



Astrophysics Milestones for Pulsar Timing Array Gravitational-wave Detection

Nihan S. Pol^{1,2,3}, Stephen R. Taylor¹, Luke Zoltan Kelley⁴, Sarah J. Vigeland⁵, Joseph Simon^{6,7}, Siyuan Chen^{8,9,10}, Zaven Arzoumanian¹¹, Paul T. Baker¹², Bence Bécsy¹³, Adam Brazier¹⁴, Paul R. Brook^{2,3}, Sarah Burke-Spolaor^{2,3,15}, Shami Chatterjee¹⁴, James M. Cordes¹⁴, Neil J. Cornish¹³, Fronefield Crawford¹⁶, H. Thankful Cromartie^{14,42}, Megan E. DeCesar^{37,40,41}, Paul B. Demorest¹⁸, Timothy Dolch^{19,20}, Elizabeth C. Ferrara^{38,39,21}, William Fiore^{2,3}, Emmanuel Fonseca²², Nathan Garver-Daniels^{2,3}, Deborah C. Good²³, Jeffrey S. Hazboun^{24,41}, Ross J. Jennings¹⁴, Megan L. Jones⁵, Andrew R. Kaiser^{2,3}, David L. Kaplan⁵, Joey Shapiro Key²⁴, Michael T. Lam^{25,26}, T. Joseph W. Lazio⁶, Jing Luo²⁷, Ryan S. Lynch²⁸, Dustin R. Madison^{2,3,41}, Alexander McEwen⁵, Maura A. McLaughlin^{2,3}, Chiara M. F. Mingarelli^{29,30}, Cherry Ng³¹, David J. Nice¹⁷, Timothy T. Pennucci^{32,33,41}, Scott M. Ransom³², Paul S. Ray³⁴, Brent J. Shapiro-Albert^{2,3}, Xavier Siemens^{5,35}, Ingrid H. Stairs²³, Daniel R. Stinebring³⁶, Joseph K. Swiggum^{17,41}, Michele Vallisneri⁶, Haley Wahl^{2,3}, and Caitlin A. Witt^{2,3}

The NANOGrav Collaboration

¹ Department of Physics and Astronomy, Vanderbilt University, 2301 Vanderbilt Place, Nashville, TN 37235, USA; nihan.pol@nanograv.org

² Department of Physics and Astronomy, West Virginia University, P.O. Box 6315, Morgantown, WV 26506, USA

³ Center for Gravitational Waves and Cosmology, West Virginia University, Chestnut Ridge Research Building, Morgantown, WV 26505, USA

⁴ Center for Interdisciplinary Exploration and Research in Astrophysics (CIERA), Northwestern University, Evanston, IL 60208, USA

⁵ Center for Gravitation, Cosmology and Astrophysics, Department of Physics, University of Wisconsin-Milwaukee, P.O. Box 413, Milwaukee, WI 53201, USA

⁶ Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109, USA

⁷ Department of Astrophysical and Planetary Sciences, University of Colorado, Boulder, CO 80309, USA

⁸ Laboratoire de Physique et Chimie de l'Environnement et de l'Espace, LPC2E UMR7328, Université d'Orléans, CNRS, F-45071 Orleans, France

⁹ Station de Radioastronomie de Nancay, Observatoire de Paris, Université PSL, CNRS, Université d'Orléans, F-18330 Nancay, France

¹⁰ FEMTO-ST Institut de recherche, Department of Time and Frequency, UBFC and CNRS, ENSMM, F-25030 Besancon, France

¹¹ X-Ray Astrophysics Laboratory, NASA Goddard Space Flight Center, Code 662, Greenbelt, MD 20771, USA

¹² Department of Physics and Astronomy, Widener University, One University Place, Chester, PA 19013, USA

¹³ Department of Physics, Montana State University, Bozeman, MT 59717, USA

¹⁴ Cornell Center for Astrophysics and Planetary Science and Department of Astronomy, Cornell University, Ithaca, NY 14853, USA

¹⁵ CIFAR Azrieli Global Scholars program, CIFAR, Toronto, Canada

¹⁶ Department of Physics and Astronomy, Franklin & Marshall College, P.O. Box 3003, Lancaster, PA 17604, USA

¹⁷ Department of Physics, Lafayette College, Easton, PA 18042, USA

¹⁸ National Radio Astronomy Observatory, 1003 Lopezville Road, Socorro, NM 87801, USA

¹⁹ Department of Physics, Hillsdale College, 33 E. College Street, Hillsdale, MI 49242, USA

²⁰ Eureka Scientific, Inc. 2452 Delmer Street, Suite 100, Oakland, CA 94602-3017, USA

²¹ NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA

²² Department of Physics, McGill University, 3600 University Street, Montreal, QC H3A 2T8, Canada

²³ Department of Physics and Astronomy, University of British Columbia, 6224 Agricultural Road, Vancouver, BC V6T 1Z1, Canada

²⁴ University of Washington Bothell, 18115 Campus Way NE, Bothell, WA 98011, USA

²⁵ School of Physics and Astronomy, Rochester Institute of Technology, Rochester, NY 14623, USA

²⁶ Laboratory for Multiwavelength Astrophysics, Rochester Institute of Technology, Rochester, NY 14623, USA

²⁷ Department of Astronomy & Astrophysics, University of Toronto, 50 Saint George Street, Toronto, ON M5S 3H4, Canada

²⁸ Green Bank Observatory, P.O. Box 2, Green Bank, WV 24944, USA

²⁹ Center for Computational Astrophysics, Flatiron Institute, 162 5th Avenue, New York, New York, 10010, USA

³⁰ Department of Physics, University of Connecticut, 196 Auditorium Road, U-3046, Storrs, CT 06269-3046, USA

³¹ Dunlap Institute for Astronomy and Astrophysics, University of Toronto, 50 St. George Street, Toronto, ON M5S 3H4, Canada

³² National Radio Astronomy Observatory, 520 Edgemont Road, Charlottesville, VA 22903, USA

³³ Institute of Physics, Eötvös Loránd University, Pázmány P.s. 1/A, 1117 Budapest, Hungary

³⁴ Space Science Division, Naval Research Laboratory, Washington, DC 20375-5352, USA

³⁵ Department of Physics, Oregon State University, Corvallis, OR 97331, USA

³⁶ Department of Physics and Astronomy, Oberlin College, Oberlin, OH 44074, USA

³⁷ Department of Physics and Astronomy, George Mason University, Fairfax, VA 22030, resident at U.S. Naval Research Laboratory, Washington, D.C. 20375, USA

³⁸ Department of Astronomy, University of Maryland, College Park, MD 20742, USA

³⁹ Center for Research and Exploration in Space Science and Technology, NASA/GSFC, Greenbelt, MD 20771, USA

Received 2021 February 19; revised 2021 March 24; accepted 2021 March 25; published 2021 April 26

Abstract

The NANOGrav Collaboration reported strong Bayesian evidence for a common-spectrum stochastic process in its 12.5 yr pulsar timing array data set, with median characteristic strain amplitude at periods of a year of $A_{\text{yr}} = 1.92^{+0.75}_{-0.55} \times 10^{-15}$. However, evidence for the quadrupolar Hellings & Downs interpulsar correlations, which are characteristic of gravitational-wave signals, was not yet significant. We emulate and extend the NANOGrav data set, injecting a wide range of stochastic gravitational-wave background (GWB) signals that encompass a variety of amplitudes and spectral shapes, and quantify three key milestones. (I) Given the amplitude

⁴⁰ AAAS, STPF, ORISE Fellow hosted by the US Department of Energy.

⁴¹ NANOGrav Physics Frontiers Center Postdoctoral Fellow.

⁴² NASA Hubble Fellowship Program Einstein Postdoctoral Fellow.

measured in the 12.5 yr analysis and assuming this signal is a GWB, we expect to accumulate robust evidence of an inter-pulsar-correlated GWB signal with 15–17 yr of data, i.e., an additional 2–5 yr from the 12.5 yr data set. (II) At the initial detection, we expect a fractional uncertainty of 40% on the power-law strain spectrum slope, which is sufficient to distinguish a GWB of supermassive black hole binary origin from some models predicting more exotic origins. (III) Similarly, the measured GWB amplitude will have an uncertainty of 44% upon initial detection, allowing us to arbitrate between some population models of supermassive black hole binaries. In addition, power-law models are distinguishable from those having low-frequency spectral turnovers once 20 yr of data are reached. Even though our study is based on the NANOGrav data, we also derive relations that allow for a generalization to other pulsar timing array data sets. Most notably, by combining the data of individual arrays into the International Pulsar Timing Array, all of these milestones can be reached significantly earlier.

