Draft version February 22, 2019
Preprint typeset using LaTeX style AASTeX6 v. 1.0

HERA MEMO #64: REFLECTION SYSTEMATICS IN HERA H1C DATA

Nicholas S. Kern[1] & Aaron R. Parsons[1]

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

ABSTRACT

We investigate the nature of reflection systematics in HERA season 1 data (H1C; 2017-2018). We first outline the mathematical formalism behind two kinds of systematics, signal chain reflections and antenna cross couplings, and show how they manifest in interferometric data products. Next we investigate their presence in data from the Second Internal Data Release (IDR2), finding that cable reflections corrupt most of the antennas at roughly the 3×10^{-3} level, and that cross couplings also corrupts many baselines at a variety of amplitudes, which at its worst can bee seen at the the 10^{-1} level. Future work will demonstrate how we can model and remove these systematics in a way that is not lossy to the EoR signal. This work is complementary to the work presented in HERA Memo #58, which investigated reflections in the cross-polarized visibilities.[a]

1. PRELIMINARIES

In this section, we describe the mathematical formalism of how signal chain reflections and voltage cross-couplings within our instrument manifest in our data products. To begin, we start with the standard two-element interferometer (Figure 1), consisting of two antennas, 1 and 2, whose feeds measure an incident electric field and convert it into a voltage. These signals travel from the feeds through each antenna's signal chain to the correlator, and along the way are amplified, digitized, channelized and fourier transformed into the frequency domain. The correlator then cross multiplies voltage spectra to form the fundamental interferometric data product: the visibility, V_{12}, between antenna 1 and 2, written as

$$V_{12}(\nu, t) = v_1(\nu, t) v_2^*(\nu, t). \tag{1}$$

Here we have chosen to define the visibility as the product of two antenna voltage spectra, rather than the correlation of voltage time streams: although the two are equivalent given the convolution theorem, the former will prove to be an easier basis when working with reflections. In addition, we have been explicit about the frequency and time dependence of each antenna's voltage spectra v and, by extension, the complex visibility V, although we may drop these throughout the text for brevity. As we've written it, V_{12} is a cross-correlation visibility between antenna 1 and 2. However, an interferometer can also produce the *auto-correlation* visibility, V_{11}, by simply cross-multiplying v_1 with itself. In this paper, we will use both the cross and auto-correlation visibilities to describe systematics.

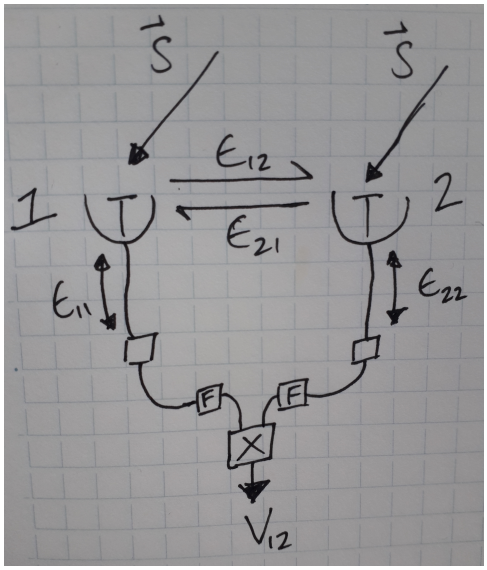

Figure 1. A schematic of a two-element interferometer with the signal path and possible sources of reflections demarcated. Sky signal ($\vec{\mathcal{S}}$) enters each antenna's feed, travels down their signal chains, is processed at a node before being directed to the F engine that digitizes and fourier transforms the signal (F) and is then sent to the correlator ($\mathcal{X}$). A possible cable reflection in antenna 1's signal chain is marked as ϵ_{11}, traversing up and down the cable connecting the feed to the node. A possible source of feed-to-feed coupling is marked as ϵ_{12}, where a signal is reflected off of antenna 1's feed or is received and then re-broadcast.

The native spaces our visibility spans are frequency and time. When fourier transforming across the frequency axis, we put the data into a time domain. To

[a] http://reionization.org/manual_uploads/HERA_Constant_Offset_Memo_Elder.pdf

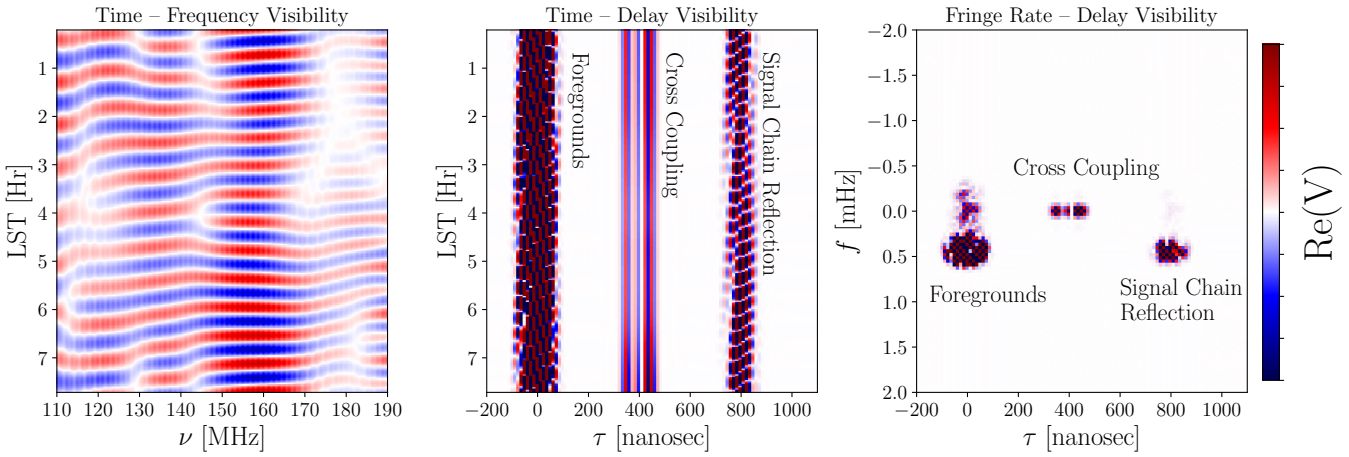

Figure 2. The real component of a simulated cross-correlation visibility with foregrounds, a signal chain reflection inserted at $\tau = 800$ ns and a cross coupling term inserted at $\tau = 400$ ns. Note that the signal chain reflection and cross-coupling have been simulated independently of each other here for visual clarity; in other words, their cross terms have been set to zero. In addition, the color scales have been artificially adjusted for visual clarity. **Left:** Visibility in time and frequency space. **Center:** Visibility in time and delay space. **Right:** Visibility in fringe-rate and delay space. Different components of the visibility–in particular systematics–are usually better separated in delay and fringe-rate space than in the original time and frequency space.

separate this from the original time domain, we refer to the fourier dual of frequency as the *delay domain* (τ). Similarly, the fourier transform of our data across time puts the data into a spectral domain, which we refer to as the *fringe-rate domain* (f) similar to Parsons et al. (2016). In the absence of explicit markers, we will use V to mean the visibility in time and frequency space, and use $\tilde{V}$ to mean the visibility in fourier space. Depending on the context, this could mean time & delay space, fringe-rate & frequency space, or fringe-rate & delay space. Which of the three we mean should be clear based on context, otherwise we will use explicit notation, as in $\tilde{V}(\nu, f)$.

The convenience of working in delay and/or fringe-rate space comes from the fact that different components of the visibilities separate better than in frequency and time space. That is to say, the covariance of our data in fourier space is generally more diagonal than in real space. Foreground signal, for example, is intrinsically spectrally smooth and will therefore occupy low delay modes, whereas a fiducial EoR model, being non-spectrally smooth, occupies low and high delay modes. This can be seen by inspecting visibility *waterfalls*, which are the visibilities plotted as a function of time and frequency. Transforming these waterfalls to the delay and/or fringe-rate domain can help us separate the systematic and better understand its behavior. Figure 2 shows simulated foreground + systematic visibilities in real and fourier space, demonstrating how systematics can be better separated in fourier space. The simulations come from the `hera_sim` mock visibility simulator, which creates a per-baseline, statistical realization of the visibility using delay and fringe-rate filters specific to the baseline orientation. In this case, the foreground visibility is generated by taking an uncorrelated white noise visibility, multiplying it with a sky amplitude spectral slope, filtering it across frequency with a truncated Gaussian window with width informed by the baseline separation, and fringe rate filtering it using the optimal fringe rate filter profile for that specific baseline. The optimal fringe-rate filters are calculated numerically following Parsons et al. (2016), whose explanation we defer to a different memo.

1.1. *Describing Signal Chain Reflections*

A reflection in the signal chain of an antenna will produce a copy of it at a specific time offset, which can also be thought of as a delay offset. An example of this is a reflection at the end of the cables in an antenna's signal chain, such that the incident signal travels back up the cable, reflects again at the start of the cable and travels back down. The reflected signal is typically only a fraction of the transmitted signal, but for EoR surveys even fractions of a foreground signal can swamp the desired EoR signal. If v_1' is antenna 1's voltage spectrum with the reflection, we can write it as

$$v_1'(\nu, t) = v_1(\nu, t) + \epsilon_{11}(\nu) v_1(\nu, t) \tag{2}$$

where ϵ is a complex quantity describing the reflection in antenna 1's signal chain, denoted as 11 because it is adding a copy of its signal to its own signal chain. The coupling coefficient, ϵ, can be broken into three constituent parameters as

$$\epsilon(\nu) = A e^{2\pi i \tau \nu + i\phi}, \tag{3}$$

where A is the amplitude, τ is the delay offset (the total time it takes to be reflected) and ϕ is the phase offset the

reflected signal may have acquired relative to the original signal. In Equation 2 we have assumed time-stability of the reflection and dropped its time dependence, although in principle the coupling parameters may have some time dependence if the physical elements driving the reflection are changing over time.

If we inserted the corrupted voltage spectra into the visibility equation (Equation 1), we get

$$V'_{12} = v_1 v_2^* + \epsilon_{11} v_1 v_2^* + v_1 \epsilon_{22}^* v_2^* + \epsilon_{11} v_1 \epsilon_{22}^* v_2^*. \quad (4)$$

We can see that in addition to the original cross-correlation term ($v_1 v_2^*$) we now also have copies of it at positive and negative delay offsets that are suppressed in amplitude by a factor of A_{11}. The time-behavior of a reflection mimics that of the original data, in that it shows the same temporal oscillation (i.e. fringing) as the foregrounds at $\tau \sim 0$, and thus also appears at the same fringe-rate modes as the foregrounds (e.g. see right of Figure 2).

Note that in the cross-correlation visibility V_{12}, first-order reflections from antenna 1 have a positive delay offset τ_1, while first-order reflections from antenna 2 have a negative delay offset $-\tau_2$ due to the complex conjugation. The second order reflection will be shifted to a delay $\tau_1 - \tau_2$, which generally does not go identically to zero for realistic cable imperfections.

The resultant auto-correlation visibility can also be computed, and is given by

$$V'_{11} = v_1 v_1^* + \epsilon_{11} v_1 v_1^* + v_1 \epsilon_{11}^* v_1^* + |\epsilon_{11}|^2 v_1 v_1^*, \quad (5)$$

where similar to the cross-correlation visibility, the first-order reflection a copy of the original visibility shifted to positive and negative delays. Note that we could have simplified this further by writing the sum of the second and third term as $2Re(\epsilon_{11} v_1 v_1^*)$, which is more intuitive given the fact that the auto-correlation visibility is a purely real quantity. However, the form we choose serves to illustrate a specific point. The auto-correlation visibility in delay space at $\tau = 0$ can be reduced to $V_{11}(\tau = 0) \approx (v_1 v_1^*)(\tau = 0)$, given that the second and third term in Equation 5 go to zero away from τ_{11} and having dropped the contribution from the second order term. Similarly, we can deduce that near $\pm\tau_{11}$, the auto-correlation visibility in delay space simplifies to

$$V'_{11}(\tau = \pm\tau_{11}) = \begin{cases} \epsilon_{11} v_1 v_1^*, & \tau = +\tau_{11} \\ v_1 \epsilon_{11}^* v_1^*, & \tau = -\tau_{11}, \end{cases}$$

which means that, in delay space, one can estimate the reflection amplitude as

$$A_{11} = \left| \frac{V_{11}(\tau = \pm\tau_{11})}{V_{11}(\tau = 0)} \right|. \quad (6)$$

If one can estimate their parameters from the data, reflections can be removed via standard antenna based calibration, which distills all antenna-based effects into a single complex per-antenna gain term. In this framework, the raw, corrupted voltage spectra of antenna 1 is related to its true, calibrated value as

$$v_1^{\rm raw} = v_1 g_1, \quad (7)$$

which when inserted into the visibility equation yields the standard antenna based calibration equation,

$$V_{12}^{\rm raw} = V_{12} g_1 g_2^* = \langle v_1 v_2^* \rangle g_1 g_2^*. \quad (8)$$

The g term is called the antenna gain, and corrects for amplitude and phase errors introduced by the various stages of the signal chain from the feed all the way to the correlator. By re-arranging Equation 2 as

$$v'_1 = v_1(1 + \epsilon_{11}) = v_1 g_1, \quad (9)$$

we can see that signal chain reflections can be completely encompassed in this gain term, and hence corrected for by estimating and applying the per-antenna gains.

