



Exploring Realistic Nanohertz Gravitational-wave Backgrounds

Bence Bécsy^{1,2} , Neil J. Cornish¹ , and Luke Zoltan Kelley^{3,4} ¹ eXtreme Gravity Institute, Department of Physics, Montana State University, Bozeman, MT 59717, USA; bencebecsy@montana.edu² Department of Physics, Oregon State University, Corvallis, OR 97331, USA³ CIERA (Center for Interdisciplinary Exploration and Research in Astrophysics), Evanston, IL 60201, USA⁴ Department of Astronomy, University of California at Berkeley, Berkeley, CA 94720, USA

Received 2022 July 4; revised 2022 November 7; accepted 2022 November 8; published 2022 December 15

Abstract

Hundreds of millions of supermassive black hole binaries are expected to contribute to the gravitational-wave signal in the nanohertz frequency band. Their signal is often approximated either as an isotropic Gaussian stochastic background with a power-law spectrum or as an individual source corresponding to the brightest binary. In reality, the signal is best described as a combination of a stochastic background and a few of the brightest binaries modeled individually. We present a method that uses this approach to efficiently create realistic pulsar timing array data sets using synthetic catalogs of binaries based on the Illustris cosmological hydrodynamic simulation. We explore three different properties of such realistic backgrounds that could help distinguish them from those formed in the early universe: (i) their characteristic strain spectrum, (ii) their statistical isotropy, and (iii) the variance of their spatial correlations. We also investigate how the presence of confusion noise from a stochastic background affects detection prospects of individual binaries. We calculate signal-to-noise ratios of the brightest binaries in different realizations for a simulated pulsar timing array based on the NANOGrav 12.5 yr data set extended to a time span of 15 yr. We find that $\sim 6\%$ of the realizations produce systems with signal-to-noise ratios larger than 5, suggesting that individual systems might soon be detected (the fraction increases to $\sim 41\%$ at 20 yr). These can be taken as a pessimistic prediction for the upcoming NANOGrav 15 yr data set, since it does not include the effect of potentially improved timing solutions and newly added pulsars.

Unified Astronomy Thesaurus concepts: Gravitational waves (678); Gravitational wave astronomy (675); Gravitational wave sources (677); Supermassive black holes (1663); Millisecond pulsars (1062)

1. Introduction

Pulsar timing arrays (PTAs) probe gravitational waves (GWs) with frequencies between a few and a few hundred nanohertz (nHz) by continuously monitoring millisecond pulsars (for a review see, e.g., Burke-Spolaor et al. 2019; Taylor 2021). The NANOGrav Collaboration recently found evidence for a low-frequency stochastic process common to all pulsars in their 12.5 yr data set (Arzoumanian et al. 2020). The three other PTAs also all found evidence for such a common red-noise process (International Pulsar Timing Array (IPTA), Antoniadis et al. 2022; EPTA, Chen et al. 2021; PPTA, Goncharov et al. 2021). It is possible that these results are the first signs of a stochastic GW background (GWB), but at this point none of the above analyses found strong evidence for Hellings–Downs (HD; Hellings & Downs 1983) spatial correlations, which are the telltale sign of a GWB. However, this is not surprising given that the power in cross-correlations is about an order of magnitude less than the total power of a GWB. This means that we can expect to detect a common process before we detect the HD correlations between pulsars (see, e.g., Pol et al. 2021; Romano et al. 2021). Upcoming PTA data sets will decide whether this red-noise process is the GWB or not.

Once a firm detection of the GWB is established, the next task will be to identify its origin. The theoretically most favored source of a GWB at nHz frequencies is an ensemble of inspiraling supermassive black hole binaries (SMBHBs; see, e.g.,

Rajagopal & Romani 1995; Enoki et al. 2004; Sesana et al. 2004; Kelley et al. 2017b). However, there are several other proposed processes that can form such a low-frequency GWB, like cosmological phase transitions (Arzoumanian et al. 2021), cosmic strings (van Haasteren et al. 2011; Blasi et al. 2021), primordial black holes (De Luca et al. 2021), inflation (Vagnozzi 2021), etc. These models all have different predictions for the spectral shape of the common signal, which presents the possibility of distinguishing these models (Kaiser et al. 2022). However, definitively identifying the source of the GWB based solely on its spectrum will be challenging given the large uncertainty in both measurement (Pol et al. 2021) and model predictions. Thus, any additional predictions of these models beyond their spectra could be useful in determining the source of the GWB.

In this paper we explore various properties of realistic GWBs from SMBHBs, some of which might help distinguish that scenario from others. As we will see, most of these properties rely on the fact that an SMBHB-based GWB is built from a finite number of sources. The statistics of individual sources in a realistic simulation were first investigated in Sesana et al. (2009), while Rosado et al. (2015) calculated detection probabilities of individual sources in the presence of a GWB. The general question of when the signal from a finite collection of individual sources becomes effectively stochastic was investigated in Cornish & Romano (2015), while the effects of the finite population on the detection prospects of the GWB have been studied in Cornish & Sampson (2016). Here we are instead focusing on how these effects can help us discern the origin of the GWB. Our simulations are using synthetic SMBHB catalogs based on the Illustris cosmological simulation (Kelley et al. 2018). We introduce and validate a method



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

that can rapidly produce simulated PTA data sets from millions of SMBHBs by modeling the large majority of the sources as a GWB and directly simulating the few brightest binaries in each frequency bin (see Section 2). We use this simulation method to investigate how the central limit theorem breaks down for such an SMBHB population (see Section 3.1), how the finite number of sources can lead to anisotropic GWBs (see Section 3.2), and how the variance of the HD correlations is affected by the properties of the SMBHB population (see Section 3.3). We also explore the detectability of individual binaries in the presence of a GWB by calculating the distribution of signal-to-noise ratios (S/Ns) of the brightest binaries in multiple realizations (see Section 4). We summarize and offer concluding remarks in Section 5.

2. Simulation Methods

2.1. SMBH Population Model

Our simulations are based on synthetic catalogs of the SMBHB population representative of the entire observable universe. These are derived from the Illustris cosmological hydrodynamic simulations (see, e.g., Vogelsberger et al. 2014), as implemented in the *holodeck* code (L. Z. Kelley 2022, in preparation), by modeling small-scale astrophysical processes in a post-processing step. This treatment accounts for environmental effects relevant to binary hardening like dynamical friction, stellar scattering, drag from a circumbinary disk, and GW emission (for more details see Kelley et al. 2017a, 2017b, 2018). The collection of mergers observed in the Illustris simulation volume can be resampled multiple times to create new realizations of a simulated catalog of SMBHBs for the entire past light cone of the observer. The standard free parameter in this implementation of SMBH binary populations is the binary lifetime. Additionally, to produce a GWB amplitude roughly consistent with the observed common process in the NANOGrav 12.5 yr data set (Arzoumanian et al. 2020), we have tuned the volume density of mergers and the distribution of SMBH masses.

Figure 1 shows the distribution of the source-frame chirp mass ($\mathcal{M}$), the luminosity distance (d_L), and the observer-frame GW frequency (f_{obs}) of such a simulated data set of SMBHBs. Here we applied a lower-frequency cutoff at $(15 \text{ yr})^{-1}$, corresponding to the observational time span we consider in this paper and the observational time span of the upcoming NANOGrav 15 yr data set (M. F. Alam et al. 2022, in preparation). This particular realization has about 115 million binaries. Note that the number density is dominated by low-frequency, low-mass, faraway systems.

