

Multi-messenger Approaches to Supermassive Black Hole Binary Detection and Parameter Estimation. II. Optimal Strategies for a Pulsar Timing Array

Tingting Liu^{1,2} , Tyler Cohen³, Casey McGrath^{4,5,6} , Paul B. Demorest⁷ , and Sarah J. Vigeland¹ 

¹ Center for Gravitation, Cosmology and Astrophysics, Department of Physics, University of Wisconsin-Milwaukee, Milwaukee, WI, 53211, USA

tingting.liu@nanograv.org

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

³ Department of Physics, New Mexico Institute of Mining and Technology, 801 Leroy Place, Socorro, NM, 87801, USA

⁴ Center for Space Sciences and Technology, University of Maryland, Baltimore County, Baltimore, MD, 21250, USA

⁵ Gravitational Astrophysics Lab, NASA Goddard Space Flight Center, Greenbelt, MD, 20771, USA

⁶ Center for Research and Exploration in Space Science and Technology II, NASA Goddard Space Flight Center, Greenbelt, MD, 20771, USA

⁷ National Radio Astronomy Observatory, 1003 Lopezville Road, Socorro, NM, 87801, USA

Received 2022 October 26; revised 2022 December 9; accepted 2023 January 11; published 2023 March 8

Abstract

Pulsar timing arrays (PTAs) are Galactic-scale gravitational wave (GW) detectors consisting of precisely timed pulsars distributed across the sky. Within the decade, PTAs are expected to detect nanohertz GWs emitted by close-separation supermassive black hole binaries (SMBHBs), thereby opening up the low-frequency end of the GW spectrum for science. Individual SMBHBs which power active galactic nuclei are also promising multi-messenger sources; they may be identified via theoretically predicted electromagnetic (EM) signatures and be followed up by PTAs for GW observations. In this work, we study the detection and parameter estimation prospects of a PTA which targets EM-selected SMBHBs. Adopting a simulated Galactic millisecond pulsar population, we envisage three different pulsar timing campaigns which observe three mock sources at different sky locations. We find that an all-sky PTA which times the best pulsars is an optimal and feasible approach to observe EM-selected SMBHBs and measure their source parameters to high precision (i.e., comparable to or better than conventional EM measurements). We discuss the implications of our findings in the context of future PTA experiments with the planned Deep Synoptic Array-2000 and the multi-messenger studies of SMBHBs such as the well-known binary candidate OJ 287.

Unified Astronomy Thesaurus concepts: [Gravitational waves \(678\)](#); [Pulsars \(1306\)](#); [Supermassive black holes \(1663\)](#)

Supporting material: machine-readable table

1. Introduction

The detection and characterization of low-frequency gravitational waves (GWs) are considered the next frontiers of GW astronomy. This frequency range can be probed by two GW experiments: the Laser Interferometer Space Antenna (Amaro-Seoane et al. 2017), which is expected to launch in the mid-2030s and will probe GW sources emitting at millihertz frequencies; and pulsar timing arrays (PTAs; e.g., Hobbs 2013; Kramer & Champion 2013; McLaughlin 2013), which have been operating since the 2000s and are expected to reach the sensitivity for nanohertz GW science in this decade.

The main astrophysical sources which emit GWs in the nanohertz range ($\sim 10^{-9}$ – 10^{-7} Hz) are supermassive black hole binaries (SMBHBs), which are thought to form following galaxy mergers. As a population, their demographics and properties encode information about, e.g., the relationship between galaxies and their central SMBHBs and the timescale of SMBHB mergers. As individual sources, SMBHBs may power active galactic nuclei (AGN) through accretion and serve as interesting and rare astrophysical laboratories to study the accretion of matter onto binary black holes. Therefore, the GW observations of individual SMBHBs will also open up new areas in multi-messenger astrophysics.

While it is anticipated that PTAs may soon detect the gravitational wave background (GWB), which is the ensemble signal from a cosmological population of SMBHBs (e.g., Taylor et al. 2016; Kelley et al. 2017; Pol et al. 2021), the best path toward detecting continuous wave (CW) signals from individual sources is less clear. By nature, CW sources are highly anisotropic, which could require different observing strategies than simply timing a large number of pulsars distributed across the sky and adding more to the array over time (which have been the main strategies toward observing the isotropic GWB). For example, should we preferentially add to the PTA pulsars near a possible SMBHB? Should the limited amount of total observing time be allocated to focus on timing a small number of pulsars with the lowest timing noise? Or is an all-sky PTA which times as many pulsars as possible the best (and a feasible) approach?

Those possible strategies have important implications for the multi-messenger studies of SMBHBs. Many binary candidates whose orbital periods would be of order months to years (i.e., corresponding to the PTA GW frequency range) have so far been reported in the literature (e.g., Graham et al. 2015a, 2015b; Liu et al. 2015; Charisi et al. 2016; Liu et al. 2016; Severgnini et al. 2018; Liu et al. 2019). Many of those candidates were identified thanks to modern synoptic surveys which monitor a large number of AGNs, which makes it possible to search systematically for intrinsically rare SMBHBs. Additionally, the typical cadence and length of a time-domain survey make it sensitive to periodic variability on order year timescales arising from mechanisms such as binary-modulated accretion (e.g., Noble et al. 2012;



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

D’Orazio et al. 2013; Farris et al. 2014), Doppler beaming (D’Orazio et al. 2015), and self-lensing (D’Orazio & Di Stefano 2018). The upcoming Rubin Observatory Legacy Survey of Space and Time (LSST; Ivezić et al. 2019) could further increase the number of known binary candidates manyfold (Kelley et al. 2018, 2021). Unfortunately, the electromagnetic (EM) follow-up or confirmation of current candidates have been less than successful (e.g., Foord et al. 2017; Guo et al. 2020; Saade et al. 2020) due to the large false positive rate in systematic searches, the lack of complementary evidence in support of the binary hypothesis, and the difficulty of distinguishing the EM emission of an SMBHB from that of a single AGN. The GW observations of EM-selected SMBHB candidates would therefore lead to at least three important breakthroughs: (1) directly and unequivocally confirming the nature of the sources via the detection of GWs, (2) studying their properties through parameter measurements via GWs and at the same time independently verifying the source parameters obtained from conventional EM observations, and (3) breaking degeneracies and extracting astrophysical information through combined GW and EM observations and connecting observational properties to theoretical models of accreting SMBHBs. Furthermore, both detection signal-to-noise (S/N) and parameter measurement uncertainties can be improved by searching for GWs at the sky locations of EM counterparts (Liu & Vigeland 2021), thereby allowing studies of low-S/N sources which would otherwise be missed in unguided all-sky searches.

While none of the reported binary candidates are within the current PTA sensitivity, it is nevertheless important to understand the capability of a given PTA experimental setup in order to understand what astrophysical information we will be able to extract from it. For example, Sesana & Vecchio (2010) have studied binary parameter measurement uncertainties as a function of, e.g., the number of pulsars in the array and their sky coverage; they have also investigated the sky localization capability which would be important for EM follow-up. Additionally, we can apply the knowledge of PTA capabilities to either fine-tune current PTAs in operation, or design next-generation PTA experiments. For instance, Burt et al. (2011) and Christy et al. (2014) have found that increasing the observing time on the best-timed pulsars would have the best effect on the PTA sensitivity volume. Based on this suggestion, the North American Nanohertz Observatory for Gravitational Waves (NANOGrav) started a high-cadence timing campaign of six pulsars beginning in 2013 (at the Green Bank Telescope) and 2015 (at the Arecibo Telescope) to increase its sensitivity to CWs⁸ (Arzoumanian et al. 2018).

In this work, we aim to study the effect of different PTA configurations on the measurement of binary parameters and explore strategies to optimize PTA experiments for observations of SMBHBs. Our study differs from previous works in several ways: first, we assume we have already identified an EM candidate and the PTA data are being searched to study the “GW counterpart” of the source. Second, we explore both the detection and parameter estimation aspects of the optimized PTA. Finally, our mock PTAs are based on realistic pulsar populations and properties, as well as observational considerations. This paper is organized as follows: in Section 2, we describe our methodology of constructing a simulated PTA and measuring binary parameters; in Section 3, we present the results of observing three mock SMBHB sources with different

PTA setups; we summarize in Section 4 and make suggestions for PTA designs which optimize single source multi-messenger observations. Throughout the paper, we adopt geometrized units: $G = c = 1$.