1.2. *Describing Antenna Cross Coupling*

We now turn our attention to another systematic that acts to couple one antenna's voltage stream with another antenna's voltage stream, which we refer to as a cross coupling or crosstalk. Note that this model for crosstalk is different than "capacitive crosstalk" created by nearby, parallel signal chains interacting with each other within cabling, receivers, analog-to-digital conversion (ADC) units etc, which is the more common usage of the term crosstalk in the radio astronomy literature. The degree to which capacitive crosstalk can be modeled with the framework we present here is deferred for future work. Our model for crosstalk simply states that one antenna's voltage is added to another antenna's voltage, with a coupling coefficient that determines how the voltage is added. Similar to the case of the signal chain reflection, we assume this coupling coefficient to be factorable into an amplitude, delay and phase.

To write down how this affects the visibility, let's consider sky emission incident on antenna 1's feed. The total measured voltage of antenna 1 is actually a sum of the measured voltage from each direction on the sky,

$$v_1 = \int_{4\pi} \mathcal{V}_1(\hat{s}) d\Omega, \quad (10)$$

where $\mathcal{V}_1$ is the contribution to v_1 along the sky direction $\hat{s}$ within $d\Omega$. Now let's assume that signal from a specific direction is reflected off the feed and received by a neighboring antenna. Given this, we can write the

coupling of v_1 to a neighboring v_2 (and vice-versa) as

$$v_1' = v_1 + \int_{4\pi} \epsilon_{21}(\hat{s})\mathcal{V}_2(\hat{s})d\Omega$$
$$v_2' = v_2 + \int_{4\pi} \epsilon_{12}(\hat{s})\mathcal{V}_1(\hat{s})d\Omega,$$

where ϵ_{12} quantifies the insertion of $\mathcal{V}_1$ signal into v_2 (and vice-versa). We see here that the feed-to-feed coupling can in principle have an angular dependence, meaning sky emission from certain directions may be preferentially coupled to specific neighboring antennas.

There are some cases where this coupling can be **independent** of the sky direction. In this case, the coupling coefficient comes out of the integral and our equations simplify to

$$v_1' = v_1 + \epsilon_{21}v_2$$
$$v_2' = v_2 + \epsilon_{12}v_1.$$

One example of this could be the coupling of one signal chain with another after the sky emission has been converted into a voltage by the feed, say if the signal from antenna 1 was received, re-broadcasted and picked up by a neighboring feed, or if the coupling happened downstream of the feeds within the signal chains at some point. Substituting these simplified equations into Equation 1, we get

$$V_{12}' = v_1v_2^* + v_1\epsilon_{12}^*v_1^* + \epsilon_{21}v_2v_2^* + \epsilon_{21}v_2\epsilon_{12}^*v_1^*, \quad (11)$$

which now contains the auto-correlation visibility terms $v_1v_1^*$ and $v_2v_2^*$ at the first-order level. Unlike the cross-correlation term, the auto-correlation terms are purely real and thus have zero phase. In the complex plane, the cross correlation term $v_1v_2^*$ winds around the origin as a function of time because its phase varies temporally. Because the autocorrelation has zero phase (and assuming ϵ_{21} is time-stable) the presence of the first-order cross coupling terms can be thought of as an additive offset from the origin of the complex plane at some arbitrary angle, about which the cross correlation rotates over time. While this offset has a constant phase, it does have a temporally varying amplitude, due to the natural variation from the auto-correlation over time. However, this variation is generally fairly slow, occurring on timescales of a beam crossing, which for HERA is roughly 40 minutes. This leads us to two critical insights on the behavior of the cross coupling terms in the cross-correlation visibility: 1) their time variability is slow, thus occupying low-fringe rate modes and 2) they have a constant phase in frequency space determined solely by the coupling coefficient, as the auto-correlation term itself has no phase. The second point also implies that the cross coupling terms have no preference for positive or negative fringe rates, meaning they are centered at the $f = 0$ fringe rate mode (e.g. see right of Figure 2).

In the more general case, where the coupling coefficient has a non-trivial angular dependence, our findings about the behavior of the systematic terms are largely the same. Imagine the scenario where the angular dependence of the coupling is fairly localized on the sky, such that only part of the full primary beam (and thus only a portion of v_1) is coupled to a neighboring antenna. In this case, we still expect an additive offset with a time-constant phase in the visibilities, but expect its amplitude to be smaller, due to only a portion of v_1 making it into v_2. Furthermore, the time dependence of the amplitude will likely be faster: given our assumption that the effective "beam" of ϵ_{12} on the sky is small compared to the primary beam, its amplitude will evolve faster than the autocorrelation, as the sum total of its perceived sky changes at a quicker rate than the full beam.

Lastly, if we were to search for these terms in the cross-correlation visibility, where would they show up? Because of the complex conjugate of ϵ_{12} in Equation 11, the $v_1v_1^*$ term will appear at a negative delay offset in V_{12}. Similarly, it can be seen that the $v_2v_2^*$ term acquires a positive delay offset due to the non-conjugation of ϵ_{21}.

We can also compute the effects of cross coupling on the measured auto-correlation visibility, V_{11}, which yields

$$V_{11}' = v_1v_1^* + v_1\epsilon_{21}^*v_2^* + \epsilon_{21}v_2v_1^* + |\epsilon_{21}|^2v_2v_2^*. \quad (12)$$

We find that the cross-correlation is inserted into the measured auto-correlation at the first order level and with a delay offset of τ_{21}. These terms are likely many order of magnitudes below the peak auto-correlation visibility amplitude, given that the cross-correlation visibilities are generally a few orders of magnitude below the auto-correlation inherently, which is further compounded by the amplitude suppression from ϵ_{21}.

We can see simply from Equation 11 that the corruption of V_{12}' by cross reflections *cannot* be factorized into antenna based gains, based simply on the presence of the ϵ_{12}-like terms, which are baseline-dependent. Removal of a cross-coupling terms in the data therefore must be done on a per-baseline basis by constructing a model of the systematic in each visibility and then subtracting it.

1.3. *Summary*

To summarize, reflections along a single antenna's signal chain produces a duplicate of the signal with suppressed amplitude and some delay offset. This is the true for both the cross and auto-correlation visibility products. Example mechanisms include cable reflections and dish-to-feed reflections within the confines of a single antenna. Reflections in the cross-correlation visibility have the same time structure as the un-reflected visibility, meaning reflected foreground signal occupies the same

fringe-rate modes as un-reflected foreground signal, but is shifted to high delays. Reflections can be removed from the raw data by creating a model of the reflections and incorporating them into the per-antenna calibration gains. If, by good engineering, reflection amplitudes can be kept under 10^{-3}, then reflections of order $\mathcal{O}(\epsilon^2)$ and higher are expected to be below fiducial EoR amplitudes and can be neglected.

Another systematic we describe is created by antenna-to-antenna cross coupling, which mixes the voltage signals between the antennas. This has the effect of introducing a copy of the auto-correlation visibility into the measured cross-correlation visibility at positive and negative delay offsets, and similarly introduce copies of the cross-correlation visibility into the measured auto-correlation visibility. In the measured cross-correlation visibility, the first-order coupling terms are slowly time variable, and occupy low fringe-rate modes centered at $f = 0$ Hz. Cross coupling terms cannot be removed via antenna based calibration, and must be modeled and subtracted at the per-baseline level.

2. HERA SEASON 1 DATA

In this section we investigate the degree to which signal chain reflections and antenna cross coupling can be found in the 1st HERA observing campaign (H1C; 2017-2018). The data comes from a single night of the second internal data release (IDR2) of the year, falling on Julian Day 2458101. The data products can be found on the NRAO Lustre system under `/lustre/aoc/projects/hera/H1C_IDR2`.

The HERA H1C array was a build-out to roughly 80 antennas throughout the observing season. The build-out was a hybrid HERA-PAPER system, utilizing the old PAPER signal chains and attaching them to new HERA antennas. The HERA antenna is a parabolic dish spanning $\sim$ 14.5 meters in diameter, with a focal height designed to minimize standing waves in the dish. Nonetheless, the standing waves that inevitably arise in parabolic dishes will appear with characteristic timescales of 50 nanoseconds (c.f. Neben et al. 2016; Thyagarajan et al. 2016; Ewall-Wice et al. 2016; DeBoer et al. 2017; Patra et al. 2018). The feeds consist of the old PAPER dipoles surrounded by a cage, which are flipped upside down and suspended 4.5 meters above the dish center, and attached to the old PAPER baluns. Once received by the feed, the voltages travel through the active balun or front-end module (FEM) and are sent through a 150 meter coaxial cable (first cable in Figure 1) to a receiver unit (RXR) in the field holding a post amplifier module (first box in Figure 1). The data in this campaign contain four receiver units, labeled Rxr3, Rxr4, Rxr5, and Rxr8 (right of Figure 3). The voltages are then sent through another coaxial cable of about 20 meters in length (second cable in Figure 1) to a container holding the analog-to-digital converters and ROACH boards that digitize and then fourier transform the signals into the frequency domain respectively (F box in Figure 1). Lastly the signals are cross multiplied by the correlator to produce the interferometric visibilities (X box in Figure 1).

The data have been processed with part of the HERA reduction and calibration pipeline (version 2.1) outlined in HERA Memo #45[1]. In particular, the RFI masks comes from the HERA reduction pipeline (version 2.1), but an added, stacked median filtering on the autocorrelations has been used to flag low-level RFI. The end-result mask applied to all visibilities is shown in Figure 4. We do not use the full calibration solution from the version 2.1 pipeline when calibrating the auto-correlation visibilities. Instead, we take the average amplitude of the calibration solution per antenna and apply it to the auto-correlations to properly set their flux scale. The reason for this, as we will see shortly in §2.1, is because we will use features in the auto-correlation delay response to diagnose systematics in the data. Calibration is a multiplicative term in frequency space, meaning it can be thought of as a convolution in the delay domain, and can therefore smooth out sharp features which would otherwise be informative for understanding the systematics. However, when calibrating the cross-correlation visibilities, we do use the full calibration solution from the version 2.1 pipeline. This is done to suppress the antenna-based systematics in the cross-correlation visibilities–which we can already learn about from the auto-correlation visibilities. In suppressing antenna-based systematics, we can hope to uncover any baseline-based systematics in the cross-correlation visibilities, if they exist.

2.1. *HERA's Signal Chain Reflections*

Whether a cable in the signal chain generates a strong reflection is conditional on the impedance match between the connected surfaces. The HERA analog system was designed to take this into consideration, but getting a tight impedance match across a wide bandpass can be difficult given that the properties of the signal chain evolve with frequency (c.f. HERA Memo #29[2]).

Many of the old PAPER signal chains could not be repurposed for the new HERA system and stop-gaps were developed in the interim. In particular, a new post amplifier model was developed and manufactured,

[1] http://reionization.org/wp-content/uploads/2018/07/IDR2.1_Memo_v2.html

[2] http://reionization.org/wp-content/uploads/2013/03/Memo29_CST_co-simulation_HERA_analogue_-system.pdf

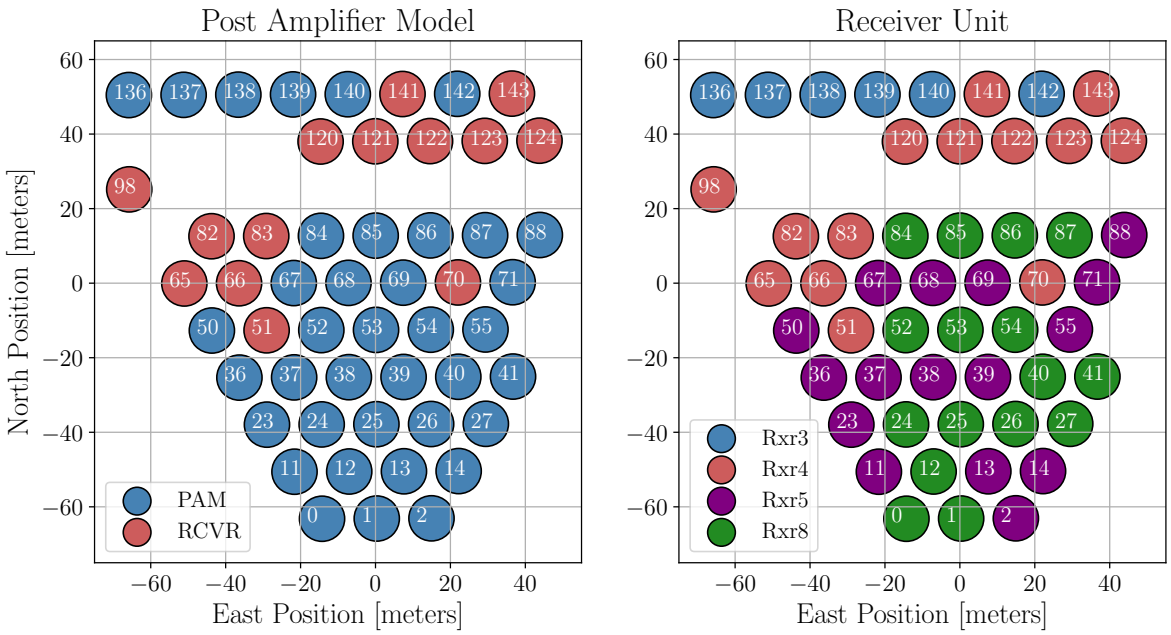

Figure 3. **Left:** Array layout in East-North-Up coordinates marking the post amplifier model for each antenna. **Right:** Array layout marking the receiver unit for each antenna. Note that all RCVR amplifiers are housed within the Rxr4 receiver unit.