2.2. Isotropic GWB plus Bright Binaries

The timing residuals in a PTA data set at observing epochs t_i can be written as

$$r(t_i) = r_n(t_i) + r_{\text{GW}}(t_i), \quad (1)$$

where r_n describes residuals due to every noise source, while r_{GW} is the contribution of GWs. In the case of a population of N SMBHBs producing GWs in the observable universe, one can write

$$r_{\text{GW}}(t_i) = \sum_{j=1}^N s_j(t_i; \theta_j), \quad (2)$$




Figure 1. Distribution of binary parameters in the simulated SMBHB catalog we use. The distribution was artificially cut off below observer-frame GW frequencies of $(15 \text{ yr})^{-1}$, which would not be observable with the 15 yr observing time span considered here. Vertical dashed lines indicate the median value of each parameter.

where s_j is the timing residual response to the GW signal from the j th binary, which is described by the parameters

$$\theta_j = \{\mathcal{M}, d_L, f_{\text{obs}}, \theta, \phi, \iota, \psi, \Phi_0, \Phi_1, \dots, \Phi_{N_{\text{PSR}}}, L_1, \dots, L_{N_{\text{PSR}}}\}, \quad (3)$$

where θ and ϕ parameterize the location of the source on the sky, ι is the inclination of the binary's orbit, ψ is the polarization angle, Φ_0 (Φ_i) is the initial phase of the GW signal at Earth (at the i th pulsar), L_i is the distance to the i th pulsar, and N_{PSR} is the number of pulsars (for an explicit expression of $s_j(t_i; \theta_j)$ see, e.g., Equation (10) in Aggarwal et al. 2019). While Equation (2) is a valid description of the total GW signal, it is impractical both for data analysis and for simulating signals. In terms of data analysis, this model has $(8 + N_{\text{PSR}})N + N_{\text{PSR}}$ parameters,⁵ which are practically impossible to explore in the realistic case, where $N \sim \mathcal{O}(10^6\text{--}10^9)$. In terms of simulating PTA data sets, individually calculating the response of millions of binaries becomes a computationally expensive task.

We will show that both of these problems can be averted by expressing the total contribution to the residuals as

$$r_{\text{GW}}(t_i) = r_{\text{GWB}}(t_i, h_c^{\text{GWB}}) + \sum_{k=1}^{N_{\text{freq}}} \sum_{j=1}^M s_j^{(k)}(t_i; \theta_j), \quad (4)$$

where r_{GWB} is an isotropic Gaussian stochastic background with the characteristic strain spectrum h_c^{GWB} , N_{freq} is the number of frequency bins considered, and $s_j^{(k)}$ is the binary with the j th-largest characteristic strain in the k th frequency bin. The second term in Equation (4) loops over all N_{freq} frequency bins

⁵ The pulsar distance parameters are the same for each GW source, but the pulsar phase parameters introduced to help convergence will be different for each binary.

and directly adds the contribution of the M brightest sources in each bin. We set h_c^{GWB} in each frequency bin based on the remaining SMBHBs as

$$h_c^{\text{GWB}}(f_k) = \sqrt{\sum_{j=M+1}^{N_k} [(h_c)_j^{(k)}]^2}, \quad (5)$$

where N_k is the total number of binaries in the k th frequency bin, f_k is the central frequency of that bin,

$$(h_c)_j^{(k)} = h_j^{(k)} \sqrt{f_j^{(k)} T_{\text{obs}}} \quad (6)$$

is the characteristic strain of the j th-brightest source in the k th frequency bin, T_{obs} is the total observation time, $f_j^{(k)}$ is the observer-frame GW frequency of that source, and $h_j^{(k)}$ is the sky- and polarization-averaged GW amplitude (Sesana et al. 2004),

$$h_j^{(k)} = \frac{8}{\sqrt{10}} \frac{(G\mathcal{M}_{\text{obs}})^{5/3} (\pi f_{\text{obs}})^{2/3}}{c^4 d_L}, \quad (7)$$

where $\mathcal{M}_{\text{obs}} = (1+z)\mathcal{M}$ is the observer-frame chirp mass. Note that Equation (5) determines h_c^{GWB} by summing up the square of the characteristic strain for all binaries except the ones directly modeled in Equation (4). This is justified by the fact that when summing up sine waves of a given frequency with random phase offsets, the expectation value of the total amplitude squared is given by the sum of the squares of individual amplitudes. This approximately still holds in this scenario, where the frequencies of sources in the same frequency bin are approximately the same.

This method of describing a realistic background by an idealized isotropic Gaussian GWB and a sum of the few brightest sources can be used for both data analysis and simulation. The *BayesHopper*⁶ algorithm uses this approach to search for multiple individual SMBHBs in the presence of a GWB (Bécsy & Cornish 2020), while in this paper we use this approach to efficiently create realistic simulated data sets. Note that for data analysis purposes the ideal number of SMBHBs to model individually is dictated by Bayesian parsimony. However, for simulating signals we want to make sure that all the relevant details are captured, so we add more binaries directly than strictly necessary and than what can be picked up by data analysis. To validate the use of this description for rapidly simulating realistic GWBs, we simulated 200 realizations of a GWB by directly simulating all binaries (as in Equation (2)) and also by simulating a GWB and 1000 outlier SMBHBs in each frequency bin (as in Equation (4)). Figure 2 shows the median power spectral density (PSD) of post-fit residuals in a given simulated pulsar for both methods. The 90% credible intervals are shown as shaded bands. Here we used a simulated noise-free data set with a single pulsar and evenly spaced observations to enable the use of fast Fourier transforms. We applied a Tukey window on the residuals before calculating the PSD to avoid spectral leakage. Both methods produce spectra with a significant variance over realizations, but both their median values and 90% credible intervals show good agreement. Note that the two peaks correspond to frequencies of 1 and 2 yr^{-1} , where the timing model fit introduces extra power. The agreement between the




Figure 2. PSD of post-fit residuals for data sets where we individually simulated all binaries (red) and where we simulated a stochastic background and 1000 SMBHBs in each frequency bin (green). We show the median and 90% credible intervals over 200 realizations as solid lines and shaded regions, respectively. The more efficient outlier method correctly reproduces both the median and the variance of the PSD from the full simulations.

PSDs produced by the two methods shows that they result in statistically equivalent frequency content. In Sections 3.1 and 3.2 we show that after removing the brightest binaries the remaining signal satisfies the assumption of Gaussianity and statistical isotropy as well.

3. Properties of a Realistic GW Background

In this section we explore how various characteristics of the background are affected by the fact that there are a finite number of sources contributing to it. We investigate the spectrum of the GWB in Section 3.1, we test for statistical isotropy in Section 3.2, and we quantify the variance of the HD correlations in Section 3.3.

3.1. GW Spectrum

In the canonical description of a GWB arising from an infinite number of circular SMBHBs evolving purely through GW emission, the characteristic strain spectrum can be described as a power law with a spectral index of $-2/3$, i.e., $h_c^{\text{GWB}} \sim f_{\text{obs}}^{-2/3}$ (Phinney 2001). The spectrum arising from a population of SMBHBs differs from this simple spectrum (see purple circles and dashed line in Figure 3). The difference is most striking at high frequencies, where the population-based spectrum has a lower amplitude than the power-law model. This is due to the fact that the power-law model counts nonphysical contributions from fractional sources (Sesana et al. 2008). The population-based spectrum also exhibits significant fluctuations at high frequencies. Figure 3 also shows the spectrum when we exclude the brightest $1/10/100/1000$ binaries in each frequency bin.⁷ We can see that as we exclude more binaries, the scatter in the spectrum decreases. This suggests that by setting $M \gtrsim 100$, the resulting spectrum is not dominated by a few bright binaries anymore. Note that as we remove more binaries, the spectra in Figure 3 terminate at

⁶ Publicly available at <https://github.com/bencebecsy/BayesHopper>.

⁷ This corresponds to setting $M = 1/10/100/1000$ in Equation (5).




Figure 3. Characteristic strain spectrum for all binaries (purple) and for various numbers of the brightest binaries removed (different shades of green). We also show a power-law spectrum for reference (dashed gray line).

lower frequencies, because at higher frequencies we removed all the binaries.