2. Methods

2.1. Constructing a Mock PTA Data Set Using a Simulated Pulsar Population

A PTA regularly monitors a number of millisecond pulsars (MSPs), which are characterized by rapid and extremely stable rotation periods. It is therefore sensitive to fluctuations in the arrival times of the radio pulses induced by GWs emitted by astrophysical sources including SMBHBs. We first simulate ten realizations of the MSP population using the pulsar population synthesis package `PsrPopPy`⁹ (Bates et al. 2014) in the snapshot mode, where simulated MSPs are generated by randomly drawing properties from distributions constrained by earlier studies. The overall population size is constrained by halting generation when the population contains the same number of MSPs (28) that would have been discovered by the Parkes Multibeam Survey (Manchester et al. 2001). We assume that the MSPs are distributed radially from the Galactic Center according to a Gaussian distribution with a standard deviation of 7.5 kpc and exponentially above and below the galactic plane with a mean scale height of 500 pc (Lorimer 2012). We assume that the MSP spin periods are distributed according to Lorimer (2012) and that their luminosities (at 1.4 GHz) are distributed according to a log-normal distribution with $\langle \log_{10} L \rangle = -1.1$ and $\sigma_{\log_{10} L} = 0.9$ (Faucher-Giguere & Kaspi 2006), where L is in units of mJy kpc². Aggarwal & Lorimer (2022) found that MSP spectral indices are best described by a broken power law with a break at 300 MHz, but since all of our simulated observations occur above 300 MHz, we simply use the high frequency part of their model: Gaussian with a mean spectral index of -1.47 and a standard deviation of 0.7. We assume that the MSP pulse widths are distributed log₁₀ normal with a mean of $0.05 P^{0.9}$ and a standard deviation of 0.3. Finally, dispersion measures (DM) and scattering timescales are computed by the NE2001 electron density model of the Milky Way (Cordes & Lazio 2002).

In order to construct a PTA from the simulated MSP population, we simulate a timing program with the proposed Deep Synoptic Array-2000 (DSA-2000; Hallinan et al. 2019), where an anticipated 25% of the time will be used for pulsar timing. We simulate timing observations of each MSP using the full DSA-2000 collecting area operating in the frequency range 0.7–2 GHz and assume an integration time of 60 minutes per pulsar per observation epoch. We assume that DSA-2000 will have a system temperature of 25 K, a gain of 10 K Jy^{-1} , and will be able to observe declinations from -30° to 90° . If a model pulsar is within the decl. limits, its pulse duty cycle is $<90\%$, and it is sufficiently bright to be detected, we then include it in our mock PTA and compute the uncertainty on its pulse time of arrival σ_{TOA} , which is dependent on both the pulsar and telescope parameters. To do so, we use the `FrequencyOptimizer`¹⁰ software package (Lam et al. 2018),

⁸ However, the program was suspended in 2020 due to the collapse of Arecibo.

⁹ <https://github.com/samb8s/PsrPopPy>

¹⁰ <https://github.com/mtlam/FrequencyOptimizer>

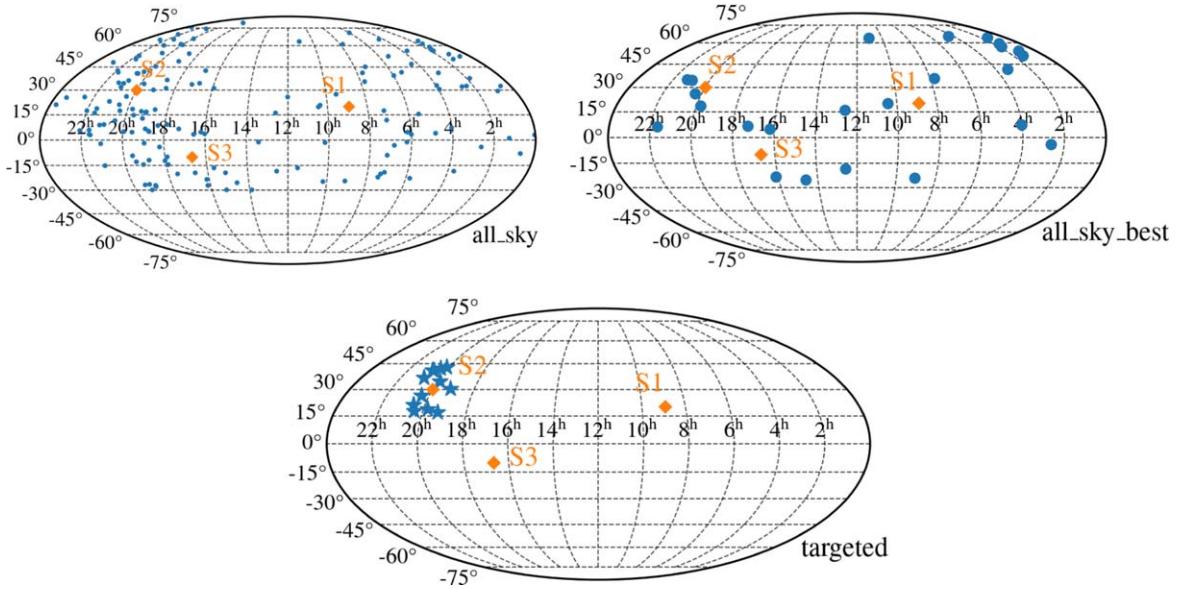


Figure 1. In the first panel, we show one realization of the simulated MSP population with $\sigma_{\text{TOA}} < 1 \mu\text{s}$ observed by DSA-2000 (small blue filled circles); those MSPs form the `all_sky` PTA. In the second panel, we show the ones with $\sigma_{\text{TOA}} < 100 \text{ ns}$ as large filled circles (`all_sky_best`). In the last panel, we show the ones which are located near S2 as stars (`targeted`). The three sources whose detection and parameter estimation prospects are examined in this work are shown as orange diamonds.

which defines the TOA uncertainty as¹¹

$$\sigma_{\text{TOA}}^2 = \sigma_{\text{DM}}^2 + \sigma_{\delta t_c}^2 + \sigma_{\text{DM}(\nu)}^2 + \sigma_{\text{tel}}^2.$$

Here, σ_{DM} is the standard error on the infinite-frequency TOA when fitting the timing residuals for the DM; it contains sources of white noise including template-fitting errors, scintillation errors, and jitter errors, where the jitter errors are computed using the scaling relationship shown in Figure 7 of Lam et al. (2019)

$$\sigma_J = 10^{-0.218} \delta^{1.216} \frac{P^{3/2}}{\sqrt{t_{\text{int}}}},$$

which is a function of the pulsar spin period P , pulse duty cycle δ , and integration time t_{int} . $\sigma_{\delta t_c}$ is the systematic error due to chromatic variations in the pulse width caused by interstellar scattering and dispersion. $\sigma_{\text{DM}(\nu)}$ is the uncertainty in the DM estimation between two frequency channels. σ_{tel} contains uncertainties due to radio frequency interference, incorrect gain calibration, and instrumental self-polarization. NANOGrav considers $\sigma_{\text{TOA}} < 1 \mu\text{s}$ to be the minimum criterion for the inclusion of an MSP in the PTA, so we adopt that cutoff here. The results are ten mock PTAs with an average of ~ 180 MSPs with $\sigma_{\text{TOA}} < 1 \mu\text{s}$. The sky distribution and distribution of σ_{TOA} for one realization are plotted in the first panel of Figure 1 and the left panel of Figure 2 (`all_sky`), respectively.