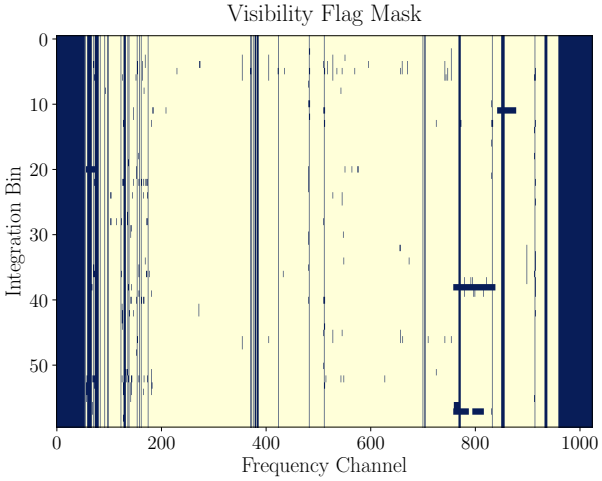

Figure 4. Visibility flagging mask used to excise RFI and bad time integrations.

which was inserted into the receiver units for most of the HERA signal chains (left of Figure 3). The new post amplifier models are denoted as (PAM), while the old PAPER post amplifier models are denoted as (RCVR).

If generated, a reflection in the first 150 m coax cable would appear at a delay of $\sim$ 1300 nanoseconds, assuming a speed of light of $c = 0.8$ in the cable. The 150 meter FEM-to-PAM cables have variable lengths at $\pm$15 meter uncertainty for most of the cables. HERA Memo #39[3] has explored the presence of reflections in the 150 m cables from H1C data and found similar results. Likewise, a reflection generated in the second 20 m coax cable would appear at a delay of $\sim$ 170 nanoseconds.

To inspect the data for reflections we take the autocorrelation visibility from each antenna and inspect it for peaks in delay space. However, the calibrated data are filled with flags due to RFI and are thus nulled at the flagged channels, which can be thought of a binary windowing function applied to the otherwise RFI-free data. Such a windowing function is not ideal for inspecting the data in fourier space, as the fourier transform of a flagged windowing function has strong sidelobes at high delays, meaning strong foreground structure at low delays will spillover and mask weaker structure at high delays. To mitigate this we employ a standard CLEAN deconvolution (Högbom 1974) on the data across frequency, with a CLEAN boundary out to $|\tau| < 2000$ ns, following the procedures outlined in Parsons & Backer (2009). This allows us to deconvolve the windowing function in fourier space, leaving us with data free of the sidelobes due to RFI. To help further suppress spectral leakage of foreground power due to a finite sampling range, we window the data in frequency space with a blackman-harris window before fourier transforming to delay space. We then average the absolute value of the visibility across time to from a single spectrum per antenna per dipole polarization. The result is shown in Figure 5, with vertical dotted white lines bounding the expected range of the 150 meter cable reflection. We can see that most of the antennas show strong peaks in

[3] http://reionization.org/wp-content/uploads/2013/03/HERA39_H1C_cable_reflections_ewall-wice.pdf

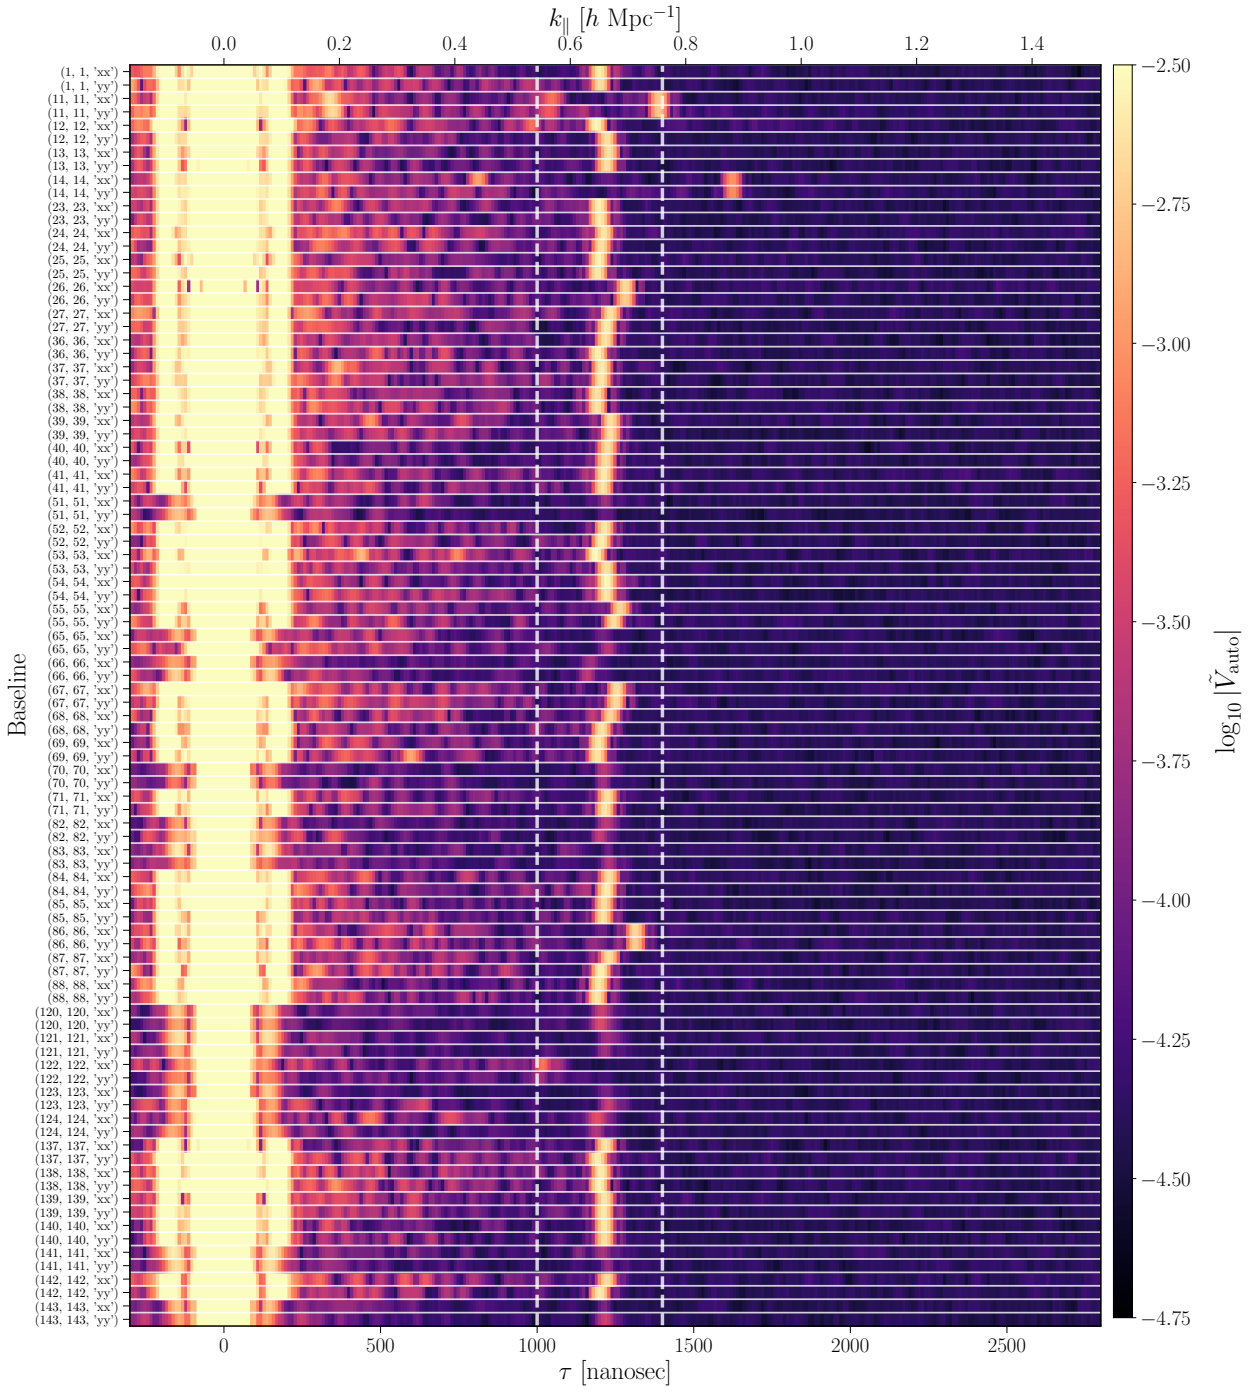

Figure 5. Auto-correlation spectra in delay space for each antenna and polarization. Each auto-correlation is CLEANed of RFI flags out to $|\tau| < 2000$ ns, transformed to delay space, averaged incoherently across LST to form a single spectrum for the entire night, and is then peak normalized. White dashed lines demarcate the expected range of delays for a reflection in the 150 m coaxial cable from the FEM to receiver unit.

the range expected for the 150 meter coaxial cable at ~ 1250 ns, with variability at about ± 100 ns. We also conclude: 1) antennas that have cable reflections on one polarization also tend to have reflections in the other polarization; 2) antennas that have spikes in the delay range of $1000 < \tau < 1500$ ns also have spikes at $\tau \sim 200$ ns and 3) there is a variety of per-antenna structure at intermediate delays of $200 < \tau < 700$ ns, that could possibly arise from cable imperfections creating reflections midway through the 150 m cable.

To try to understand why some antennas have stronger reflections than others, and why the bump at 200 ns is correlated to the bump at 1250 ns, we plot each of the visibility spectra according to their post amplifier

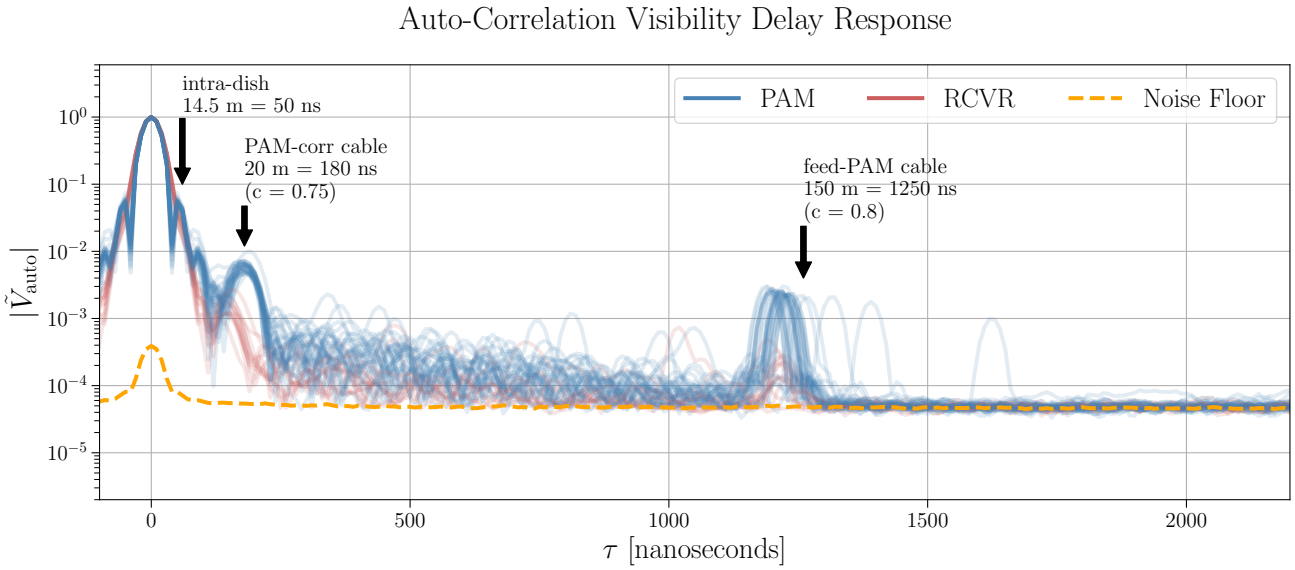

Figure 6. The same auto-correlation spectra in delay space from Figure 5, but separated by the post amplifier model PAM (blue) and RCVR (red). Strong reflections from the 150 m cable at ~ 1200 ns are correlated with reflections in the 20 m cable at ~ 180 ns. Structure at delays of ~ 50 ns could be explained by intra-dish reflections, while the structure farther out at ~ 100 ns is harder to explain, and was not seen in early simulations of the HERA feed and dish delay response.

