

Absolving the SSINS of Precision Interferometric Radio Data: A New Technique for Mitigating Faint Radio Frequency Interference

MICHAEL J. WILENSKY,¹ MIGUEL F. MORALES,^{1,2} BRYNA J. HAZELTON,^{1,3} NICHOLE BARRY,^{4,5} RUBY BYRNE,¹ AND SUMIT ROY⁶

¹*Department of Physics, University of Washington, Seattle, WA 98195, USA*

²*Dark Universe Science Center, University of Washington, Seattle, 98195, USA*

³*eScience Institute, University of Washington, Seattle, WA 98195, USA*

⁴*School of Physics, The University of Melbourne, Parkville, VIC 3010, Australia*

⁵*ARC Centre of Excellence for All Sky Astrophysics in 3 Dimensions (ASTRO 3D), Australia*

⁶*Department of Electrical and Computer Engineering, University of Washington, Seattle, WA 98195, USA*

Submitted to Publications of the Astronomical Society of the Pacific June 2, 2019, Accepted August 19, 2019

ABSTRACT

We introduce a new pipeline for analyzing and mitigating radio frequency interference (RFI), which we call Sky-Subtracted Incoherent Noise Spectra (SSINS). SSINS is designed to identify and remove faint RFI below the single baseline thermal noise by employing a frequency-matched detection algorithm on baseline-averaged amplitudes of time-differenced visibilities. We demonstrate the capabilities of SSINS using the Murchison Widefield Array (MWA) in Western Australia. We successfully image aircraft flying over the array via digital television (DTV) reflection detected using SSINS and summarize an RFI occupancy survey of MWA Epoch of Reionization data. We describe how to use SSINS with new data using a documented, publicly available implementation with comprehensive usage tutorials.

Keywords: methods: data analysis—methods: statistical—techniques: interferometric—telescopes, dark ages, reionization, first stars

1. INTRODUCTION

Radio frequency interference (RFI) is a ubiquitous problem in radio astronomy. The omnipresence of anthropogenic radio emission and natural emissions from sources such as lightning make it exceedingly difficult to place a serviceable radio telescope in a perfectly radio-quiet location. For instance, the Murchison Widefield Array (MWA) (Tingay et al. 2013; Wayth et al. 2018), which is located in the extremely radio-quiet Murchison Radio Observatory (MRO) in Western Australia, regularly observes ORBCOMM satellite transmissions (Sokolowski et al. 2016), lightning (Sokolowski et al. 2016), as well as FM radio signals reflecting off of satellites (Zhang et al. 2018), meteor trails (Zhang et al. 2018), and even the moon (McKinley et al. 2013, 2018). The MWA also observes digital television (DTV) trans-

missions (Offringa et al. 2015; Sokolowski et al. 2016), which we show several instances of in this work.

The brightness of observed RFI signals can vary by orders of magnitude. A survey of the RFI environment of the MRO, including brightness distribution, is presented in Sokolowski et al. (2016), while an MWA RFI occupancy study is presented in Offringa et al. (2015). In Offringa et al. (2013), LOFAR (van Haarlem et al. 2013) is used to study RFI brightness for cases where a uniform spatial distribution of RFI emitters is appropriate to assume. Fortunately, a remote telescope such as the MWA does not experience a spatially uniform distribution of RFI emitters¹. However, faint RFI contamination that is below the single baseline thermal noise is still an undeniable reality of such a telescope.

¹ This can be seen for the MRO by reconciling the observed brightness distribution in Sokolowski et al. (2016) with the theoretical brightness distribution of spatially uniform emitters in Offringa et al. (2013).

Using the MWA, we study Epoch of Reionization (EoR) cosmology via redshifted 21-cm emission from atomic Hydrogen. Reviews on the subject of 21-cm radiation and its role in cosmology are presented in Furlanetto et al. (2006) and Morales & Wyithe (2010). The EoR signal is extremely faint relative to the astrophysical foregrounds and RFI signals. Even observed RFI signals that are fainter than the single baseline thermal noise are orders of magnitude brighter than the expected EoR signal. Hence, subthermal RFI excision is a priority in our analysis.

In general, RFI mitigation strategies can be deployed at different stages of the interferometric pipeline. Reviews of RFI mitigation schemes are presented in An et al. (2017), and Baan (2019). Given the vast data rates of modern radio telescopes, it is common to perform RFI detection offline in the post-correlation stage of the pipeline. Post-correlation mitigation strategies include neural network approaches (Burd et al. 2018; Kerrigan et al. 2019), filters in the uv-plane (Sekhar & Athreya 2018), variants of fringe filters (Athreya 2009; Offringa et al. 2012a), as well as methods that take advantage of the cyclostationarity of common RFI signals (Bretteil & Weber 2005; Hellbourg et al. 2012). Several post-correlation detection methods are compared in Offringa et al. (2010), including surface fitting, singular value decomposition, and combinatorial thresholding algorithms such as SUMTHRESHOLD. Successful implementations of the SUMTHRESHOLD method are presented in Offringa et al. (2015) and Peck & Fenech (2013) using the AOFLAGGER² and SERPent³ software packages, respectively. We use an RFI detection algorithm in this work similar to SUMTHRESHOLD, with slight differences.

Currently, all MWA EoR observations are flagged for RFI during pre-processing using AOFLAGGER. Like many of the state-of-the-art methods referred to above, AOFLAGGER flags a single baseline at a time. Clearly, single baseline flaggers are fundamentally limited to the sensitivity attainable by a single baseline. As we will show, there exists RFI observed by the MWA that evades detection by AOFLAGGER because it is fainter than the single baseline thermal noise. We offer Sky-Subtracted Incoherent Noise Spectra (SSINS) as a flexible, sensitive, and statistically simple RFI analysis pipeline to detect, classify, and excise such faint RFI. SSINS was developed for the MWA, and it has also undergone limited testing on the Hydrogen Epoch of Reion-

ization Array (HERA) (DeBoer et al. 2017) and the Long Wavelength Array (LWA) (Taylor et al. 2012).

After explaining the SSINS method in detail (§2), we show several common cases of RFI that are successfully flagged by SSINS in the MWA EoR highband, which extends from 167.1 Mhz to 197.7 Mhz (§3). We also image RFI and hypothesize about the source (§4). We then discuss the implementation of SSINS, explain how it can be customized for different RFI environments, and summarize RFI mitigation efforts for data used in an EoR power spectrum limit in Barry et al. (2019b) (§5). Finally, we discuss possibilities for future work (§6).

2. THE METHOD

The measurements from an interferometer can be broken down into three components: the astrophysical sky signals, thermal noise, and RFI. We will first separate the slowly varying sky signal from the modulated RFI (§2.1), then increase our sensitivity to the faint RFI signal through successive integrations and frequency-matched flaggers (§2.2–2.4). In Figures 1–3 we compare an observation free of RFI (left) with an observation contaminated by faint RFI (right).

2.1. Sky Subtraction

The first step, sky subtraction, relies on the time variation of the sky signal being slow relative to the visibility cadence. We denote the visibility belonging to the ij antenna pair, time integration, t_n , frequency, ν , and polarization, p , as $V_{ij}(t_n, \nu, p)$. By subtracting data from subsequent time integrations, the majority of the astrophysical sky signal will be removed while leaving much of the modulated RFI and thermal noise. We write the sky-subtracted visibilities as

$$\Delta V_{ij}(t_n, \nu, p) = V_{ij}(t_{n+1}, \nu, p) - V_{ij}(t_n, \nu, p). \quad (1)$$

What remains has a noise-like component and potentially an RFI-like component.