Unified Astronomy Thesaurus concepts: Gravitational waves (678); Gravitational wave astronomy (675); Millisecond pulsars (1062); Pulsar timing method (1305); Astronomical simulations (1857)

1. Introduction

Supermassive black hole binaries (SMBHBs) are expected to form following the mergers of massive galaxies during the buildup of hierarchical structure in Λ CDM (Burke-Spolaor et al. 2019, and references therein). No observationally confirmed SMBHBs at subparsec separations are known, but a growing number of strong candidates have been identified (Graham et al. 2015; Charisi et al. 2016; Liu et al. 2019; Chen et al. 2020). If SMBHBs are indeed able to form, and also to reach \sim milliparsec separations, then they will become the sources of the strongest gravitational waves (GWs) in the universe. The stochastic superposition of GWs produced by SMBHBs across cosmic time is expected to produce an aggregate background signal with an approximate characteristic strain spectrum of the form $h_c(f) = A_{\text{GWB}}(f/f_{\text{yr}})^\alpha$ with $\alpha = -2/3$ for a continuous population of circular and purely GW-driven binary systems (Rajagopal & Romani 1995; Jaffe & Backer 2003; Sesana et al. 2008). At low frequencies ($f \lesssim 10^{-8}$ Hz), interactions between SMBHBs and their ambient galactic environments can turn over or flatten the spectrum ($\alpha > -2/3$; Sampson et al. 2015; Kelley et al. 2017b; Taylor et al. 2017). At high frequencies ($f \gtrsim 10^{-8}$ Hz), rapid binary hardening can lead to a low occupation of spectral bins that steepens the spectrum ($\alpha < -2/3$) and may introduce spikes from individual sources (Sesana et al. 2008). Despite these effects, the stochastic GW background (GWB) from SMBHBs is expected to be the dominant source of GWs in the nanohertz frequency regime.

Pulsar Timing Arrays (PTAs; Sazhin 1978; Detweiler 1979; Foster & Backer 1990) measure the times of arrival (TOAs) of radio pulses from millisecond pulsars as a means of measuring the local spacetime curvature, and thus signs of passing GWs, analogous to how the LIGO-Virgo Collaboration (Harry & LIGO Scientific Collaboration 2010; Accadia et al. 2012) uses pairs of perpendicular laser arms. As PTAs accrue larger data sets over time, the increasing number of recorded pulses leads to improved sensitivity, especially as new pulsars are added to the arrays (Siemens et al. 2013; Taylor et al. 2016). The NANOGrav 12.5 yr data set (Alam et al. 2021) shows conclusive evidence (with Bayesian odds of $\sim 10^4:1$) that a common-spectrum low-frequency stochastic process is present across a large fraction of the 45 ms pulsars included in the analysis (Arzoumanian et al. 2020). Its properties are consistent with expectations of the strain spectrum from a stochastic GWB, but it does not yet show statistically significant quadrupolar inter-pulsar correlations, which are widely considered the definitive evidence for the GWB (Hellings & Downs 1983; Tiburzi et al. 2016). The latter are the so-called

“Hellings & Downs” (HD) correlations of GWB-induced timing delays between pairs of pulsars, which are a function only of the angular separation between the pulsars on the sky. Future data sets, such as the in-preparation NANOGrav 15 yr data set, in addition to data from the European PTA (EPTA; Desvignes et al. 2016), Australian Parkes PTA (PPTA; Kerr et al. 2020), and the International Pulsar Timing Array (IPTA; Verbiest et al. 2016; Perera et al. 2019), will reveal the nature of the signal. These data will either show spatial correlations consistent with a GWB, or other correlation signatures that suggest unmodeled noise processes affecting the PTA data.

In this paper, we forecast the evolution of GWB inference into the future, under the assumption that the signal is produced by cosmological SMBHBs. There may be alternate sources of the GWB, and there have been many recent suggestions of such interpretations for the NANOGrav 12.5 yr results (see, e.g., Blasi et al. 2021; De Luca et al. 2021; Ellis & Lewicki 2021; Vaskonen & Veermäe 2021). The scaling results that we derive here are generalizable to such other sources of the GWB. Similar analyses performed in the past were focused only on the time to detection of the GWB (e.g., Taylor et al. 2016; Rosado et al. 2015; Vigeland & Siemens 2016; Kelley et al. 2017b), and while this remains important (especially given the results in Arzoumanian et al. 2020), we consider here for the first time the evolution of the GWB parameter measurement uncertainties. The evolution of these properties is of crucial importance in determining the source and underlying physics that produces the GWB. For example, while a GWB from SMBHBs produces a spectral index of $\alpha = -2/3$, a GWB produced by primordial GWs can have a spectral index, $\alpha = -2$ or -1 (Grishchuk 2005; Lasky et al. 2016), while a GWB produced by some models of cosmic strings have $\alpha = -7/6$ (Ölmez et al. 2010). A distinction between these source models can only be made once the spectral index measurement uncertainty is small enough to exclude one (or more) of the predicted spectral indices. Similarly, different models of a given source can produce different amplitudes of the GWB, and thus knowing when we can distinguish between them is dependent on the evolution of the measurement uncertainties on the GWB amplitude.

However, for a background produced by inspiraling SMBHBs, the canonical $\alpha = -2/3$ spectral index GW strain power law emerges when assuming a continuous distribution of circular binary sources evolving purely due to GW emission (Phinney 2001). Following galaxy mergers, two SMBHBs can only become bound and reach the small separations ($\lesssim 10^{-2}$ pc) required to produce detectable GWs through environmental interactions, such as dynamical friction and stellar “slingshot”

scattering (Begelman et al. 1980). These interactions make binaries harden faster than they would due purely to GW emission, and thus the resulting GW strain can be lower than in the $-2/3$ power law, particularly at low frequencies ($f \ll 1 \text{ yr}^{-1}$). At higher frequencies, finite number effects (Sesana et al. 2008) can also cause significant deviations from a pure power law. We explore the effects of these variations in the SMBHB GWB spectrum by injecting and recovering “realistic” spectra that have been produced through full mock populations of SMBH binaries.

The paper is structured as follows. In Section 2, we describe the methods used to simulate and analyze the PTA data sets. In Section 3, we show that the Bayesian model odds ratio can be related to the frequentist signal-to-noise ratio (S/N) through a simple analytic expression for PTAs, while in Section 4, we derive a new statistic, the total S/N, that encapsulates the information in both the auto- and cross-correlation components of the PTA data. In Section 5, we describe the evolution of the GWB detection significance, parameter uncertainties, and spectral characterization of the GWB, and define three key milestones that PTAs should achieve along the way. Finally, in Section 6, we show how the timeline for achieving these milestones described can be accelerated by combining the data from individual PTAs into the IPTA.

2. Methods

2.1. Simulation Framework

We have developed a framework for realistic pulsar timing data simulation that employs the NANOGrav 12.5 yr data set (Alam et al. 2021) as its foundation, deriving observational time stamps and TOA uncertainties from all 45 pulsars therein, as well as useful metadata such as observing radio frequency and telescope. We account for pulse-phase jitter noise, as well as the usual TOA uncertainties from radiometer noise, using the maximum likelihood pulsar noise parameters measured by NANOGrav in the 12.5 yr data preparation and analysis. We epoch-average the TOA uncertainties and time stamps, which reduces the effective number of TOAs by almost an order of magnitude. We also inject intrinsic red noise in each pulsar according to its measured values in the NANOGrav 12.5 yr data set, with the caveat that we filter out a GWB-like $\alpha = -2/3$ red process from each pulsar by modeling it alongside the intrinsic pulsar noise. This is important in isolating the true intrinsic red noise in each pulsar, which may otherwise be conflated with a common process in single-pulsar noise estimation. These simulations are the first to inject red noise in each pulsar that reflects the measured NANOGrav common red noise being appropriately filtered out.

In order to forecast data accumulation beyond the existing 12.5 yr baseline, we derive distributions of the observational cadences and TOA uncertainties from the final year’s worth of data in each pulsar. We then adopt a status quo approach of drawing new observation times and uncertainties from these distributions until we reach a 20 yr array baseline. We retain the 45 pulsars from the 12.5 yr analysis and do not add any new pulsars as a function of time. This latter choice implies that the temporal growth rate of the detection statistics that we derive in this work will be relatively conservative because the addition of pulsars to the PTA is known to increase the detection significance of the GWB (Siemens et al. 2013).