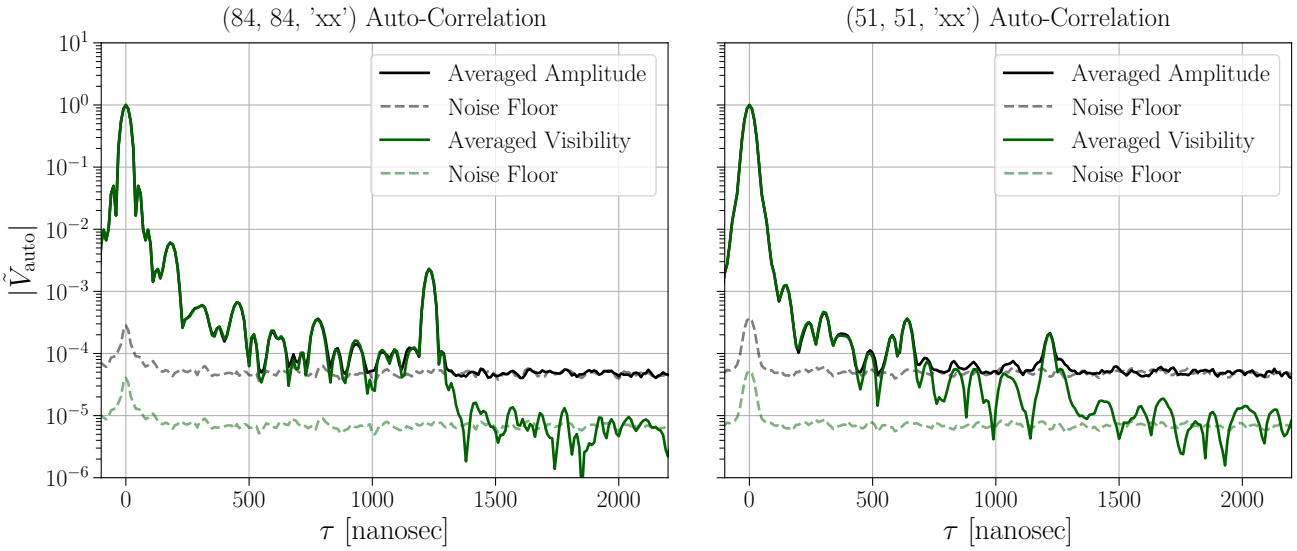

Figure 7. Auto-correlation delay response. Black shows a per-integration CLEAN and then a time average of the visibility amplitude in fourier space. Green shows a complex visibility average, a CLEAN on the averaged integrations and then a transformation to fourier space. Dashed shows the noise floor for the averaged and un-averaged data. The spike in the noise curve at low delay is spurious and not real, and is a consequence of foreground structure leaking through the imperfect differencing scheme for estimating the noise. **Left:** Auto-correlation for antenna 84 and 'XX' polarization. This antenna is connected to a PAM amplifier, and therefore shows a strong reflection at 190 and 1200 ns. **Right:** Auto-correlation for antenna 51 and 'XX' polarization, which is not connected to a PAM amplifier. In both antennas, we can see a clear systematic tail at low delays–that is higher for the PAM signal chain–and which doesn't integrate down like the structure at $\tau > 1200$ ns.

model (PAM or RCVR), shown in Figure 6. Now we can clearly see that antennas connected to the newer PAM-model post amplifiers (blue) have significantly stronger peaks at ~ 1200 nanoseconds than those connected to the older RCVR-model amplifiers (red). In addition, we can see this to be the case not only for the 150 m cable but also the 20 m cable. The 150 m cable reflection show reflection amplitudes of roughly 3×10^{-3} for the PAM signal chains, and roughly 3×10^{-4} for the RCVR signal chains. Similarly, the 20 m cables show reflection amplitudes at the 8×10^{-3} level for the PAM signal chains.

We also see evidence for structure at the $50 - 100$ nanosecond level, which is more dramatic in antennas

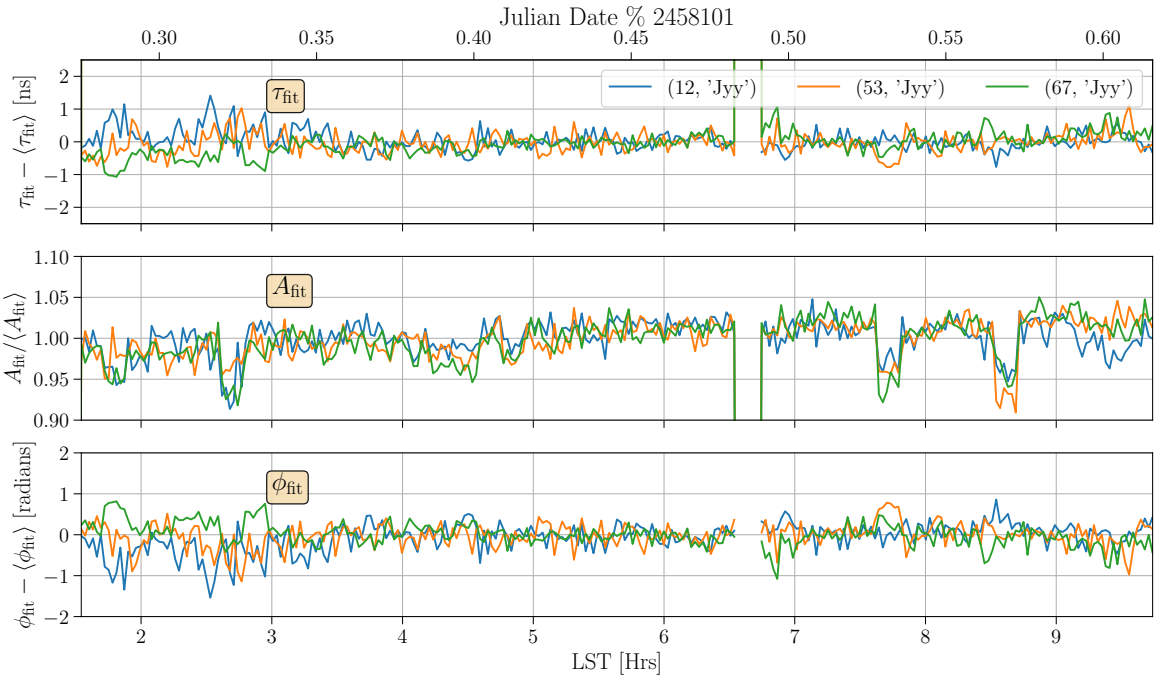

Figure 8. 150 meter cable reflection parameter fits over the course of a night. For the three antennas shown, systematic delay and phase fluctuations are kept to within 1 nanosecond and 1 radian respectively, while reflection amplitudes drift by a few percent throughout the night.

connected to the the PAM post amplifiers. A reflection across the diameter of the dish at 50 nanoseconds (marked) could explain the first bump with a suppression of 10^{-1}, which isn't terribly inconsistent with simulations of the HERA dish and feed (Figure 5 of Ewall-Wice et al. 2016). The second bump at ~ 90 ns, however, is harder to explain: if it was a harmonic of the 50 ns feature it would appear another factor of 3 or 4 more suppressed. Furthermore, simulations of the HERA feed and dish do not predict structure at a delay of 100 ns with a suppression of only 10^{-2} of the peak auto-correlation.

Lastly, we see evidence in Figure 6 of a slowly tapering systematic tail in the auto-correlations that extends to fairly high delays of $\tau > 500$ ns. The noise floor estimated from time differencing is shown in orange, and we see for most antennas some structure that rumbles above the noise at intermediate delays. If this is indeed a systematic floor and not an artifact of the fourier transform and CLEANing, it would be higher in amplitude by a factor of 10 at $\tau = 300$ ns and have a shallower slope as a function of delay than expected given simulations of the HERA antenna's chromaticity (Ewall-Wice et al. 2016).

The reflection parameters of cable reflections in principle need not be time stable. If the cables, for example, are exposed to the ambient temperature in the field, then the natural temperature drift throughout a night could cause the cables to expand and contract, leading to shifts in the light-travel distance throughout a nightly observation. To get a rough sense of the variation of the 150 m cable reflection parameters, we fit for the parameters of the 150 m cable reflection about a dozen times per hour across the 2458101 observation. They way we fit for the reflection will be described in future work, but suffice it to say we fit for the peak of the bump in the fourier transformed auto-correlations (e.g. **??** as a function of time. Figure 8 shows the result for a few antennas, showing fairly stable reflection delays, phases and amplitudes throughout the entirety of the night. This result suggests that, depending on the desired accuracy, one may be able to remove reflections sufficiently using only a handful of reflection models per night for each antenna.

2.2. *HERA's Voltage Cross Coupling*

Next we inspect the cross-correlation visibilities to look for evidence of the cross coupling systematic outlined in §1. Figure 9 shows cross-correlation visibility spectra sharing antenna 82 in common for a few baseline orientations. The inset plot in the upper left shows the array layout near antenna 82. Numbers on the upper left of each circle marks the antenna number, while the number on the upper right marks the light travel delay to the antenna from antenna 82. The line within each circle center shows the orientation of the dipole for the particular polarization: East-West for XX polarization and North-South for YY polarization.

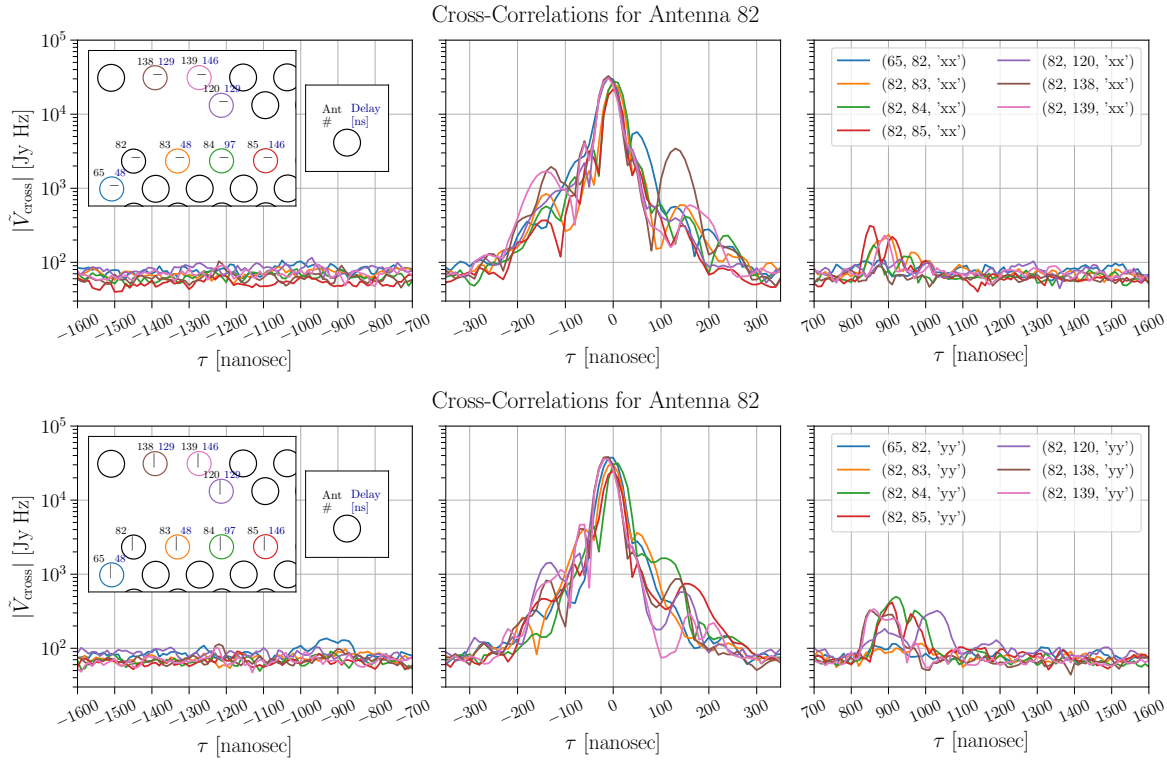

Figure 9. Cross-correlation visibilities in delay space tied to antenna 82. The inset plot shows the array layout near antenna 82 and marks the number of each antenna (upper left #), the light-travel delay in nanoseconds from antenna 82 (upper right #) and the orientation of the dipole (circle center). **Top:** Visibilities for the XX polarization, whose dipole is oriented along the East-West direction. **Bottom:** Visibilities for the YY polarization, whose dipole is oriented along the North-South direction.