To separate the thermal noise from the RFI it is helpful to understand the statistical properties of the thermal noise. For visibility data, the noise is a circular complex Gaussian process. In other words, for each visibility, the real and imaginary components of the noise are independently and identically distributed Gaussian random variables with mean equal to zero. Following standard derivations such as in Zwillinger & Kokoska (2000), the amplitudes, X , of a circular Gaussian process are Rayleigh distributed:

$$f_X(x; \sigma) = \frac{x}{\sigma^2} e^{-x^2/2\sigma^2}, \quad (2)$$

where σ^2 is the variance of the Gaussian which describes the real (or imaginary) component of the process. The

² <https://sourceforge.net/p/aoflagger/wiki/Home/>

³ <https://github.com/daniellefenech/SERPent>

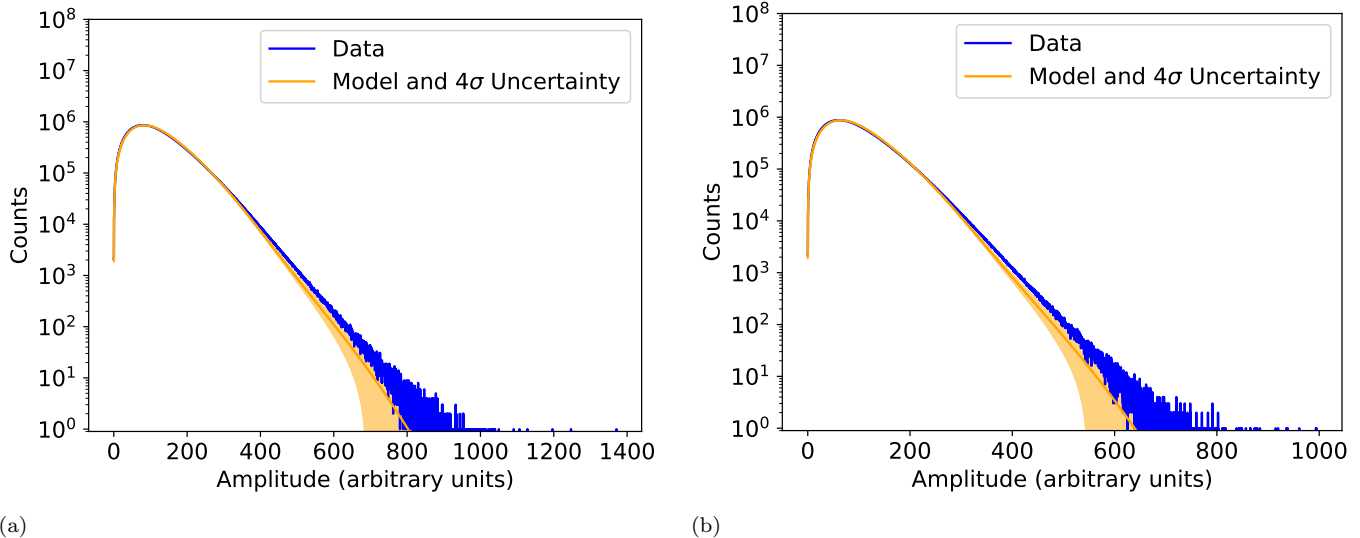


Figure 1. Histograms of the amplitudes of time-differenced visibilities for two different two-minute MWA observations, similar to Figure 8.2 of Barry (2018). (a) belongs to an observation deemed clean by the methods shown in this paper, while (b) belongs to an observation that is shown to have some digital television contamination (§2.4). The measurements are shown in blue, while an accompanying Rayleigh-mixture fit is shown in orange with 4σ error bars. These figures are nearly indistinguishable, despite differences in contamination.

mean of the Rayleigh distribution is

$$E[X] = \sqrt{\frac{\pi}{2}}\sigma, \quad (3)$$

and the variance is

$$E[X^2] - E[X]^2 = \frac{4 - \pi}{2}\sigma^2. \quad (4)$$

Figure 1 shows the Rayleigh-like amplitude distributions of the sky-subtracted visibilities for our two example observations along with a Rayleigh-mixture fit. Because the observed noise has a frequency dependence, each frequency channel was fit independently by maximum likelihood estimation to form the final model fit and 4σ errors. The similarity of Figures 1(a) and (b) shows that the sensitivity of a single visibility is not sufficient to mitigate the faint RFI that we are interested in.

2.2. The Incoherent Average

We boost our RFI sensitivity by averaging the amplitudes of the sky-subtracted visibilities over all of the baselines in the array, leaving a single dynamic spectrum per polarization. This is not an entirely novel idea. We refer the reader to the last two paragraphs of §4.2 in Offringa et al. (2015). Formally, for an array with N_A antennas, the incoherent average, Y , is

$$Y(t_n, \nu, p) = \frac{2}{N_A(N_A - 1)} \sum_{i=1}^{N_A} \sum_{j>i}^{N_A} |\Delta V_{ij}(t_n, \nu, p)|. \quad (5)$$

Since we discard the phase of the sky-subtracted visibilities in the average, we call this an incoherent average. Note that we do not include autocorrelations in the average. We refer to the remaining spectra as sky-subtracted incoherent noise spectra (SSINS), incoherent noise spectra, or just noise spectra.

Examples of incoherent noise spectra for the same observations as in Figure 1 are shown in Figure 2. The spectra in Figure 2(b) actually show some RFI from digital television broadcasting, although it is extremely difficult to discern from the surrounding noise even in this very sensitive space. In order to determine the nature of such features, we boost the contrast in a way that allows for the rigorous application of statistics.

2.3. Mean Subtraction

Mean subtraction transforms the data so that the data of a clean observation will be standardized: it will appear as if it were sampled from a zero-mean, unit width Gaussian probability distribution. We describe this process formally using the central limit theorem.

In our context, we take the central limit theorem to say the following (see Billingsley (1995)). Let (X_1, X_2, \dots, X_N) be independent and identically distributed random variables with finite mean, μ , and variance, Σ^2 . Then, as N grows large, the sample means given by

$$S_N = \frac{1}{N} \sum_{k=1}^N X_k \quad (6)$$




Figure 2. The incoherent noise spectrum for the E-W polarization corresponding to the histograms in Figure 1, on a two second cadence. The MWA employs a two-stage polyphase filter bank, involving a coarse channelization and then a fine channelization of each coarse channel. The periodic banding in frequency that is seen in these spectra is a result of that filter. The spectrum on the left is clean. However, if one examines the spectrum on the right extremely closely, they may notice a smudge between 174 and 181 Mhz in the first twenty seconds of the observation. This smudge is made much more obvious after the mean-subtraction transformation, shown in Figure 3, and it is DTV interference.

converge in distribution to a normally distributed random variable of mean, μ , and variance, Σ^2/N . Let us assume the thermal noise to be independent between baselines and ignore baseline-to-baseline noise variation⁴. Having averaged over so many baselines⁵ in Equation 5, the central limit theorem states that the thermal background in the incoherent noise spectrum will be very nearly normally distributed at each frequency with a mean described by Equation 3 and variance described by Equation 4, divided by the number of baselines in the array, which we denote N_{bl} :

$$N_{bl} = \frac{N_A(N_A - 1)}{2}. \quad (7)$$

Note that the Rayleigh distribution has only a single parameter, so we can actually relate these two quantities. Using Equations 3 and 4, and writing the underlying mean for each frequency as μ_ν , we write the underlying thermal background distribution for each frequency as

$$f_Y(y; \mu_\nu, N_{bl}) = \sqrt{\frac{N_{bl}}{2\pi C \mu_\nu^2}} \exp\left[-\frac{N_{bl}}{2C \mu_\nu^2}(y - \mu_\nu)^2\right], \quad (8)$$

where

$$C = \frac{4}{\pi} - 1 \quad (9)$$

⁴ Statements of the central limit theorem exist for non-identically distributed sequences, which would be important here if baseline-to-baseline variation in noise levels were a dominant effect. See Billingsley (1995).

⁵ Over 8000 for the MWA

is the ratio of the Rayleigh variance to the square of its mean.

At this stage, the background distribution for the entire spectrum is a mixture distribution of all the frequencies. If for each frequency we subtract the mean and normalize with respect to the standard deviation, then the background for the entire spectrum will be described by a single distribution (the standard normal distribution), thus simplifying the problem.

If there are sufficiently many time samples in an observation, then an estimation of the means per frequency can be obtained simply by taking the mean of the spectrum in time:

$$\hat{\mu}(\nu, p) = \frac{1}{N_t - 1} \sum_{n=1}^{N_t-1} Y(t_n, \nu, p). \quad (10)$$

Writing the mean only as a function of frequency and polarization of course assumes that the thermal process is at least wide-sense (weakly) stationary in time. In the event that there is some drift, then a trend line or trend polynomial can be calculated for each frequency and polarization. For now, we will work with the simplest, most prevalent case.