Next, we consider two variants of our mock timing program: (1) an all-sky PTA consisting of only pulsars with $\sigma_{\text{TOA}} < 100 \text{ ns}$ (`all_sky_best`) and (2) a targeted campaign in an area of the sky consisting of ~ 10 pulsars within a $\sim 10^\circ$ radius, where

¹¹ The total σ_{TOA} in Lam et al. (2018) originally contains an additional term, σ_w , containing sources of white noise. We ultimately determined that since σ_{DM} is derived from the white noise covariance matrix, adding σ_w in quadrature would double-count these contributions so we omit that term in our analysis.

the GW source is chosen to be located within that sky area (`targeted`). Since the antenna pattern functions increase with a closer (but not perfect) alignment between the source and the pulsar (see the definitions in Ellis et al. 2012), a PTA campaign which targets pulsars near an EM-selected source may be more advantageous for detecting its GW signal. We plot the pulsars in one realization of `all_sky_best` as large filled circles in the second panel of Figure 1, and those in `targeted` as stars in the third panel. Their respective σ_{TOA} distributions are shown in the last two panels in Figure 2. The properties of a sample of simulated MSPs are shown in Table 1, and ten realizations of the mock population are available in machine-readable form online.

Then, for each pulsar sample in each realization, we construct mock PTA observations which have a baseline of 15 yr, which is approximately the length of the current PTAs. The cadence of observations is chosen to be either monthly or weekly, to imitate the current “monthly” and “high-cadence” campaigns of NANOGrav (Alam et al. 2021).

2.2. Parameters and GW Signals of an SMBHB

We consider a binary system with component masses m_1 and m_2 . The so-called chirp mass \mathcal{M} is given by $\mathcal{M} \equiv (m_1 m_2)^{3/5} / (m_1 + m_2)^{1/5}$. We assume the system is in a circular orbit whose GW frequency is related to its orbital frequency $f_{\text{GW}} = 2f_{\text{orb}}$ and is often expressed as $\omega = \pi f_{\text{GW}}$. The source is located at a luminosity distance d_L ; since in this work we will only consider nearby sources, we ignore the effect of redshift: $(1+z) \approx 1$. For simplicity, we only consider the “Earth term” in the timing residual and do not consider the “pulsar term,”¹² and hence only the combination $\mathcal{M}^{5/3} / d_L$ can

¹² Our justification for only considering the Earth term is mainly twofold: (1) our methodology (see Section 2.3) is only concerned with the parameter measurement uncertainty and does not inform the possible bias in the posterior distribution introduced by dropping the pulsar term; (2) extracting source information from the pulsar term would require precise measurements of the pulsar distance which we do not assume in this work.

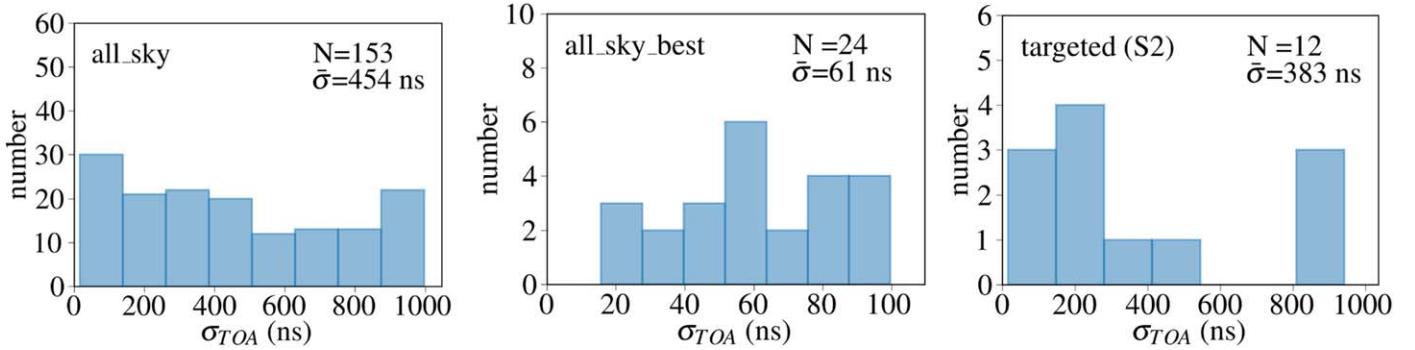


Figure 2. We show the respective σ_{TOA} distributions, sample sizes, and mean σ_{TOA} of the three pulsar samples shown in Figure 1.

be constrained. It is often convenient to redefine the amplitude as $\alpha = \mathcal{M}^{5/3}/d_{\text{L}}/\omega^{1/3}$.

The sky position of the source is written as polar and azimuthal angles ϕ and θ , respectively, (corresponding to R.A. and decl. in units of radians: $\phi = \text{R.A.}$, $\theta = \pi/2 - \text{decl.}$). We assume the frequency evolution of the binary over the course of ~ 15 yr is negligible (i.e., a monochromatic signal). We further assume the BHs are non-spinning, since PTAs are largely insensitive to the effects of spin in an SMBHB system (see Sesana & Vecchio 2010).

Such a binary system is then described by the following parameters: $p = \{\alpha, \omega, \phi, \theta, i, \psi, \Phi_0\}$, where i , ψ , and Φ_0 are the inclination, GW polarization angle, and initial orbital phase, respectively. For simplicity, we adopt fixed values of $\psi = \pi/2$, $\Phi_0 = \pi/2$, and an intermediate inclination of $i = \pi/4$ for each of the sources, since these are the less astrophysically interesting parameters for our study. We then follow the usual definitions in Ellis et al. (2012) and Aggarwal et al. (2019) to compute the Earth term timing residual Res corresponding to each pulsar which has coordinates ϕ_{psr} and θ_{psr} .

The first source (S1) which will be “observed” by our mock PTA is the circular and non-spinning analog of the well-known SMBHB candidate OJ 287 (e.g., Sillanpaa et al. 1988; Lehto & Valtonen 1996; see, e.g., Dey et al. 2019; Valtonen et al. 2021 for recent reviews). Adopting the binary parameters in Dey et al. (2018), we estimate a GW amplitude of $\alpha \approx 3$ ns. We note that while it is possible to compute a more accurate waveform for this complex binary candidate (see, e.g., Valtonen et al. 2021), in this work we only seek to obtain estimates of its detection and parameter measurement prospects.

As can be seen in Figure 1, S1 is located in an area of the sky where few mock pulsars are nearby (e.g., within a ~ 10 deg radius), thereby making a future targeted pulsar timing campaign with DSA-2000 unlikely. Prompted by this, we consider a mock source (S2) which (1) has similar binary parameters but is located at the part of the sky where it may be observed by a targeted campaign and (2) represents a generic SMBHB candidate, whereas S1 is based on an actual candidate. We additionally consider S3, which shares the same binary parameters as S2 but is at a less advantageous location for pulsar searching and timing (and hence the possibility of a targeted campaign). On the other hand, it is located outside the Galactic plane and suffers from less reddening and extinction effects. Thus, S3 represents sources which could be easily detected or observed by most EM observatories. The parameters of the three sources are listed in Table 2, where the S1 parameters are based on Dey et al. (2018).

To summarize, we observe the three sources with two all-sky PTAs with either weekly or monthly cadence over 15 yr. However, an *all_sky* campaign at a weekly cadence is not under consideration here since the total observing time would exceed the 25% available for pulsar timing. For S2, we additionally consider a targeted campaign at both cadences. Those mock programs are also summarized in Table 3.

2.3. Estimating Binary Parameters Using the Fisher Matrix Formalism

The Fisher matrix is often used to assess the parameter estimation performance of an experiment in fields such as cosmology (e.g., Albrecht et al. 2006) and GW astrophysics (e.g., Sesana & Vecchio 2010; McGrath & Creighton 2021). While it can lead to a less accurate prediction of the performance of the experiment (see, e.g., Vallisneri 2008) than the mock data analysis approach, it is significantly less computationally expensive and thus can be applied to a large number of experimental setups and realizations. In this section, we describe how we apply the Fisher matrix formalism to predict the prospects of mock PTAs as described in Section 2.1 of measuring the parameters of the three SMBHBs as depicted in Section 2.2.