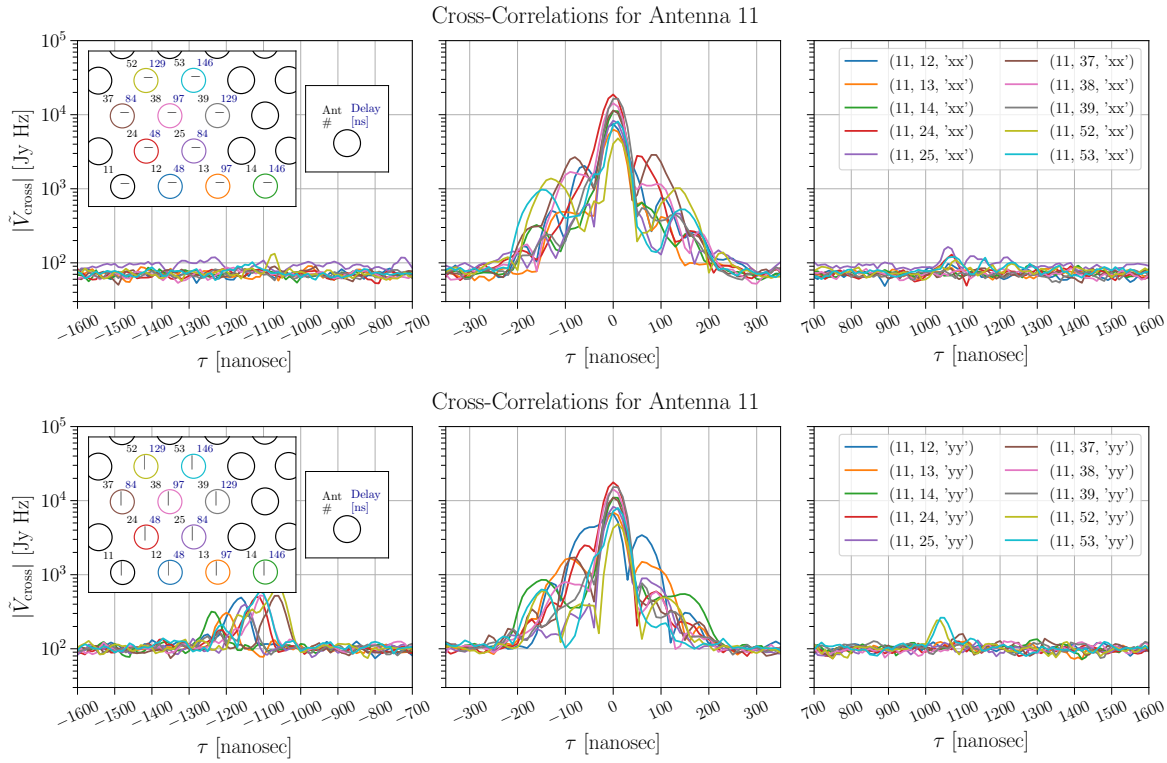

Figure 10. Cross-correlation visibilities in delay space tied to antenna 11.

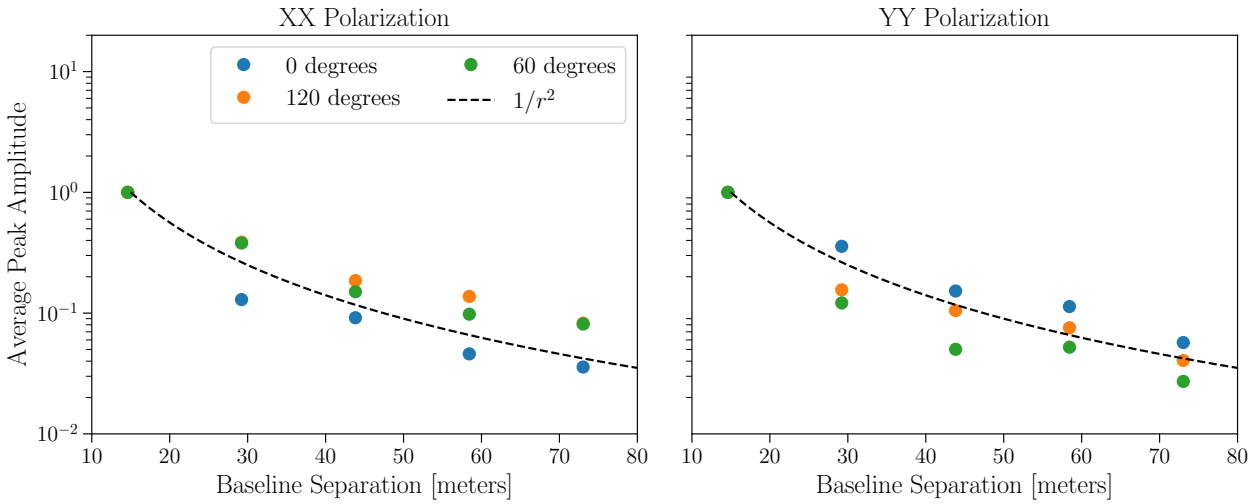

Figure 11. The dependence of the low-delay spike amplitude on baseline separation. For a given angle in ENU coordinates, we take all 1-unit, 2-unit, 3-unit, 4-unit and 5-unit baselines and incoherently average all baselines within each group. We then estimate the peak amplitude at the appropriate horizon delay of each group, and normalize all groups by the 1-unit peak amplitude. We perform this for baselines oriented along 0 degrees, 60 degrees, and 120 degrees in ENU coordintes. The dashed line is the $1/r^2$ amplitude dependance on baseline separation expected from over-the-air reflections between feeds.

There are a few salient features in Figure 9. First are the spikes at low delay ($|\tau| < 300$ ns) that seem to correspond to the baseline horizon, i.e. the light-travel delay between antennas. We can see this by cross-referencing the delay to each antenna in the inset array plot and matching it to the delay of the spikes for each baseline. The fact that the low-delay spikes match the geometric baseline horizon is indicative of either the "pitchfork" effect (Thyagarajan et al. 2015) or feed-to-feed cross coupling (i.e. reflections of incident sky emission between feeds). The pitchfork effect is explained as the boosting of measured signal collected at the horizon (along the baseline axis) due to the shortening of the projected baseline, making it more sensitive to diffuse foreground emission (in effect mimicking the autocorrelation). Diffuse foreground emission from the horizon, inherently spectrally smooth, shows up at the delay of the geometric baseline separation (Parsons et al. 2012), meaning that we expect to see peaks in the cross correlation visibility at the geometric baseline delay. This effect is most pronounced for wide-field arrays, whose weaker beam attenuation at low declinations does not completely null the incident emission from the horizon. While HERA has a more compact beam response than most wide-field arrays, simulations done by Thyagarajan et al. (2016) nonetheless predicted the presence of the pitchfork effect in HERA visibilities.

Not surprisingly, the low-delay spikes in Figure 9 are dipole-dependent, in the sense that the baselines that lie along the dipole axis–and thus at its null–are less pronounced (e.g. compare baseline (82, 83) in the XX and YY polarization). However, a spike in power at the geometric baseline horizon would also be expected of feed-to-feed reflections, so this fact doesn't really help us discriminate one mechanism from another.

Also highlighted in Figure 9 is the non-noise-like structure at high delay ($|\tau| > 700$ ns), which is of great concern because it occupies the delay modes we hope to use to measure the EoR. Its structure is much more complex, but in general, it seems to peak at a certain delay and then taper off in amplitude towards higher delays. The fact that it shows up at large delays is hard to explain: the baseline horizon is many times smaller, and the only element in the signal chain that could produce correlations in the voltage streams at delays larger than a few hundred nanoseconds is the 150 m coaxial cable. However, as we saw in Figure 6, these tend to produce structure at delays of ~ 1200 nanoseconds, so why and how this structure gets inserted at delays of $700 < \tau < 1300$ ns is hard to explain. Furthermore, the high-delay spikes do not seem to be correlated with the low-delay spikes, suggestive of a different physical mechanism for its origin than feed-to-feed reflections or the pitchfork effect. We also present a similar figure for antenna 11 (Figure 10), demonstrating the diversity of the low-delay and high-delay spikes for different baseline orientations and at different locations within the array.

In Figure 11 we show the dependance of the low-delay spike amplitude on physical baseline separation. We do this by specifying an angle in the East-North-Up (ENU) coordinates of the array–say 0 degrees North–and then collate all baselines separated by 1-hex-unit (14.5 meters), 2-hex-units (29 meters), 3-hex-units etc into separate groups. Each group represents a set of baselines that should be nominally redundant and thus measure the same modes on the sky. For each group we aver-

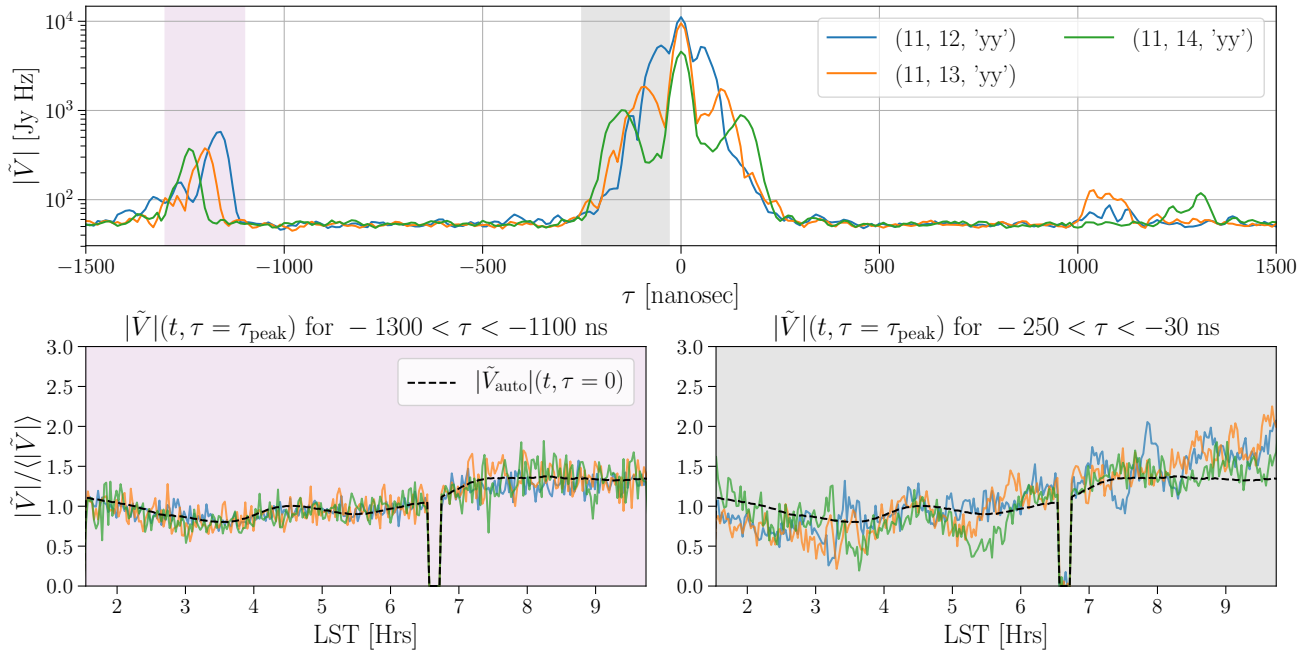

Figure 12. Time dependance of the spike amplitude. **Top:** Cross-correlation spectra showing low-delay ($|\tau| < 300$ ns) spikes at each baseline's horizon, and high-delay ($|\tau| > 700$ ns) spikes. **Bottom-Left:** Time dependence of the peak amplitude of the high-delay spikes for $\tau < -700$, normalized by their time average. The dashed line shows the time dependance of the auto-correlation $\tau = 0$ mode, which shows a close match to the time structure of the crosstalk at high delay. **Bottom-Right:** Time dependance of the peak amplitude of the baseline horizon spikes for $-250 < \tau < -30$ ns normalized by their time average, showing worse agreement with the auto-correlation $\tau = 0$ mode.

age the absolute value of the visibilities in delay space, and then take the peak amplitude at the baseline geometric horizon. We then normalize each group by the 1-unit peak for the given ENU angle. Figure 11 shows the result for the baseline types: 0 degrees North, 60 degrees North and 120 degrees North. A $1/r^2$ curve is also plotted (dashed), which shows the expected amplitude dependance of an over-the-air reflection between feeds given the inverse square law of radiation propagation. This figure seems to suggest that the low-delay spikes do indeed follow a $1/r^2$ dependence quite well. However, this does not indicate conclusively that the spikes are due to feed-to-feed reflections, as the pitchfork effect should also show a decreasing amplitude with increasing baseline separation, however its exact dependence with r depends on the structure of the sky emission at the horizon. Visibility simulations of the diffuse foreground sky at the horizon could make a prediction for the pitchfork r dependance at these LSTs, and thus help to discriminate the origin of the spikes at the baseline horizon delays in the cross correlation visibilities.

