The Detectability of FRBs with HERA

Vighnesh Nagpal[1] and Joshua S. Dillon[1]

[1]*Department of Astronomy, University of California, Berkeley*

April 21, 2022

1 Introduction

We are currently in a period of rapid advancement in Fast Radio Burst (FRB) science, with telescopes such as CHIME, ASKAP and FAST discovering thousands of new FRBs in the past few years. However, FRB science at the low frequencies HERA probes has only just started to develop, with LOFAR's recent detection FRBs (Pleunis et al., 2021a) from the known repeating source FRB 20180916B marking the very first FRB detection in the frequency range accessible to HERA. This gives rise to an exciting question: could HERA detect FRBs as well?

Detecting FRBs using HERA presents a unique challenge, as it is optimised for cosmological power spectrum studies rather than transient detection. HERA's long integration time relative to the pulse width of most FRBs makes detecting them using HERA more challenging. However, HERA's large size may overcome this suboptimal integration time to potentially enable a detection or at least useful upper limits on FBR rates at low frequency. Additionally, surveys conducted by telescopes in the Southern Sky, such as the Parkes Telescope (Yang et al., 2021) and ASKAP (Shannon et al., 2018) have found numerous FRB sources in the patch of sky accessible to HERA, which could be potential targets for us. In this memo, we simulate HERA observations of FRBs, develop methods to compute their signal-to-noise ratios (SNR) and in doing so forecast what sort of FRBs HERA could potentially detect.

2 FRB Modeling and Parameterization

For simplicity, our simulations assume that the burst follows a Gaussian profile in time, though we note that recent studies (Pleunis et al., 2021b) in the field have found bursts with varying morphologies. The Gaussian model is given by:

$$I(t,\nu) = \frac{I_0}{\sigma\sqrt{2\pi}} \exp\left[-\frac{(t - t_{\rm FRB}(\nu))^2}{2\sigma^2}\right]. \tag{1}$$

Here $t_{FRB}(\nu)$ refers to the time at which the component of the burst at frequency ν arrives at the detector, and where $\sigma = \mathrm{FWHM}/(2\sqrt{2\ln 2})$, where FWHM is the full-width at half-maximum of the Gaussian time profile. Also for simplicity, we assume that I_0 is constant with frequency, which is probably too optimistic, in the sense that no observed LOFAR FRBs span the full $\sim$200 MHz of HERA's band.

The arrival time of the burst is a function of ν, due to dispersion along the line of sight due to passage through an ionised medium. The equation for the frequency dependent arrival time follows from the definition of the dispersion measure:

$$t_{\rm FRB}(\nu) = t_{\rm ref} + k_{DM} DM \left(v^{-2} - v_{\rm ref}^{-2}\right) \tag{2}$$

Here, t_{ref} is the reference time at which the burst peaks at frequency ν_{ref}. For the purposes of this work, we adopt 150 MHz as the reference frequency. The amount of dispersion is controlled by two constants: the dispersion measure (DM) of the burst, and kDM, a measure of the free electron density along the line of sight, the value of which we fix to $4.1488 \times 10^3\,\mathrm{MHz^2\,cm^3\,pc^{-1}\,s}$.

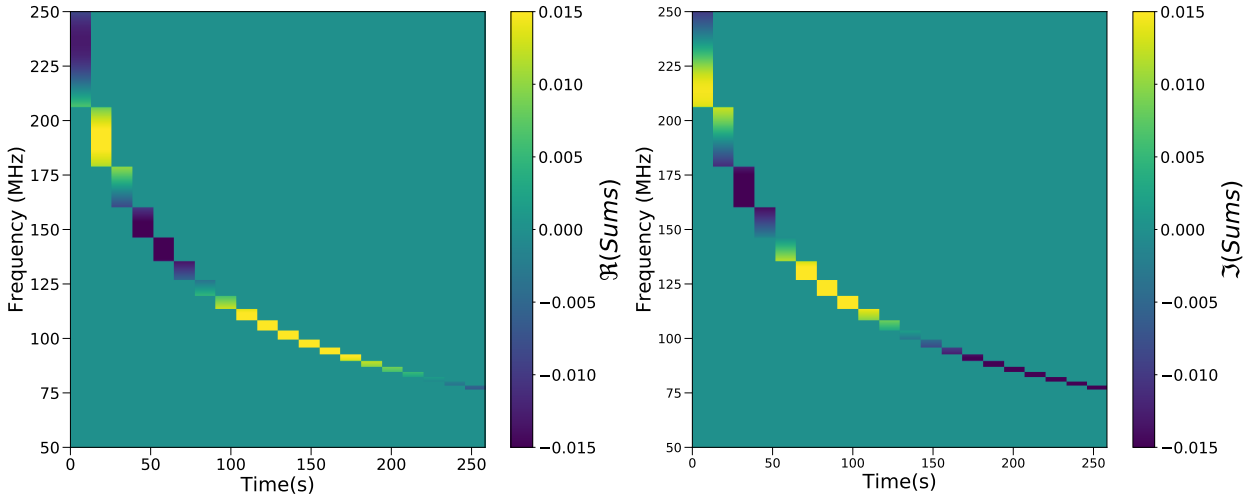

Figure 1: Simulated HERA visibility of an FRB with DM=300, as seen by a 51 m E, 37 m S baseline.The panel on the left shows the real part of the FRB visibility, while the panel on the right shows the imaginary component. The phasing observed in both panels is a consequence of the fact that the FRB moves across the sky over the ∼9 minutes it takes to disperse across the 50-250 MHz frequency range.

When simulating bursts, we start off with a simulation time resolution of 0.01 seconds in order to resolve the burst profile. To compute visibility sums and differences, we average down to 0.1 s even and odd sub-integrations, from which we then compute sums and differences at HERA's 9.6 s cadence.

3 Visibility Simulation and SNR Calculation

The visibility we would expect for a single point source like an FRB as observed by a single baseline is

$$V_{ij}^{\mathrm{FRB}}(t,\nu) = I(t,\nu)B(\hat{\mathbf{r}}(t))\exp\left[\frac{-2\pi i\nu}{c}\mathbf{b}_{ij}\cdot\hat{\mathbf{r}}(t)\right] \tag{3}$$

where $I(t,v)$ is the FRB flux in frequency v at time t, $\hat{\mathbf{r}}(t)$ is the unit vector pointed toward the FRB at time t, $B(\hat{\mathbf{r}}(t))$ is the beam power at the FRB position at time t, and $\mathbf{b}_{ij}$ the baseline vector between antennae i and j. Figure 1 shows the real and imaginary components of a simulated FRB visibility.