Finally, we inject ten different realizations of noise and isotropic GWB signals (with spectral index $\alpha = -2/3$) for each of eight different amplitudes in the range $10^{-16} \leq A_{\text{GWB}} \leq 5 \times 10^{-15}$ into our synthesized 20 yr baseline array. We then study how the signal evolves in our data by analyzing increasingly longer slices of each simulation, from 10 yr up to the final 20 yr. Put together, our PTA simulation technique is the most sophisticated that has ever been presented in the literature.

2.2. Bayesian Data Analysis Methods

The methods used for modeling the injected GWB are the same as those used in the NANOGrav 11 yr (Arzoumanian et al. 2018) and 12.5 yr (Arzoumanian et al. 2020) analyses. Thus, we only briefly summarize the probabilistic framework for estimating the GWB parameters. In the standard detection pipeline, the GWB is modeled as a power law across frequency for the characteristic strain,

$$h_c(f) = A_{\text{GWB}} \left(\frac{f}{1 \text{ yr}^{-1}} \right)^\alpha, \quad (1)$$

where A_{GWB} is the amplitude of the GWB, f is the frequency, and $\alpha = -2/3$ is the spectral index expected for a population of inspiraling SMBHBs whose inspiral is dominated by GWs. This GWB spectrum can be expressed in terms of the timing-residual cross-spectral density between pulsars a and b as,

$$S_{\text{ab}}(f) = \Gamma_{\text{ab}} \frac{A_{\text{GWB}}^2}{12\pi^2} \left(\frac{f}{1 \text{ yr}^{-1}} \right)^{-\gamma} \text{yr}^3, \quad (2)$$

where $\gamma \equiv 3 - 2\alpha$, and Γ_{ab} is the overlap reduction function that encodes the interpulsar spatial correlations, such as the HD spatial correlations (Hellings & Downs 1983) expected for an isotropic GWB.

In addition to the GWB, we also need to model the intrinsic noise for each pulsar in the PTA. For each pulsar, we model uncorrelated instrumental Gaussian noise in the multiplicative EFAC and quadrature-additive EQUAD parameters. Noise that is correlated across observing frequencies but uncorrelated across TOAs is modeled using the ECORR parameter. These parameters are collectively referred to as “white noise.” We include an additional power-law process per pulsar to model any intrinsic low-frequency noise (i.e., “spin red noise”) in the given pulsar. The aforementioned white-noise parameters that correct the reported TOA uncertainties were folded into the simulated uncertainties in the preparation of our data sets, so we do not separately model these white-noise parameters in our analysis. The red noise for the pulsars is allowed to vary when we search for the GWB.

The GWB parameter constraints are derived by modeling the GWB as a time-correlated process that is statistically common (i.e., with a common spectrum) but uncorrelated among all pulsars in the PTA. This ignores the influence of interpulsar correlations on GWB parameter constraints, which should be minimal as these are dominated by autocorrelation information in the pulsars. Constraints on the GWB spectral index, α , are reported as marginalized over all parameters, including A_{GWB} , while constraints on A_{GWB} , are reported as marginalized over all noise parameters but conditioned on the fiducial spectral index $\alpha = -2/3$. The inclusion of HD correlations in the search process is computationally expensive, so we only use it

to discriminate whether the observed common red process is from a GWB or other spatially correlated signals like terrestrial clock errors (monopolar spatial correlations) or solar system ephemeris errors (dipolar spatial correlations). These spatial correlations are encoded in the overlap reduction function in Equation (2).

To calculate the Bayesian detection significance, we use a product-space sampling method that computes the probabilistic preference for one signal model over another (Arzoumanian et al. 2018; Taylor et al. 2020, and references therein). We construct a super-model consisting of the union of multiple GWB signal models, where an indexing variable determines which model is “active” and used to calculate the corresponding likelihood. The ratio of the posterior probabilities of the model indices can then be used to estimate the Bayesian odds for the preference of one model over the other. Thus, to calculate the Bayesian odds for an HD-correlated process in the data set, \mathcal{O}_{HD} , we create a super-model with the same GWB parameters in both models, while the index parameter toggles the presence of HD correlations.

In addition to power-law strain spectrum models, we also measure the GWB strain spectrum agnostically with an independent amplitude at each frequency. Our priors correspond to independent log-uniform constraints on the characteristic strain amplitude for each Fourier-basis frequency, given by k/T , where $k=1, 2, \dots, 30$, and T is the time span between the first and last TOA in the given data set. Unless explicitly specified, all of the results presented in this paper have been derived using 30 frequencies to model the GWB (and intrinsic red noise) and employing the recent JPL solar system ephemeris, DE438 (Folkner & Park 2018). However, we also explore the effect of using a different number of frequencies and BayesEphem (Vallisneri et al. 2020) in modeling the GWB.

2.3. Frequentist Data Analysis Methods

In addition to the Bayesian pipeline, we also use a frequentist approach to model the GWB (Arzoumanian et al. 2016, 2018). This uses an optimal statistic, \hat{A}^2 , to measure the amplitude of the GWB, corresponding to the sum of the correlations between pulsar pairs weighted by the intrinsic pulsar noise, and conditioned on the assumed spectral index of the GWB. The optimal statistic can also be used to calculate the S/N , ρ_{template} , for a given interpulsar correlation template, e.g., monopolar or dipolar for clock and ephemeris errors, respectively, and HD (or quadrupolar) for an astrophysical GWB. To calculate the template S/N , the optimal statistic first calculates the amount of cross-correlated power between different pulsar pairs in the PTA. This pairwise cross-correlated power, suitably binned in angular separation space, can be used as a visual test for the presence of quadrupolar correlations in the data, or as intermediate products for further study of the spatial correlations.

In this work, we use the noise-marginalized optimal statistic (Vigeland et al. 2018), which is more robust when pulsars have intrinsic red noise. We use posterior samples from Bayesian MCMC analyses to marginalize over the pulsar intrinsic red noise and calculate corresponding distributions for \hat{A}^2 and ρ_{template} .

These Bayesian and frequentist methods reflect production-level analyses that are employed in PTA detection and parameter estimation analyses. Thus, our results are a crucial




Figure 1. The Bayesian odds ratio preferring a GWB over a spatially uncorrelated process (\mathcal{O}_{HD}) is plotted against the cross-correlation signal-to-noise ratio (S/N , ρ). These quantities were calculated using the injected simulated data sets across all realizations and injected amplitudes. The theoretical prediction (dashed blue line) for the relation between these two measures of confidence agrees with the empirical best-fit model (solid blue line).

validation of the efficacy of the pipelines that will be performing the important business of SGWB detection in the next several years.

3. Bridging Bayesian Odds and Frequentist S/N s

Under certain circumstances, it is possible to relate Bayesian model selection to frequentist hypothesis testing. When data is informative such that the likelihood is strongly peaked, the Bayesian evidence can be computed under the Laplace approximation, such that $Z_{\mathcal{H}} \equiv \int d\theta p(d|\theta, \mathcal{H})p(\theta|\mathcal{H}) \approx p(\theta_{\text{ML}}|\mathcal{H})\Delta V_{\mathcal{H}}/V_{\mathcal{H}}$, where $p(d|\theta_{\text{ML}}, \mathcal{H})$ maximizes the likelihood with parameters θ given data d under model \mathcal{H} (Romano & Cornish 2017). $\Delta V_{\mathcal{H}}/V_{\mathcal{H}}$ measures the compactness of the parameter space volume occupied by the likelihood with respect to the total prior volume, incorporating the Bayesian notion of model parsimony. Taking the ratio of the Bayesian evidence between two models, labeled 1 and 2, and assuming equal prior odds, allows the Bayesian odds ratio, \mathcal{O}_{12} , to be written as $\ln \mathcal{O}_{12} \approx \ln \Lambda_{\text{ML}}(d) + \ln[(\Delta V_1/V_1)/(\Delta V_2/V_2)]$, where $\Lambda_{\text{ML}}(d)$ is the maximum likelihood ratio. The relevant maximum likelihood statistic for PTA GWB detection is the aforementioned optimal statistic: a noise-weighted two-point correlation statistic between all unique pulsar pairs, which compares a model with HD interpulsar correlations versus one with no correlations between pulsars (Anholm et al. 2009; Demorest et al. 2013; Chamberlin et al. 2015). The signal-to-noise ratio (S/N , ρ) of such GWB-induced correlations can be written as $\ln \Lambda_{\text{ML}} = \rho^2/2$. While the likelihood may be marginally more compact under the model with HD correlations, there is no difference in parameter dimensionality; hence, we ignore the likelihood compactness terms. The relationship between the Bayesian odds ratio in favor of HD correlations and the frequentist S/N of such correlations can then be written as

$$\ln \mathcal{O}_{\text{HD}} \approx \rho_{\text{HD}}^2/2. \quad (3)$$

In a bid to assess the validity of this relationship for PTA GWB searches, we used our simulated data sets to calculate both of these statistics across all realizations, injected amplitudes, and time slices. These values, and the theoretical relationship in Equation (3), are shown in Figure 1. We also allowed the prefactor of $1/2$ to vary in an empirical fit, finding a value of 0.52 ± 0.01 , which agrees well with Equation (3).