We begin by first computing the Fisher matrix for each pulsar (indexed by a)

$$\mathcal{F}_{ij}^a = \sum_b \frac{1}{\sigma^2} \frac{\partial \text{Res}_b}{\partial p_i} \frac{\partial \text{Res}_b}{\partial p_j},$$

where i and j are the indices of the binary parameters, b denotes each observation, σ represents the σ_{TOA} of the a -th pulsar as discussed in Section 2.1¹³ and Res is the Earth term timing residual corresponding to the a -th pulsar as described in Section 2.2. Here we fix the values of ϕ and θ to mimic the search for a CW signal where the source has been pinpointed (which is a reasonable assumption if the candidate is identified electromagnetically), which means we do not compute the partial derivatives with respect to ϕ or θ , and hence $i, j = 1 \dots 5$ which correspond to the parameters $\{\alpha, \omega, i, \psi, \Phi_0\}$.

We then compute the Fisher matrix for multiple pulsars at various locations $\{\phi_{\text{psr}}, \theta_{\text{psr}}\}$ (i.e., a PTA) by summing the

¹³ We have assumed white noise for the Fisher matrix formalism. Properly including pulsar red noise would require nontrivial modifications to the Fisher matrix, which is beyond the scope of this paper. However, for a discussion of the potential effect of red noise on detection and parameter estimation, see Section 3.4.

Table 1
Sample of Simulated Pulsars

<i>P</i> (ms)	\dot{P} (s^{-1})	DM (pc cm^{-3})	<i>t</i> _{scatter} (ms)	<i>W</i> ₅₀ (ms)	<i>l</i> (deg)	<i>b</i> (deg)	R.A. (deg)	Decl. (deg)	<i>S</i> ₁₄₀₀ (mJy)	<i>L</i> ₁₄₀₀ (mJy kpc ²)	α	<i>d</i> _{true} (kpc)	<i>X</i> (kpc)	<i>Y</i> (kpc)	<i>Z</i> (kpc)	σ_{TOA} (μs)	σ_{tel} (μs)	$\sigma_{\delta\text{DM}}$ (μs)	σ_{W} (μs)	σ_{J} (μs)	<i>A</i> _{red} ($\mu\text{s yr}^{1/2}$)	γ_{red}	<i>A</i> _{eff} / <i>A</i> _{eff,tot}
6.64	2.19E-21	615.09	7.42	0.80	-0.39	1.85	264.39	-28.29	9.11E-04	0.08	-0.91	9.18	-0.06	-0.67	0.30	-1.00	-1.00	-1.00	1.09	6.50E-04	-5.60	1.00	
3.19	1.86E-19	177.42	2.01	0.06	6.82	-9.80	279.87	-27.63	0.004	0.15	-2.08	5.86	0.69	2.77	-1.00	1757.16	0.01	1757.16	5.68	0.05	0.09	-5.60	1.00
3.27	9.73E-21	336.38	3.86	0.11	-15.20	-7.19	264.06	-45.61	0.002	0.33	-2.45	14.57	-3.79	-5.45	-1.82	-2.00	-2.00	-2.00	0.11	0.02	-5.60	1.00	
4.42	1.37E-19	112.95	0.00	0.48	33.95	15.35	269.57	7.82	8.00E-04	0.03	-2.70	6.26	3.37	3.49	1.66	-1.00	-1.00	-1.00	0.53	0.02	-5.60	1.00	
14.04	4.74E-21	638.35	10.42	0.50	-3.47	0.52	263.76	-31.60	0.012	0.71	-3.04	7.63	-0.46	0.89	0.07	-2.00	-2.00	-2.00	1.00	8.17E-04	-5.60	1.00	
25.52	7.16E-20	1415.46	1143.50	0.79	12.99	-0.46	273.89	-17.89	2.26E-04	0.08	-0.18	19.35	4.35	-10.36	-0.16	-1.00	-1.00	-1.00	2.12	1.95E-04	-5.60	1.00	
7.00	7.94E-21	552.88	1514.56	0.21	54.88	0.25	293.04	19.55	2.54E-06	7.14E-04	-2.13	16.76	13.71	-1.14	0.07	-1.00	-1.00	-1.00	0.29	0.00	-5.60	1.00	
14.44	2.79E-19	943.11	500.37	0.29	-43.43	-0.23	221.07	-60.08	0.002	0.54	-1.69	15.14	-10.41	-2.50	-0.06	-2.00	-2.00	-2.00	0.58	0.02	-5.60	1.00	
28.56	5.91E-21	215.86	0.04	1.29	4.94	10.38	259.74	-19.19	0.006	0.52	-0.80	9.23	0.78	-0.54	1.66	101.11	0.04	101.11	27.34	3.62	3.04E-04	-5.60	1.00
1.98	6.71E-21	372.40	1.68	0.34	-2.35	4.21	260.91	-28.64	0.002	0.16	-0.64	8.46	-0.35	0.07	0.62	-1.00	-1.00	-1.00	0.25	0.01	-5.60	1.00	

Notes. *P*—pulse period; \dot{P} —period derivative; DM—NE2001 DM; *t*_{scatter}—scattering timescale; *W*₅₀—intrinsic pulse FWHM; and *l*, *b*—galactic longitude and latitude. Note that in *PsrPopPy*, $-180^\circ < l < +180^\circ$. R.A., decl.—R.A. and decl.; *S*₁₄₀₀—flux at 1.4 GHz; *L*₁₄₀₀—luminosity at 1.4 GHz; α —spectral index; *d*_{true}—true distance from Earth; and *X*, *Y*, *Z*—galactic Cartesian positions. The Sun is at *X* = 0, *Y* = 8.5 kpc, *Z* = 0 in this coordinate system. σ_{TOA} —total TOA uncertainty; σ_{tel} —telescope noise; $\sigma_{\delta\text{DM}}$ —DM estimation uncertainty: $\sigma_{\delta\text{DM}} = \sqrt{\sigma_{\text{DM}}^2 + \sigma_{\delta\text{IC}}^2 + \sigma_{\text{DM}(\nu)}^2}$; σ_{W} —white noise errors; σ_{J} —jitter noise; *A*_{red}, γ_{red} —red noise power spectrum amplitude and spectral index, respectively; and *A*_{eff}/*A*_{eff, tot}—fraction of the total collecting area used in observation.

(This table is available in its entirety in machine-readable form.)

Table 2
SMBHB Parameters

Source	R.A. deg	decl. deg	\mathcal{M} $\log M_{\odot}$	f_{GW} $\log \text{Hz}$	d_{L} $\log \text{Mpc}$	α ns
S1	133.7036	20.1085	9	-8.28	3.2	3.4
S2	300	30	9	-8	2.5	13.9
S3	250	-10	9	-8	2.5	13.9

Note. $\Psi = \pi/2$, $\Phi_0 = \pi/2$, and $i = \pi/4$ for each source. S1 represents an OJ 287-like SMBHB, whereas S2 and S3 are mock SMBHBs.

Table 3
Mock PTA Campaigns

Source	PTA	Campaign
S1	all_sky	15 yr, monthly
	all_sky_best	15 yr, monthly
	all_sky_best	15 yr, weekly
S2	all_sky	15 yr, monthly
	all_sky_best	15 yr, monthly
	all_sky_best	15 yr, weekly
	targeted	15 yr, monthly
	targeted	15 yr, weekly
S3	all_sky	15 yr, monthly
	all_sky_best	15 yr, monthly
	all_sky_best	15 yr, weekly

Note. Properties of the sources are listed in Table 2.

matrices

$$\mathcal{F} = \sum_a \mathcal{F}^a.$$

Instead of directly inverting the matrix \mathcal{F} , we first ensure it is well-conditioned, by normalizing the matrix by the factor $\sqrt{\mathcal{F}_{ii}\mathcal{F}_{jj}}$ so that the diagonal elements are 1 and the off-diagonal elements are of order unity. We then use singular value decomposition to invert the normalized matrix: $\mathcal{C}_N = \mathcal{F}_N^{-1}$. We then divide \mathcal{C}_N by the normalization factor to obtain the final covariance matrix \mathcal{C} , where $\mathcal{C} = \mathcal{F}^{-1}$. Finally, we check the accuracy of the inversion by multiplying the covariance matrix by the original Fisher matrix and confirm that its difference from the identity matrix is lower than an error threshold: $\max(|\mathbb{I} - \mathcal{C}\mathcal{F}|) < 10^{-3}$. We can then obtain the measurement uncertainty σ_i of parameter p_i from the i -th diagonal element $\sigma_i^2 = \mathcal{C}_{ii}$. In this work, we focus on the uncertainties of the two more astrophysically interesting parameters, α and ω .