To quantify how much h_c^{GWB} is dominated by a few sources at a given frequency, we look at the minimum number of binaries contributing to at least 90% of h_c^{GWB} at a given frequency (N_{90}). Figure 4 shows the median N_{90} value over 20 realizations. We show this for the total population and also for populations where we removed the top $M=1/10/100/1000$ binaries. The total number of binaries in each frequency bin is also indicated. While the bulk of the GWB signal is built up from $\mathcal{O}(10^3)$ binaries at the lowest frequency bins, it is dominated by less than 10 binaries above about 100 nHz. This is expected from our qualitative assessment of the spectra in Figure 3, where we have seen an increased scatter at high frequencies. We can also see that as we remove the brightest binaries from the population, the effective number of binaries contributing to the GWB increases. Note that as we remove more and more binaries, we can reach a limit where we are removing a significant fraction of the total number of binaries, and thus removing more binaries actually decreases N_{90} . (see, e.g., “-1000” and “-100” lines around 40 nHz in Figure 4). We also show the frequencies where the total numbers of binaries reach 100 and 1000 (vertical dashed lines). Note that the “-1000” and “-100” lines approach these horizontal lines, indicating the point where we remove all the binaries in the given bin.

Figure 4 illustrates why modeling a realistic background as a GWB and a few individual sources is well motivated. Modeling all the binaries as a GWB is not justified, since the background level for the full population is dominated by a few sources at high frequencies, so the central limit theorem breaks down. However, if we only consider the population with hundreds of the brightest binaries removed, the background level is determined by a large ($\gtrsim 100$ –1000) number of binaries, so treating it as a Gaussian background is a good approximation. In this case the non-Gaussian nature of the background is modeled by adding the brightest binaries as individual sources to the full signal (see Equation (4)).




Figure 4. Number of binaries contributing 90% of the total strain in each frequency bin (N_{90}). We show this for the whole data set (purple) and for data sets with the top 1/10/100/1000 binaries removed (different shades of green). We also show the total number of binaries in each bin (N_{tot} ; gray histogram). Vertical dashed lines indicate frequencies where the total numbers of binaries reach 1000 and 100. Note that the lines for the data sets without the top 1000/100 asymptote to the corresponding vertical lines, indicating frequency bins where we remove a large fraction of the binaries.

3.2. Statistical Isotropy

In the simplest description, it is also usually assumed that the GWB exhibits statistical isotropy. However, this is not expected to be true for realistic backgrounds. SMBHBs are hosted by galaxies, which cluster into galaxy clusters, so nearby galaxy clusters can result in an overabundance of nHz GW sources in a particular sky location. In addition, SMBHBs can produce anisotropy even when they are uniformly distributed in volume, if a small number of them dominate the nHz GW sky. This is in stark contrast with a GWB of primordial cosmological origin, which would be perfectly statistically isotropic, so the origin of the GWB can be firmly established by detecting anisotropy (Mingarelli et al. 2013). Indeed, there are numerous methods to search for signatures of anisotropy in the GWB (Mingarelli et al. 2013; Taylor & Gair 2013; Cornish & van Haasteren 2014; Hotinli et al. 2019; Ali-Haïmoud et al. 2020; Taylor et al. 2020), and by some estimates PTAs might be able to detect anisotropy just a few years after detecting the isotropic component of the GWB (Pol et al. 2022).

In our modeling we assume a uniform-in-volume distribution of sources. This means that the following results on anisotropy can be taken as a lower limit, since the effect of local inhomogeneities would further increase the level of anisotropy. To model the angular distribution of GW power, we pixelate the sky using the HEALPix framework (Górski et al. 2005). Since this produces equal-area pixels, randomly assigning each

SMBHB to a pixel is equivalent to distributing them uniformly on the sky. We carry out our analysis independently for each frequency bin. For each pixel, we sum up h_c^2 values of all the binaries assigned to that pixel, resulting in a discretized function of sky location $h_c^2(\Omega)$.

Figure 5 shows that distribution for a particular realization at the frequency bin between 66.5 and 68.7 nHz. This frequency bin has a total number of 9801 binaries. The different rows show maps with different numbers of the brightest binaries removed. For each of these we show both the raw map of $h_c(\Omega)$ (left) and a reconstructed map with angular scales restricted to those used for the anisotropy test described below (right). Note that the mean amplitude has been subtracted to better show the fluctuations. We allow the color scales to be different for each sky map, but we force it to be symmetric around 0. This particular example shows a highly anisotropic distribution due to a few brightest binaries dominating the sky, but as more and more of the binaries are excised, the remaining GWB gets more and more isotropic. By the time we remove the top 10 binaries, the GWB sky map restricted to large angular scales is visibly indistinguishable from an isotropic distribution. This is similar to the findings of previous studies reported in Taylor & Gair (2013) and Mingarelli et al. (2017), where authors found that small angular scale power can be reduced (and thus the map can be made more isotropic) by removing the brightest source (see Figure 2 in both references).

To quantitatively assess the statistical isotropy of maps like those shown in Figure 5, we follow the formalism introduced in Hajian & Souradeep (2003) and Hajian & Souradeep (2004). This method was used to test the statistical isotropy of the cosmic microwave background using data from the Wilkinson Microwave Anisotropy Probe (Hajian et al. 2005). We start by decomposing our maps in terms of $Y_{\ell m}$ spherical harmonics:

$$h_c(\Omega) = h_c(\theta, \varphi) = \sum_{\ell=0}^{\ell_{\max}} \sum_{m=-\ell}^{\ell} a_{\ell m} Y_{\ell m}(\theta, \varphi). \quad (8)$$

The $a_{\ell m}$ spherical harmonic coefficients can be used to compute the bipolar spherical harmonic coefficients:

$$A_{ll'}^{\ell M} = \sum_{mm'} a_{\ell m} a_{\ell' m'} \mathfrak{C}_{lm'l'm'}^{\ell M}, \quad (9)$$

where $\mathfrak{C}_{lm'l'm'}^{\ell M}$ are Clebsch–Gordan coefficients. The bipolar power spectrum is defined as

$$\kappa_{\ell} = \sum_{ll'M} |A_{ll'}^{\ell M}|^2. \quad (10)$$

Note that for statistically isotropic maps the usual angular power spectrum $C_{\ell} = 1/(2\ell + 1) \sum_m |a_{\ell m}|^2$ contains all information about the map, and we have $A_{ll'}^{\ell M} = (-1)^l C_{\ell} \sqrt{2\ell + 1} \delta_{ll'} \delta_{\ell 0} \delta_{M0}$ and $\kappa_{\ell} = \kappa_0 \delta_{\ell 0}$. After correcting for the estimation bias (Equation (14) of Hajian & Souradeep 2003), κ_{ℓ} has an expectation value of 0 for all $\ell \geq 1$ under statistical isotropy. The variance can also be analytically computed (Equation (17) of Hajian & Souradeep 2003), so one can test for statistical isotropy by checking whether κ_{ℓ} are consistent with zero within their theoretical variance.

We calculate κ_{ℓ} for $1 \leq \ell \leq 6$, which corresponds to probing angular scales larger than 30° . The cutoff at $\ell = 6$ is motivated by the angular resolution sensitivity forecasts from Pol et al. (2022). We apply a Gaussian low-pass filter on the a_{ℓ}^m values as defined in Equation (5) of Hajian et al. (2005) with $l_s = 20$.