Now we consider the quantity

$$Z(t_n, \nu, p) = \frac{Y - \hat{\mu}}{\sqrt{C \hat{\mu}^2 / N_{bl}}} \quad (11)$$

We call this the mean-subtracted incoherent noise spectrum. Here, for each data point, we have subtracted the estimated mean and normalized with respect to the

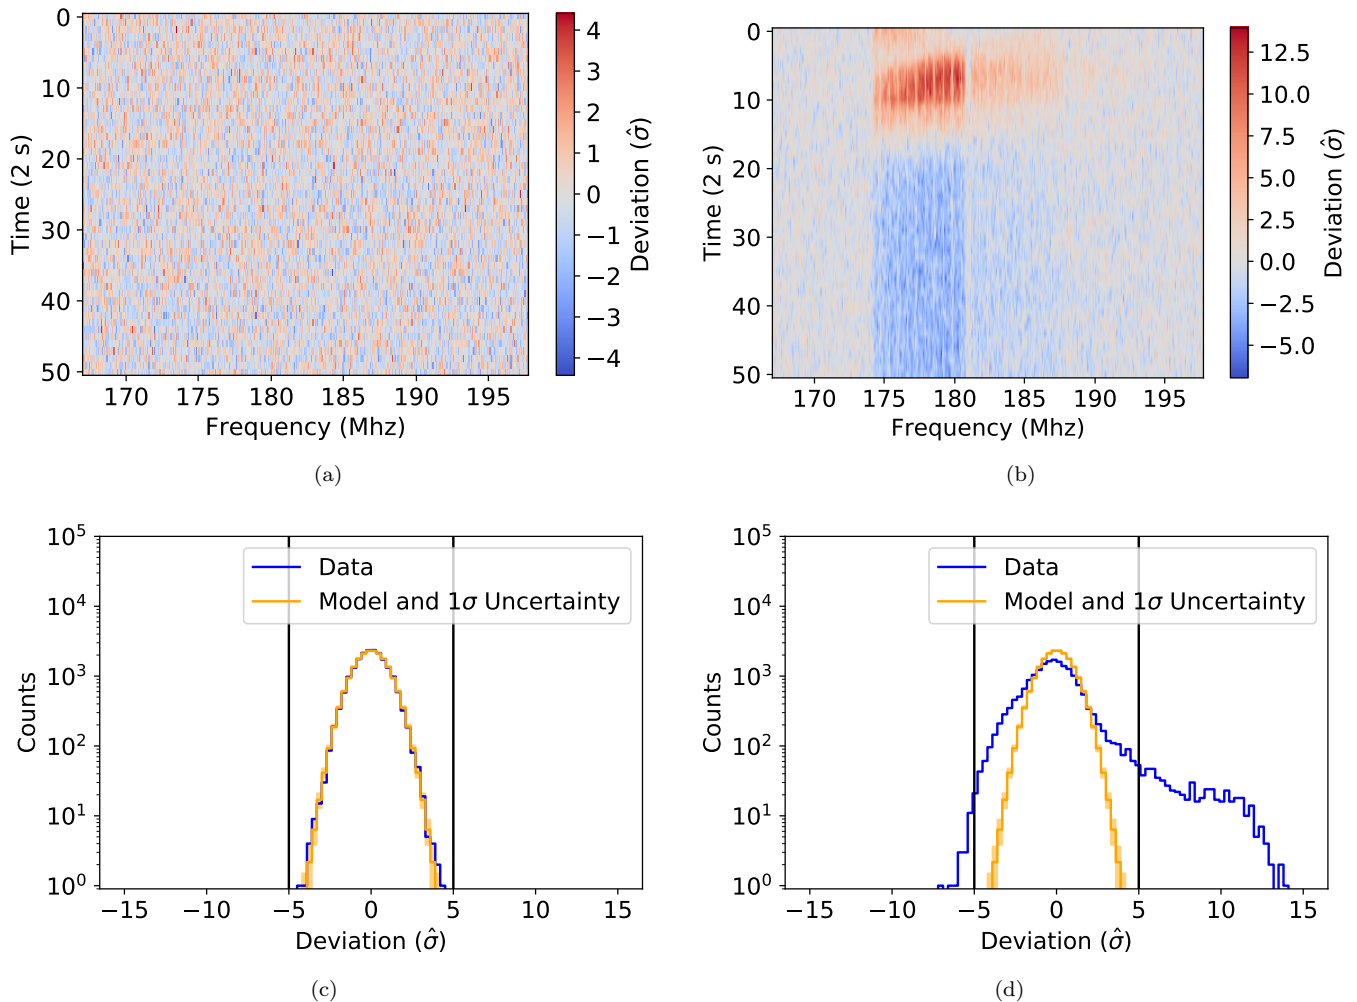


Figure 3. Top Row: the mean-subtracted incoherent noise spectra for a single polarization, corresponding to the data in Figure 2. Bottom Row: Histograms for the mean-subtracted incoherent noise spectra in the top plots, shown with vertical lines that demarcate a point beyond which very few outliers are expected at this data volume. The spectrum in (a) is clean. There are no discernible features such as clustering in time or frequency and the extent of the data is within the range of outliers expected for a data volume of this size. The spectrum in (b), however, features a noticeable cluster of positive outliers that extends from 174-181 Mhz in the beginning of the observation, and a slightly less noticeable one from 181-188 Mhz. Not only is this clustering antithetical to stationary noise, but the data in the brightest feature has outliers as strong as $14\hat{\sigma}$, which is not expected at this data volume. In correspondence with this reasoning, the purportedly clean observation in (c) looks highly Gaussian, while the contaminated observation in (d) is clearly deviating from the thermal model.

estimated standard deviation in the corresponding frequency channels and polarizations. Should the observation be totally clean of RFI and if sky-subtraction indeed fully removed the sky from the data, then the remaining data, which will be purely thermal, ought to be very nearly distributed according to

$$f_Z(z) = \frac{1}{\sqrt{2\pi}} e^{-z^2/2}, \quad (12)$$

otherwise known as the standard normal distribution. In other words, the quantity in Equation 11 is a z-score for each sample in the spectrum.

An example of mean-subtraction results following from Figures 1 and 2 is shown in Figure 3. The accompanying histograms are plotted along with a standard normal distribution. The clean observation in Figure 3(c) conforms to the standard normal distribution exceedingly well, while the contaminated observation in Figure 3(d) has outliers well beyond what is expected, indicating highly non-thermal behavior. Recall that in Figure 1, it was nearly impossible to discern any difference between these two observations.

The mean-subtracted spectra are very useful for highlighting RFI that is only marginally brighter than the surrounding thermal noise after incoherently averaging

over all baselines. In order to programmatically identify and flag RFI, we deploy a frequency-matched flagger in the mean-subtracted spectrum.

2.4. The Frequency-Matched Flagger

We extend the probability theory from the previous section to develop a frequency-matched flagger. We introduce this development by showing the results of thresholding on a single-sample basis, and then we show how information from multiple samples can be incorporated.

If the total data volume is M , then the expected number of outliers in a clean mean subtracted spectrum, N_{out} , beyond a certain threshold, τ , is

$$N_{out} = M \operatorname{erfc} \left(\frac{\tau}{\sqrt{2}} \right), \quad (13)$$

where erfc is the complementary error function:

$$\operatorname{erfc}(x) = \frac{2}{\sqrt{\pi}} \int_x^\infty dt e^{-t^2}. \quad (14)$$

This includes both positive and negative outliers. For a typical two-minute MWA EoR highband observation, the number of expected outliers when $\tau = 5$ is about 0.05. In other words, about one in twenty observations ought to have just a single sample of that strength or greater. Statistically speaking, samples of that strength are exceedingly unlikely to be thermal and so we can be confident that those samples are contaminated.

It is tempting to flag all data beyond the significance threshold initially, but this can be problematic when there is RFI contamination. Recall that the estimated mean and standard deviation are determined by a time-average in each frequency channel, as described in Equation 11. If RFI that is brighter than the thermal noise contaminates some times, then the mean estimate for the contaminated channels will be skewed upward relative to an estimate drawn from only the clean data. As a result, some clean data may appear to be outlying in the negative direction beyond the significance threshold. For example, consider the bright red cluster followed by the blue trough in Figure 3(b). As shown in Figure 3(d), some data in this observation lies beyond the threshold in the negative direction when the z-scores are initially calculated. Flagging everything beyond the threshold initially will inevitably flag some clean data in this case. We can decrease this type of overflagging by taking an iterative approach to flagging, wherein the RFI contamination of the mean estimate is progressively removed in each iteration. This also allows us to probe deeper into the observation for fainter RFI in each iteration.

