometers: a MITEoR Northern Sky Map from 128 to 175 MHz.” 2017, MNRAS.