This ensures that there is no spectral leakage affecting the analysis. A key difference between our analysis and the one presented in Hajian et al. (2005) is that we do not use the theoretical variance of κ_{ℓ} to decide whether a map is isotropic or not. Instead, we compare the measured κ_{ℓ} values to κ_{ℓ} values calculated from N_{iso} realizations of a simulated isotropic map with the same angular power spectrum. We call a map statistically isotropic if the measured κ_{ℓ} values are within the maximum and minimum κ_{ℓ} from the truly isotropic maps for all ℓ . This procedure allows us to set the false positive and false negative probabilities of our test by choosing the value of N_{iso} . As discussed below, we perform multiple trials of this test, so we want to make sure that anisotropic maps rarely pass the test by chance. We choose $N_{\text{iso}} = 5$, which results in the false positive and false negative rates reported in Table 1. Note that increasing N_{iso} would further decrease the probability of anisotropic maps passing the test at the cost of fewer truly isotropic maps passing it.

We performed κ_{ℓ} -tests on all four maps shown in the right column of Figure 5. The maps with all binaries and with the single brightest binaries removed failed the test and are thus not consistent with being statistically isotropic. The maps with 10 and 100 binaries removed passed the test. It is instructive to compare the maps in the left and right columns. Based on the raw maps, we might say that even the map with 10 binaries removed is not isotropic. However, the maps on the right correspond to what the κ_{ℓ} -test actually sees. It is evident even in those large angular scale features that the first two maps are dominated by a few sources. Note in particular that these maps are highly asymmetric around their mean values, resulting in the lack of dark-blue colors. However, once we remove the 10 brightest binaries, the map is consistent with statistical isotropy up to $\ell = 6$. This example reinforces our previous qualitative assessment that removing the few brightest binaries makes the GWB more isotropic.

We can ask how many binaries one needs to remove to get a background consistent with statistical isotropy. To answer that, we start with a κ_{ℓ} -test on a map with all binaries and remove binaries one at a time until we reach a statistically isotropic map. We call the number of binaries removed to achieve this N_{rem} . We repeat this procedure over all frequency bins and multiple realizations (see Figure 6). We can see that there is a trend of needing to remove more binaries as we move to higher frequencies. Figure 6 also shows the number of binaries with an amplitude within a factor of three of the brightest binary, $N_{\text{top}}(33.3\%)$, after removing the N_{rem} brightest binaries. We can see that we typically reach isotropy when there are more than ~ 10 binaries with comparable amplitudes. This is consistent with the fact that at $\ell = 6$ there are 13 different spherical harmonics, so once there are $\gtrsim 10$ comparable sources randomly placed on the sky, we cannot distinguish that from an isotropic sky. Note that at the highest frequencies $N_{\text{rem}} \simeq N_{\text{tot}}$, where N_{tot} is the total number of binaries in each frequency bin. This is because in our procedure we set $N_{\text{rem}} = N_{\text{tot}}$ if we cannot achieve statistical isotropy with any value of N_{rem} . This is a similar feature to the one in Figure 4, where removing more binaries does not get us closer to a true stochastic background when N_{tot} is small to start with. We also show the expected value of N_{rem} assuming that all maps are isotropic, $E[N_{\text{rem}}]_{\text{isotropy}} = \alpha/(1 - \alpha) = 0.75$, which is based on N_{rem} following a negative binomial distribution with the success probability of $\alpha = 0.43$ corresponding to the false negative probability of the κ_{ℓ} -test reported in Table 1. We can see that at




Figure 5. Angular distribution of the characteristic amplitude in the frequency bin with $66.5 \text{ nHz} < f_{\text{obs}} < 68.7 \text{ nHz}$. We show sky maps for all binaries (first row) and with 1/10/100 brightest binaries removed (second/third/fourth rows). The left column shows the raw maps, while the right column shows the maps reconstructed with the same range of angular scales as was used for the anisotropy test ($1 \leq \ell \leq 6$). Only the maps with 10 and 100 binaries removed passed the isotropy test in this example. The monopole term was subtracted from each map. Note the changing scale of the color maps, which we forced to be symmetric around 0.




Figure 6. Minimum number of sources removed to achieve statistical isotropy based on the κ_ℓ -test with ℓ between 1 and 6, $N_{\text{rem}}(1 \leq \ell \leq 6)$. We show the median and 68% credible interval of N_{rem} over 30 GWB realizations. We also show the number of binaries with an h_c within a factor of 3 of the brightest binary in the frequency bin, $N_{\text{top}}(33.3\%)$. Note how statistical isotropy is reached when $N_{\text{top}}(33.3\%) \sim 10$. The black line shows the total number of sources in each frequency bin. The dashed (dotted) horizontal line shows the expected value of N_{rem} due to the false negative (positive) probability of our test (see Table 1) if all maps are isotropic (anisotropic).

Table 1

False Negative and False Positive Probabilities of the Anisotropy Test with $N_{\text{iso}} = 5$

Isotropy Test Result	Pass	Fail
Truly isotropic map	57%	43%
Truly anisotropic map ^a	12%	88%

Note.

^a Based on randomly selected map dominated by a single source.

the lowest frequencies the median value of N_{rem} is only slightly above this expectation value, suggesting that a fraction of these maps are isotropic even without removing any bright binaries. We also show the expected value of N_{rem} assuming that every map is anisotropic, $E[N_{\text{rem}}]_{\text{anisotropy}} = 7.33$, which is higher than the median N_{rem} at any frequency and thus not expected to significantly affect the results.

3.3. Variance of the Hellings–Downs Curve

It can be shown that an isotropic, unpolarized, time-stationary, Gaussian GWB produces residuals that are correlated between pulsar pairs, and the amount of correlation between pulsar i and pulsar j is described by the HD curve (Hellings & Downs 1983):⁸

⁸ Note that we follow the normalization convention where $\Gamma_{ii} = 1$.

$$\Gamma_{ij} = \frac{1}{2} - \frac{1}{4}x_{ij} + \frac{3}{2}x_{ij}\ln(x_{ij}) + \frac{1}{2}\delta_{ij}, \quad (11)$$

where the term with δ_{ij} encodes correlations due to the pulsar terms and vanishes when looking at different pulsars, and x_{ij} is defined as

$$x_{ij} = \frac{1 - \cos \gamma_{ij}}{2}, \quad (12)$$

where γ_{ij} is the angular separation between pulsars i and j . It has been shown in Cornish & Sesana (2013) that Equation (11) holds not only for GWBs but also for a so-called continuous-wave (CW) signal with a quasi-constant frequency and amplitude from an individual SMBHB.

While Equation (11) predicts the expectation value of the correlations as a function of γ , individual correlations can have a significant variance (see Allen 2022 and references therein). It has been suggested that the way the variance varies with angle might be used to distinguish between primordial and astrophysical sources of the GWB (Allen 2022).

There are two sources of that variance that are present even in idealized noise-free observations, which we can call pulsar variance and cosmic variance, following the nomenclature of Allen (2022). Pulsar variance arises from the fact that individual correlations can depend not only on γ but also on three other angles describing the position of the two pulsars on the sky. Having a large number of pulsar pairs, one can in principle average out the pulsar variance. Cosmic variance refers to statistical deviations from Equation (11) in a given realization of the GWB. Unlike pulsar variance, we cannot average out cosmic variance, since we only have one universe to make observations of.

To quantify the variance of the HD correlations, we performed three different simulations, each with 100 fully independent realizations:

- (i) GWB: an isotropic stationary Gaussian stochastic background with a characteristic strain amplitude of 1×10^{-15} measured at $f_{\text{obs}} = 1 \text{ yr}^{-1}$ and spectral index of $-2/3$.
- (ii) CW: GW signal of an individual SMBHB with fixed amplitude, $\mathcal{M} = 10^9 M_\odot$, $f_{\text{obs}} = 10 \text{ nHz}$, and nuisance parameters drawn randomly from uniform distributions.
- (iii) Population: realistic GWB simulated as in Equation (4).