We use the Fagnoni Vivaldi beam model.[1] When computing visibilities using this beam, we take into account the movement of the FRB across the sky over the course of the observing window–calculating the altitude and azimuth at each measurement time. We then use this time dependent position of the FRB to calculate the beam power at a given time and frequency, which is incorporated into the visibilities.

To calculate the observability of the FRB, we need a model for noise in a given observation given a time- and frequency-dependent system temperature. To calculate T_{sys} we use a constant receiver temperature of 100 K:

$$T_{sys}(t,\nu) = T_{sky}(t,\nu) + 100\,\mathrm{K}. \tag{4}$$

Our sky temperature is derived from simulated autocorrelations using RIMEz, which includes both diffuse emission and point sources.[2] We convert autocorrelations to T_{sky} via the following expression.

$$T_{sky}(t,\nu) = \left|V_{ii}^{\mathrm{foregrounds}}(\nu,t)\right|\left(\frac{10^{-26}\,\mathrm{W}}{\mathrm{Jy\,m^2\,Hz}}\right)\frac{c^2}{2k_B\nu^2\omega_p(\nu)}. \tag{5}$$

[1] Found at `/lustre/aoc/projects/hera/H4C/beams/NF_HERA_Vivaldi_efield_beam_healpix.fits`
[2] `/lustre/aoc/projects/hera/zmartino/hera_calib_model/H4C_1/abscal_files_unique_baselines`

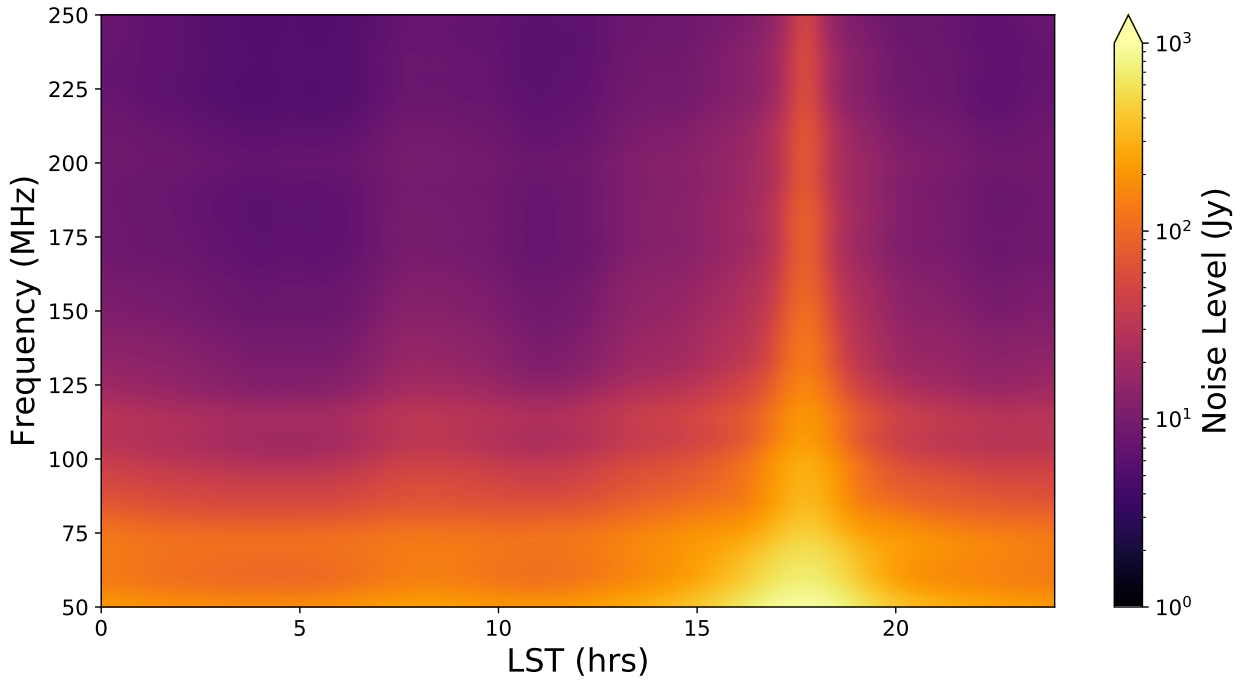

Figure 2: Noise levels estimated from autocorrelations simulated using RIMEz. Because HERA is sky noise dominated for most frequencies and LSTs, the noise peaks when the galaxy is overhead at ∼18 hrs. The process for estimating these noise levels is outlined in Equations 4-6.

Using T_{sys} and assuming that all antennas have the same noise level, we can infer the thermal noise on any given visibility measurement as

$$\sigma_{ij}(\nu) = \left(\frac{T_{sys}(\nu,t)}{\sqrt{t_{int}\Delta\nu}}\right)\left(\frac{\mathrm{Jy\,m^2\,Hz}}{10^{-26}\,\mathrm{W}}\right)\frac{2k_B\nu^2\omega_p(\nu)}{c^2}. \tag{6}$$

We show that thermal noise level in Figure 2 Thus, the total SNR for a given visibility to a single FRB can be calculated by summing over all times and frequencies in quadrature:

$$SNR_{ij} = \sqrt{2\sum_{\nu}\sum_{t}\left|\frac{V_{ij}(t,\nu)}{\sigma_{ij}(\nu,t)}\right|^2} \tag{7}$$

Here the factor of 2 inside the square root comes from the fact that noise is a 2D-Gaussian in the complex plane, but the signal of interest is fundamentally a real quantity that is the same in every visibility, up to a baseline-dependent phase factor. Thus, only part of the noise contributes to the denominator in the SNR.[3] Note that most V_{ij} in this sum are zero or numerically very close to zero; most times and frequencies do not have appreciable FRB flux.

Using the single-visibility SNR defined above, we compute the signal-to-noise ratio for the whole array for simulated FRBs, SNR_{array}, as:

$$SNR_{array} = \sqrt{\sum_{\ell}\left(SNR_{ij}\right)^2 n_\ell} \tag{8}$$

where ℓ is an index that runs over all unique baseline groups and n_ℓ is the number of total baselines in each redundant group.

[3]This assumes that after the FRB is detected, it can be sufficiently localized to coherently phase up all visibilities together before averaging. Hence, while we might consider searching for FRBs in $|V_{ij}|$, that quantity can be biased in the low-SNR regime. Thus, we will likely need to follow up candidates detected in this way by examining the full, complex visibilities.