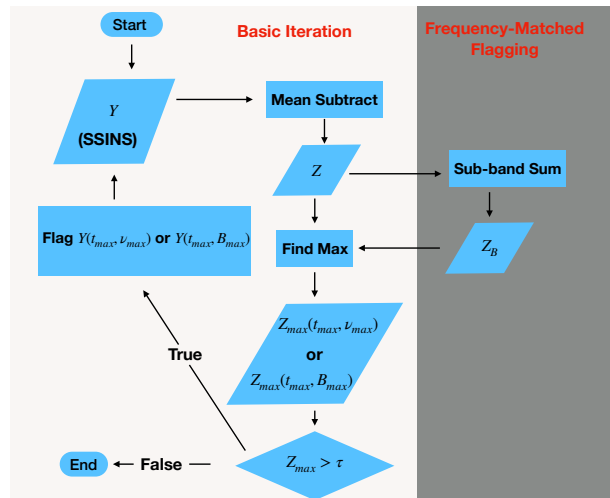


Figure 4. A flowchart of the iterative flagging procedure detailed in §2.4. The white area shows the basic iteration performed on the z-scores calculated in Equation 11, while the grey shows two additional steps that are included when frequency-matched flagging is implemented with the z-scores calculated in Equation 15.

We describe the basic iterative flagging procedure below and present a flow chart in Figure 4. First, we calculate the z-scores of each noise spectrum sample using mean-subtraction. Next, we identify the time and frequency of the strongest outlier in the mean-subtracted spectrum that is beyond the threshold in either the positive or negative direction. Then, we flag this time and frequency in all polarizations regardless of the polarization in which the outlier was found. This is done in case that RFI is present in other polarizations at extremely faint levels; it is often the case that RFI is not exactly polarized along the directions measured by the antennas. We then repeat this process until no more outliers exist beyond the threshold. The results of this process using $\tau = 5$ on the contaminated observation shown in previous figures is shown in Figure 5(a). In this figure, we deviate from the previous convention in Figures 1-3. Now, the same observation is in both panels. We are comparing differences in flagging between the single-sample outlier iteration detailed above and a frequency-matched flagger, detailed in the main text below. We can see that there appears to be incomplete flagging in the feature noticed in Figure 3(b), and that another feature immediately adjacent to it of similar width persists in the observation at fainter levels. Any individual remaining sample in these features lies beneath the significance threshold, and so cannot be caught by the method above at the threshold set. To boost our sensitivity to faint features such as these, we combine samples of the mean-subtracted spectrum across frequencies.




Figure 5. A comparison of single-sample thresholding and frequency-matched flagging on the contaminated observation of the previous figures. Black data is flagged. In (a), we show the results of iteratively flagging single-sample outliers beyond the $\tau = 5$ threshold. Leftover RFI persists, since some RFI is beneath this threshold on a single-sample basis. In (b), we show the results of frequency-matched flagging, where we have used bandwidth information about Western Australian DTV to inform the flagger. No remaining interference can be seen in the spectrum. A movie of each flagging iteration is available in the ancillary files.

So far, we have only relied on knowledge of the thermal background without including any specific information about the particular source of the RFI. In the examples shown, data are taken in Western Australia over the frequencies 167-197 MHz. A subset of the Western Australian digital television channels are broadcast in this observing band and are broadcast 7 MHz wide, adjacent to one another⁶. The broad features shown in the previous figures correspond to DTV channels 6 and 7. The MWA also sometimes observes DTV channel 8. We can hunt for these particular contaminants by summing the mean-subtracted spectrum over the frequencies belonging to a particular type of contaminant. A similar targeted sub-band summing method was employed in §4 of [Offringa et al. \(2015\)](#) exactly for DTV interference.

Formally, if an RFI signal occupies a sub-band, B , that spans N_B frequency channels of the instrument between the signal’s lower frequency, ν_L , and upper frequency, ν_U , we take

$$Z_B(t_n, p) = \frac{1}{\sqrt{N_B}} \sum_{\nu=\nu_L}^{\nu_U} Z(t_n, \nu, p). \quad (15)$$

This sum is precisely constructed so that a clean Z_B will be standard normal distributed just like a clean Z , so that the same significance threshold can be applied to any sub-band as is applied to a single sample. When RFI in a sub-band is observed for only part of

the observation, even a sequence of very weak positive (or negative) outliers at a given time can sum to be greater in absolute value than the desired significance threshold, thereby increasing the sensitivity to RFI that occupies that sub-band. We can then iterate similarly as in the single-sample case to flag the observation for sub-band outliers in addition to single-sample outliers, where now we identify the strongest outlier among all suspected RFI sub-bands and times in the observation and then proceed as before. We show an example of the frequency-matched flagger in Figure 5(b), where it successfully excised DTV RFI belonging to multiple broadcasting channels. A movie showing the status of the observation in each iteration is available in the ancillary files.

The SSINS frequency-matched flagger can be adapted to search for any RFI contaminant within the observing band, including RFI that occupies the entire band. In §3, we illustrate these adaptive capabilities using specific examples of common RFI signals that we observe with the MWA.

3. SOME COMMON RFI OCCUPANTS IN MWA DATA

The sensitivity boost afforded by the baseline averaging allows SSINS to identify a wealth of faint RFI that goes undetected by single baseline algorithms. However, RFI that persists throughout the entire observing time can still evade SSINS due to the fact that such RFI will always contaminate the mean estimation no matter how deeply it is flagged. We have successfully implemented

6

https://www.acma.gov.au/~media/Licence-Issue-and-Allocation/Publication/pdf/TVRadio-Handbook_Electronic_edition-pdf?la=en

methods in SSINS to compensate for this behavior. Below, we explore the behavior of SSINS through examples for three different RFI classifications: broadband streaks, narrowband RFI, and DTV. In each case, the data was passed through AOFLAGGER before making the incoherent noise spectra, except where indicated in §3.2.

3.1. Faint Broadband Streaks

First, we show an observation with faint streaks that are not band limited. See Figure 6. These streaks are too faint to be detected by the standard single-baseline AOFLAGGER implementation and appear to be quite common in our brief survey of MWA EoR High-band data. The exact nature of these streaks is so far unknown. Since they are not band-limited, we cannot directly appeal to our knowledge of the RFI environment in the same way that we can with the DTV interference shown earlier. It is possible that they are instrumental in origin. Despite not knowing the physical origin of such events, these streaks still occupy a programmable sub-band for the frequency-matched flagger, that is, the whole observing band, and so we can still adequately flag them using the algorithm in this paper.

3.2. Narrowband Interference

Next we show an example of narrowband interference observed by the MWA that only occupies one or two fine frequency channels (Figure 7). Though this RFI is present in the observation at dramatically fainter levels after AOFLAGGER, we show the pre-AOFLAGGER spectrum in this case because of the interference’s proximity to a coarse band edge, which makes viewing the event difficult after routine flagging of the edges.

In this case, there are two interference events that are present through the entirety of the observation. The higher frequency line is a full order of magnitude brighter than the lower frequency line, and they are separated by exactly one coarse channel. We hypothesize that the dimmer line is caused by the brighter line’s proximity to a coarse band edge, which are prone to aliasing due to the cascaded fourier transform in the digital signal path⁷.