Figure 2. The evolution of the total S/N, $\hat{\rho}$, as a function of the HD cross-correlation S/N, ρ_{HD} . As we can see, HD template S/N values of 1, 3, and 7 correspond to $\hat{\rho}$ values of 5, 25, and 105 respectively.

This is the first time that this relationship has been validated for PTA GWB detection.

4. A Statistic to Quantify PTA Milestones

There are many factors that influence the detectability of a GWB signal in a given PTA configuration. Some of these are related to the signal itself, i.e., the amplitude of the characteristic strain spectrum at $f = 1/\text{yr}$, A_{GWB} , as well as the frequency-dependent shape of the spectrum. The quality of our detector is parameterized through the overall observational timing baseline of the array, the number of pulsars, and the timing quality of each pulsar (given by its radio-frequency-dependent noise characteristics). We need a statistic that accounts for all of these factors, acting as a fiducial scaling parameter in terms of which we can track the evolution of parameter uncertainties and signal detectability. The optimal statistic is not appropriate here, because it only considers cross-correlations and thus underestimates the information content in the full signal (especially the autocorrelations) that dominate parameter estimation.

The full signal likelihood (modeling auto- and cross-correlations) is a sufficient statistic for this purpose. Hence, we form a total-signal S/N, $\hat{\rho}$, through the log-likelihood ratio of the GWB+noise versus noise-only models, where noise is treated as uncorrelated between pulsars. Therefore,

$$\hat{\rho} = \sqrt{2 \ln \Lambda} \equiv \sqrt{2 [\ln p(d|\theta_{\text{signal}}, \theta_{\text{noise}}) - \ln p(d|\theta_{\text{noise}})]}. \quad (4)$$

This differs from the HD cross-correlation S/N, ρ_{HD} (or an equivalent cross-correlation statistic) in considering all distinct pairings of the pulsars, including autocorrelations.

We would like to note that this new statistic is not meant to be a replacement for the cross-correlation S/N. Because an astrophysical GWB is characterized by the unique HD cross-correlation signature, the corresponding cross-correlation S/N, ρ_{HD} , will be the primary arbiter for confirming a signal as an astrophysical GWB.

As PTAs move closer to detection, the power in the GWB at low frequencies is expected to dominate the intrinsic white and red noise in some of the pulsars in the PTA. In this “intermediate-signal regime,” the total S/N, $\hat{\rho}$, is expected to be much larger than the corresponding HD S/N (ρ_{HD} , Romano et al. 2020). We can directly connect them by calculating the two statistics on the same simulated injected data sets. The binned values for the two statistics calculated this way are

shown in Figure 2 and related by the empirical scaling relation

$$\hat{\rho} = 25 \left(\frac{\rho_{\text{HD}}}{3} \right)^{1.7}. \quad (5)$$

Using this scaling relation, ρ_{HD} values of [1, 3, 7] map to $\hat{\rho}$ values of approximately [5, 25, 105] respectively. Because the power in autocorrelations is significantly larger than that in cross-correlations, the corresponding S/N values for the total SN, $\hat{\rho}$, are higher than those for HD cross-correlation, ρ_{HD} (Romano et al. 2020).

From existing scaling laws for ρ_{HD} (Siemens et al. 2013), we can generalize $\hat{\rho}$ to any PTA configuration. For a PTA in the weak- and intermediate-signal regime (Siemens et al. 2013), the total S/N can be written as

$$\hat{\rho} \propto \left[M c \frac{A_{\text{GWB}}^2}{\sigma^2} T^\gamma \right]^{1.7}, \quad (6a)$$

$$\hat{\rho} \propto \left[M \left(\frac{A_{\text{GWB}}}{\sigma \sqrt{c}} \right)^{1/\gamma} T^{1/2} \right]^{1.7}, \quad (6b)$$

respectively, where M is the number of pulsars in the array, $c = 1/\Delta t$ is the inverse of the observational cadence, A_{GWB} is the amplitude of the GWB, σ is the white-noise timing rms, T is the timing baseline of the PTA, and γ is as previously defined ($\gamma \equiv 3 - 2\alpha = 13/3$ for a population of SMBHBs).

5. PTA Milestones

5.1. Milestone I: Detection

The upper panel of Figure 3 shows the evolution of the median $\hat{\rho}$ as a function of data baseline and injected GWB amplitude, the middle panels show the cross-correlated power in different angular separation bins for a GWB with an amplitude of $A_{\text{GWB}} = 2 \times 10^{-15}$ for three different data baselines (12, 15, and 20 yr), and the lower panel shows the evolution of the S/N for HD, monopole, and dipole cross-correlations as a function of data baseline for a GWB with amplitude between $A_{\text{GWB}} = [1-2] \times 10^{-15}$. With an increase in the length of the PTA data set, $\hat{\rho}$ for the GWB increases, as well as the S/N for HD correlations, ρ_{HD} . The bottom panel of Figure 3 shows that, starting at the 15 yr slice, the median HD correlations are preferred over other types of spatial correlation for a GWB with strength $A_{\text{GWB}} = [1-2] \times 10^{-15}$. Only with HD correlations clearly favored over monopolar or dipolar signals (which happens near the 18 yr baseline in our simulations) can a signal be confidently attributed to GWBs instead of uncharacterized noise sources.

To test the influence of the number of frequencies used to model the GWB, we repeat the analysis in the lower panel of Figure 3 for a GWB modeled with only the five lowest frequencies. The choice of these five lowest frequencies is designed to avoid GWB significance and parameter estimation being contaminated by poorly modeled excess noise in the data set that dominates at higher frequencies (Arzoumanian et al. 2020). We find that the initial detection of a GWB is not inhibited or delayed by using only the five lowest frequency bins, which only reduces the S/N for HD correlations by a few percent. We also explored the influence that accounting for systematics in solar system ephemeris modeling has on our GWB detection significance. We do so by including