Finally, we estimate the total S/N by summing the contribution from each pulsar: $\text{S/N}^2 = \sum_a \text{S/N}_a^2$, where $\text{SNR}_a = \sqrt{\sum_b \left(\frac{\text{Res}_b}{\sigma}\right)^2}$. In this work, we focus our analysis in the strong signal regime (defined in this work as $\text{S/N} > 5$) so that σ_i , instead of being a lower limit, approaches the actual measurement uncertainty. Therefore we only record the value of σ_i if S/N is greater than 5 in a realization.

3. Results

3.1. PTA Observations of an OJ 287-like Binary Candidate

We first investigate whether S1 can be observed by a future all-sky PTA at a either weekly or monthly cadence, i.e., our

simulated all_sky and all_sky_best, by performing a Fisher matrix analysis for each simulated PTA. At a monthly cadence, only 2/10 of the realizations of an all_sky PTA exceed our S/N threshold of 5. Their S/Ns slightly decrease in all_sky_best, which is expected because the pulsars in all_sky_best are a subset of the ones in all_sky (solid line in Figure 3). If only the best pulsars (all_sky_best) are monitored at a weekly cadence, 9/10 of the realizations have $\text{S/N} > 5$, and the mean value increases to ~ 9 (dashed line).

We proceed to compare the measurement uncertainties between the GW- and EM-based methods, where the GW-based uncertainties are computed using the Fisher matrix method in Section 2.3, and the EM-based uncertainties are based on the reported uncertainties of m_1 , m_2 , and P_{orb} from Dey et al. (2018). The measurement uncertainty of the GW amplitude (quantified by $\Delta\alpha/\alpha$) is at the $\sim 100\%$ level (middle panel), far exceeding the EM-based measurement uncertainty of $\sim 0.2\%$ for this binary candidate. By contrast, ω can be constrained at the few percent level despite the modest S/N (right panel); it is however still greater than the EM-based uncertainty of $\sim 0.1\%$.

Thus, we tentatively conclude that a future PTA with DSA-2000 which resembles our all_sky or all_sky_best could detect the GW signal from OJ 287 with ~ 15 yr of data at $>$ monthly cadence. However, evidence for its binarity based only on GW observations will be modest, unless GW data (especially the more precise measurement of f_{GW}) are interpreted alongside EM observations. Further, the stochastic GWB has a fiducial power-law spectral shape which increases in amplitude at low GW frequencies; therefore the GWB would have a significant effect on the detectability of sources emitting at (particularly) lower frequencies. See Section 3.4 for a discussion of this effect on the three mock sources considered in this paper.

It is worth noting that the all_sky campaign only increases the S/N by < 1 compared to all_sky_best, despite monitoring ~ 6 times the number of pulsars. This suggests that if the total number of observations is a limited resource, it may be best spent on the highest quality pulsars, either at the same cadence or a higher cadence. We explore this further in the following sections.

3.2. PTA Observations of a Mock Binary Candidate

In this section, we consider the detectability and parameter measurement uncertainties of the mock source S2, which has a larger GW amplitude ($\alpha \approx 14$ ns) and is located in a part of the sky where more pulsars may be found and monitored. We first compute the expected results of a monthly, 15 year-long campaign. As we show in the left panel of Figure 4, the highest S/N is reached in the all_sky campaign; it decreases in all_sky_best and decreases further in targeted. This can be understood by the fact that targeted has comparable pulsar timing noise as all_sky, but ~ 10 times fewer pulsars in the array. Similarly, all_sky_best has much fewer pulsars than all_sky and hence a lower S/N, despite a higher median pulsar timing quality.

As we also show in the last two panels of Figure 4, targeted results in poorer constraints on α and ω , consistent with its lowest S/N. Nonetheless, these GW-based parameter measurements are better than the constraints obtained from typical EM observations for the following reasons. (Recall that S2 represents a generic EM binary candidate and therefore its

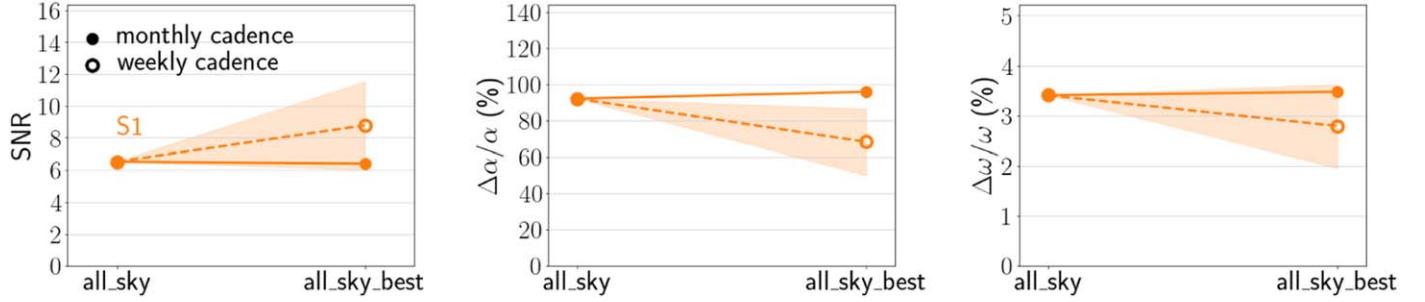


Figure 3. In the left panel, we plot the S/Ns of S1 (based on the binary candidate OJ 287) as observed by two types of PTAs, `all_sky` and `all_sky_best`, both with a 15 year-long baseline. Recall that S1 is not observed by a targeted campaign (see the text for details). In the middle and right panels, we show the measurement uncertainties of α and ω , respectively. The solid line compares the mean S/Ns or parameter uncertainties for the two PTAs at the same, monthly observing cadence, whereas the dashed line compares `all_sky` at a monthly cadence vs. `all_sky_best` at a weekly cadence. The shaded bands represent 1σ uncertainties from the realizations, except for the monthly cadence case which does not have sufficient $S/N > 5$ realizations to compute uncertainties.

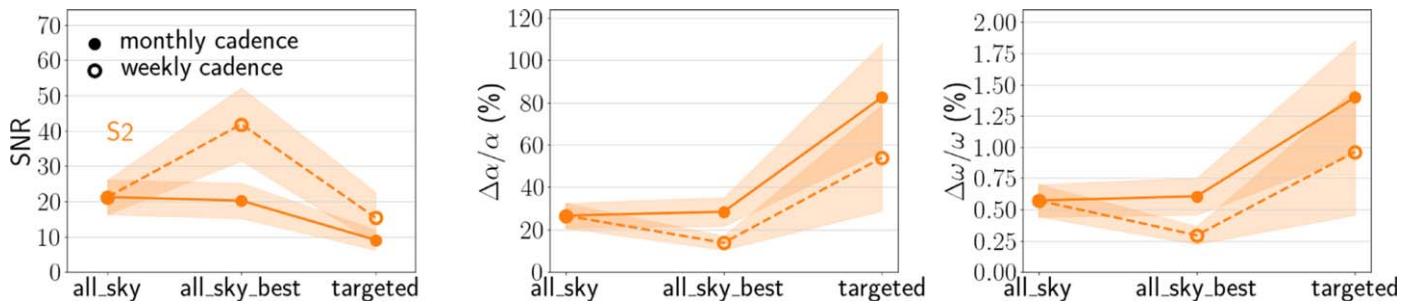


Figure 4. Same as Figure 3, but for the mock source S2. Additionally, the results of the `targeted` campaign are included.