We can also learn about the systematics from their time behavior (i.e. across LST). In §1 we made a prediction that voltage cross coupling systematics (i.e. crosstalk) should be an overall additive offset in the complex visibilities with constant phase and slowly variable amplitude. If the cross coupling occurs downstream of the feed, the time variation should mimic that of the auto-correlation. If the coupling occurs upstream of the feed (i.e. inter-feed reflections of sky radiation), then depending on the angular dependence of the reflection coefficient on the sky, the crosstalk amplitude may have a faster time evolution than the auto-correlation timescale.

We can test these predictions by looking at the time-dependence of the low-delay and high-delay spikes for a few baselines. Figure 12 shows this for baselines anchored at antenna 11. The top plot has the same curves from Figure 10, showing the low-delay spikes at each baseline's horizon delay and the high-delay structure popping out from the noise. If we take the peak value of each spike and plot the time dependence of its amplitude, we get the bottom row of Figure 12: the pink-shaded plot (left) showing the time dependence of the high-delay spikes, and the grey-shaded plot (right) showing the time dependence of the low-delay spikes, where each curve has been normalized by its time-averaged value. Overplotted is the time dependence of the $\tau = 0$ mode of antenna 11's auto-correlation (dashed), which shows good agreement with the high-delay structure and not as good agreement with the low-delay structure.

From this plot, we've learned that the high-delay spikes have a slowly variable amplitude that tracks the amplitude of the average auto-correlation quite well,

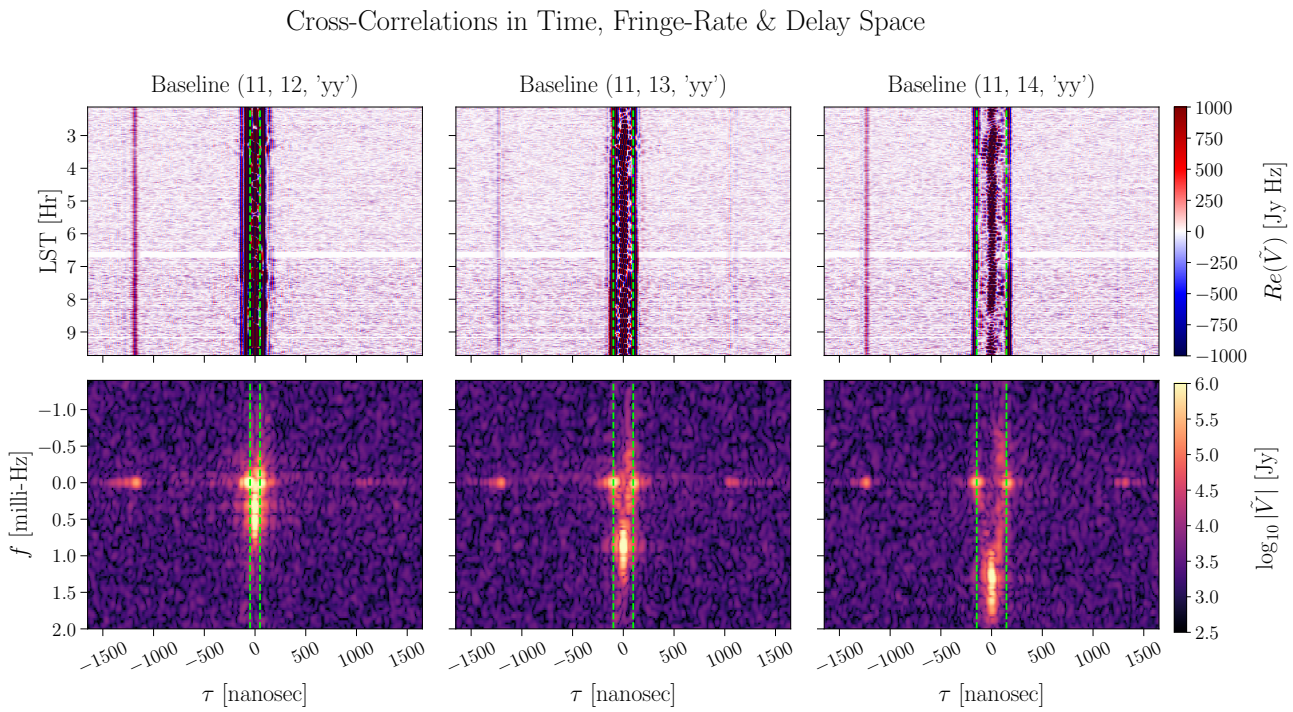

Figure 13. The same baselines from Figure 14 but plotted in time & delay space (**Top**), and fringe-rate & delay space (**Bottom**) without any time averaging. The top row shows the real component of the visibilities, which can be useful for assess how rapidly a mode is wrapping around the complex plane (i.e. oscillation from red to blue, aka "fringing"). The bottom row shows the absolute value of the data in fringe-rate and delay space. The dashed green lines demarcate the baseline's geometric horizon. Modes that occupy low fringe-rate modes in the bottom plots (e.g. $\tau \sim -1300$ ns) also have a stable position in the complex plane in the top plots. The main lobe of the foregrounds within the baseline horizon moves to progressively higher fringe-rates for longer baselines, as expected, but the low-delay spikes at the horizon remain centered at $f = 0$ mHz.

which is exactly what we would expect of voltage cross-coupling downstream of the feed. For the low-delay spikes, we still can't conclusively discriminate their origin from this plot. The coarse time structure of the low-delay spikes seems to somewhat agree with the auto-correlation, but in particular we can see lots of extra small-scale time structure (bottom-right of Figure 12). However, this doesn't mean that the spikes cannot be cross coupling in origin, because as we described in §1, if feed-to-feed reflections (i.e. cross coupling) have a non-trivial and localized angular dependence on the sky, the time evolution of the reflection terms in the cross-correlation visibility can have faster time evolution than the auto-correlation visibility.

Another way of looking at this is to plot the full fringe-rate and delay dependence of the visibilities. In Figure 13, we show visibility waterfalls in time and delay space (top row) and fringe-rate and delay space (bottom row) for the same three baselines, with the green dashed lines representing each baseline's horizon delay. The top row plots the real component of the visibility, which is useful to assessing how fast the delay mode is wrapping around the complex plane over time (i.e. how quickly its color-scale changes across time). In the top-row plots, we can see that delay modes within the horizon of each baseline show the expected oscillatory structure of foregrounds "fringing" as they track through the primary beam. However, delay modes at each baseline's horizon exhibit much slower time variation comparatively. Similarly, the high-delay modes we already know to be slowly variable in amplitude also exhibit slow time variation in the real component. We might then ask ourselves, exactly how slow do they evolve?

To quantify this, the bottom-row plots of Figure 13 show the absolute values of the visibilities in fringe-rate & delay space, which recall is just the equivalent of taking the top-row and fourier transforming across LST. Here we can clearly see that the spike structures at high-delay and low-delay are indeed centered at the zero-fringe rate mode, and both contain some spillover to positive and negative fringe-rates. The low-delay spikes contain more spillover to higher fringe-rates, which is another way of saying that they have some faster time-variable modes, indeed as we already observed in Figure 13. The main foreground lobes within each baseline's horizon shifts to increasingly higher fringe rates with longer baseline separations as expected given that larger East-West baseline separations produce visibilities with faster time structure (Parsons et al. 2016).

One might be tempted to look at the bottom-row of

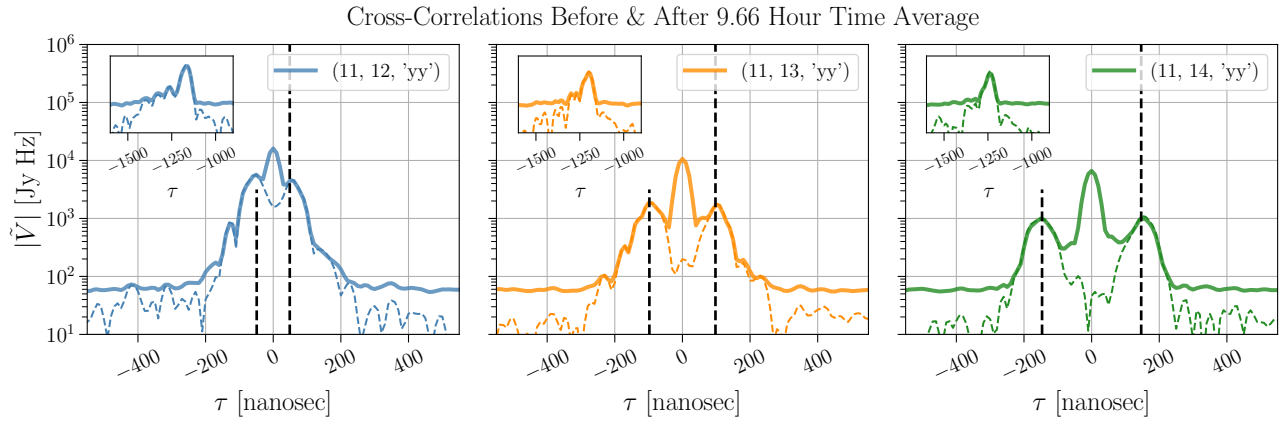

Figure 14. Cross-correlation visibility spectra in delay space before (bold) and after (dashed) taking a long time average of the complex visibilities. The modes that remain (dashed) are modes the modes that occupy low fringe rates. The inset shows the same curves at negative delays where the other non-fringing systematic appears. Dashed vertical lines mark the geometric baseline horizon for each visibility.

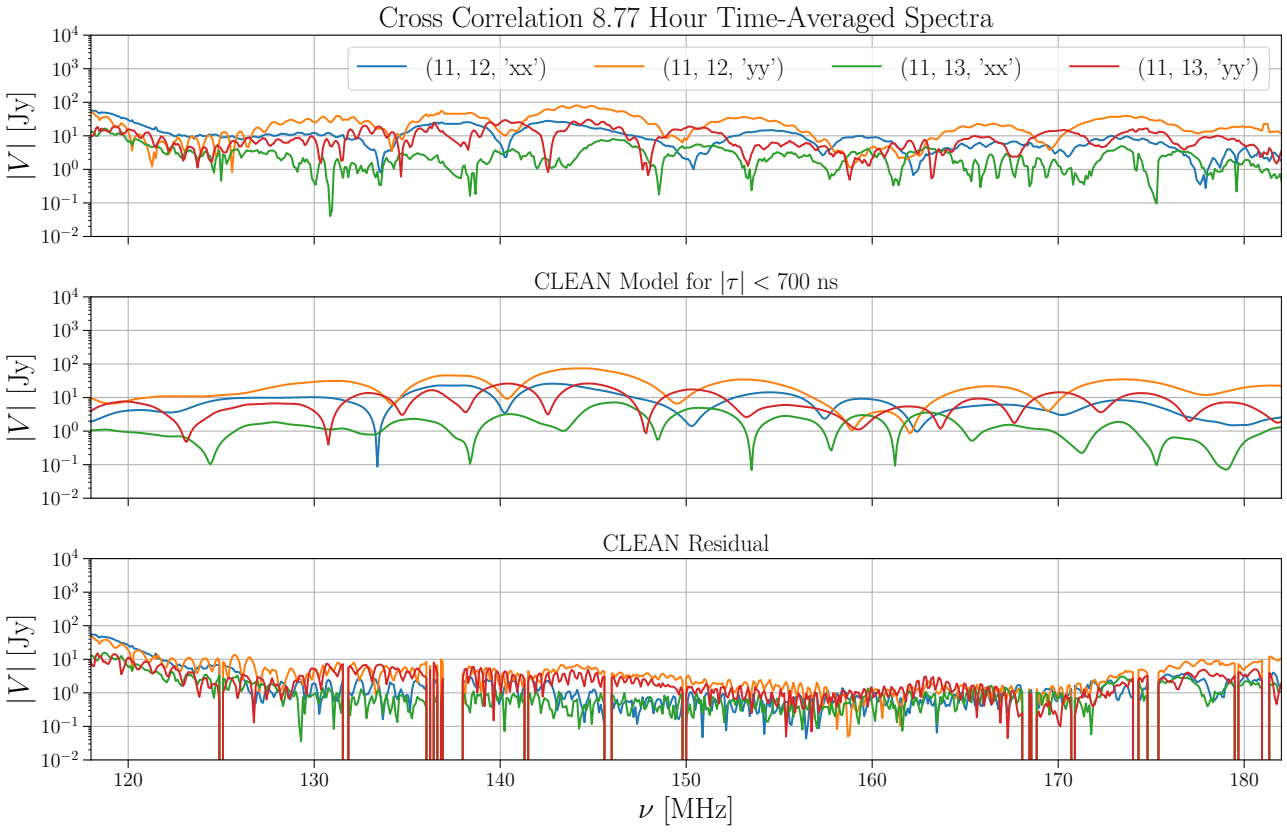

Figure 15. Delay CLEANed and time-averaged cross correlation visibilities in the frequency domain. **Top:** Time-average of the CLEAN data, containing structure at all delays. **Middle:** Time-average of the CLEAN model, with an upper delay limits of $|\tau| < 700$ ns, meaning most of the power comes from the non-fringing structure at the geometric horizon. **Bottom:** Time-average of the CLEAN residual, meaning most of the power comes from the non-fringing structure at high delays.