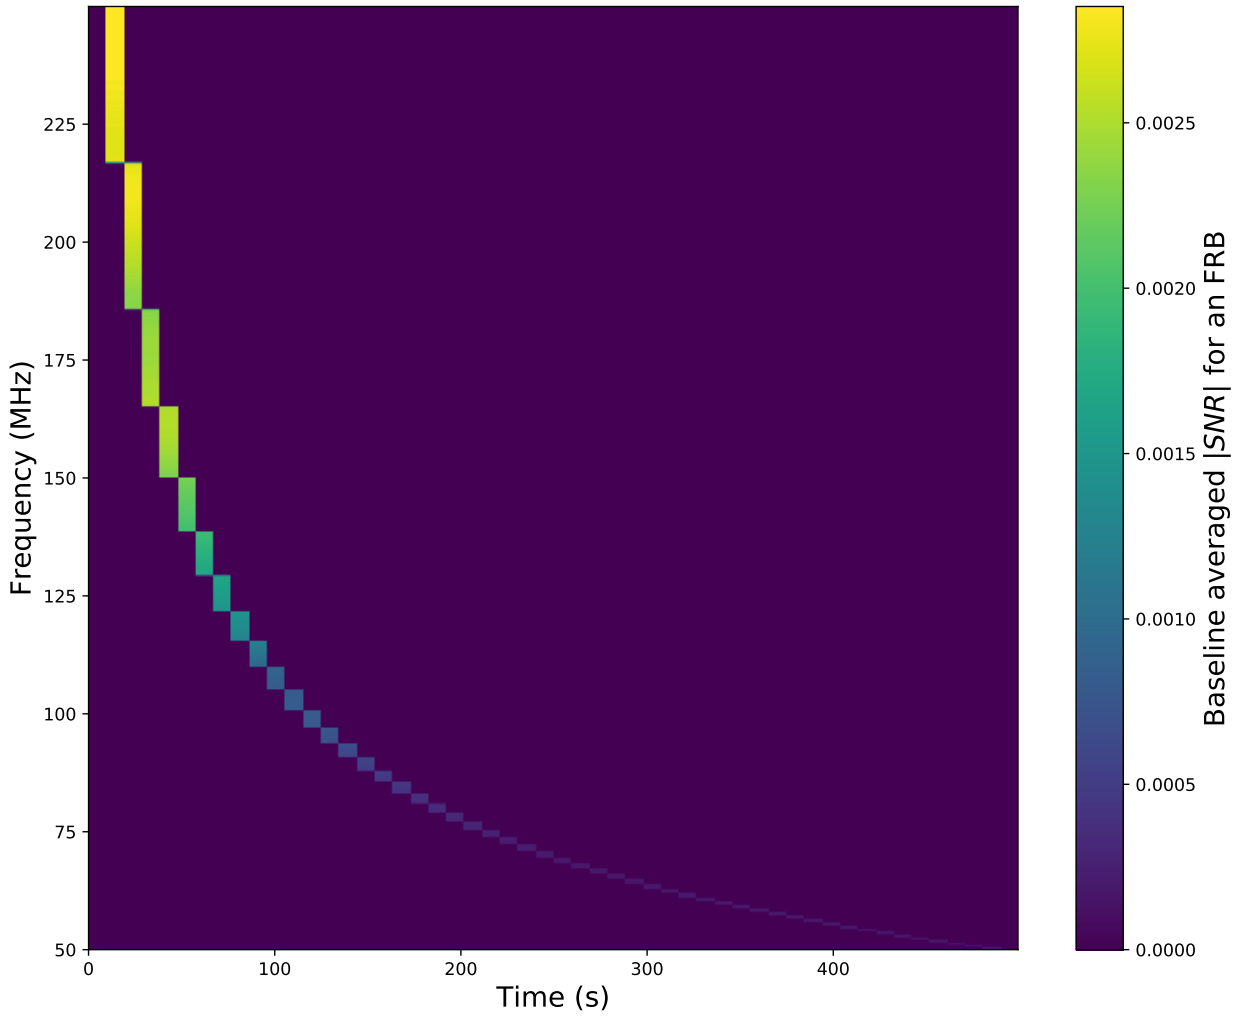

Figure 3: Inverse variance weighted average of the $|SNR|$ for a given FRB with DM=300 and N=11, per channel and per integration, but summed over all redundant baselines. Before carrying out this average, we exploited the redundancy of baselines to drive down the noise by a factor dependent on the number of redundant baselines sharing the same vector. As can be observed, the SNR decreases as we move to lower frequencies due to the corresponding increase in the noise level. However, it is important to note that this effect may not always hold exactly, since the movement of the FRB across the sky can lead to SNR peaking at lower frequencies that arrival time corresponds to near-zenith passage, especially for high-DM FRBs.

4 Statistical simulation pipeline

We can now simulate a series of 1 Jy FRBs distributed evenly on the celestial sphere and in JD. The JD is picked uniformly randomly in time. The random draws on the sphere are done as follows:

$$\alpha = 2\pi u,\ \delta = cos^{-1}(2v - 1) \tag{9}$$

In this, u, v are random variables distributed evenly on $[0, 1]$ and θ, ϕ are spherical angular coordinates. We then compute the horizontal coordinates (with respect to HERA) for each such FRB, for the entire time the FRB disperses from 250 MHz to 50 MHz. Because HERA's beam falls off steeply away from zenith, we adopt an altitude cutoff of 1.2 radians = 68.7° for the maximum altitude the FRB reaches during its

dispersion. For FRBs which do not reach at least this altitude over their time in HERA's frequency range, we assign an SNR of 0 and do not simulate any visbilities for the sake of computational expediency.

For FRBs with maximum altitudes higher than this cutoff, we simulate full HERA array visibilities and use these to estimate the predicted SNRs. For such FRBs, we use the times at which the source passes the altitude cutoff to define our observational window. We then simulate sums and diffs for each redundant baseline in the array using the method outlined in section 3. In Figure 4, we plot the results of simulating a population of FRBs using this process. Owing to the fact that the location and arrival times of bursts in the simulated population are randomly drawn, the vast majority of FRBs on the sky will not be observable by HERA. However, the small fraction that land close to zenith will have high SNR and are thus potentially detectable using HERA.