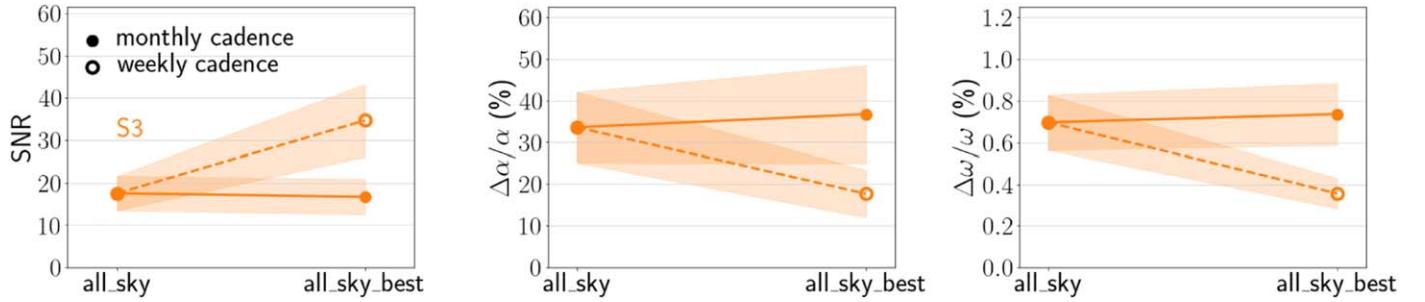


Figure 5. Same as Figure 3, but for S3. Recall that S3 is also not observed by a `targeted` campaign.

parameters are assumed to be measured from standard methods.) First, an observed variability period caused by modulated accretion in the binary may not directly correspond to its intrinsic orbital period, since this relationship may be mass ratio-dependent and the observed period may be a few times the orbital period (see, e.g., D’Orazio et al. 2013; Farris et al. 2014; Noble et al. 2021); therefore, an EM-based frequency measurement would have an (underestimated) uncertainty of 100%, which is ~ 2 orders of magnitude larger than the GW-basement measurement (right panel in Figure 4). Second, standard BH mass measurements suffer from a systematic uncertainty of ~ 0.3 dex (e.g., Shen 2013); combined with the aforementioned uncertainty on ω , this translates to a $\Delta\alpha/\alpha$ of approximately 120%, which is a few times higher than the GW-based measurement (middle panel in Figure 4). These precise GW parameter measurements can allow meaningful comparisons with EM-based measurements as a robust test of the binary model, breaking the degeneracies arising from EM-only observations, and testing the predictions of SMBHB theory.

It is however important to note that `all_sky` boasts the highest S/N simply due to the fact that it is timing ~ 6 times more pulsars (than `all_sky_best`) and that the combined contribution of the bottom 85% of the pulsars to the S/N only increases the S/N by ~ 1 . Therefore it would be interesting to investigate whether `all_sky_best` or `targeted` is a more efficient approach to achieve a similar scientific outcome.

To do so, we compare the results of the high-cadence, weekly `all_sky_best` and `targeted` campaigns with the (monthly) `all_sky` campaign. As shown in Figure 4, we predict that `all_sky_best` results in the highest S/N and best parameter measurements; this is despite the fact that `all_sky_best` costs $\sim 40\%$ less total telescope time than `all_sky`. Likewise, a high-cadence `targeted` campaign requires only of order a third of the time as `all_sky`, but achieves a similar S/N. Furthermore, both high-cadence campaigns can measure α and ω at precise levels, which are better than typical EM measurements as discussed previously. Therefore, we conclude that `all_sky_best` and `targeted` are both possible, if not preferable, alternatives to `all_sky`.

3.3. PTA Observations of a Second Mock Binary Candidate

Finally, we examine the detection and parameter estimation prospects of S3. As we show in Figure 5, the expected outcomes of `all_sky` and `all_sky_best` follow the same trend; namely, `all_sky` results in a higher S/N for the same cadence, but the high-cadence `all_sky_best` campaign outperforms `all_sky` at a lower cost of total telescope time. Further, despite the slightly lower S/N than S2 (consistent with the difference in sky location), the parameters of S3 can still be measured to moderate to high precision.

As discussed earlier, the sky location of S3 may be more representative of that of a likely EM candidate: it is located where it is more likely to be found by EM observatories, including all-sky surveys; on the other hand, it is not at such a fortuitous location that it happens to be in close proximity to a large number of (high quality) MSPs. Fortunately, its detection prospects are nonetheless good in both all-sky campaigns and are only mildly impacted by its “inferior” location. This suggests that, contrary to conventional wisdom, dedicated timing campaigns targeting pulsars near likely sources may not be necessary, efficient, or possible for CW observations.

3.4. Effects of the GWB and Pulsar Red Noise on the Detectability of S1–S3

Recently, a red noise process which has a common spectrum across pulsars has been observed in regional PTAs and the combined International Pulsar Timing Array data set (Arzoumanian et al. 2020; Chen et al. 2021; Goncharov et al. 2021; Antoniadis et al. 2022). While the PTAs have not confirmed this common red noise process as the GWB, it already needs to be accounted for in more recent CW searches (e.g., Arzoumanian et al. 2023) and will have an effect on PTA’s sensitivity to CWs, especially at lower frequencies. Further, some level of red noise is present in many, if not most, MSPs, which would further compromise the observability of CW sources.

While our Fisher matrix methodology does not handle red noise, in this section, we seek to understand the potential effect of the GWB and pulsar red noise on the detectability of CW sources using simulated PTA sensitivity curves and speculate their implications for binary parameter estimation.

The GWB has a power-law power spectrum of the form

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

where we adopt the median value of the amplitude from the NANOGrav 12.5 yr analysis: $A_{\text{GWB}} = 1.92 \times 10^{-15}$ (Arzoumanian et al. 2020) and $\alpha = -2/3$ (Phinney 2001).

The power spectrum of the pulsar red noise is given by (e.g., Alam et al. 2021)

$$P(f) = A_{\text{RN}}^2 \left(\frac{f}{\text{yr}^{-1}} \right)^{\gamma_{\text{RN}}},$$

where A_{RN} and γ_{RN} are the amplitude and the index, respectively. We simulate A_{RN} and γ_{RN} for each mock pulsar using the model developed by Shannon & Cordes (2010). They determined that the natural log of the measured, red timing noise after a second-order polynomial fit is normally distributed as $\ln(\sigma_{\text{TN},2}) \sim \mathcal{N}(\ln(\hat{\sigma}_{\text{TN},2}), \delta)$. $\hat{\sigma}_{\text{TN},2}$ is the post-fit red noise

scaled to the spin period P and period derivative \dot{P} given by

$$\hat{\sigma}_{\text{TN},2} = 10^{15} C_2 P^{-\alpha} \left(\frac{\dot{P}}{P^2} \right)^\beta T_{\text{yr}}^\gamma \text{ ns}, \quad (1)$$

where C_2 , α , and β are estimated over the pulsar population and T_{yr} is the timescale of the residuals in years. We assume that the spin period derivatives are uncorrelated with spin period and are distributed log-normally with $\langle \log_{10} \dot{P} \rangle = -19.9$ and $\sigma_{\log_{10} \dot{P}} = 0.5$. We adopt the updated values for C_2 , α , β , δ , and γ , estimated for a larger sample of MSP red noise measurements, in Table 4 of Lam et al. (2017), row “MSP_{10,PPTA} + NANO.” The red noise index, $\gamma_{\text{RN}} = -(2\gamma + 1)$, is assumed to be -5.6 for all MSPs. By combining Equation (1) and Equation (15) of Lam et al. (2017) for the measured red noise in terms of A_{RN} and γ_{RN} , we can directly draw A_{RN} for each model pulsar where

$$f(A_{\text{RN}}) = \frac{1}{\sqrt{2\pi\delta^2}} \exp \left[- \left(\frac{(\ln(\hat{A}_{\text{RN}}) - \ln(A_{\text{RN}}))^2}{2\delta^2} \right) \right],$$

and

$$\hat{A}_{\text{RN}} = (-(1 + \gamma_{\text{RN}}))^{1/2} 10^{15} C_2 P^{-\alpha} \left(\frac{\dot{P}}{P^2} \right)^\beta \mu\text{s yr}^{1/2}.$$