Narrowband RFI such as this that persists throughout the entire observation poses a detection problem for the SSINS algorithm, since the estimation of the thermal parameters will always be contaminated by the RFI. The statistical properties of the RFI usually do not perfectly resemble the thermal noise, but even so,

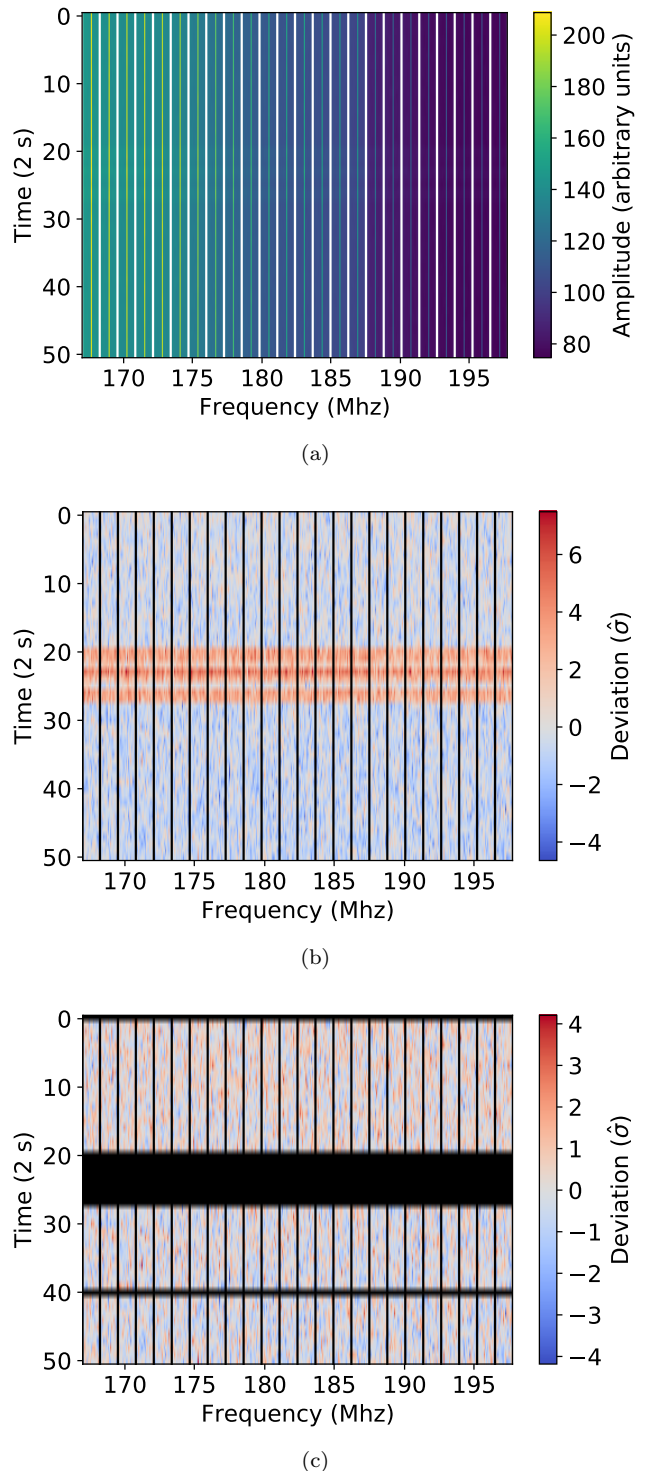


Figure 6. The SSINS of an observation which show faint broadband streaks. The pre-existing flags in (a) and (b) are over the coarse band edges, which are routinely flagged in pre-processing due to systematic difficulties. The streaks are barely noticeable to the eye in (a), but the mean-subtracted spectrum in (b) shows them prominently. The frequency-matched flagger excises these streaks neatly, as shown in (c). These exceedingly faint events are often missed by AOFLAGGER and are not uncommon. They can appear isolated within an observation, or in series like this.

⁷ See "Aliases" in https://wiki.mwatelescope.org/display/MP/Memos?preview=/14156367/18481185/MEMO_CascadedFT_2012_05_25.pdf

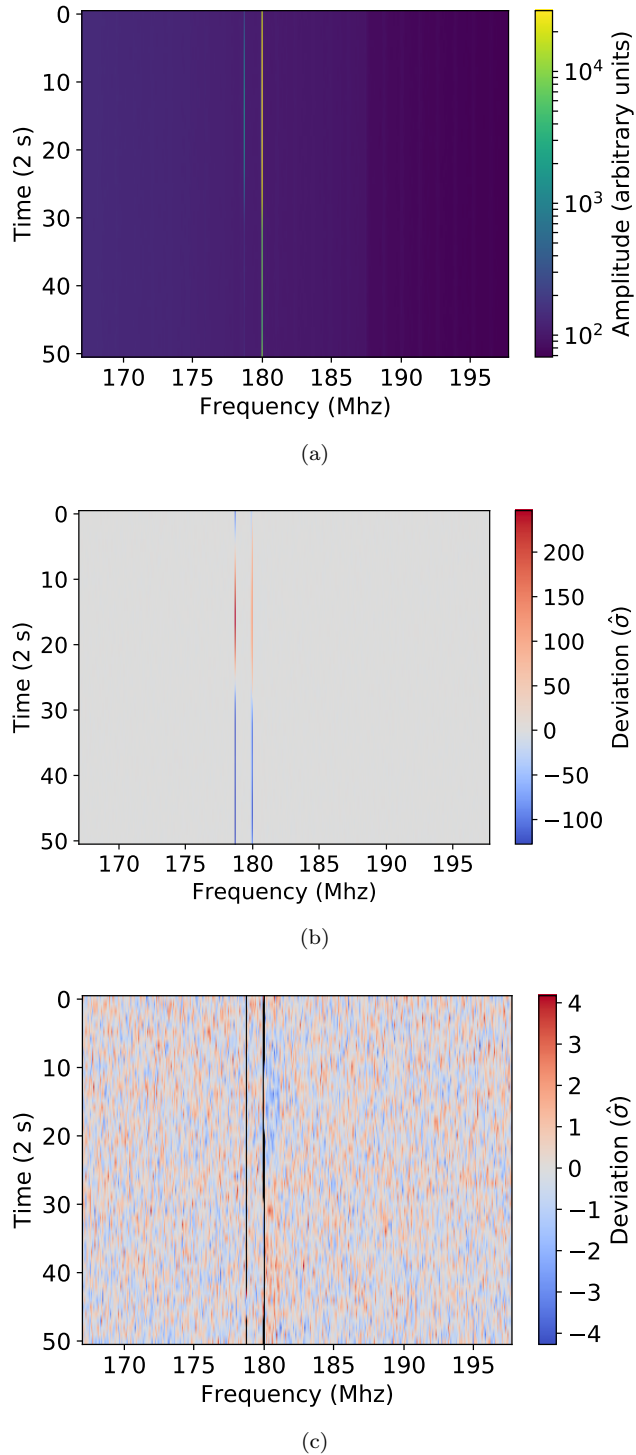


Figure 7. Pictured is an example of narrowband RFI that lasts the entirety of an observation. Note the logarithmic color scale in (a). This extremely bright event is a full two orders of magnitude brighter than the typical clean spectrum, and so the finer structure typically seen in an MWA incoherent noise spectrum is washed out in this colormap. We also note a disturbance precisely one coarse channel wide to the right of the main interference event in (c), characterized by a blue trough followed by a red excess. This is the only known instance of this feature in MWA data so far. Its exact nature is unknown. It may be related to the sheer brightness of this RFI event.

such RFI is almost always incompletely flagged by the SSINS frequency-matched flagger. Rarely, there is even RFI whose brightness varies so little over the course of the observation that it is totally camouflaged by the mean-subtraction step. Clearly, narrowband RFI that only occupies some fraction of the observation does not have this problem. In Figure 7(c), we show the mean-subtracted spectrum after frequency-matched flagging. We can see that the fainter event (alias) was flagged almost completely. If we examine the brighter event at the higher coarse channel, we notice the flag mask is thicker at some times compared to others. The RFI seems to occupy two adjacent fine frequency channels. The times indexed 10 to 20 are missing flags in the higher fine frequency channel, while the times indexed 30 to 48 are missing flags in the lower fine frequency channel.

We handle incomplete flagging in the following way. First, note that as we flag more times in a channel, fewer samples enter into the estimation of the mean. Eventually, the uncertainty in the estimator will be so high that the estimate will be untrustworthy. Furthermore, as more and more data within a channel are flagged, the chance that clean data remain in that channel diminishes. Combining these ideas, we set a threshold where once the amount of remaining unflagged data falls below the threshold, the entire remainder of the channel is flagged. The exact value that this threshold should take depends on the RFI environment of the telescope. While many narrowband events in the MWA survey shown in §5.2 would demand a threshold as aggressive as 0.7 (flag if less than 70% remains), an extremely brief survey of HERA data shows that a threshold of 0.25 successfully flags many of HERA’s persistent narrowband occupants.

3.3. A DTV Signal

Finally, we show an observation with DTV interference that was partially caught by AOFLAGGER. In this case, the DTV was caught on some baselines, but not all of them, so we can see a leftover DTV footprint in the noise spectrum. See Figure 8.