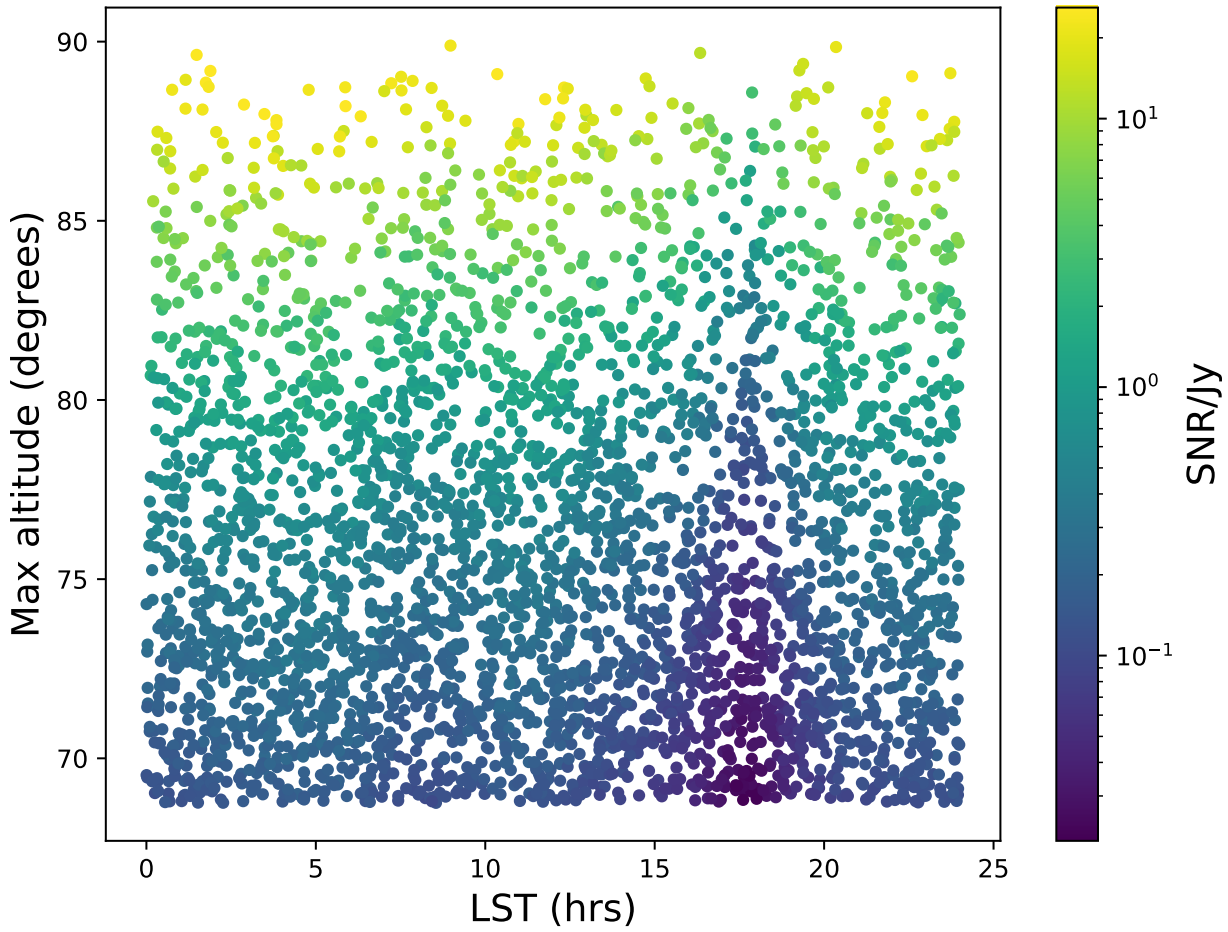

Figure 4: A population of 10^5 1 Jy FRBs with DM=300 and FWHM=0.1s with randomly drawn burst times and locations on the celestial sphere. Each point is a member of the sub-sample of FRBs that pass our maximum altitude cutoff (as these are the only ones for which we calculate SNRs) plotted with respect to the LST at which its 150 MHz component arrives (x-axis), the maximum altitude their source reaches during the burst (y-axis), and its SNR (colour). As can be seen, the number of FRBs that reach a given altitude decreases as we approach zenith, and the SNR conversely increases. In addition, we observe that the population of bursts close to LST=18 hrs have significantly lower SNRs than other times, an effect we attribute to high sky temperature from the galaxy.

We conducted a parameter exploration of the following four parameters in order to assess the impact on burst SNRs for HERA: DM, FWHM[ms], frequency range[MHz], and the size of the array, which is parametrised by the hexagonal N. In Table 1, we list the values (with the fiducial choice in bold) we test for

each of these parameters. In our simulations, all parameters were fixed at the fiducial value except the one being varied.

DM	FWHM [ms]	Hexagonal N (Number of Antennae)	Frequency Range [MHz]
100	10	5 (56)	**50-250**
300	**100**	7 (120)	50-100
900	1000	9 (208)	100-150
2700	10000	**11 (320)**	150-200
			200-250
			110-190

Table 1: Parameter values tested in our simulations. Bolded entries indicate fiducial values. In each column, we simulated all values while keeping all other columns fixed at the fiducial value.

5 Results

In this section, we show the results of the statistical simulation methodology outlined in section 4. With so few detected FRBs at HERA frequencies, it is difficult to directly predict the number of FRBs that HERA might see without assuming a flux-dependent rate of FRBs. Instead, we choose to simply calculate the fraction of 1 Jy FRBs observable with at least a given SNR. Since SNR scales linearly with flux, we can thus interpret our results as a function of SNR/Jy. In general, dimmer FRBs will only be detectable at some SNR threshold near the center of the primary beam while brighter FRBs will be detectable over a larger solid angle and will thus be detected more often.

We begin by examining the impact that the size of the array—quantified by the hexagonal array number N—has on FRB detectability. As can be seen in Figure 5, the detection prospects improve as N increases. Assuming a fiducial 800 FRBs per sky per day (or 400 night^{-1}sky^{-1} assuming 12 hours of observing per night; (CHIME/FRB Collaboration et al., 2021)) with $N = 11$, we would expect to observe around 1 FRB per night with an SNR/Jy > 5. This expectation drops to 0.4 and 0.08 for N=7 and N=5 respectively.