To simulate the PTA sensitivity to CWs in the presence of these noise terms, we use the software package `hasasia`¹⁴ (Hazboun et al. 2019). We first compute the sensitivity curve by using the sky positions of simulated `all_sky` pulsars and only include white timing noise (see Section 2.1). Following the previous sections, we assume a 15 yr PTA with a monthly cadence. The resulting sensitivity is shown as the solid curve in Figure 6. For a monochromatic source (which we assume throughout the paper) emitting at f_0 with a strain amplitude h_0 , the expected S/N is $h_0 \sqrt{T_{\text{obs}}/S_{\text{eff}}(f_0)}$, where S_{eff} is the effective strain–noise power spectral density for the PTA and is related to the sensitivity curve $h_c(f) = \sqrt{f S_{\text{eff}}(f)}$ (Hazboun et al. 2019). Therefore, the resulting S/Ns for S1 and S2/S3 are 2.2 and 9.3, respectively, consistent with our results in Section 3 that the sources are detectable at the intermediate to high S/N level at monthly cadence.¹⁵

Next, we recompute the sensitivity in the presence of the GWB. The result is shown as the dashed curve. As expected, the GWB severely decreases the PTA sensitivity to CWs, especially at frequencies lower than $\sim 2 \times 10^{-8}$ Hz where it decreases by up to an order of magnitude. Consequently, S1 would not be detectable above the GWB (S/N = 0.2), while S2 and S3 may still be detectable, albeit with a lower S/N (2.7).

Finally, we include the pulsar red noise and show the resulting sensitivity curve as the dashed–dotted curve. The effect is less pronounced, with the sensitivity decreasing by less than a factor of a few at all frequencies, resulting in an S/N = 0.2 (2.1) for S1 (S2/S3).

We further anticipate that PTA’s ability to estimate binary parameters would also decrease in the presence of red noise. This is based on the fact that at a given frequency, source

¹⁴ <https://github.com/Hazboun6/hasasia>

¹⁵ Note however that the sensitivity curve is averaged over the sky and is not directly applicable to targeted searches. Therefore, the sensitivity curves in Figure 6 are meant for comparison with each other for the purpose of understanding the effect of red noise.

detectability and parameter measurability closely track each other (see Figures 3–5), i.e., a lower S/N corresponds to a larger measurement uncertainty.

While our Fisher matrix method cannot handle red noise and hence our conclusions in this work regarding source detectability and observing strategies are drawn assuming only white noise (i.e., the solid curve in Figure 6), future work would be able to make more accurate predictions of the detectability of CW sources by taking into account additional noise, especially the GWB. In practice, the (relatively small) effect of pulsar red noise could be mitigated by only including pulsars with low red

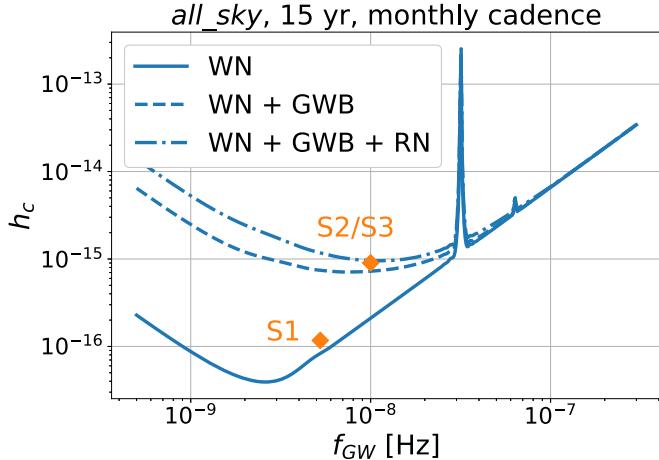


Figure 6. We show the expected decrease in sensitivity to CWs by the PTA in the presence of (1) only pulsar white noise (solid curve), (2) pulsar white noise and the GWB (dashed curve), and (3) pulsar white noise, GWB, and pulsar red noise (dashed-dotted curve). We superimpose the GW strains and frequencies of the three mock sources (orange diamonds) for visual purposes.

noise levels in the PTA, while the effect of the astrophysical GWB cannot be removed.

4. Summary and Conclusions

In this work, we have applied the Fisher matrix formalism to predict the SMBHB parameter estimation prospects of a future PTA with DSA-2000. We focus on the case where the sky location of the source is known via the prior identification of a possible EM counterpart, motivated by (1) the boosted detection S/N and parameter measurements by searching for the GW signal at a fixed sky location, and (2) the possibility of studying SMBHBs in a true multi-messenger fashion. Compared to previous work, we have taken a more realistic approach, taking into consideration factors such as practical and observational constraints of the PTA, the Galactic pulsar population, the number and timing precision of MSPs detectable by a telescope, and their spatial distribution.

As a case study, we first examined the detection prospects of the well-known SMBHB candidate OJ 287 (using S1 as a surrogate). We estimate that a future PTA with DSA-2000 may be able to detect the source with 15 yr of data (assuming only pulsar white noise). While the measurement uncertainty of the GW amplitude is large, which translates to a poor constraint on the system mass, the PTA constraint on the GW frequency is at a level which is more comparable to the EM equivalent, permitting the possibility of robustly testing its binary model through a joint interpretation of GW and EM data.

We then generalized this approach to a mock source, and our main results can be summarized as follows:

1. We have considered three types of PTAs: an all-sky PTA which observes all pulsars in our mock pulsar sample (`all_sky`), an all-sky PTA which only observes the best quality pulsars (`all_sky_best`), and a PTA which targets pulsars near an EM-selected binary candidate (`targeted`). At equal observing cadence, `all_sky` naturally results in the highest S/N and best parameter measurements. However, even with less total observing time, the high-cadence `all_sky_best` outperforms the other ones.
2. We have examined the detection and parameter estimation prospects of a mock source placed at different sky locations. If it is at a fortuitous location where a higher-cadence targeted campaign is possible (i.e., S2), the targeted campaign yields results similar to a lower-cadence `all_sky` campaign. More realistic is a sky location where `targeted` is not feasible due to the dearth of pulsars near the source (i.e., S3); this scenario may be more likely for EM-selected binary candidates. Fortunately, at such a location, its detection prospects with `all_sky` or `all_sky_best` may only decrease slightly compared to S2.
3. For a more typical EM candidate (represented by S2 and S3), the GW amplitude and frequency parameters may be constrained at a level which is better than standard EM measurements. Particularly, the GW frequency can be easily constrained within a few percent; this level of precision is consistent with similar estimates from previous work using both mock data analysis (e.g., Liu & Vigeland 2021) and Fisher matrix-based approaches (e.g., Sesana & Vecchio 2010).

Those results suggest the following strategies for a PTA to achieve the best science outcome at the lowest cost:

1. Ideally, an all-sky PTA which times of all pulsars detectable by pulsar surveys at a high cadence yields the best outcome. In reality, such a PTA costs more time than what is available on a single telescope or array. A PTA which only times the best pulsars can achieve similar results at a lower cost of total telescope time, even if observing at a higher cadence.
2. Applying the principle of choosing quality over quantity, an upgraded and downsized version of present-day PTAs would be capable of detecting EM-selected SMBHBs with a variety of source parameters and measuring their parameters at precision levels which are better than conventional EM measurements. Such a PTA can operate as a true GW observatory of SMBHBs and complement EM searches and observations with profound implications for the multi-messenger studies of SMBHBs.
3. A dedicated PTA which targets a small number of MSPs near an EM-selected candidate may not be feasible, since the sky locations where MSPs can be found and where EM binary candidates may be observed likely do not coincide. Fortunately, this may not be necessary either, because an all-sky PTA consisting of a small number of high quality pulsars is capable of observing a source at any location with good results.

We end with a caveat about our preference for `all_sky_best` over `all_sky`. We have assumed known sky locations and therefore have not considered the PTA measurement uncertainties of ϕ and θ in this work. However, it has been known that sky localization improves with the number pulsars in the PTA (Sesana & Vecchio 2010). Therefore, it cannot be determined from our study whether `all_sky_best` remains the best strategy for an unguided search over the whole sky without an EM counterpart beforehand. It would therefore be interesting to investigate the source localization capability of `all_sky_best` (with its more realistic pulsar spatial distribution and σ distribution) and to reexamine its overall single source detection capability.