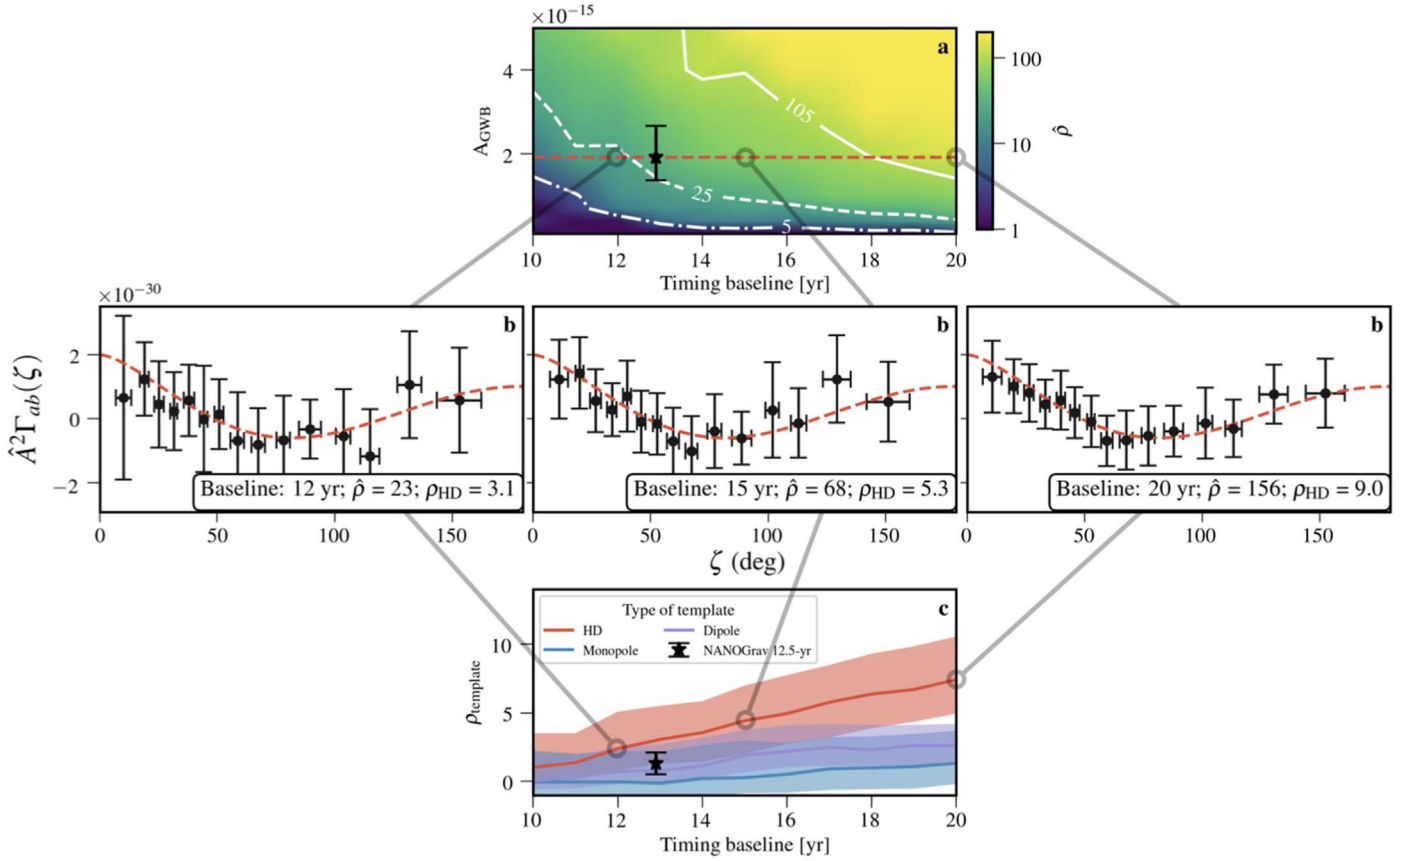


Figure 3. Top: the evolution of the median total S/N, $\hat{\rho}$, over 10 realizations as a function of time. The dashed-dotted, dashed, and solid contours represent total S/N values of 5, 25, and 105 corresponding to HD S/Ns, ρ_{HD} , of 1, 3, and 7 respectively. The black data point represents the amplitude of the GWB from the NANOGrav 12.5 yr analysis. Middle: the median binned cross-correlated power over 10 realizations for increasing S/N from left to right. The uncertainties in this panel represent the spread of the cross-correlated power over the 10 realizations. The theoretical HD curve for a power-law GWB with an amplitude of $A_{\text{GWB}} = 2 \times 10^{-15}$ is shown by the red dashed line. Bottom: the cross-correlation S/N ratio, ρ_{template} , for different types of spatial correlations with an injected power-law GWB amplitude of $A_{\text{GWB}} = [1-2] \times 10^{-15}$. The shaded regions represent the median 95% credible intervals across ten realizations of the PTA data set. The black data point represents the significance of the GWB from the NANOGrav 12.5 yr analysis.

BAYESEPHM (Vallisneri et al. 2020) in our model, which marginalizes over uncertainties in Jupiter orbital elements, gas giant masses, and the celestial reference frame. We find that BAYESEPHM reduces detection significance by only ~ 4 units of total S/N, $\hat{\rho}$, consistent with earlier results reported in Vallisneri et al. (2020). Thus, the presence of solar system ephemeris uncertainties is unlikely to inhibit the detection of a GWB by PTAs.

As shown in Figure 3 as a black data point, detection statistics for the NANOGrav 12.5 yr data set are consistent with our simulations, and we expect the significance of a GWB to increase in the real data set at least at the rate shown here, if not faster due to the addition of more pulsars to the array (Siemens et al. 2013) and new advanced noise mitigate schemes that should allow the GWB spectrum to be modeled to higher frequencies.

5.2. Milestone II: Source of The GWB

Figure 4 shows the evolution of the fractional parameter measurement uncertainty, $\Delta X/X$, as a function of total S/N, where X is the median measured value of the GWB amplitude (A_{GWB}) and spectral index (α), and ΔX is the corresponding 95% credible interval parameter uncertainty. These

uncertainties are well fit by the relations,

$$\frac{\Delta A_{\text{GWB}}}{A_{\text{GWB}}} = 44\% \times \left(\frac{\hat{\rho}}{25}\right)^{-2/5}, \quad (7a)$$

$$\frac{\Delta \alpha}{\alpha} = 40\% \times \left(\frac{\hat{\rho}}{25}\right)^{-1/2}. \quad (7b)$$

The choice of using the five lowest frequencies to model the GWB does not affect the fractional parameter uncertainties until the signal has total S/N $\hat{\rho} \gtrsim 60$, whereupon using 30 frequencies improves their recovery. At such high detection significance, the GWB begins to dominate more than just the five lowest frequency bins, accruing greater significance and more informative constraints from higher frequencies. We also find that the inclusion of BAYESEPHM has a minimal effect on the fractional uncertainty of the measured GWB parameters.

The measured spectral index of the GWB will be the primary arbiter of the source of the GWB. A spectral index of $\alpha = -2/3$, corresponding to our injected signal, is expected for a purely GW-driven population of inspiraling SMBHB systems, while other more exotic sources of the GWB are predicted to have different spectral indices. At initial detection ($\hat{\rho} = 25$ or $\rho_{\text{HD}} = 3$), the spectral index measurement should have a fractional uncertainty of 40% (Equation 7(b)). In our suite of simulations, this precision is already sufficient to disfavor (at

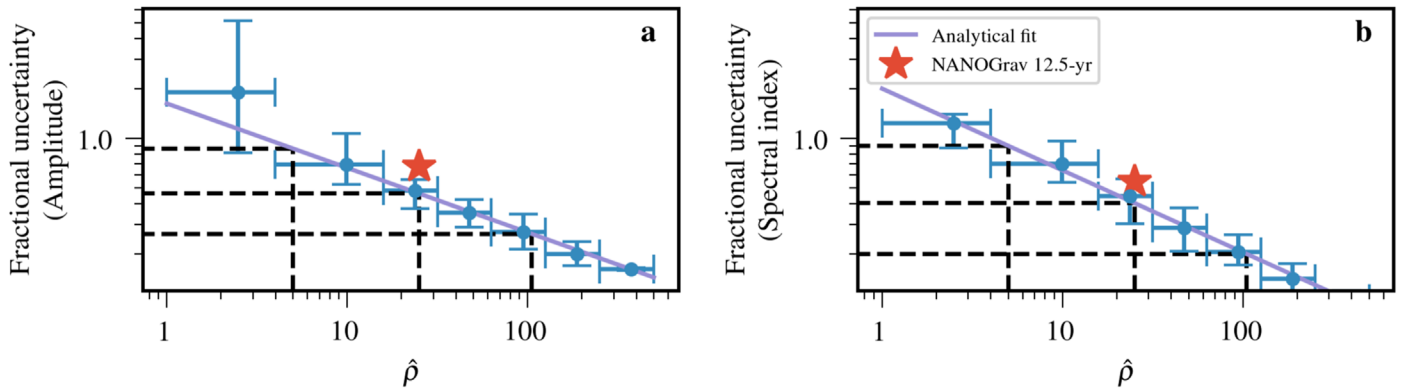


Figure 4. The evolution of the median fractional uncertainty on the measured amplitude, A , and spectral index, α , as a function of total S/N , $\hat{\rho}$, is shown in panels (a) and (b), respectively. The evolution of the fractional uncertainty for the amplitude and spectral index with the S/N is best fit by a power law as shown (parameterized in Equation (7)). For the amplitude and spectral index, the fractional uncertainties corresponding to $\hat{\rho}$ values of 5, 25, and 105 are 84%; 44% and 25%; and 90%, 40%, and 20%, respectively, and are shown by the dashed black lines. The star shows the fractional uncertainty from the NANOGrav 12.5 yr analysis.

95% credibility) models of primordial GWs with matter ($\alpha = -2$; Grishchuk 2005) and radiation-dominated ($\alpha = -1$; Lasky et al. 2016) equations of state.⁴³ This precision is also sufficient to disfavor some models of the GWB produced by cosmic strings, such as those from kinks and cusps in the string loops, which are predicted to have a spectral index, $\alpha = -7/6$ (Ölmez et al. 2010). As the significance of the GWB signal grows, the fractional uncertainty on the measured spectral index will continue to decrease and allow us to test other sources of the GWB. As we show later, an increase in the timing baseline of the data set allows more comprehensive spectral modeling of the GWB, thereby reducing the need to rely on power-law fits.