Figure 6 shows results analogous to those shown in Figure 5 but for populations of FRBs differing in pulse FWHMs rather than N. As expected, FRB detectability rises strongly with increasing FWHM because the burst duration gets more comprable to the 9.6 s integration time. For the fiducial case of FWHM=100 ms, we once again find a 0.3% detection rate for 5.0 SNR/Jy FRBs. This number rises to $\sim$1% and $\sim$2% for FWHM=1000 ms and 10000ms respectively. In terms of the actual number of detections per night, we expect that full HERA may be able to detect 2, 8, and 16 FRBs per night at the 5.0 SNR/Jy level for the FWHM=100, 1000, and 10000 ms cases respectively. Conversely, detection prospects for a 10 ms are bleak: we expect to detect virtually no FRBs with FWHM=10 ms (a timescale typical of higher frequency detections of FRBs) with SNR/Jy > 2.5.

Figure 7 shows the results for when the observing frequency range is varied. This is meant to help assess the possibility that the FRB is more narrow-band than our observations, as was often the case in Pleunis et al. (2021a). As we would expect, the widest range (50-250 MHz) has the highest detectability, as SNR is additive in quadrature. Detectability decreases as we shift to lower frequency ranges, as evidenced by the steep drop-offs in the 50-100 MHz and 100-150 MHz curves as SNR/Jy increases. This is because of the higher noise at low frequencies (Figure 2).

Finally, we show the results of varying the bursts' DM in Figure 8. Most notable here is the observation that DM does not seem to significantly impact FRB detectability. Though the values we test span a wide range (DM=100,300,900, and 2700) the detection curves for each are very similar. On a given day, these curves imply that irrespective of DM, HERA could detect 0.3% of FRBs anywhere on the sky with a SNR/Jy of at least 5.0. As before, if we assume 12 hours of observation per day and 800 1 Jy FRBs per sky per day (i.e. 400 per night), this translates to HERA detecting 1 FRB per night at the 5σ threshold.

Owing to the higher noise at lower frequencies, the dominant contribution to the SNR of an FRB is from the higher frequency component. However, the apparent movement of the source on the sky is relatively

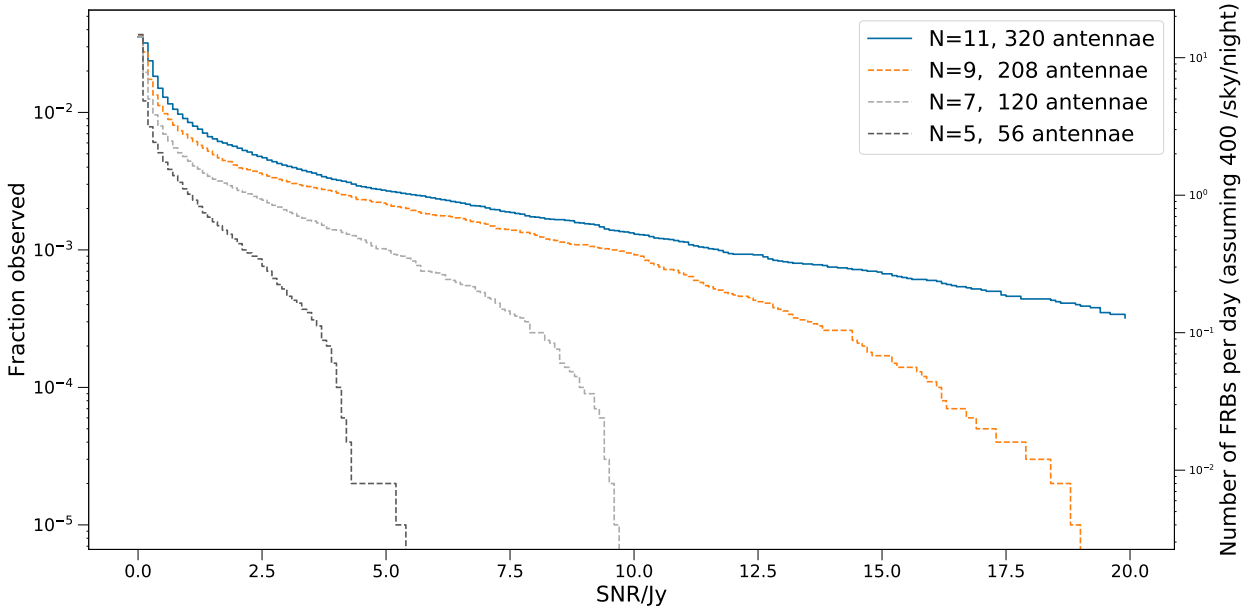

Figure 5: Detection statistics for four populations of simulated FRBs, in which the varied parameter is the array size, which itself is parametrised using hexagonal N. We test four values for N: 5,7, 9, 11 (full HERA), corresponding to 56, 120, 208, and 320 antennae respectively. Points on the curves for the three populations can be interpreted as follows: the y-coordinate is the fraction of FRBs anywhere on the sky observable with an SNR/Jy greater than the x-coordinate. The second y-axis shows these fractions converted into an expectation for the number of FRBs observed above a given SNR/Jy threshold per day using an occurrence rate of 400 FRB night^{-1}sky^{-1}. As might be expected, detectability increases alongside N. These detection curves imply that for N=11, we would expect to observe around 1 FRB day^{-1} with an SNR/Jy>5. This expectation drops to 0.4 and 0.08 for N=7 and N=5 respectively. Even higher SNRs/Jy occur more rarely because they require for the FRB to be located very near to the peak of the primary beam.

small during the arrival time range for the higher frequency components of the burst, owing to the nature of the dispersion relation. However, there are a few additional effects that impact the SNR. Low DM FRBs that fall into the primary beam should spend a higher fraction of time in the primary beam owing to their lower dispersion time. On the other hand, high DM FRBs are more likely to have at least some part of their burst's spectrum fall in the primary beam. This effect can be seen in Figure 9, in which each panel is analogous to Figure 4 but for populations of FRBs divided based on their DM. In the DM=100 case, we observe fewer FRBs that have their max altitude very near zenith. However, the SNR/Jy for this sample is very high. This contrasts with the trends seen in the DM=2700 population, within which a larger proportion of FRBs reach zenith, an effect largely balanced out by the lower average SNR/Jy of this sample.

6 Conclusion and Next Steps

We have performed a preliminary forecasting of HERA's ability to detect FRBs without a dedicated FRB detection system with high temporal resolution. Our results indicate that full HERA should be able to see on the order of 1 FRB per day with SNR/Jy > 5. Though the advent of full HERA will significantly improve our expected capability to detect such FRBs, our work indicates that it is possible that there may exist some undetected FRBs lurking in already existing HERA data. Motivated by this, we aim to search for FRBs in already existing archival data and develop the infrastructure to perform such searches in HERA data yet to come.