For these simulations we use the sky location of the 67 pulsars in the NANOGrav 15 yr data set (M. F. Alam et al. 2022, in preparation) and simulate evenly sampled observations over 15 yr with a 30-day cadence and with no noise. From the resulting $r_k(t_i)$ residuals at times t_i in pulsar k we calculate the zero-lag correlation between each pulsar pair as

$$\mathcal{C}_{kl} = \frac{1}{N_{\text{obs}}} \sum_{i=1}^{N_{\text{obs}}} r_k(t_i) r_l(t_i), \quad (13)$$

where N_{obs} is the number of observations per pulsar. We also introduce

$$\hat{\mathcal{C}}_{kl} = \mathcal{C}_{kl} \left(\frac{1}{N_{\text{PSR}}} \sum_{k=1}^{N_{\text{PSR}}} \mathcal{C}_{kk} \right)^{-1}, \quad (14)$$

which normalizes $\mathcal{C}_{kl}$ with the autocorrelation terms, so that it follows the same convention as Equation (11).




Figure 7. Mean and variance of interpulsar correlations. We show results for an idealized GWB, a single CW source, and a realistic population-based GWB. We also show the theoretically expected mean and variance of the correlations under different assumptions. The small discrepancy between the theoretical and simulated CW curves is due to the anisotropic distribution of pulsars used for the simulation (see the [Appendix](#)). Note that we show the results in 20 angular bins, and the results were averaged over 100 realizations to reduce the scatter of these curves.

Figure 7 shows the mean and 1σ region of $\hat{c}_{kl}$ as a function of γ . The means follow the expected HD curve for all three simulations.⁹ We also show the theoretical variance for a purely stochastic background from Allen (2022) and for a single CW with pulsar terms from the [Appendix](#) based on results from Allen (2022). We can see that our results are in good agreement with these predictions. In particular, the GWB curve follows the prediction for an isotropic Gaussian background perfectly. The Population simulation also follows the shape of the same theoretical curve, but with a slightly lower variance at all angles. This constant shift is most likely due to its different spectral properties, which influence the normalization between the mean and the variance of the correlations (i.e., Equations (C30) and (C32) of Allen 2022). This also seems to suggest that our realistic GWB behaves more like an idealized GWB than a single CW in this regard. The CW simulations also show some agreement with the corresponding theoretical predictions. In particular, these correctly predict that a single source has less variance than a GWB at large angles but more at small angles. The deviation can be explained by the anisotropic distribution of pulsars on the sky, since they agree very well for a simulation with isotropically distributed pulsars (see the [Appendix](#)).

The 1σ bounds we see in Figure 7 include both pulsar and cosmic variance. To reduce the pulsar variance, one can take the mean correlation in a set of angular separation bins. In the limit of a large number of pulsars, such an averaging should

⁹ This is expected since we are showing the mean over many realizations, which removes the cosmic variance from the mean curves. For a given realization, we expect the mean to be different from the HD curve by an amount predicted by the cosmic variance.




Figure 8. Variance of the mean correlations over 100 realizations. We show results for an idealized GWB, a single CW source, and a realistic population-based GWB. Besides the same theoretical curves as in Figure 7, we also show the theoretical variance from cosmic variance only (Allen 2022). Note the different vertical scale compared to Figure 7.

remove the pulsar variance, so the only source of variance is the cosmic variance (assuming noiseless observations). To test this scenario, we analyzed the same simulations as above, but instead of calculating the variance of individual correlations, we first took the mean of the correlations in each realization and then calculated the variance of those means over all realizations. The results are shown in Figure 8, where we used 20 equal-sized angular bins for the averaging. We also show the same theoretically expected variances as in Figure 7. In addition, solid gray lines indicate the cosmic variance (Equation (G11) in Allen 2022). As expected, the variance of the correlations is drastically reduced by the binning and averaging procedure. However, the variance is still larger than the cosmic variance at some angular separations. This is due to the finite number of pulsar pairs available for averaging in our simulations. Increasing the bin sizes used for the averaging brings these lines closer to the theoretical curves at the cost of losing angular resolution. Also note that the results from our three simulations are all consistent with each other within their statistical uncertainty in this averaged representation. In this exercise we weighted all pulsar pairs equally in the averaging. Allen & Romano (2022) derived the optimal way of reducing the pulsar variance, which does not use equal weights for different pulsar pairs.

Figures 7 and 8 only show the first and second moments of the distributions of correlations. However, one advantage of our direct simulation approach over analytical calculations is that we can easily produce the full distribution of correlations. To that end, we show the distribution of the deviation from the HD curve both with and without binning in Figure 9. We can see that the binned correlations show a consistent distribution for all three of our simulations. This is not surprising given that their means and variances are indistinguishable, as we have seen in Figure 8. As expected, the individual correlations show




Figure 9. Histogram of deviations from the HD curve for all three simulations we performed. The bottom panel shows the mean and 90% credible interval for the distribution of all individual correlation values. The top panel shows the distribution of mean correlation values calculated in each angular bin and each realization.

a significantly wider distribution. We can also see that while the population-based and idealized GWB simulations show almost identical distributions, the histogram is significantly different for the simulations with an individual CW signal. As was previously pointed out in Cornish & Sesana (2013), the overall distribution of correlation deviations has a heavier tail for a stochastic background than for a single CW source. Also note that the distribution of individual correlations shows significant non-Gaussian features, as is expected based on analytical calculations (see footnote 40 in Allen 2022).

4. Prospects for Distinguishing Individual Binaries

Our method of modeling realistic backgrounds is also uniquely suitable to study detection prospects of individual binaries in the presence of a stochastic background. By setting $M = 1$ in Equations (4) and (5), we can separate the single brightest binary and the rest of the GWB in each bin. Thus, we can take into account the confusion noise coming from the GWB itself, along with the white and red noise in each pulsar. Then, we can calculate the expected S/Ns for those outlier sources as (see, e.g., Equation (219) of Di Matteo et al. 2019)

$$S/N = (s(t; \theta) | s(t; \theta))^{1/2}, \quad (15)$$

where $s(t; \theta)$ is the CW signal in question, and we define the inner product on residuals as

$$(a|b) = a^T C^{-1} b, \quad (16)$$

where C is the noise covariance matrix that takes into account white and red noise in each pulsar and a stochastic background




Figure 10. GWB spectrum in a particular realization (blue histogram) and the brightest binary in each bin (blue circles). The binary we focus our attention on is indicated by the orange circle.

based on Equation (5) with $M = 1$. We calculate inner products using the ENTERPRISE software package (Ellis et al. 2020).

4.1. Variance of S/N over Extrinsic Parameters

The S/N can have a significant variance even when we fix the intrinsic parameters of the source. These are due to effects of extrinsic parameters: inclination angle, polarization angle, initial phase, and two angles describing the sky location of the source. To explore the effect of these on S/N, we focus on a particular binary with fixed amplitude/distance, chirp mass, and GW frequency and calculate the distribution of S/Ns over extrinsic parameters. Figure 10 shows this binary on the h_c - f_{obs} plane (orange circle). We also show the GWB spectrum for this particular realization (blue histogram) and the single brightest binary in each frequency bin (blue circles). Note that we chose our binary to be the one that produces the highest S/Ns in this realization. We use a simulated data set based on the properties of pulsars in the NANOGrav 12.5 yr data set (Alam et al. 2021), and we extend the observation time to 15 yr, as was done in Pol et al. (2021).

Figure 11 shows the distribution of the median S/N over the sky for this particular source. We also show the location of the pulsars in our simulated array as red stars. The sizes of the stars are scaled by a rough estimate of the sensitivity of the pulsars defined as

$$\text{Sensitivity} \propto \left(\sum_{i=1}^{N_{\text{obs}}} \frac{1}{(\Delta t_i)^2} \right)^{1/2}, \quad (17)$$

where Δt_i is the nominal time of arrival (TOA) error of the i th observation. Note that this simple expression reproduces the expected scaling both with the number of observations ($\sim \sqrt{N_{\text{obs}}}$) and with the TOA errors ($\sim 1/\Delta t_i$). We can see in Figure 11 that the sensitivity on the sky shows a dipolar structure, where one gets S/Ns almost a factor of four higher in one direction than the antipodal direction. This is in agreement with the highly anisotropic upper limits found by the search for individual binaries in the NANOGrav 11 yr data set (Aggarwal et al. 2019). The reason for this is the highly anisotropic distribution of NANOGrav pulsars, which tend to be concentrated around the galactic center.