5.3. Milestone III: Properties of The GWB

Once there is sufficient evidence that the GWB is most likely due to a population of merging SMBHBs, we can use our spectral characterization to probe the astrophysics of the underlying SMBHB population. For example, the amplitude of a GWB produced by SMBHBs is set primarily by the mass distribution of SMBHBs, often parameterized relative to their host galaxies (e.g., in the $M-\sigma$ relation; Ferrarese & Merritt 2000; Gebhardt et al. 2000), and the efficiency by which they reach subparsec separations. The recovered amplitude can thus be used to distinguish between different population models of SMBHBs.

Here we consider three such models (McWilliams et al. 2014; Sesana et al. 2016; Simon & Burke-Spolaor 2016) with GWB amplitudes representative of typical values predicted in the literature: $A_{\text{GWB}} = [0.4_{-0.16}^{+0.26}, 1.5_{-0.3}^{+0.3}, 4.0_{-1.8}^{+3.3}] \times 10^{-15}$, which were also examined in the NANOGrav 11 yr GWB analysis (Arzoumanian et al. 2018). At initial detection, the amplitude measurement will have a fractional uncertainty of 44%, which will be sufficient to distinguish the first model from the other two. Models two and three are distinguishable with a fractional uncertainty of $\sim 37\%$ on the amplitude, occurring near the 17 yr slice of our simulated data. Thus, almost immediately after the initial detection of the GWB, we

expect to be able to clearly distinguish between typical models in the literature.

5.3.1. Constraining Dynamical Influences on the GWB Spectral Shape

A realistic, astrophysical GWB from SMBHBs is, however, expected to deviate noticeably from the pure power-law spectrum we have assumed so far. To quantify the influence of these deviations on our detection prospects, we use astrophysically motivated simulations of SMBHB populations to inject three non-power-law GWB spectra into the synthesized data.

To calculate plausible and self-consistent GWB spectra, we generate full mock populations of SMBH binaries. We employ two entirely independent approaches to simulate SMBHB populations. In the first, we anchor the population on the galaxy–galaxy merger rate derived from the Illustris cosmological simulations (Genel et al. 2014; Vogelsberger et al. 2014a, 2014b). While the Illustris population of SMBH mergers has been used extensively (Kelley et al. 2017a, 2017b, 2018), we have generalized our approach to allow for a more flexible range of binary properties and thus resulting GWB spectra. In particular, we have chosen binary hardening rates and masses so as to produce (i) a GWB spectrum that is very nearly a power law with a relatively steep “negative” spectral index at low frequencies (Figure 5: gray), and (ii) a strain spectrum that is nearly “flat” at low frequencies (Figure 5: yellow).

The second method uses spectra originally generated for the astrophysical inference performed on the NANOGrav 11 yr data set (Arzoumanian et al. 2018). This method uses a semi-analytic model (Simon & Burke-Spolaor 2016) to simulate a binary population using observational-based measurements of the galaxy stellar mass function, galaxy merger rate, and SMBH mass–host galaxy relation. The eccentricity distribution and binary hardening rate is incorporated over a wide range of parameters (Taylor et al. 2017). Model parameters are chosen to produce (iii) a GW strain spectrum that is attenuated strongly enough to turn over, producing a “positive” spectral index at low frequencies (Figure 5: blue).

Overall, the amplitude of the simulated spectra is calibrated primarily by the distribution of SMBH masses, and the low-frequency spectral index is increased from the fiducial $-2/3$ by increasing the environmental hardening rate, mostly due to coupling with the nuclear stellar environment (Simon & Burke-

⁴³ But note that such a broadband primordial GWB signal with similar amplitude to the common-spectrum process found in the NANOGrav 12.5 yr analysis could be constrained at higher frequencies by ground-based GW interferometers and big bang nucleosynthesis (e.g., Kuroyanagi et al. 2021).




Figure 5. Per-frequency spectral modeling of the GWB spectrum for the 20 yr slice. Each frequency bin shows the median strain and uncertainty across 10 realizations, while the injected GWB spectrum is shown by the solid lines. The few lowest frequency bins closely track the shape of the injected GWB, with the lowest frequency bin being the strongest discriminator between the three types of injected spectra. The high-frequency bins are dominated by the white noise in the PTA and thus are insensitive to the GWB.

Spolaor 2016; Kelley et al. 2017b). For both methodologies, we have calibrated all three realistic spectra to have an amplitude of $A_{\text{GWB}} = 1 \times 10^{-15}$ at a frequency of 1 yr^{-1} and are shown by the solid lines in Figure 5.

The total S/N recovered using our standard detection pipeline (assuming a pure power-law strain spectrum) is shown for these realistically modeled GWB spectra in Figure 6. Even in the “positive” model, where a significant amount of GW strain is attenuated at low frequencies, the measured S/N is unaffected at the 15–17 yr slices (where initial detection is predicted to happen) and is only decreased by $\sim 50\%$ at the 20 yr slice. This point is crucial because it demonstrates that the assumption of a pure power-law strain spectrum allows for a high S/N detection, despite omitting the detailed spectral structure. Additionally, even significant attenuation of the GWB signal at low frequencies does not represent a significant obstacle to detection. Any eventual reduction in the recovered S/N compared to the predicted S/N can then be used to interpret the presence of a turnover in the GWB spectrum.

However, many of the astrophysical details about SMBHB environments and demographics are missed by power-law analyses. The additional encoded information can be extracted using per-frequency modeling of the GWB (Lentati et al. 2013; Taylor et al. 2013). In this approach, the strain in each frequency bin is modeled independently, allowing us to measure the spectral shape of the underlying GWB. A per-frequency spectral analysis of 20 yr of data is shown in Figure 5 for our three astrophysically motivated injected GWB spectra, with points and uncertainties representing strain measured independently in each frequency bin. We see that the few lowest bins are dominated by the injected GWB and closely track their spectral shapes. In a follow-up analysis, we will explore full astrophysical model inference to determine how accurately astrophysical parameters can be extracted from realistic GWBs. However, the results shown here are very encouraging in that even models with some degeneracy (i.e., in the overall signal amplitude, A_{GWB}) can be disentangled within the next several years.

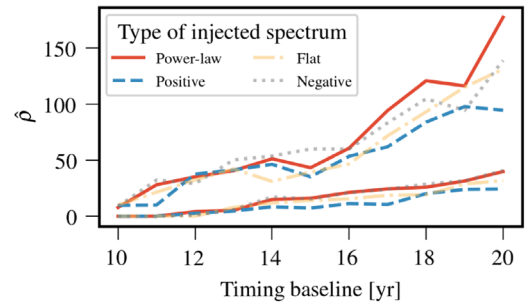


Figure 6. Detection significance for a variety of injected GWB spectra that are all modeled using a pure power-law spectrum in the detection pipeline. The two sets of curves represent the median 95% credible interval upper and lower bounds on the S/N. We can see that power-law modeling of the GWB is robust for making a detection of the GWB.

6. Accelerating PTA Milestones

The S/N for a GWB signal in PTA data can be raised by increasing either the time span of the data set or the number of monitored pulsars (Siemens et al. 2013). However, new pulsars usually cannot be added into a PTA immediately upon discovery and generally require multiple years of timing data to assess their appropriateness for PTA analysis. A solution is to combine the data already collected by PTAs around the world into a single data set and analyze this joint data set for a GWB signal. Beyond NANOGrav, there are currently two other PTAs apart that have decades’ worth of timing data on millisecond pulsars: the Parkes Pulsar Timing Array (PPTA; Hobbs 2013) and the European Pulsar Timing Array (EPTA; Kramer & Champion 2013). Together, these PTAs, along with emerging PTAs in India (InPTA; Joshi et al. 2018), China (CPTA, Lee 2016) and South Africa (SAPTA; Bailes et al. 2016), form the International Pulsar Timing Array (IPTA; Hobbs et al. 2010). To date, the IPTA has published two data releases (Verbiest et al. 2016; Perera et al. 2019), which result in an increase in the number of pulsars relative to individual regional PTAs at that time, while also increasing the time span of data for pulsars that are common across the three PTAs. The IPTA has analyzed one of these data sets for GW signals (Lentati et al. 2016; Verbiest et al. 2016), while the analysis of the latest data release (Perera et al. 2019) is currently in progress.