This work is supported by the NANOGrav National Science Foundation Physics Frontiers Center award no. 2020265. The material is based upon work supported by NASA under award number 80GSFC21M0002. S.J.V. acknowledges support from NSF grant PHY-2011772. T.L. thanks Jolien Creighton, Jeff Hazboun, and David Kaplan for helpful suggestions. We thank the referee for useful comments that improved the paper.

Software: `astropy` (Astropy Collaboration et al. 2018), `FrequencyOptimizer` (Lam et al. 2018), `hasasia` (Hazboun et al. 2019), `matplotlib` (Hunter 2007), `PsrPopPy` (Bates et al. 2014).

ORCID iDs

Tingting Liu  <https://orcid.org/0000-0001-5766-4287>
Casey McGrath  <https://orcid.org/0000-0002-6155-3501>
Paul B. Demorest  <https://orcid.org/0000-0002-6664-965X>
Sarah J. Vigeland  <https://orcid.org/0000-0003-4700-9072>

References

Aggarwal, K., & Lorimer, D. R. 2022, arXiv:[2203.05560](https://arxiv.org/abs/2203.05560)

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

Albrecht, A., Bernstein, G., Cahn, R., et al. 2006, arXiv:[astro-ph/0609591](https://arxiv.org/abs/astro-ph/0609591)

Amaro-Seoane, P., Audley, H., Babak, S., et al. 2017, arXiv:[1702.00786](https://arxiv.org/abs/1702.00786)

Antoniadis, J., Arzoumanian, Z., Babak, S., et al. 2022, *MNRAS*, 510, 4873

Arzoumanian, Z., Brazier, A., Burke-Spolaor, S., et al. 2018, *ApJS*, 235, 37

Arzoumanian, Z., Baker, P. T., Blumer, H., et al. 2020, *ApJL*, 905, L34

Arzoumanian, Z., Baker, P. T., & Blecha, L. 2023, arXiv:[2301.03608](https://arxiv.org/abs/2301.03608)

Astropy Collaboration, Price-Whelan, A. M., Sipócz, B. M., et al. 2018, *AJ*, 156, 123

Bates, S. D., Lorimer, D. R., Rane, A., & Swiggum, J. 2014, *MNRAS*, 439, 2893

Burt, B. J., Lommen, A. N., & Finn, L. S. 2011, *ApJ*, 730, 17

Charisi, M., Bartos, I., Haiman, Z., et al. 2016, *MNRAS*, 463, 2145

Chen, S., Caballero, R. N., Guo, Y. J., et al. 2021, *MNRAS*, 508, 4970

Christy, B., Anella, R., Lommen, A., et al. 2014, *ApJ*, 794, 163

Cordes, J. M., & Lazio, T. J. W. 2002, arXiv:[astro-ph/0207156](https://arxiv.org/abs/astro-ph/0207156)

Dey, L., Gopakumar, A., Valtonen, M., et al. 2019, *Univ*, 5, 108

Dey, L., Valtonen, M. J., Gopakumar, A., et al. 2018, *ApJ*, 866, 11

D’Orazio, D. J., & Di Stefano, R. 2018, *MNRAS*, 474, 2975

D’Orazio, D. J., Haiman, Z., & MacFadyen, A. 2013, *MNRAS*, 436, 2997

D’Orazio, D. J., Haiman, Z., & Schiminovich, D. 2015, *Natur*, 525, 351

Ellis, J. A., Siemens, X., & Creighton, J. D. E. 2012, *ApJ*, 756, 175

Farris, B. D., Duffell, P., MacFadyen, A. I., & Haiman, Z. 2014, *ApJ*, 783, 134

Faucher-Giguere, C.-A., & Kaspi, V. M. 2006, *ApJ*, 643, 332

Foord, A., Gultekin, K., Reynolds, M., et al. 2017, *ApJ*, 851, 106

Goncharov, B., Shannon, R. M., Reardon, D. J., et al. 2021, *ApJL*, 917, L19

Graham, M. J., Djorgovski, S. G., Stern, D., et al. 2015a, *Natur*, 518, 74

Graham, M. J., Djorgovski, S. G., Stern, D., et al. 2015b, *MNRAS*, 453, 1562

Guo, H., Liu, X., Zafar, T., & Liao, W.-T. 2020, *MNRAS*, 492, 2910

Hallinan, G., Ravi, V., Weinreb, S., et al. 2019, *BAAS*, 51, 255

Hazboun, J. S., Romano, J. D., & Smith, T. L. 2019, *PhRvD*, 100, 104028

Hobbs, G. 2013, *CQGra*, 30, 224007

Hunter, J. D. 2007, *CSE*, 9, 90

Ivezic, Z., Kahn, S. M., Tyson, J. A., et al. 2019, *ApJ*, 873, 111

Kelley, L. Z., Blecha, L., Hernquist, L., Sesana, A., & Taylor, S. R. 2017, *MNRAS*, 471, 4508

Kelley, L. Z., Blecha, L., Hernquist, L., Sesana, A., & Taylor, S. R. 2018, *MNRAS*, 477, 964

Kelley, L. Z., D’Orazio, D. J., & Di Stefano, R. 2021, *MNRAS*, 508, 2524

Kramer, M., & Champion, D. J. 2013, *CQGra*, 30, 224009

Lam, M. T., Cordes, J. M., Chatterjee, S., et al. 2017, *ApJ*, 834, 35

Lam, M. T., McLaughlin, M. A., Arzoumanian, Z., et al. 2019, *ApJ*, 872, 193

Lam, M. T., McLaughlin, M. A., Cordes, J. M., Chatterjee, S., & Lazio, T. J. W. 2018, *ApJ*, 861, 12

Lehto, H. J., & Valtonen, M. J. 1996, *ApJ*, 460, 207

Liu, T., Gezari, S., Ayers, M., et al. 2019, *ApJ*, 884, 36

Liu, T., Gezari, S., Burgett, W., et al. 2016, *ApJ*, 833, 6

Liu, T., Gezari, S., Heinis, S., et al. 2015, *ApJL*, 803, L16

Liu, T., & Vigeland, S. J. 2021, *ApJ*, 921, 178

Lorimer, D. R. 2012, in IAU Proc. 8, Symp. S291, Neutron Stars and Pulsars: Challenges and Opportunities after 80 years (Cambridge: Cambridge Univ. Press), 237

Manchester, R. N., Lyne, A. G., Camilo, F., et al. 2001, *MNRAS*, 328, 17

McGrath, C., & Creighton, J. 2021, *MNRAS*, 505, 4531

McLaughlin, M. A. 2013, *CQGra*, 30, 224008

Noble, S. C., Krolik, J. H., Campanelli, M., et al. 2021, *ApJ*, 922, 175

Noble, S. C., Mundim, B. C., Nakano, H., et al. 2012, *ApJ*, 755, 51

Phinney, E. S. 2001, arXiv:[astro-ph/0108028](https://arxiv.org/abs/astro-ph/0108028)

Pol, N. S., Taylor, S. R., Kelley, L. Z., et al. 2021, *ApJL*, 911, L34

Saade, M. L., Stern, D., Brightman, M., et al. 2020, *ApJ*, 900, 148

Sesana, A., & Vecchio, A. 2010, *PhRvD*, 81, 104008

Severgnini, P., Cicone, C., Della Ceca, R., et al. 2018, *MNRAS*, 479, 3804

Shannon, R. M., & Cordes, J. M. 2010, *ApJ*, 725, 1607

Shen, Y. 2013, *BAI*, 41, 61

Sillanpaa, A., Haarala, S., Valtonen, M. J., Sundelius, B., & Byrd, G. G. 1988, *ApJ*, 325, 628

Taylor, S. R., Vallisneri, M., Ellis, J. A., et al. 2016, *ApJL*, 819, L6

Vallisneri, M. 2008, *PhRvD*, 77, 042001

Valtonen, M. J., Dey, L., Gopakumar, A., et al. 2021, *Galax*, 10, 1