Figure 11. Distribution of the median S/N over the sky for a particular source. Red stars indicate the locations of pulsars in our array (based on the NANOGrav 12.5 yr data set; see Alam et al. 2021), with their sizes corresponding to their sensitivity as defined in Equation (17). The black cross indicates the sky pixel with the highest median S/N.




Figure 12. Distribution of S/Ns for a particular source. The top panel shows the overall distribution (gray), along with histograms restricted to the half of the sky centered on the most sensitive sky location (red) and the antipodal point (blue). The bottom panel shows how S/N correlates with the orbital inclination (ι). The coloring indicates each source’s angular distance from the most sensitive sky location (see black cross in Figure 11).

The top panel of Figure 12 shows the distribution of S/Ns for this particular source, marginalized over all external parameters. There is a significant variance in the S/N, with the 95% credible interval ranging from 0.8 to 5.1. We also show the distribution restricted to the more and less sensitive hemisphere, which we define relative to the most sensitive sky pixel marked with a black




Figure 13. Location of the loudest sources on the h_c - f_{obs} plane over 500 realizations of the binary population with 100 random extrinsic parameters each (50,000 in total). The color of each circle indicates the S/N of the loudest source in the particular realization. Note that the color scale was capped at S/N = 5 for better visibility. We also show the GWB spectra for each realization in blue. Sources can have significant S/Ns even if they are below the characteristic strain spectrum for the GWB if they have favorable extrinsic parameters (e.g., sky locations and inclinations).

cross in Figure 11. This shows that the sky location significantly contributes to the variance of the overall distribution. While most other extrinsic parameters do not show a clear correlation with S/N, the well-known effect of the inclination angle (ι) appears here as well. This is evident from the strong dependence of the GW amplitude on ι (see, e.g., Equation (13)–(14) of Aggarwal et al. 2019). The bottom panel of Figure 12 shows the two-dimensional distribution of $\cos \iota$ and S/N. Sources with face-on orientation ($\cos \iota = \pm 1$) produce significantly higher S/N values compared to edge-on systems ($\cos \iota = 0$). We also color-code the points on this panel with the source’s angular distance from the most sensitive sky location marked in Figure 11. We can see that orbital inclination and sky location represent the majority of the S/N variance. This means that PTAs are more likely to first see an individual SMBHB with approximately face-on orientation, located in the part of the sky toward which our PTA is most sensitive.

4.2. Expected Properties of the Highest-S/N SMBHB

We are interested in the expected properties of the CW signal that will first be detectable. To find those, we create 50,000 realizations of a realistic GWB made up from 500 realizations of the SMBHB population with 100 different random extrinsic parameters each. We use a simulated PTA based on the NANOGrav 12.5 yr data set extended to a 15 yr time span (Pol et al. 2021). Thus, the following results can be treated as pessimistic predictions for the upcoming NANOGrav 15 yr data set (M. F. Alam et al. 2022, in preparation), since we do not take into account the improved timing precision and the addition of new pulsars. For each realization we find the individual binary with the highest S/N looking through all the frequency bins. Figure 13 shows the location of these loudest sources on the h_c - f_{obs} plane,




Figure 14. Distribution of the detector-frame GW frequency and chirp mass, luminosity distance, inclination angle, and S/N of the loudest source over 500 binary population realizations with 100 realizations each (50,000 in total). Dashed lines and quoted values correspond to the 5th/50th/95th percentiles. Gray histograms show all realizations, while red histograms show only those where the S/N is larger than 5 ($\sim 8\%$ of all realizations).

color-coded with their corresponding S/N. We also show the spectrum of the GWB for all realizations. We can see that the loudest sources are concentrated at moderate frequencies. The lack of high-S/N sources at higher frequencies is due to the white noise in the data set, which has $h_c \sim f_{\text{obs}}^{3/2}$. Note that a large fraction of the loudest sources lie below the characteristic strain spectrum of the GWB, which acts as a noise source here. This is possible owing to the fact that the S/N depends not only on h_c but also on extrinsic parameters as we have seen in Section 4.1. Selecting the loudest source over all frequencies preferentially selects sources with the most favorable sky locations, inclinations, etc., thus resulting in significant S/Ns even if the source has $h_c < h_c^{\text{GWB}}$.

Figure 14 shows the one- and two-dimensional marginal distributions of the detector-frame GW frequency and chirp mass, luminosity distance, inclination angle, and S/N for the loudest sources over 50,000 realizations. The median f_{obs} is ~ 6 nHz, the median chirp mass is $\sim 5 \times 10^9 M_\odot$, and the median luminosity distance is ~ 1.5 Gpc. This suggests that we will most likely first see a CW source at moderate frequencies with a very high chirp mass at a considerable distance. This is in agreement with the findings of Rosado et al. (2015) and Kelley et al. (2017b). The distribution of $\cos i$ is significantly different from the flat distribution corresponding to isotropic inclination distribution. This is a well-known selection effect

due to the increased S/N for face-on systems (see Figure 12). The probability of the loudest source being within 30° of face-on is $\sim 26\%$ for the whole population and $\sim 38\%$ for the $S/N > 5$ subpopulation (compare with $\sim 13\%$ for the isotropic distribution).

We can also see in Figure 14 that the median S/N is 1.8, which is not expected to be detectable. However, there is about a 6% chance to get a source with $S/N > 5$. Depending on the details of the detection algorithm used, those might be detectable, which is an interesting prospect for the search for individual binaries in the upcoming NANOGrav 15 yr data set (M. F. Alam et al. 2022, in preparation). We also show histograms for these high-S/N sources in red. We can see that they largely follow the same distribution for f_{obs} , $\mathcal{M}$, and d_L as all realizations. The same analysis extended to a 20 yr time span yields similar f_{obs} , $\mathcal{M}$, and d_L values; a median S/N of 4.5; and $S/N > 5$ in about 41% of the realizations.

Note that another reason why these predictions are pessimistic is that we only check the binary with the highest h_c in each bin. In principle, binaries with lower h_c can end up producing the highest S/N signals. This means that by not taking those into account, our S/N distributions are biased low. On the other hand, these results do not marginalize over the uncertainty in the astrophysical models used for the SMBHB populations, which could have a large effect on these results. We will incorporate these additional details in a future study to provide more robust predictions for future PTA detection prospects.

5. Conclusion and Future Work

In this paper we presented a new approach to efficiently simulating realistic PTA data sets by modeling the signal as a combination of an isotropic Gaussian GWB and a few of the brightest binaries modeled individually. This produces data sets consistent with the naive method of directly modeling all binaries in a fraction of the time. We used this simulation technique to explore various properties of realistic PTA data sets. We show that the data sets are dominated by a small number of binaries at high frequencies, resulting in the well-known lack of GW power compared to a simple power-law model (see, e.g., Sesana et al. 2008). We test these data sets for statistical isotropy and find that they can be made isotropic by removing the few brightest binaries at all except the highest frequencies. We calculate the mean and the variance of the spatial correlations in our simulated data sets and find good agreement with analytical results presented in Allen (2022).