A unique aspect of HERA that may aid this effort is the phasing pattern an FRB should show in HERA diffs (Figure 10); a consequence of the frequency-dependent arrival time of the FRB. We hope to exploit this

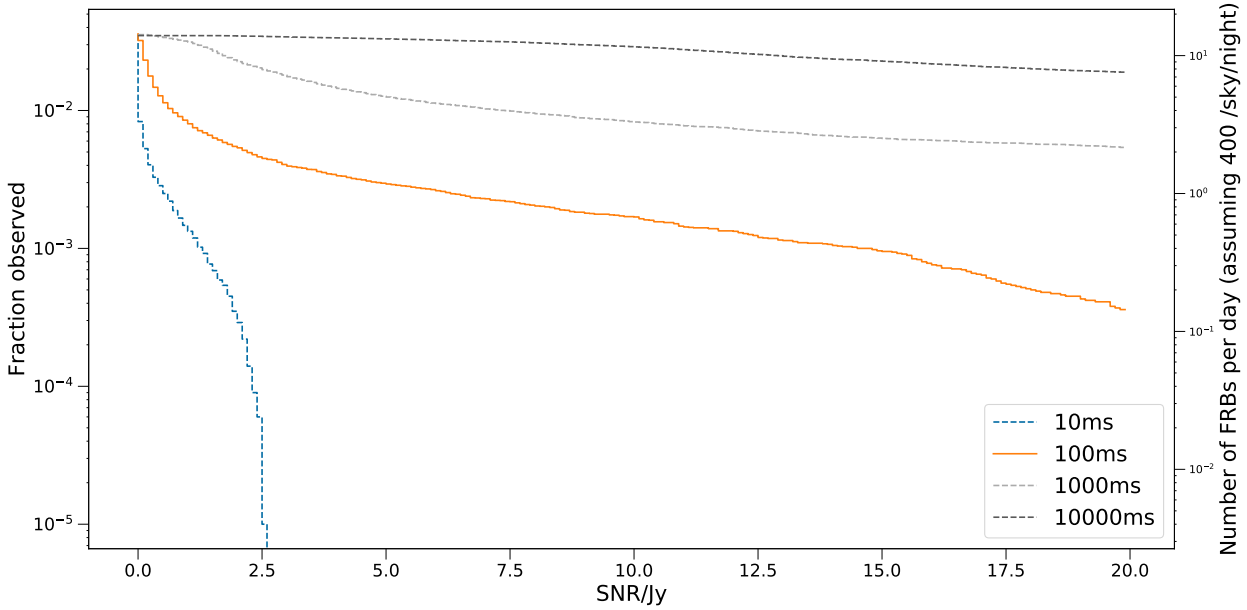

Figure 6: Detection statistics analogous to those shown in Figure 5, but with the FWHM of the burst being the parameter varied. We test four values for the FWHM: 10 ms, 100 ms, 1000 ms, and 10000 ms. As illustrated in the plot, virtually none of the simulated FRBs with a FWHM of 10 ms exceed a SNR/Jy of 2.5. On the other hand, the number detected at SNR/Jy=20 (assuming the 400 FRB night^{-1}sky^{-1} occurrence rate) rises to 5 and 10 for the 100 ms and 1000 ms pulses respectively.

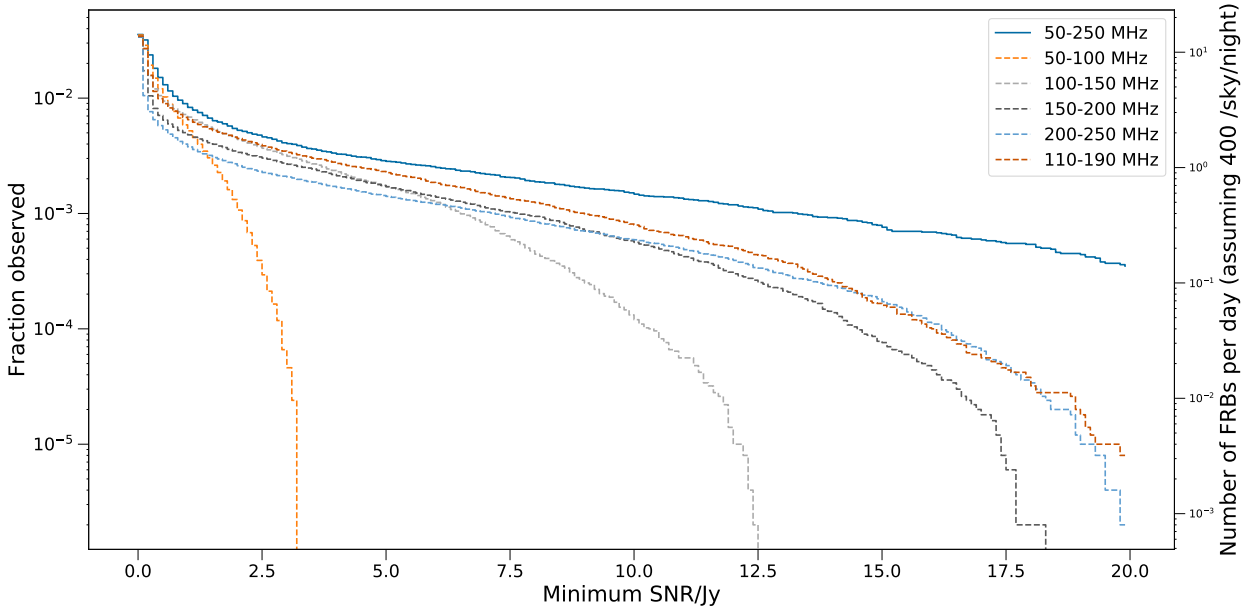

Figure 7: Detection statistics analogous to those shown in Figure 5, but here we focus on the effects of varying the frequency window being observed. Due to the fact that the noise level generally rises with decreasing frequency (Figure 2) while our FRB model has a frequency-independent flux, we see that the low frequency windows have lower SNRs. Of particular interest here is the 110-190 MHz window, which corresponds to the frequency range in archival data from HERA Phase I (Abdurashidova et al., 2022).