To quantify the improvement offered by the IPTA over NANOGrav-only data, we use the IPTA’s second data release (Perera et al. 2019), DR2, and compare its sensitivity to that of the NANOGrav subset of the DR2 data set, i.e., the NANOGrav 9 yr data set (NG9; NANOGrav Collaboration et al. 2015). We estimate the sensitivity using HASASIA (Hazboun et al. 2019a, 2019b), a software package that can compute the per-frequency sensitivity of any PTA given the white- and red-noise properties of the pulsars. We use the noise properties of the pulsars provided with the public DR2 data, but do not include intrinsic red noise because these parameters are highly covariant with a common process and would require a full GWB detection analysis for accurate calculation.

The per-frequency sensitivities of IPTA DR2 and NG9 are shown in Figure 7. IPTA DR2 is more sensitive across all frequencies than NG9. Consequently, IPTA DR2 achieves a total-signal significance that is a factor of 30 higher than NG9. This corresponds to a decrease in the fractional uncertainty on the GWB amplitude and spectral index by factors of 4 and 5.5, respectively. Note that these are only estimates because (as




Figure 7. The sensitivity curve for IPTA DR2 and the NANOGrav 9 yr data set, which was one of the components going into the construction of DR2. These sensitivity curves were constructed using HASASIA and only used the white-noise parameters released with IPTA DR2. The IPTA DR2 is more sensitive at all frequencies than the NANOGrav 9 yr data set.

previously mentioned) no intrinsic red noise was accounted for in the sensitivity curves, which would have the effect of reducing each of the S/N values. However, given these estimates, an analysis of a combined IPTA data set offers the opportunity to detect a GWB and extract the underlying astrophysics of the source population at a faster pace and with higher significance than can be achieved with any individual PTA.

7. Conclusion

The recent NANOGrav evidence of a common-spectrum process is a tantalizing hint of the first signs of a GWB with growing significance in PTA data. As such, we have developed a road map for the next several years that includes three key PTA scientific milestones: (I) robust detection of the GWB through measuring HD-correlated timing delays across the PTA, (II) discriminating the origin of the GWB as being from either a population of SMBHBs or exotica, and (III) unveiling the demographic and dynamical properties of SMBHBs as encoded in the shape of the recovered GWB strain spectrum. All of these milestones could be achievable by the middle of this decade under conservative assumptions, and even sooner under the auspices of the IPTA.

This work has been carried out by the NANOGrav collaboration, which is part of the International Pulsar Timing Array. The NANOGrav project receives support from National Science Foundation (NSF) Physics Frontiers Center award number 1430284. This work made use of the Super Computing System (Spruce Knob) at West Virginia University (WVU), which are funded in part by the National Science Foundation EPSCoR Research Infrastructure Improvement Cooperative Agreement #1003907, the state of West Virginia (WVEPS-CoR via the Higher Education Policy Commission), and WVU. Some of our models utilized the IllustrisTNG simulations public data release (<http://www.tng-project.org/>; Nelson et al. 2019) and sublink merger trees (Rodriguez-Gomez et al. 2015).

The Flatiron Institute is supported by the Simons Foundation. Pulsar research at UBC is supported by an NSERC Discovery Grant and by the Canadian Institute for Advanced Research. J.S. and M.V. acknowledge support from the JPL RTD program. S.B.S. acknowledges support for this work from NSF grants #1458952 and #1815664. S.B.S. is a CIFAR Azrieli Global Scholar in the Gravity and the Extreme Universe

program. S.R.T. acknowledges support from NSF grant AST-#2007993 and a Dean’s Faculty Fellowship from Vanderbilt University’s College of Arts & Science. T.T.P. acknowledges support from the MTA-ELTE Extragalactic Astrophysics Research Group, funded by the Hungarian Academy of Sciences (Magyar Tudományos Akadémia), which was used during the development of this research. T.D. and M.T.L. acknowledge NSF AAG award number 2009468. Support for Dr. Cromartie was provided by NASA through the NASA Hubble Fellowship Program grant #HST-HF2-51453.001 awarded by the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., for NASA, under contract NAS5-26555. This work is supported in part by NASA under award number 80GSFC17M0002.

Software: ENTERPRISE (Ellis et al. 2019), enterpri-se_extensions (Taylor et al. 2018), hasasia (Hazboun et al. 2019a), libstempo (Vallisneri 2020), matplotlib (Hunter 2007), PTMCMC (Ellis & van Haasteren 2017).

ORCID iDs

Nihan S. Pol <https://orcid.org/0000-0002-8826-1285>
 Stephen R. Taylor <https://orcid.org/0000-0003-0264-1453>
 Luke Zoltan Kelley <https://orcid.org/0000-0002-6625-6450>
 Sarah J. Vigeland <https://orcid.org/0000-0003-4700-9072>
 Joseph Simon <https://orcid.org/0000-0003-1407-6607>
 Siyuan Chen <https://orcid.org/0000-0002-3118-5963>
 Paul T. Baker <https://orcid.org/0000-0003-2745-753X>
 Bence Bécsy <https://orcid.org/0000-0003-0909-5563>
 Paul R. Brook <https://orcid.org/0000-0003-3053-6538>
 Sarah Burke-Spolaor <https://orcid.org/0000-0003-4052-7838>
 Shami Chatterjee <https://orcid.org/0000-0002-2878-1502>
 James M. Cordes <https://orcid.org/0000-0002-4049-1882>
 Neil J. Cornish <https://orcid.org/0000-0002-7435-0869>
 Fronefield Crawford <https://orcid.org/0000-0002-2578-0360>
 H. Thankful Cromartie <https://orcid.org/0000-0002-6039-692X>
 Megan E. DeCesar <https://orcid.org/0000-0002-2185-1790>
 Paul B. Demorest <https://orcid.org/0000-0002-6664-965X>
 Timothy Dolch <https://orcid.org/0000-0001-8885-6388>
 Elizabeth C. Ferrara <https://orcid.org/0000-0001-7828-7708>
 William Fiore <https://orcid.org/0000-0001-5645-5336>
 Emmanuel Fonseca <https://orcid.org/0000-0001-8384-5049>
 Nathan Garver-Daniels <https://orcid.org/0000-0001-6166-9646>
 Deborah C. Good <https://orcid.org/0000-0003-1884-348X>
 Jeffrey S. Hazboun <https://orcid.org/0000-0003-2742-3321>
 Ross J. Jennings <https://orcid.org/0000-0003-1082-2342>
 Megan L. Jones <https://orcid.org/0000-0001-6607-3710>
 Andrew R. Kaiser <https://orcid.org/0000-0002-3654-980X>
 David L. Kaplan <https://orcid.org/0000-0001-6295-2881>
 Joey Shapiro Key <https://orcid.org/0000-0003-0123-7600>
 Michael T. Lam <https://orcid.org/0000-0003-0721-651X>
 Jing Luo <https://orcid.org/0000-0001-5373-5914>
 Ryan S. Lynch <https://orcid.org/0000-0001-5229-7430>
 Dustin R. Madison <https://orcid.org/0000-0003-2285-0404>
 Alexander McEwen <https://orcid.org/0000-0001-5481-7559>
 Maura A. McLaughlin <https://orcid.org/0000-0001-7697-7422>