Our methodology also allows for calculating S/Ns of any binary in our data sets. We use that to calculate the distribution of the S/N of the brightest source over realizations for a simulated PTA based on the NANOGrav 12.5 yr pulsars with time spans extended to 15 yr. These calculations account for both pulsar noise and the confusion noise from the GWB. We find that the brightest binary tends to have moderate GW frequency (few times the inverse of the observational time span), high chirp mass, large distance, and nearly face-on orientation. The median S/N of the brightest source is 1.8, and about 6% of the realizations produce a source with S/N higher than 5. These S/N values are pessimistic estimates for the upcoming NANOGrav 15 yr data set (M. F. Alam et al. 2022, in preparation), since they do not take into account any potential improvements of the timing model solutions or the addition of new pulsars relative to the 12.5 yr data set.

However, they rely on a particular model of the SMBHB population, which might introduce biases. If we further increase the observing time span to 20 yr, the fraction of realizations producing a binary with S/N greater than 5 increases to about 41%, and the median S/N increases to 4.5. We also calculate the S/N distribution for a given binary over different values of their extrinsic parameters (sky location, inclination and polarization angle, initial phase). We find that these parameters can have a significant effect on the S/N. The sky location and the inclination angle are particularly impactful parameters, as they can change the S/N by a factor of 2–4.

The methods presented in this paper will continue to be useful tools to understand the interplay between the stochastic background and individual sources. In the future, we plan to incorporate different models of the SMBHB population, allowing us to produce simulated data sets using different assumptions about SMBHB formation and evolution. The software used for this paper was made to be compatible with the `holodeck`¹⁰ software package (L. Z. Kelley 2022, in preparation), allowing us to use any SMBHB population model developed there. That will also allow us to make our detection prospect predictions marginalized over astrophysical uncertainties. In addition, we also plan to run PTA detection pipelines on realistic data sets produced by the methods presented here. In particular, analyzing such data sets with the `QuickCW`¹¹ pipeline (Bécsy et al. 2022) would allow us to make more accurate predictions than the simple S/N calculations presented in this paper. We also plan to analyze such data sets with `BayesHopper` to investigate how PTAs will be able to detect multiple individual sources in the future (Bécsy & Cornish 2020).

The authors thank Stephen Taylor and Nihan Pol for providing access to software used in Pol et al. (2021) to create simulated data sets based on measured noise properties of NANOGrav pulsars. We thank Chiara Mingarelli and Alberto Sesana for feedback on the manuscript. We also thank Bruce Allen and Joseph Romano for useful discussions about the variance of HD correlations. We appreciate the support of the NSF Physics Frontiers Center award PFC-1430284 and the NSF Physics Frontiers Center award PFC-2020265.

Software: `healpy` (Zonca et al. 2019), `HEALPix` (Górski et al. 2005), `ENTERPRISE` (Ellis et al. 2020), `corner` (Foreman-Mackey 2016), `Numba` (Lam et al. 2015, 2022).

Appendix Variance of Spatial Correlations with Different Normalizations

In this appendix we explore how the variance of the HD correlations from a single binary depends on the way we normalize the correlation values. We reproduce some results of Allen (2022) and derive formulae for the case with a different normalization used in Section 3.3. We compare these results with simulations for a CW signal in a simulated PTA with an isotropic distribution of pulsars. Note that we follow the convention that the HD curve at zero separation is $1/2$, while Allen (2022) normalizes to $1/3$. Thus, our results need to be multiplied by $2/3$ when comparing with results in Allen (2022).

¹⁰ Publicly available at <https://github.com/nanograv/holodeck>.

¹¹ Publicly available at <https://github.com/bencebecsy/QuickCW>.

A.1. Without Pulsar Terms

From Equation (A31) of Allen (2022) the mean HD correlation for a binary with a given inclination (ι) and GW amplitude ($\mathcal{A}$) is

$$\mu = \frac{\mathcal{A}^2}{16} [1 + 6 \cos^2 \iota + \cos^4 \iota] \mu_u(\gamma), \quad (\text{A1})$$

where $\mu_u(\gamma)$ is the mean HD curve (see Equation (D29) in Allen 2022). Neglecting pulsar terms, the variance is

$$\begin{aligned} \sigma^2 &= \frac{\mathcal{A}^4}{256} [1 + 6 \cos^2 \iota + \cos^4 \iota]^2 \sigma_u^2(\gamma) \\ &+ \frac{\mathcal{A}^4}{512} [\sin^8 \iota] \sigma_c^2(\gamma), \end{aligned} \quad (\text{A2})$$

where σ_u and σ_c are defined in Equations (D37) and (F2) in Allen (2022), respectively.

We can average these over $\cos \iota \in [-1, 1]$ to get (see Equation (A32) of Allen 2022)

$$\mu_{\text{average}} = \frac{\mathcal{A}^2}{5} \mu_u(\gamma) \quad (\text{A3})$$

and

$$\sigma_{\text{average}}^2 = \frac{71\mathcal{A}^4}{1260} \sigma_u^2(\gamma) + \frac{\mathcal{A}^4}{1260} \sigma_c^2(\gamma). \quad (\text{A4})$$

We normalize these so that $\mu_{\text{average}} = \mu_u$, which implies $\mathcal{A}^2 = 5$ and

$$\sigma_{\text{average}} = \sqrt{\frac{355}{252} \sigma_u^2(\gamma) + \frac{5}{252} \sigma_c^2(\gamma)}. \quad (\text{A5})$$

This is appropriate if we normalize the correlation values globally over all realizations. However, if we normalize the correlations realization by realization, then we need to normalize Equation (A1) before averaging. We thus require $\mathcal{A}^2/16[1 + 6 \cos^2 \iota + \cos^4 \iota] = 1$, which implies

$$\sigma^2 = \sigma_u^2(\gamma) + \frac{\sin^8 \iota}{2[1 + 6 \cos^2 \iota + \cos^4 \iota]^2} \sigma_c^2(\gamma). \quad (\text{A6})$$

If we average this over $\cos \iota \in [-1, 1]$, we get

$$\sigma_{\text{average}} = \sqrt{\sigma_u^2(\gamma) + \frac{1}{4} \alpha \sigma_c^2(\gamma)}, \quad (\text{A7})$$

where $\alpha = 4 + \pi - 3\pi/\sqrt{2}$.

Figure 15 shows the comparison of simulation results with Equations (A5) and (A7). The simulation normalized per realization seems to match Equation (A7) really well. The one normalized globally generally follows Equation (A5) but deviates at small and large angles.

A.2. With Pulsar Terms

If we include pulsar terms, the mean correlation is unchanged, but the variance becomes (see Equation (B10) of




Figure 15. Variance of HD correlations from a CW simulation without pulsar terms and a simulated PTA with isotropically distributed pulsars. We also show the theoretical variance with two different averaging procedures described in Section A.1.

Allen 2022)

$$\begin{aligned} \sigma^2 &= \frac{\mathcal{A}^4}{256} [1 + 6 \cos^2 \iota + \cos^4 \iota]^2 \sigma_u^2(\gamma) \\ &- \frac{3\mathcal{A}^4}{512} [\sin^8 \iota] \sigma_p^2(\gamma) \\ &+ \mathcal{A}^4 \left(\frac{3}{512} [1 + 6 \cos^2 \iota + \cos^4 \iota]^2 + \frac{5}{1024} [\sin^8 \iota] \right) \\ &\times \sigma_c^2(\gamma), \end{aligned} \quad (\text{A8})$$

where σ_p is defined in Equation (E8) of Allen (2022).