We perform a similar demonstration as with the fainter DTV events shown in §2.4 by first flagging only bright single-sample outliers and then checking to see if fainter occupants lie beneath the single-sample threshold (Figure 8(c)). A notable feature that emerges from single-sample flagging is that the cluster of significant outliers around the DTV event seem to span more than 7 MHz. The excess width of this event is not seen in the pre-AOFLAGGER spectrum (not shown), indicating that this is a pre-processing artifact. As described in Offringa et al. (2012b), AOFLAGGER deploys a morphological detection algorithm that can overflag broad

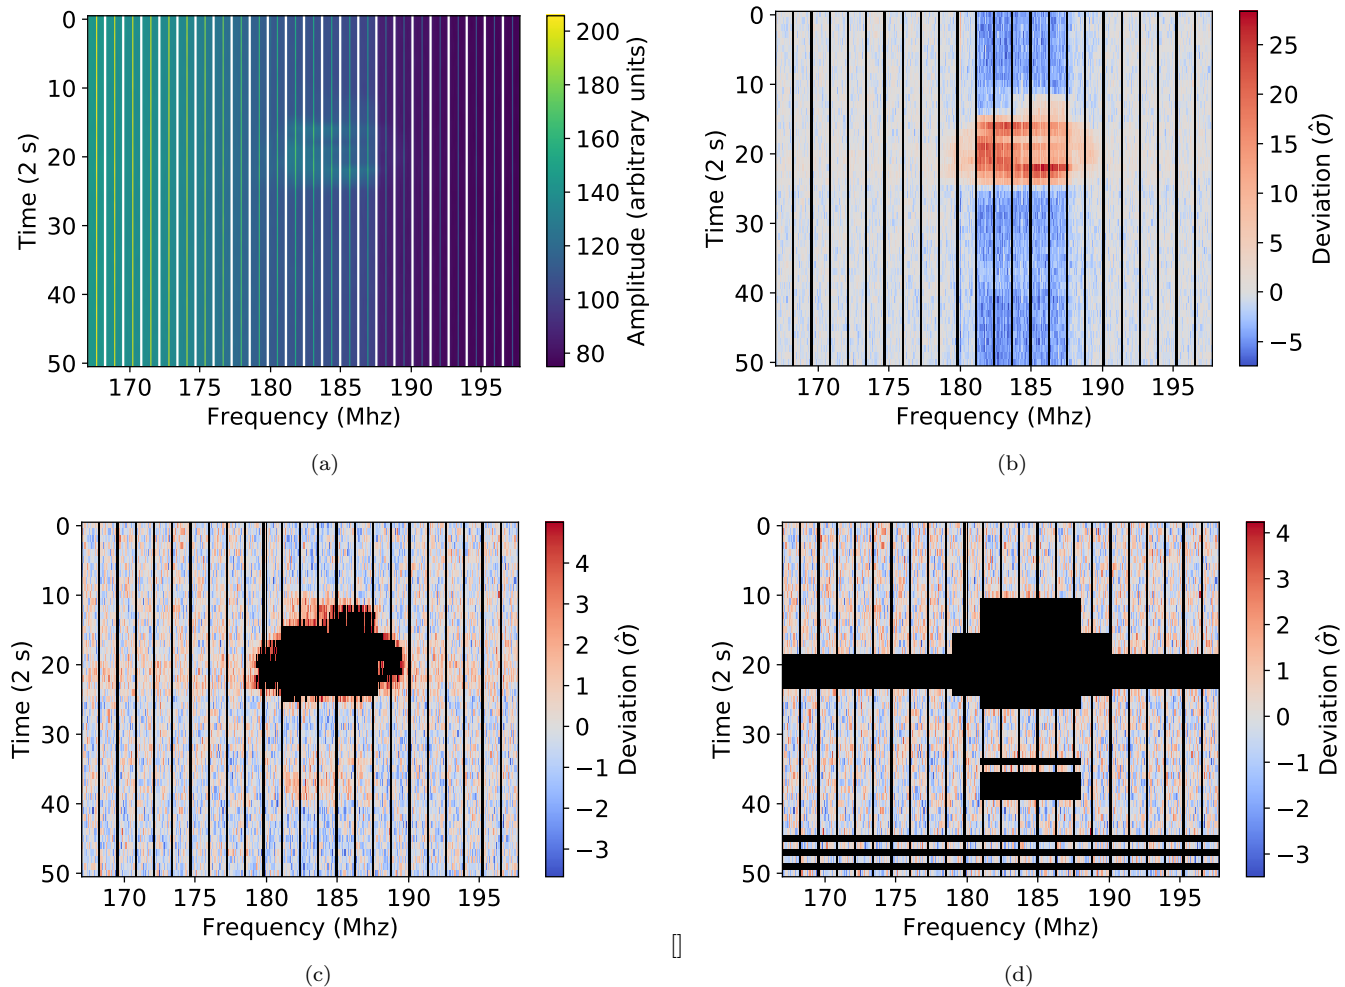


Figure 8. The SSINS of a DTV example after applying AOFLAGGER (a), alongside the mean-subtracted spectrum (b), as well as the results of single-sample iterative flagging (c), and frequency-matched flagging (d). From the flags reported by AOFLAGGER, we know it was caught on some baselines, however there is clearly leftover DTV corresponding to DTV channel 7. Several features emerge after single-sample flagging in (c). First, it appears that there is a broadband streak (§3.1) simultaneous with the DTV interference. Second, it appears there may have been a second DTV interference event later in the observation and much fainter. Third, the DTV interference seems to have associated outliers that span more than the advertised 7 MHz, which is a pre-processing artifact described in the main text. With the frequency-matched flagger we can search for DTV, broadband streaks, single-sample outliers, and the pre-processing artifact simultaneously. The results are shown in (d), where all notable features are excised.

contaminants in a baseline by an amount proportional to the algorithm’s user-set aggression threshold. Indeed, summing the AOFLAGGER flags for this observation over the set of baselines does show overflagging of this event in frequency by an amount similar to the feature we see in the incoherent noise spectrum. A feature of this size manifests in the incoherent noise spectrum due to the fact that visibilities have been averaged in time and frequency relative to the operating time-frequency resolution of AOFLAGGER. Flags are applied before averaging, so time-frequency bins with fewer samples entering them will have noise that has not been averaged down as much compared to those bins in which all

possible contributing samples were averaged together. The overflagged bins contributing to the incoherent noise spectrum are then brighter than the surrounding uncontaminated ones, thus appearing like RFI to our statistical test. We can adapt the frequency-matched flagger with a custom sub-band to identify this pre-processing artifact. The results of frequency-matched flagging are shown in Figure 8(d).

This type of DTV interference is extremely common in MWA EoR Highband Observations. Roughly one third of the observations included in the EoR limit in [Beardsley et al. \(2016\)](#) had some trace of DTV RFI according to the frequency-matched flagger. These observations




Figure 9. A full-sky horizon-to-horizon image of the DTV channel 6 event in Figure 3(b), with 52000 GLEAM sources removed using FHD. Other than the shape of the MWA beam and some diffuse structure not included in the calibration/subtraction model, the image is largely featureless. However, all the way in the second southern sidelobe lies a faint streak belonging to the DTV6 event (annotated with a bright red arrow). We hypothesize that the DTV signal is reflecting off of an aircraft on the southern horizon into the array. This is but one example of a collection of observations that show a similar feature in the second southern sidelobe, some of which are at the same time of night but on different dates, as one might expect from scheduled flights to and from Perth.

were subsequently removed in a reduction of the same data set in Barry et al. (2019b), as discussed in §5.2. While in this case AOFLAGGER was able to remove some of the interference, it is often the case that the contamination is beneath the level where AOFLAGGER can make an appreciable difference.

The MWA’s extremely remote location makes it unlikely that it would directly observe a DTV transmission. However, it is clear from noise spectra such as those in the previous sections that the MWA is observing DTV. There are several hypotheses to explain this, including tropospheric ducting and reflection off of aircraft or satellites. We explore these possibilities by imaging DTV events found using SSINS.