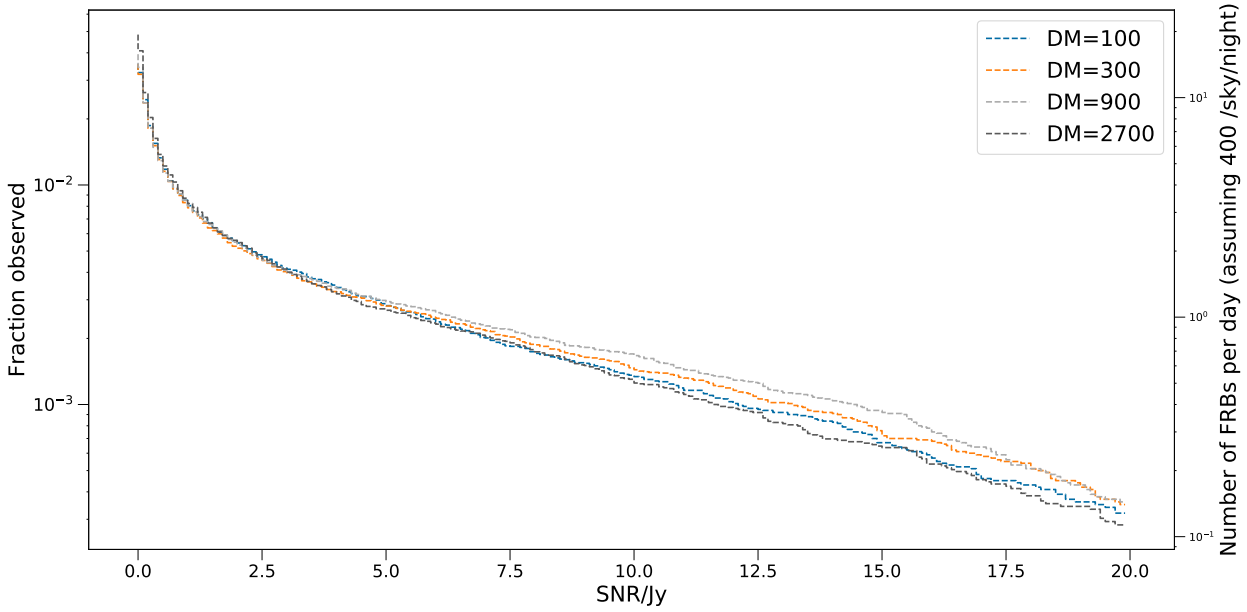

Figure 8: Detection statistics for four populations of simulated FRBs with differing dispersion measures: 100, 300, 900, 2700. These detection curves imply that the DM of an FRB source does not seem to strongly impact its detectability with HERA, an effect that persists over the wide range of DMs from 100 to 2700.

distinctive pattern for the purposes of FRB detection—perhaps by training a machine learning algorithm to detect it—inspired by success in recent studies (Zhang et al., 2018).

References

Abdurashidova, Zara, James E. Aguirre, Paul Alexander, et al. (2022). "First Results from HERA Phase I: Upper Limits on the Epoch of Reionization 21 cm Power Spectrum". In: *The Astrophysical Journal* 925.2, p. 221. DOI: 10.3847/1538-4357/ac1c78. URL: https://doi.org/10.3847%2F1538-4357%2Fac1c78.

CHIME/FRB Collaboration, Mandana Amiri, Bridget C. Andersen, et al. (Dec. 2021). "The First CHIME/FRB Fast Radio Burst Catalog". In: ApJS 257.2, 59, p. 59. DOI: 10.3847/1538-4365/ac33ab. arXiv: 2106.04352 [astro-ph.HE].

Pleunis, Z., D. Michilli, C. G. Bassa, et al. (Apr. 2021a). "LOFAR Detection of 110-188 MHz Emission and Frequency-dependent Activity from FRB 20180916B". In: ApJ 911.1, L3, p. L3. DOI: 10.3847/2041-8213/abec72. arXiv: 2012.08372 [astro-ph.HE].

Pleunis, Ziggy, Deborah C. Good, Victoria M. Kaspi, et al. (Dec. 2021b). "Fast Radio Burst Morphology in the First CHIME/FRB Catalog". In: ApJ 923.1, 1, p. 1. DOI: 10.3847/1538-4357/ac33ac. arXiv: 2106.04356 [astro-ph.HE].

Shannon, R. M., J. P. Macquart, K. W. Bannister, et al. (Oct. 2018). "The dispersion-brightness relation for fast radio bursts from a wide-field survey". In: Nature 562.7727, pp. 386–390. DOI: 10.1038/s41586-018-0588-y.

Yang, X., S. B. Zhang, J. S. Wang, et al. (Nov. 2021). "81 New candidate fast radio bursts in Parkes archive". In: MNRAS 507.3, pp. 3238–3245. DOI: 10.1093/mnras/stab2275. arXiv: 2108.00609 [astro-ph.HE].

Zhang, Yunfan Gerry, Vishal Gajjar, Griffin Foster, et al. (Oct. 2018). "Fast Radio Burst 121102 Pulse Detection and Periodicity: A Machine Learning Approach". In: ApJ 866.2, 149, p. 149. DOI: 10.3847/1538-4357/aadf31. arXiv: 1809.03043 [astro-ph.HE].