Chiara M. F. Mingarelli  <https://orcid.org/0000-0002-4307-1322>

Cherry Ng  <https://orcid.org/0000-0002-3616-5160>

David J. Nice  <https://orcid.org/0000-0002-6709-2566>

Timothy T. Pennucci  <https://orcid.org/0000-0001-5465-2889>


Scott M. Ransom  <https://orcid.org/0000-0001-5799-9714>

Paul S. Ray  <https://orcid.org/0000-0002-5297-5278>

Brent J. Shapiro-Albert  <https://orcid.org/0000-0002-7283-1124>

Xavier Siemens  <https://orcid.org/0000-0002-7778-2990>

Ingrid H. Stairs  <https://orcid.org/0000-0001-9784-8670>

Daniel R. Stinebring  <https://orcid.org/0000-0002-1797-3277>

Joseph K. Swiggum  <https://orcid.org/0000-0002-1075-3837>

Michele Vallisneri  <https://orcid.org/0000-0002-4162-0033>

Haley Wahl  <https://orcid.org/0000-0001-9678-0299>

Caitlin A. Witt  <https://orcid.org/0000-0002-6020-9274>

References

- Accadia, T., Acernese, F., Alshourbagy, M., et al. 2012, *Jinst*, **7**, 3012
- Alam, M. F., Arzoumanian, Z., Baker, P. T., et al. 2021, *ApJS*, **252**, 4
- Anholm, M., Ballmer, S., Creighton, J. D. E., Price, L. R., & Siemens, X. 2009, *PhRvD*, **79**, 084030
- Arzoumanian, Z., Baker, P. T., Blumer, H., et al. 2020, *ApJL*, **905**, L34
- Arzoumanian, Z., Baker, P. T., Brazier, A., et al. 2018, *ApJ*, **859**, 47
- Arzoumanian, Z., Brazier, A., Burke-Spolaor, S., et al. 2016, *ApJ*, **821**, 13
- Bailes, M., Barr, E., Bhat, N. D. R., et al. 2016, arXiv:1803.07424
- Begelman, M. C., Blandford, R. D., & Rees, M. J. 1980, *Natur*, **287**, 307
- 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
- Chamberlin, S. J., Creighton, J. D. E., Siemens, X., et al. 2015, *PhRvD*, **91**, 044048
- Charisi, M., Bartos, I., Haiman, Z., et al. 2016, *MNRAS*, **463**, 2145
- Chen, Y.-C., Liu, X., Liao, W.-T., et al. 2020, arXiv:2008.12329
- De Luca, V., Franciolini, G., & Riotto, A. 2021, *PhRvL*, **126**, 041303
- Demorest, P. B., Ferdman, R. D., Gonzalez, M. E., et al. 2013, *ApJ*, **762**, 94
- Desvignes, G., Caballero, R. N., Lentati, L., et al. 2016, *MNRAS*, **458**, 3341
- Detweiler, S. 1979, *ApJ*, **234**, 1100
- Ellis, J., & Lewicki, M. 2021, *PhRvL*, **126**, 041304
- Ellis, J., & van Haasteren, R. 2017, jellis18/PTMCMCSampler: Official Release, Zenodo, doi:10.5281/zenodo.1037579
- Ellis, J. A., Vallisneri, M., Taylor, S. R., & Baker, P. T. 2019, ENTERPRISE: Enhanced Numerical Toolbox Enabling a Robust Pulsar Inference Suite ascl:1912.015
- Ferrarese, L., & Merritt, D. 2000, *ApJL*, **539**, L9
- Folkner, W. M., & Park, R. S. 2018, Planetary ephemeris DE438 for Juno, Tech. Rep. IOM 392R-18-004, Jet Propulsion Laboratory, Pasadena, CA
- Foster, R. S., & Backer, D. C. 1990, *ApJ*, **361**, 300
- Gebhardt, K., Bender, R., Bower, G., et al. 2000, *ApJL*, **539**, L13
- Genel, S., Vogelsberger, M., Springel, V., et al. 2014, *MNRAS*, **445**, 175
- Graham, M. J., Djorgovski, S. G., Stern, D., et al. 2015, *MNRAS*, **453**, 1562
- Grishchuk, L. P. 2005, *PhyU*, **48**, 1235
- Harry, G. M. & LIGO Scientific Collaboration 2010, *CQGra*, **27**, 084006
- Hazboun, J., Romano, J., & Smith, T. 2019a, *JOSS*, **4**, 1775
- Hazboun, J. S., Romano, J. D., & Smith, T. L. 2019b, *PhRvD*, **100**, 104028
- Hellings, R. W., & Downs, G. S. 1983, *ApJL*, **265**, L39
- Hobbs, G. 2013, *CQGra*, **30**, 224007
- Hobbs, G., Archibald, A., Arzoumanian, Z., et al. 2010, *CQGra*, **27**, 084013
- Hunter, J. D. 2007, *CSE*, **9**, 90
- Jaffe, A. H., & Backer, D. C. 2003, *ApJ*, **583**, 616
- Joshi, B. C., Arumugasamy, P., Bagchi, M., et al. 2018, *JApA*, **39**, 51
- 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
- Kerr, M., Reardon, D. J., Hobbs, G., et al. 2020, *PASA*, **37**, e020
- Kramer, M., & Champion, D. J. 2013, *CQGra*, **30**, 224009
- Kuroyanagi, S., Takahashi, T., & Yokoyama, S. 2021, *JCAP*, 2021, 071
- Lasky, P. D., Mingarelli, C. M. F., Smith, T. L., et al. 2016, *PhRvX*, **6**, 011035
- Lee, K. J. 2016, in ASP Conf. Ser. 502, *Frontiers in Radio Astronomy and FAST Early Sciences Symp.* 2015, ed. L. Qain & D. Li (San Francisco, CA: ASP), 19
- Lentati, L., Alexander, P., Hobson, M. P., et al. 2013, *PhRvD*, **87**, 104021
- Lentati, L., Shannon, R. M., Coles, W. A., et al. 2016, *MNRAS*, **458**, 2161
- Liu, T., Gezari, S., Ayers, M., et al. 2019, *ApJ*, **884**, 36
- McWilliams, S. T., Ostriker, J. P., & Pretorius, F. 2014, *ApJ*, **789**, 156
- NANOGrav Collaboration, Arzoumanian, Z., Brazier, A., et al. 2015, *ApJ*, **813**, 65
- Nelson, D., Springel, V., Pillepich, A., et al. 2019, *ComAC*, **6**, 2
- Ölmez, S., Mandic, V., & Siemens, X. 2010, *PhRvD*, **81**, 104028
- Perera, B. B. P., DeCesar, M. E., Demorest, P. B., et al. 2019, *MNRAS*, **490**, 4666
- Phinney, E. S. 2001, arXiv:astro-ph/0108028
- Rajagopal, M., & Romani, R. W. 1995, *ApJ*, **446**, 543
- Rodriguez-Gomez, V., Genel, S., Vogelsberger, M., et al. 2015, *MNRAS*, **449**, 49
- Romano, J. D., & Cornish, N. J. 2017, *LRR*, **20**, 2
- Romano, J. D., Hazboun, J. S., Siemens, X., & Archibald, A. M. 2020, arXiv:2012.03804
- Rosado, P. A., Sesana, A., & Gair, J. 2015, *MNRAS*, **451**, 2417
- Sampson, L., Cornish, N. J., & McWilliams, S. T. 2015, *PhRvD*, **91**, 084055
- Sazhin, M. V. 1978, *SvA*, **22**, 36
- Sesana, A., Shankar, F., Bernardi, M., & Sheth, R. K. 2016, *MNRAS*, **463**, L6
- Sesana, A., Vecchio, A., & Colacino, C. N. 2008, *MNRAS*, **390**, 192
- Siemens, X., Ellis, J., Jenet, F., & Romano, J. D. 2013, *CQGra*, **30**, 224015
- Simon, J., & Burke-Spolaor, S. 2016, *ApJ*, **826**, 11
- Taylor, S. R., Baker, P. T., Hazboun, J. S., Simon, J. J., & Vigeland, S. J. 2018, Enterprise Extensions, https://github.com/nanograv/enterprise_extensions
- Taylor, S. R., Gair, J. R., & Lentati, L. 2013, *PhRvD*, **87**, 044035
- Taylor, S. R., Simon, J., & Sampson, L. 2017, *PhRvL*, **118**, 181102
- Taylor, S. R., Vallisneri, M., Ellis, J. A., et al. 2016, *ApJL*, **819**, L6
- Taylor, S. R., van Haasteren, R., & Sesana, A. 2020, *PhRvD*, **102**, 084039
- Tiburzi, C., Hobbs, G., Kerr, M., et al. 2016, *MNRAS*, **455**, 4339
- Vallisneri, M. 2020, Libstempo: Python Wrapper for Tempo2, Astrophysics Source Code Library, ascl:2002.017
- Vallisneri, M., Taylor, S. R., Simon, J., et al. 2020, *ApJ*, **893**, 112
- Vaskonen, V., & Veermäe, H. 2021, *PhRvL*, **126**, 051303
- Verbiest, J. P. W., Lentati, L., Hobbs, G., et al. 2016, *MNRAS*, **458**, 1267
- Vigeland, S. J., Islo, K., Taylor, S. R., & Ellis, J. A. 2018, *PhRvD*, **98**, 044003
- Vigeland, S. J., & Siemens, X. 2016, *PhRvD*, **94**, 123003
- Vogelsberger, M., Genel, S., Springel, V., et al. 2014a, *Natur*, **509**, 177
- Vogelsberger, M., Genel, S., Springel, V., et al. 2014b, *MNRAS*, **444**, 1518