Figure 13 and conclude: the low-delay spikes at the baseline horizon have slowly variable time-structure with power centered at $f = 0$ mHz, therefore it must be due to voltage cross coupling, which we expect to exhibit such behavior. Unfortunately this is complicated by the fact that we also expect the pitchfork effect to produce similar results. The time dependence of sky emission incident from the horizon depends on the fringe-rates observed by the baseline in that direction. It just so happens that the fringe-rates observed by a antenna-pari at the horizon along the baseline axis is centered at fringe-rates of zero, meaning that if the emission were due to a pitchfork effect, we'd expect it to be slowly time variable. This means we still cannot conclusively discriminate the origin of the low-delay spike structure from either the pitchfork effect or feed-to-feed reflections using the evidence gathered above.

To further demonstrate the fact that the spike structure at low-delay and high-delays are slowly time variable, we plot the cross-correlation spectra in delay space before and after a time average of the complex visibilities (Figure 14). The time-averaged spectrum is equivalent to taking the $f = 0$ mHz bin from Figure 13 and plotting it as a function of delay, however, we can more clearly see in Figure 14 that most of the power in the spike structure at low and high delays comes from a nearly time-constant term in the data, while components that are noise-like or foreground-like (in the main lobe of the primary beam), average down after a time average.

The time-averaged visibilities (dashed of Figure 14) can also be inspected in the frequency domain. Figure 15 shows the frequency structure of the time-averaged visibility amplitude for a few baselines. In the top plot we show the full frequency spectra, and then in the middle and bottom plot separate the structure arising from non-fringing power at low delays and non-fringing power at high delays, respectively. We can see there are certain frequencies where the systematics are worse than other frequencies, in particular at the band edges and in the range of 130 – 150 MHz for the high-delay systematics.

3. CROSS COUPLING MODELS

Here we propose and test various physical mechanisms for the non-fringing structure found at high delays in Figure 12, Figure 13 & Figure 14. To summarize, we don't find a single compelling physical model that can explain all of the observed behavior of the systematic, but have evidence to rule out a handful of simple and not so-simple physical models.

3.1. *Single Noise Source in the Field*

A single noise source in the field can create a non-fringing systematic in the data, but we can rule out this model simply based on the fact that there is no way to boost the signal to high delays. A noise source in the field should appear at anywhere within a baseline's geometric horizon (depending on the orientation of the baseline with respect to the noise source). The fact that the systematics appear at high delay rules this out. In addition, one might expect the amplitude of the noise source to not vary with LST, however, in Figure 12 we saw that indeed the amplitude of the high-delay systematic does vary with LST and furthermore tracks the total sky amplitude quite well, suggesting that the systematic is somehow tied to the sky.

3.2. *Inter-Feed Then Cable Reflections*

One mechanism that might generate high-delay, non-fringing structure is to have feed-to-feed reflections between antenna 1 & 2, which then gets reflected in the cables of each antenna's signal chain. Assuming we only have coupling of antenna 1 voltage into antenna 2 for simplicity, we can write the antenna voltages as

$$\begin{aligned} v_1' &= v_1(1+\epsilon_{11}) \\ v_2' &= (v_2 + \epsilon_{12}v_1)(1+\epsilon_{22}), \end{aligned}$$

which when inserted into the visibility equation yields

$$\begin{aligned} V_{12}' = v_1(&v_2^* + \epsilon_{22}^* v_2^* + \epsilon_{12}^* v_1^* + \epsilon_{12}^* v_1^* \epsilon_{22}^*) \\ &+ v_1\epsilon_{11}(v_2^* + \epsilon_{22}^* v_2^* + \epsilon_{12}^* v_1^* + \epsilon_{12}^* v_1^* \epsilon_{22}^*). \end{aligned}$$

The relevant term at large negative delay is $v_1 v_1^* \epsilon_{12}^* \epsilon_{22}^*$, which shows up at a delay of $\tau = -\tau_{22} - \tau_{12}$, and the relevant term at large positive delay is $v_1 v_1^* \epsilon_{11} \epsilon_{12}^*$, which shows up at a delay of $\tau = \tau_{11} - \tau_{12}$.

There are two predictions this model makes that we can rule out with evidence gathered above. 1) The delay with which the systematic appears is the cable reflection delay ± the geometric baseline separation. Take antenna 82 for example: antenna 82 has a weak cable reflection near $\tau \sim 1250$ ns (Figure 5), meaning cross coupling systematics with its nearest neighbors should appear at characteristic delays of 1250 ns ±50 ns. The systematics we see in the data in Figure 5 are not even close to the expected value, peaking near $\tau \sim 900$ ns. Furthermore, 2) we can see that a cable reflection in 82 should boost systematics to positive delays, but 82 has a very weak reflection compared to, say antenna 84, whose reflection should boost the cross coupling systematic to negative delays in the $V_{82,84}$ visibility. However, we see no evidence of the systematic at negative delays in $V_{82,84}$ (Figure 9): in other words, the amplitude of the systematic doesn't seem dependent on the amplitude of the cable reflection as predicted by the model above (assuming $|\epsilon_{12}| \sim |\epsilon_{21}|$).

3.3. *Cable Reflection Then Signal Broadcasting*

Let's swap the order of our previous model. In this case, we have a signal that enters the feed of antenna 1,

travels down the 150 m cable and is reflected back up the signal chain, exits the feed and is received by antenna 2. In this case we can write the antenna voltages as

$$v_1' = v_1(1+\epsilon_{11})$$
$$v_2' = (v_2 + \epsilon_{11}\epsilon_{12}v_1)(1+\epsilon_{22}).$$

If we insert this into the visibility equation and keep only non-fringing terms (i.e. auto-correlation terms) we get

$$V_{12}' = v_1(\epsilon_{11}^*\epsilon_{12}^* v_1^* + \epsilon_{11}\epsilon_{11}^*\epsilon_{12}^* v_1^* + \epsilon_{11}^*\epsilon_{12}^* v_1^*\epsilon_{22}^* + \epsilon_{11}\epsilon_{11}^*\epsilon_{12}^* v_1^*\epsilon_{22}^*) + \ldots,$$

with the first term being the least suppressed, and therefore likely the strongest in the data. However, we can see that we face a similar problem as before: the systematics appear at integer values of the cable delay and geometric delay summed. It is hard to produce systematics at delays of $\sim$ 900 nanoseconds, like we see in Figure 9, when we can see that the cables produce reflections on the order of $\sim$ 1200 ns. Furthermore, one might reason that signal traversing backwards through the FEM and out of the feed is an unlikely scenario due to the reverse isolation in the low-noise amplifier (LNA) in the FEM.

3.4. *A Single Broadcasting Antenna*

One way to get the structure to show up at the intermediate delays of $\tau \sim 900$ ns is to postulate that the excess signal comes from a single antenna that has a shortened cable and is actively broadcasting. Consider a single antenna that collects radiation which hits some stage in the FEM and is reflected back out the feed, such that the signal is broadcasted to neighboring antennas with only a geometric delay. Now consider the part of the signal that made in through the FEM: suppose it traveled down this antenna's cable, was reflected and traveled back up the cable and made it through the FEM somehow and was also broadcasted, such that it acquires a cable delay in addition to a geometric delay when received by neighboring antennas. If this antenna's cable was somehow cut shorter than the others, then we can have a scenario where the antenna is broadcasting two coherent signals into the field with a relative delay of roughly 900 ns.

In this case, the measured voltage of antenna 1 and antenna 2 can be written as

$$v_1' = v_1 + \epsilon_{31}\epsilon_{33}v_3 + \epsilon_{31}v_3$$
$$v_2' = v_2 + \epsilon_{32}\epsilon_{33}v_3 + \epsilon_{32}v_3,$$

in which case their measured visibilities (again keeping only non-fringing terms) can be written as

$$V_{12}' = \epsilon_{31}\epsilon_{33}v_3\epsilon_{32}^*\epsilon_{33}^* v_3^* + \epsilon_{31}\epsilon_{33}v_3\epsilon_{32}^* v_3^* + \epsilon_{31}v_3\epsilon_{32}^*\epsilon_{33}^* v_3^* + \epsilon_{31}v_3\epsilon_{32}^* v_3^* + \ldots,$$

out of which the terms that would show up at large positive and negative delays are the second and third terms, respectively. However, one problem with this model is its implied symmetry: the second and third term are identical other than the fact that one term shows up at $\tau_{33}+\tau_{12}$ and the other at $-\tau_{33}+\tau_{12}$. This means that in our case of a single broadcasting antenna, we would expect to see systematics at *both* negative and positive delays in the cross-correlation visibilities, where Figure 9 and Figure 10 shows that we distinctly don't see in the data. Additionally, we would also expect these terms to come through in the measured auto-correlation, however, we don't see evidence of such structure in the auto-correlation visibilities of the data.

3.5. *Diagnostic Plots*

While we haven't been able to find a single physical model that explains all of the behavior of the non-fringing systematics at high-delay, we have been able to reject a handful of simple models. In order to facilitate further discussion, in this section we provide a few more diagnostic plots of the high-delay systematics in the data.

We begin by asking the question, can we somehow distill the baseline-dependent behavior of the systematic into an antenna-based space? To do this, we iterate over all baselines in the data and for each one take the peak amplitude of the cross correlation visibility in the regime $\tau < -700$ ns and append it to a list for antenna 1 and repeat for $\tau > 700$ ns and append that peak value to a list for antenna 2. We do this for all baselines and take the average of each antenna's list to get the mean of the peak crosstalk amplitude for each antenna. If we refer back to Equation 11, we can see that mathematically, this process is assigning the ϵ_{12}^* term to antenna 1 and the ϵ_{21} term to antenna 2. Put into words, this is some kind of measurement of the "reflectivity" of each antenna, or the amount of radiation leaked from one antenna into any neighboring antenna.

The top row of Figure 16 shows this for each antenna and linear visibility polarization in the array, while the bottom row shows the average of the peak delay for each antenna, weighted by the peak amplitude in each antenna's list. We can immediately see a clear preference for high systematic amplitudes in the North Western corner of the array, and what is most interesting about this trend is that it is independent of the post amplifier model and receiver model (c.f. Figure 3). The average delay–weighted by the systematic amplitude peak–for each antenna does not show a clear preference for any particular location of the array, however.