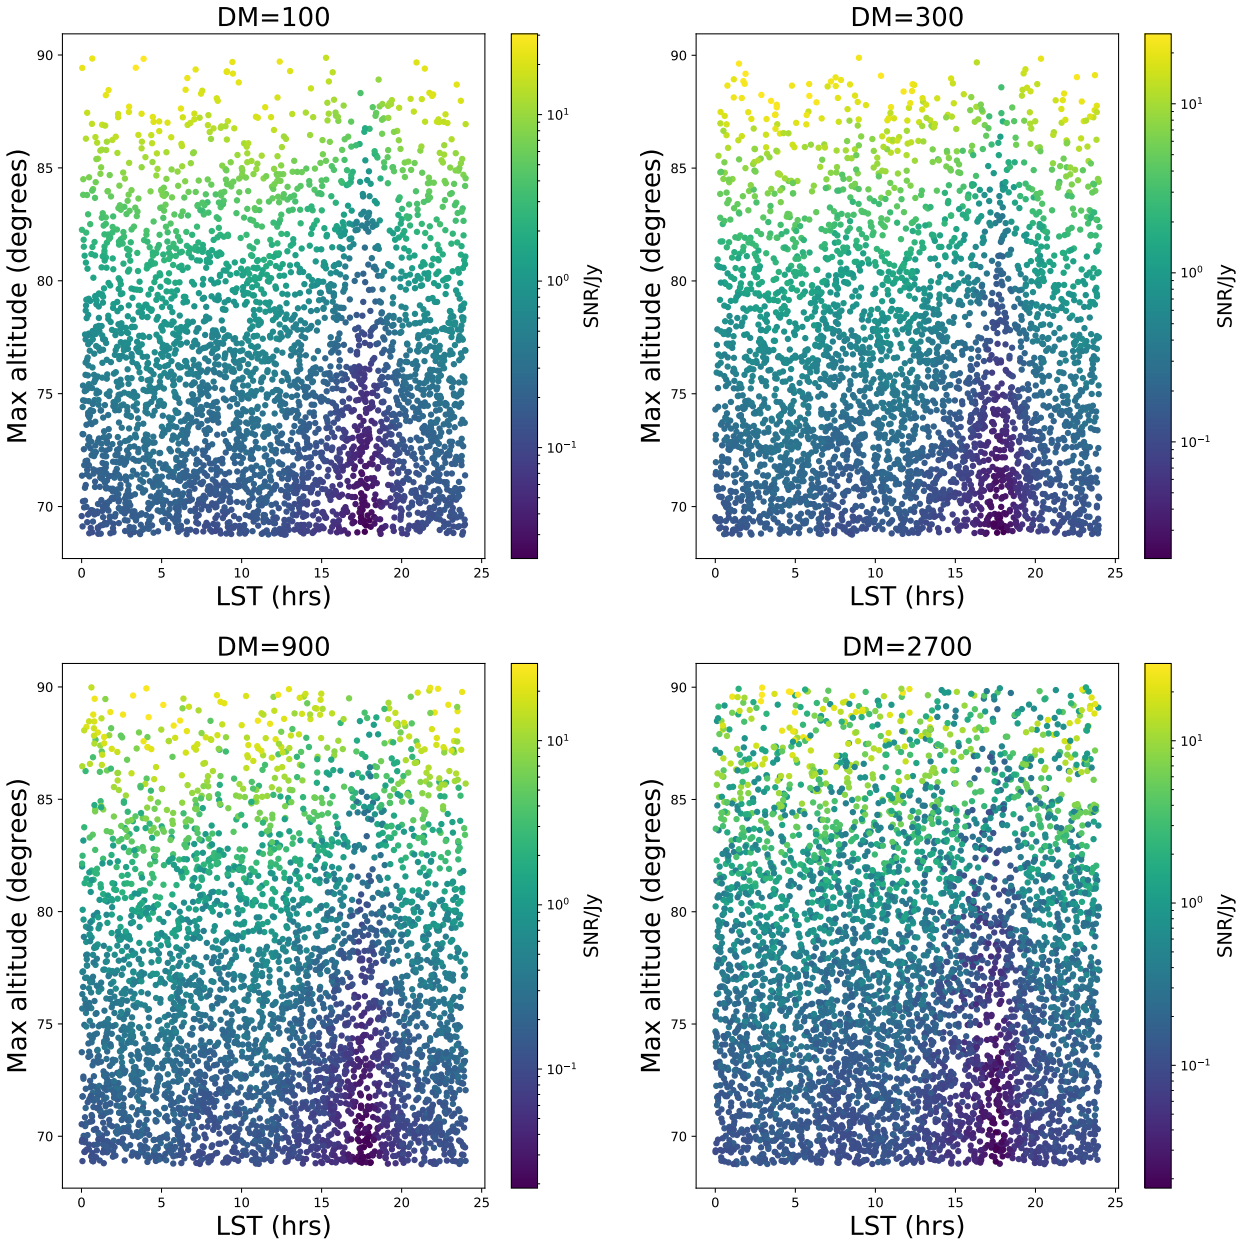

Figure 9: SNR/Jy statistics for populations of simulated FRBs plotted with respect to their central LST and the source's max altitude during the duration of the burst, with SNR shown by the colour. Here, each panel is analogous to Figure 4, with the panels differing in the DM of the FRB population displayed. As DM increases, we observe two competing effects. The number of bursts reaching a given maximum altitude increases with DM, as allowed by the longer dispersion time. Conversely, the average SNR of bursts with a given max altitude decreases with increasing DM, as the long dispersion time means that these bursts spend a longer portion of time deeper in the beam, which drives down SNR. The competing nature of these effects may be responsible for how the detectability of an FRB does not strongly depend on its DM, as observed in Figure 8.

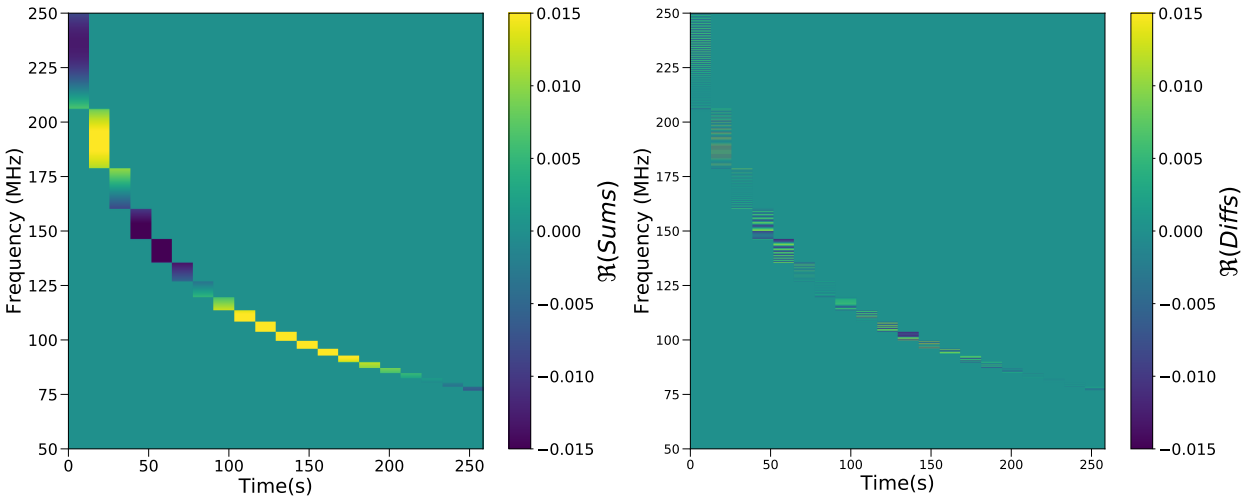

Figure 10: Sums and Diffs for a unit flux FRB with DM=300, N=11, and FHWM=100 ms, for the same 51 m E, 38 m S baseline as in Figure 1. Of particular interest here is the interweaving pattern that can be seen in the diffs (right panel). This occurs because the arrival time of the FRB shifts back and forth between the even and odd interleaved 100 ms sub-integrations as it disperses as a function of frequency. This characteristic signature may prove a smoking gun in burst detection efforts.