Averaging this over $\cos \iota \in [-1, 1]$ gives (see Equation (B11) of Allen 2022)

$$\sigma^2 = \frac{71\mathcal{A}^4}{1260} \sigma_u^2(\gamma) - \frac{3\mathcal{A}^4}{1260} \sigma_p^2(\gamma) + \frac{109\mathcal{A}^4}{1260} \sigma_c^2(\gamma). \quad (\text{A9})$$

Plugging in $\mathcal{A}^4 = 25$ gives

$$\sigma_{\text{average}} = \sqrt{\frac{355}{252} \sigma_u^2(\gamma) - \frac{5}{84} \sigma_p^2(\gamma) + \frac{545}{252} \sigma_c^2(\gamma)}. \quad (\text{A10})$$

This is the average standard deviation if we normalize correlations globally, over many realizations. Alternatively, we can express $\mathcal{A}$ from Equation (A1) and plug into Equation (A8) before averaging. This gives

$$\begin{aligned} \sigma^2 &= \sigma_u^2(\gamma) - \frac{3 \sin^8 \iota}{2[1 + 6 \cos^2 \iota + \cos^4 \iota]^2} \sigma_p^2(\gamma) \\ &+ \left(\frac{3}{2} + \frac{5 \sin^8 \iota}{4[1 + 6 \cos^2 \iota + \cos^4 \iota]^2} \right) \sigma_c^2(\gamma). \end{aligned} \quad (\text{A11})$$

Averaging over $\cos \iota \in [-1, 1]$ gives

$$\sigma_{\text{average}} = \sqrt{\sigma_u^2(\gamma) - \frac{3}{4} \alpha \sigma_p^2(\gamma) + \left(\frac{3}{2} + \frac{5}{8} \alpha \right) \sigma_c^2(\gamma)}. \quad (\text{A12})$$

Figure 16 shows simulation results (both with global normalization and per realization normalization) and




Figure 16. Variance of HD correlations from a CW simulation with pulsar terms and a simulated PTA with isotropically distributed pulsars. We also show the theoretical variance with two different averaging procedures described in Section A.2.

Equations (A10) and (A12). We can see that both match well with the appropriate simulation.

ORCID iDs

Bence Bécsy <https://orcid.org/0000-0003-0909-5563>

Neil J. Cornish <https://orcid.org/0000-0002-7435-0869>

Luke Zoltan Kelley <https://orcid.org/0000-0002-6625-6450>

References

Aggarwal, K., Arzoumanian, Z., Baker, P. T., et al. 2019, *ApJ*, **880**, 116
 Alam, M. F., Arzoumanian, Z., Baker, P. T., et al. 2021, *ApJS*, **252**, 4
 Ali-Haïmoud, Y., Smith, T. L., & Mingarelli, C. M. F. 2020, *PhRvD*, **102**, 122005
 Allen, B. 2022, *PhRvD*, in press (arXiv:2205.05637)
 Allen, B., & Romano, J. D. 2022, arXiv:2208.07230
 Antoniadis, J., Arzoumanian, Z., Babak, S., et al. 2022, *MNRAS*, **510**, 4873
 Arzoumanian, Z., Baker, P. T., Blumer, H., et al. 2020, *ApJL*, **905**, L34
 Arzoumanian, Z., Baker, P. T., Blumer, H., et al. 2021, *PhRvL*, **127**, 251302
 Bécsy, B., & Cornish, N. J. 2020, *CQGrA*, **37**, 135011
 Bécsy, B., Cornish, N. J., & Digman, M. C. 2022, *PhRvD*, **105**, 122003
 Blasi, S., Brdar, V., & Schmitz, K. 2021, *PhRvL*, **126**, 041305

Burke-Spolaor, S., Taylor, S. R., Charisi, M., et al. 2019, *A&ARv*, **27**, 5
 Chen, S., Caballero, R. N., Guo, Y. J., et al. 2021, *MNRAS*, **508**, 4970
 Cornish, N. J., & Romano, J. D. 2015, *PhRvD*, **92**, 042001
 Cornish, N. J., & Sampson, L. 2016, *PhRvD*, **93**, 104047
 Cornish, N. J., & Sesana, A. 2013, *CQGrA*, **30**, 224005
 Cornish, N. J., & van Haasteren, R. 2014, arXiv:1406.4511
 De Luca, V., Franciolini, G., & Riotto, A. 2021, *PhRvL*, **126**, 041303
 Di Matteo, T., King, A., & Cornish, N. J. 2019, *Black Hole Formation and Growth* (Berlin: Springer)
 Ellis, J. A., Vallisneri, M., Taylor, S. R., & Baker, P. T. 2020, ENTERPRISE: Enhanced Numerical Toolbox Enabling a Robust Pulsar Inference Suite, Zenodo, doi:10.5281/zenodo.4059814
 Enoki, M., Inoue, K. T., Nagashima, M., & Sugiyama, N. 2004, *ApJ*, **615**, 19
 Foreman-Mackey, D. 2016, *JOSS*, **1**, 24
 Goncharov, B., Shannon, R. M., Reardon, D. J., et al. 2021, *ApJL*, **917**, L19
 Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, **622**, 759
 Hajian, A., & Souradeep, T. 2003, *ApJL*, **597**, L5
 Hajian, A., & Souradeep, T. 2004, arXiv:astro-ph/0501001
 Hajian, A., Souradeep, T., & Cornish, N. 2005, *ApJL*, **618**, L63
 Hellings, R. W., & Downs, G. S. 1983, *ApJL*, **265**, L39
 Hotinli, S. C., Kamionkowski, M., & Jaffe, A. H. 2019, *OJAp*, **2**, 8
 Kaiser, A. R., Pol, N. S., McLaughlin, M. A., et al. 2022, *ApJ*, **938**, 115
 Kelley, L. Z., Blecha, L., & Hernquist, L. 2017a, *MNRAS*, **464**, 3131
 Kelley, L. Z., Blecha, L., Hernquist, L., Sesana, A., & Taylor, S. R. 2017b, *MNRAS*, **471**, 4508
 Kelley, L. Z., Blecha, L., Hernquist, L., Sesana, A., & Taylor, S. R. 2018, *MNRAS*, **477**, 964
 Lam, S. K., Pitrou, A., & Seibert, S. 2015, in Proc. Second Workshop on the LLVM Compiler Infrastructure in HPC, LLVM '15 (New York: ACM)
 Lam, S. K., stuartarchibald, Pitrou, A., et al. 2022, numba/numba: Version 0.55.0, Zenodo, doi:10.5281/zenodo.5847553
 Mingarelli, C. M. F., Lazio, T. J. W., Sesana, A., et al. 2017, *NatAs*, **1**, 886
 Mingarelli, C. M. F., Sidery, T., Mandel, I., & Vecchio, A. 2013, *PhRvD*, **88**, 062005
 Phinney, E. S. 2001, arXiv:astro-ph/0108028
 Pol, N., Taylor, S. R., & Romano, J. D. 2022, *ApJ*, **940**, 173
 Pol, N. S., Taylor, S. R., Zoltan Kelley, L., et al. 2021, *ApJL*, **911**, L34
 Rajagopal, M., & Romani, R. W. 1995, *ApJ*, **446**, 543
 Romano, J. D., Hazboun, J. S., Siemens, X., & Archibald, A. M. 2021, *PhRvD*, **103**, 063027
 Rosado, P. A., Sesana, A., & Gair, J. 2015, *MNRAS*, **451**, 2417
 Sesana, A., Haardt, F., Madau, P., & Volonteri, M. 2004, *ApJ*, **611**, 623
 Sesana, A., Vecchio, A., & Colacino, C. N. 2008, *MNRAS*, **390**, 192
 Sesana, A., Vecchio, A., & Volonteri, M. 2009, *MNRAS*, **394**, 2255
 Taylor, S. R. 2021, arXiv:2105.13270
 Taylor, S. R., & Gair, J. R. 2013, *PhRvD*, **88**, 084001
 Taylor, S. R., van Haasteren, R., & Sesana, A. 2020, *PhRvD*, **102**, 084039
 Vagnozzi, S. 2021, *MNRAS*, **502**, L11
 van Haasteren, R., Levin, Y., Janssen, G. H., et al. 2011, *MNRAS*, **414**, 3117
 Vogelsberger, M., Genel, S., Springel, V., et al. 2014, *MNRAS*, **444**, 1518
 Zonca, A., Singer, L., Lenz, D., et al. 2019, *JOSS*, **4**, 1298