We can also reverse the process and, for each cross-correlation visibility, assign the peak systematic amplitude at negative delays (i.e. the ϵ_{12}^* term) to antenna

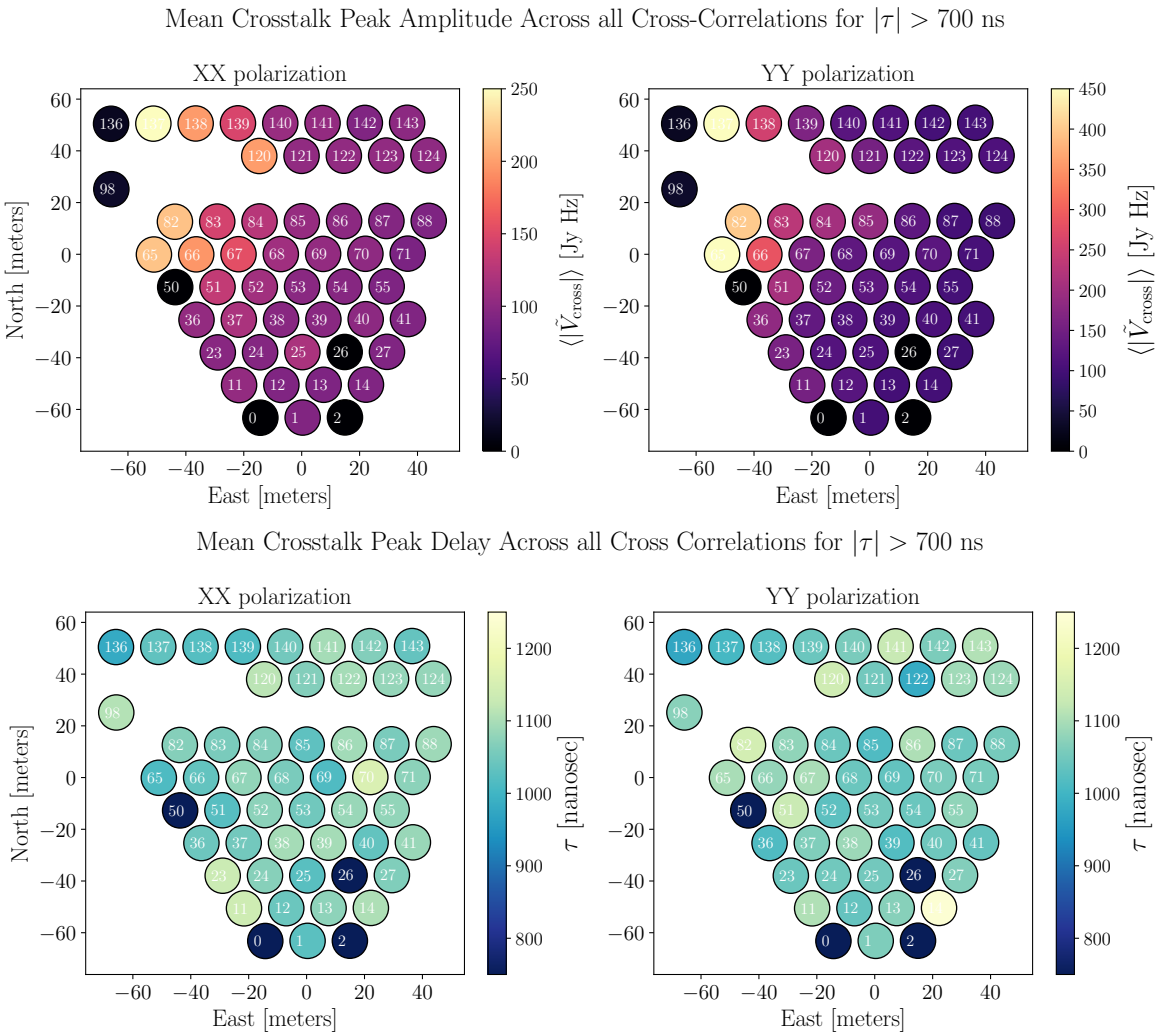

Figure 16. Antenna-based plots of the high-delay ($|\tau| > 700$ ns) systematic, assigning $\epsilon_{12}^* \to 1$ and $\epsilon_{21} \to 2$. **Top:** The average systematic peak amplitude across all cross-correlation visibilities for each antenna and auto-visibility polarization. **Bottom:** The average delay of the peaks, weighted by the peak amplitude.

2 and the peak systematic amplitude at positive delays to antenna 1. This is now some kind of measurement of the amount of leaked radiation an antenna collects from any neighboring antenna. Similar to before, the top row of Figure 17 shows the average systematic amplitude for each antenna and polarization, while the bottom row shows the average delay of the peak weighted by the peak amplitude. Interestingly, the amplitude plots show less of a preference for any particular location in the array, and the delay plots now show a clear trend for the North West corner of the array. Furthermore, the minimum delay of the antennas in that corner reads out at about 850 nanoseconds, which happens to coincide with some structure in the auto-correlations of antenna 98 and 136 at similar delays.

It's not clear if the evidence presented above proves that the high-delay systematics are tied to any particular antenna, however, there is some evidence that in trying to distill the baseline-dependent systematic into an antenna-based space, some trends pop out that seem to indicate some preference for the North Western corner of the array, and in particular the flagged antennas 98 and 136. More work is needed in trying to develop a physical model that can explain all of the observed behavior of the systematic, including the range of delays it appears at, its variable amplitudes and odd amplitude shape as a function of delay (something of a decaying sawtooth Figure 9).

4. SUMMARY

This memo outlined the phenomenology of two kinds of instrument systematics: signal chain reflections and antenna cross coupling (i.e. crosstalk). We showed in §1 how signal chain reflections and cross coupling terms show up in the measured cross-correlation and auto-correlation visibilities, and made predictions for their frequency and time structure. In §2 we investigated data

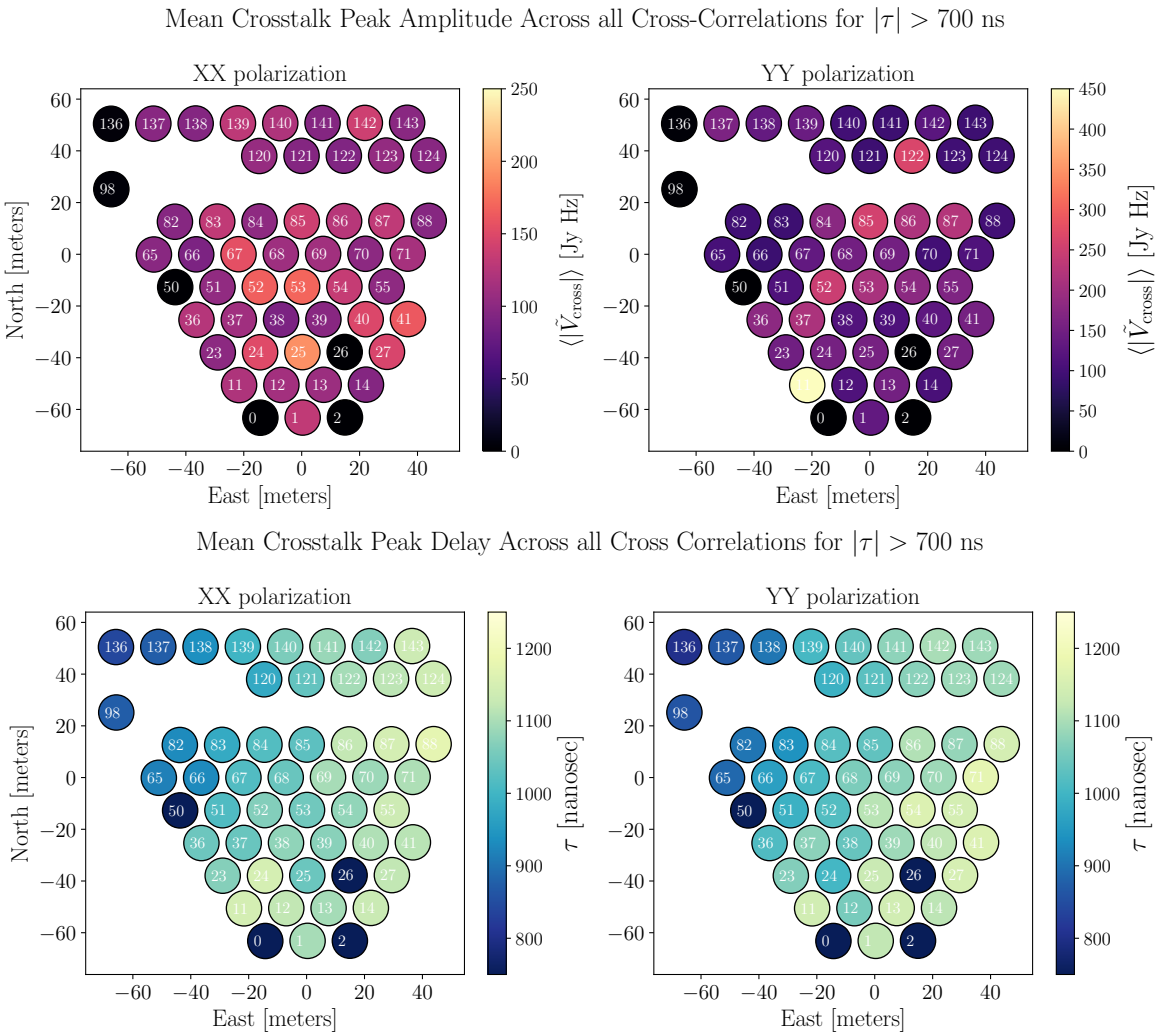

Figure 17. Antenna-based plots of the high-delay ($|\tau| > 700$ ns) systematic, assigning $\epsilon_{12}^* \to 2$ and $\epsilon_{21} \to 1$. **Top:** The average systematic peak amplitude across all cross-correlation visibilities for each antenna and auto-visibility polarization. **Bottom:** The average delay of the peaks, weighted by the peak amplitude.

from HERA's first observing season (H1C) on night JD = 2458101, and find the following

1) Cable reflections can be found in the auto-correlation visibilities at high significance. The 150 m cables from FEM to receiver are found to have reflections with delays of $1200 < \tau < 1400$ ns, and amplitudes of roughly 2×10^{-3} from signal chains with PAM post amplifiers, and 2×10^{-4} for those with RCVR post amplifiers. Signal chains with PAM amplifiers also show reflections in the 20 m cable from receiver to correlator at delays of roughly 190 ns, with amplitudes of roughly 8×10^{-3}.

2) The 150 m cable reflection shows fairly stable reflection parameters over the course of a single night, and removal to high fidelity will likely require a per-antenna reflection model a couple times per night.

3) All antennas show a systematic tail extending out to $\tau \sim 1000$ ns at a dynamic range of $\sim 10^{-4}$. This tail does not average down after LST integration, but structure outside of $\tau > 1200$ ns seems to integrate down like noise, at least down to a dynamic range of $\sim 10^{-5}$.

4) The CLEANed cross correlation visibilities across the full band clearly show evidence of excess power at $\pm$ each baseline's horizon delay that has non-negligible bleed-out into the EoR window. These spikes could be due to either a pitchfork effect (Thyagarajan et al. 2015) or due to feed-to-feed reflections of incident sky emission. We provide a series of diagnostic plots that show the time-dependence of these spikes have a somewhat slow time-dependence, with their power centered at $f = 0$ mHz but extending out to non-zero fringe rate modes enough such that their full time de-

pendence is faster than the time evolution of the auto-correlations. Their amplitude dependence as a function of baseline separation follows a $1/r^2$ quite well, which is expected for over-the-air feed reflections, but could also be the case for a pitchfork effect. At this time, we cannot conclusively say what the origin of the excess power at the baseline horizon is (pitchfork or feed-to-feed reflections), but we do have a better understanding of its time and frequency behavior, which can be used to separate this structure from EoR structure if necessary.

5) We also see evidence for significant excess power at high delays $|\tau| > 700$ ns. We show that the time dependence of these features closely follow that of the auto-correlation $\tau = 0$ mode, which is what we expect for cross coupling that occurs downstream of the feed (i.e. after the feeds convert the incident electric field to a voltage). These systematics are also tightly confined to low fringe rate modes centered at $f = 0$ mHz, which can be exploited to separate them from EoR if necessary. We still do not understand why and how the structure appears at such high delays: their appearance is not directly correlated with the cable reflections or the excess power at the baseline horizon. However, to get to such high delays, it is likely that one or multiple 150 meter cables in the array are somehow involved: no other differential path length exists in the signal path that could boost the structure to such high delays, unless there is something much more complex happening within the analog system that is creating these systematics.

6) The gathered evidence seems to rule out a few simplistic physical models of the high-delay non-fringing systematic, presented in §3. An effort to distill the baseline-dependent systematic into an antenna-based space seems to show a correlation of the systematic to the North Western corner of the array, but it is still uncertain whether this suggests that one or a few antennas are responsible for the systematic.

Future work will demonstrate how to model and remove these systematics without attenuating the desired EoR signal in the data.

REFERENCES

DeBoer, D., et al. 2017, PASP, 129, 45001

Ewall-Wice, A., et al. 2016, ApJ, 831, 196

Högbom, J. A. 1974, A&AS, 15, 417

Neben, A. R., et al. 2016, ApJ, 826, 199

Parsons, A., Pober, J., Aguirre, J., Carilli, C., Jacobs, D., & Moore, D. 2012, ApJ, 756, 165

Parsons, A. R., & Backer, D. C. 2009, AJ, 138, 219

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

Patra, N., et al. 2018, Experimental Astronomy, 45, 177

Thyagarajan, N., Parsons, A. R., DeBoer, D. R., Bowman, J. D., Ewall-Wice, A. M., Neben, A. R., & Patra, N. 2016, ApJ, 825, 9

Thyagarajan, N., et al. 2015, ApJ, 804, 14