4. IMAGING DTV INTERFERENCE




Figure 10. The Stokes I residual image of the DTV event in Figure 8, showing just the primary beam (center lobe in Figure 9). This is 34 seconds of data. The broad North-South streak in the Eastern edge of the primary beam (annotated with a bright red arrow) suggests possible motion of a source - likely a reflective aircraft. A snapshot-by-snapshot movie is available in the ancillary files.

The images were made using the Fast Holographic Deconvolution (FHD) software⁸ (Sullivan et al. 2012; Barry et al. 2019a). So that RFI could not drive the calibration solution, we calibrated on a clean part of each observation and transferred that calibration to the contaminated part of the respective observation. In order to see these faint DTV events in the images, it was necessary to subtract out 52000 sources from the GaLactic and Extragalactic All-sky MWA (GLEAM) catalog (Hurley-Walker et al. 2017) using FHD. Example images using the observations from Figures 3(b) and 8 are shown in Figures 9 and 10, respectively. We separated Figure 10 into 2s-snapshots and made a movie (available in the ancillary files), where it is clear that some sort of flying object moving nearly due North-South is reflecting DTV into the array.

Due to the nonzero extent of the moving source, we can rule out the possibility of satellites and near-Earth asteroids, which would appear point-like. Aircraft are reflective and fly low enough to appear extended in the image. The speed (extent) of the object is atypically

slow (large) for a commercial jet such as a Boeing 737 flying at a typical cruising altitude of 10 km. A parallaxic estimation of the object’s altitude⁹ ultimately proved inconclusive due to the faintness of the RFI (Xiang Zhang, personal communication). As an alternative candidate, we suggest this could be a slow, low-flying bush plane whose apparent angular size may be due to its low height and near-field effects.

5. FREQUENCY-MATCHED FLAGGER CUSTOMIZATION

In this section, we describe the process of developing a customized SSINS frequency-matched flagger for a radio telescope. This is an exploratory process wherein the user becomes familiar with the incoherent noise spectra of their telescope. Some of the spectra features will be due to RFI, while other features may point to subtleties of the instrument. Once equipped with a thorough catalog of sub-bands to search, the user may process large amounts of data in an automated way. To exemplify

⁸ <https://github.com/EoRImaging/FHD>

⁹ See Loi et al. (2016) and citations therein for other uses of this type of altitude estimation

what can be attained from this process, we summarize the RFI occupancy analysis for the data used in the EoR limit presented in Barry et al. (2019b).

5.1. *Exploring the Data and Developing a Frequency-Matched Flagger*

The first step to applying SSINS to a dataset is acquiring the software. It is implemented in PYTHON along with a comprehensive set of unit tests. It is publicly available on GitHub¹⁰, where the user will also find installation instructions including the list of dependencies: PYUVDATA¹¹, NUMPY, SCIPY, SIX, H5PY, PYYAML, ASTROPY, and optionally MATPLOTLIB. There are tutorials and other documentation available¹² with simple usage examples for getting started.

After becoming familiar with the software, the next step is to examine the data, typically by hand, without attempting to use the frequency-matched flagger. While going through a survey of observations baseline by baseline in a modern radio telescope would be untenable, SSINS compresses the data of each observation into two dynamic spectra per polarization: a raw incoherent noise spectrum and its mean-subtracted form. It is useful to look at both forms of the spectra side-by-side, particularly if narrow RFI is expected. In this way, an entire season of observations can be examined in short order.

While features in the incoherent noise spectra may correspond well with known factors in the RFI environment, it is also likely that SSINS will reveal RFI and even other instrumental features that had not been previously observed or considered. As the user catalogs the various characteristics of the data, they can build a dictionary of different sub-bands or single frequencies that appear commonly occupied. These occupations, regardless of their physical cause, can be input into the frequency-matched flagger settings. For example, some faint features in HERA incoherent noise spectra were later found to correspond with correlator malfunctions, some of which had been previously identified in other cases through other means. While technically this was not RFI per se, it was a feature that could be identified and removed with SSINS.

Once the set of identifiable contaminants has been collected, the frequency-matched flagger can be deployed for a first round of flagging. The frequency-matched flagger is very quick. A 2-minute, 30 Mhz MWA spec-

trum can be flagged in less than a second, while a 10-minute, 100 Mhz HERA spectrum that is significantly more contaminated typically takes between 20 and 60 seconds. As with the initial inspection, the results of the frequency-matched flagger can be quickly ascertained by eye. Oftentimes, fainter RFI than could be initially seen in the noise spectra is unearthed as the frequency-matched flagger iterates. Of course, the user will wish to add these to the dictionary of occupants. If no more identifiable sub-band occupancies are revealed, one may wish to adjust the parameters of the flagger, such as the significance threshold. This iterative process of data examination and deep cleaning is often illuminating.

5.2. *Results with a Season of MWA Data*

The SSINS package itself was developed and tested in the manner described above using a season’s worth of data from 2013. This same dataset was used for the EoR power spectrum limit featured in Beardsley et al. (2016) as well as Barry et al. (in review). Below, we describe the filter settings for an RFI analysis of this dataset and also summarize occupancy levels.

The data in this analysis was pre-processed with COTTER, which uses AOFLAGGER to identify and flag bright RFI. Then, SSINS was run on the data in order to identify and catalog leftover faint RFI. We used a significance threshold of 5 and sought single sample outliers, broadband streaks, DTV interference, as well as the broader DTV sub-band discussed in §3.3, all of which are occupants detailed in §3. We also completely flagged a fine frequency channel if it ever reached an occupancy fraction of 0.7 during frequency-matched flagging. We present RFI occupancy for the season in Figure 11, summarizing total RFI occupancy as seen by SSINS after AOFLAGGER as well as DTV occupancy. Such figures can be used to pick out particularly bad days from the season and search for patterns in time.

The occupancy data from SSINS was used to make data cuts for the sake of improving the EoR limit in Barry et al. (2019b). We did not reflag the data using SSINS, but instead just used its outputs to develop quality metrics for observations. Ultimately, we cut all observations with any trace of DTV as well as all observations with greater than 40% occupancy. We found that this improved the limit despite removing roughly 1/3 of the observations originally included.

6. DISCUSSIONS AND CONCLUSIONS

We have described the SSINS RFI detection algorithm in detail. The substantial sensitivity boost afforded by the incoherent average over the baselines allows for detection of faint RFI that escapes other high-performing

¹⁰ <https://github.com/mwilensky768/SSINS>

¹¹ <https://github.com/RadioAstronomySoftwareGroup/pyuvdata>

¹² <https://ssins.readthedocs.io/en/latest>




Figure 11. (a) A scatter plot of the total RFI occupation in a season of data as seen by SSINS. (b) A scatter plot of the DTV occupancy as seen by SSINS. In each panel, each circle represents a single two-minute observation. Each line of circles is a night of observations on the Julian Date (JD) shown on the vertical axis, while the horizontal axis gives that observation’s local sidereal time (LST) in hours. Black circles were not found to have additional RFI after AOFLAGGER. All observations outside of the plotted LST range and all missing circles in a night are observations that were removed by a previous jackknife test, detailed in [Beardsley et al. \(2016\)](#). In (a), the color shows the fraction of samples in the noise spectra found to be contaminated, disregarding coarse band edges, while in (b), the color shows the fraction of times contaminated by DTV interference, regardless of which broadcasting channels were identified. These plots provide an occupancy overview of the season, letting one pick out particularly bad days or times of night by eye, such as the line of high occupancy observations nearest JD 2456600, which was a day that also had many observations removed by the previous jackknife test.

single-baseline algorithms. We demonstrated its effectiveness on several different types of RFI found in the MWA EoR highband, including DTV interference. We have implemented countermeasures for persistent RFI with some success, but we plan to make improvements to these detection efforts. Overall, we observe that the increased sensitivity afforded by SSINS helps us understand the general pervasiveness of faint RFI.

Developing a custom SSINS frequency-matched flagger for a new RFI environment is an iterative process that is often quite instructive. We confirmed the quality of our own EoR highband frequency-matched flagger by successfully imaging faint DTV reflected from aircraft as well as successfully improving the EoR limit in [Barry et al. \(2019b\)](#). To ease the development process, we have written tutorials for basic software usage and provided an overview of how SSINS can be used in practice.

Though SSINS is capable of detecting fainter RFI than current state-of-the-art algorithms and is a general improvement to the field, we plan to further improve the methodology to exploit its full capabilities. We propose the following items as future work:

- **Improved Narrowband RFI Detection:** The current implementation tends to underflag narrowband RFI that persists through the entire observation. An algorithm that preliminarily identifies entirely contaminated channels and interpolates over

these channels during mean estimation will allow for more complete flagging.

- **Blind Sub-band Hunting:** Presently, the user must input specific sub-bands into the frequency-matched flagger that are established from knowledge of the RFI environment or a hand grading of the data. An algorithm that searches for likely occupied sub-band candidates would ease analysis of extremely large data sets.
- **EoR RFI Power Spectrum Shape Characterization:** The effect of faint RFI on an EoR power spectrum has not yet been precisely characterized. We aim to theoretically predict particular RFI footprints in the power spectrum and then subsequently confirm the presence of these footprints in power spectra made from MWA data.
- **Data Retention for EoR Limit:** In [Barry et al. \(2019b\)](#), we used SSINS to entirely remove poor-quality observations from the limit calculation, rather than propagating flags from the frequency-matched flagger to the visibilities and keeping any remaining unflagged data within these observations. We expect the EoR limit analysis could be more sensitive if we include remaining clean data from previously cut observations.

As SSINS is applied to more radio projects in the future, more possibilities for enhancements will be found. In this light, we look forward to exploring the full capabilities of the SSINS framework.

7. ACKNOWLEDGEMENTS

MW would like to acknowledge Xiang Zhang for helpful discussions. We would like to thank the development teams of PYUVDATA, NUMPY, SCIPY, SIX, H5PY, PYYAML, ASTROPY, and MATPLOTLIB, which enabled this work. This work was directly supported by NSF grants AST-1643011, AST-1613855, and OAC-1835421. Computation on the Amazon Web Services public cloud was supported by the University of Washington student-led Research Computing Club with funding provided by

the University of Washington Student Technology Fee Committee. NB is supported by the Australian Research Council Centre of Excellence for All Sky Astrophysics in 3 Dimensions (ASTRO 3D), through project number CE170100013. This scientific work makes use of the Murchison Radio- astronomy Observatory, operated by CSIRO. We acknowledge the Wajarri Yamatji people as the traditional owners of the Observatory site. Support for the operation of the MWA is provided by the Australian Government (NCRIS), under a contract to Curtin University administered by Astronomy Australia Limited. We acknowledge the Pawsey Supercomputing Centre which is supported by the Western Australian and Australian Governments.

REFERENCES

- An, T., Chen, X., Mohan, P., & Lao, B. Q. 2017, *Acta Astronomica Sinica*, 58, 43.
<https://arxiv.org/abs/1711.01978>
- Athreya, R. 2009, *The Astrophysical Journal*, 696, 885, doi: 10.1088/0004-637x/696/1/885
- Baan, W. A. 2019, *Journal of Astronomical Instrumentation*, 8, 1940010, doi: 10.1142/S2251171719400105
- Barry, N. 2018, PhD thesis, University of Washington
- Barry, N., Beardsley, A. P., Byrne, R., et al. 2019a, *Publications of the Astronomical Society of Australia*, 36, e026, doi: 10.1017/pasa.2019.21
- Barry, N., Wilensky, M., Trott, C. M., et al. 2019b, *The Astrophysical Journal*, 884, 1, doi: 10.3847/1538-4357/ab40a8
- Beardsley, A. P., Hazelton, B. J., Sullivan, I. S., et al. 2016, *ApJ*, 833, 102, doi: 10.3847/1538-4357/833/1/102
- Billingsley, P. 1995, *Probability and Measure*, 3rd edn. (New York: John Wiley & Sons), 357–362
- Bretteil, S., & Weber, R. 2005, *Radio Science*, 40, RS5S15, doi: 10.1029/2004RS003124
- Burd, P. R., Mannheim, K., Mrz, T., et al. 2018, *Astronomische Nachrichten*, 339, 358, doi: 10.1002/asna.201813505
- DeBoer, D. R., Parsons, A. R., Aguirre, J. E., et al. 2017, *PASP*, 129, 045001, doi: 10.1088/1538-3873/129/974/045001
- Furlanetto, S. R., Oh, S. P., & Briggs, F. H. 2006, *Physics Reports*, 433, 181 , doi: <https://doi.org/10.1016/j.physrep.2006.08.002>
- Hellbourg, G., Weber, R., Capdessus, C., & Boonstra, A.-J. 2012, *Comptes Rendus Physique*, 13, 71 , doi: <https://doi.org/10.1016/j.crhy.2011.10.010>
- Hurley-Walker, N., Callingham, J. R., Hancock, P. J., et al. 2017, *MNRAS*, 464, 1146, doi: 10.1093/mnras/stw2337
- Kerrigan, J., Plante, P. L., Kohn, S., et al. 2019, *Monthly Notices of the Royal Astronomical Society*, doi: 10.1093/mnras/stz1865
- Loi, S. T., Cairns, I. H., Murphy, T., et al. 2016, *Journal of Geophysical Research (Space Physics)*, 121, 1569, doi: 10.1002/2015JA022052
- McKinley, B., Briggs, F., Kaplan, D. L., et al. 2013, *AJ*, 145, 23, doi: 10.1088/0004-6256/145/1/23
- McKinley, B., Bernardi, G., Trott, C. M., et al. 2018, *MNRAS*, 481, 5034, doi: 10.1093/mnras/sty2437
- Morales, M. F., & Wyithe, J. S. B. 2010, *ARA&A*, 48, 127, doi: 10.1146/annurev-astro-081309-130936
- Offringa, A. R., de Bruyn, A. G., Biehl, M., et al. 2010, *MNRAS*, 405, 155, doi: 10.1111/j.1365-2966.2010.16471.x
- Offringa, A. R., de Bruyn, A. G., & Zaroubi, S. 2012a, *Monthly Notices of the Royal Astronomical Society*, 422, 563, doi: 10.1111/j.1365-2966.2012.20633.x
- Offringa, A. R., van de Gronde, J. J., & Roerdink, J. B. T. M. 2012b, *A&A*, 539, A95, doi: 10.1051/0004-6361/201118497
- Offringa, A. R., de Bruyn, A. G., Zaroubi, S., et al. 2013, *MNRAS*, 435, 584, doi: 10.1093/mnras/stt1337
- Offringa, A. R., Wayth, R. B., Hurley-Walker, N., et al. 2015, *PASA*, 32, e008, doi: 10.1017/pasa.2015.7
- Peck, L. W., & Fenech, D. M. 2013, *Astronomy and Computing*, 2, 54 , doi: <https://doi.org/10.1016/j.ascom.2013.09.001>
- Sekhar, S., & Athreya, R. 2018, *AJ*, 156, 9, doi: 10.3847/1538-3881/aac16e

- Sokolowski, M., Wayth, R. B., & Lewis, M. 2016, arXiv e-prints, arXiv:1610.04696.
<https://arxiv.org/abs/1610.04696>
- Sullivan, I. S., Morales, M. F., Hazelton, B. J., et al. 2012, ApJ, 759, 17, doi: [10.1088/0004-637X/759/1/17](https://doi.org/10.1088/0004-637X/759/1/17)
- Taylor, G. B., Ellingson, S. W., Kassim, N. E., et al. 2012, Journal of Astronomical Instrumentation, 1, 1250004, doi: [10.1142/S2251171712500043](https://doi.org/10.1142/S2251171712500043)
- Tingay, S. J., Goeke, R., Bowman, J. D., et al. 2013, PASA, 30, e007, doi: [10.1017/pasa.2012.007](https://doi.org/10.1017/pasa.2012.007)
- van Haarlem, M. P., Wise, M. W., Gunst, A. W., et al. 2013, A&A, 556, A2, doi: [10.1051/0004-6361/201220873](https://doi.org/10.1051/0004-6361/201220873)
- Wayth, R. B., Tingay, S. J., Trott, C. M., et al. 2018, PASA, 35, 33, doi: [10.1017/pasa.2018.37](https://doi.org/10.1017/pasa.2018.37)
- Zhang, X., Hancock, P., Devillepoix, H. A. R., et al. 2018, MNRAS, 477, 5167, doi: [10.1093/mnras/sty930](https://doi.org/10.1093/mnras/sty930)
- Zwillinger, D., & Kokoska, S. 2000, Standard Probability and Statistics Tables and Formulae (Boca Raton: CRC Press LLC), 55