



Chinese Pulsar Timing Array Upper Limits on Microhertz Gravitational Waves from Supermassive Black-hole Binaries Using PSR J1713+0747 FAST Data

R. Nicolas Caballero^{1,2} , Heng Xu¹ , Kejia Lee^{1,3,4}, Siyuan Chen^{2,3,5} , Yanjun Guo^{1,6}, Jinchun Jiang¹ , Bojun Wang¹,
Jiangwei Xu^{1,2,3}, and Zihan Xue^{1,2,3}

¹ National Astronomical Observatories, Chinese Academy of Sciences, Beijing 100101, China; caballero.astro@gmail.com, hengxu@bao.ac.cn, kjlee@pku.edu.cn

² Kavli Institute for Astronomy and Astrophysics, Peking University, Beijing 100871, China

³ Department of Astronomy, School of Physics, Peking University, Beijing 100871, China

⁴ Beijing Laser Acceleration Innovation Center, Beijing 101400, China

⁵ Shanghai Astronomical Observatory, Chinese Academy of Sciences, Shanghai 200030, China

⁶ Max-Planck-Institut für Radioastronomie, Auf dem Hügel 69, 53121 Bonn, Germany

Received 2024 March 13; revised 2024 December 29; accepted 2025 February 11; published 2025 March 26

Abstract

We derive the gravitational-wave (GW) strain upper limits from resolvable supermassive black-hole binaries using the data from the Five-hundred-meter Aperture Spherical radio Telescope, in the context of the Chinese Pulsar Timing Array project. We focus on circular orbits in the μHz GW frequency band between 10^{-7} and 3×10^{-6} Hz. This frequency band is higher than the traditional pulsar timing array band and is less explored. We used the data of the millisecond pulsar PSR J1713+5307 observed between 2019 August and 2021 April. A dense observation campaign was carried out in 2020 September to allow for the μHz band coverage. Our sky-averaged continuous source upper limit at the 95% confidence level at 1 μHz is 1.26×10^{-12} , while the same limit in the direction of the pulsar is 4.77×10^{-13} .

Key words: (stars:) pulsars: general – gravitational waves – methods: statistical – methods: observational – methods: data analysis

1. Introduction

The primary work on the search for gravitational waves (GWs) at nHz frequencies over the last decades has been conducted using radio pulsars. In particular, the work is based on the employment of pulsar timing arrays (PTAs; Foster & Backer 1990), which are ensembles of millisecond pulsars (MSPs) at different sky locations. MSPs are chosen for the work of GW searches due to their remarkably stable rotations (see, e.g., Verbiest et al. 2009). At the same time, they allow for the highest-precision recording of pulse arrival times (times-of-arrival; ToAs), which can be of the order of tens of nanoseconds, as is the case with data in this paper. By comparison, the canonical pulsar population, which comprises the younger pulsars, exhibits a substantial amount of rotational irregularities (Hobbs et al. 2010; Parthasarathy et al. 2019). These manifest in the form of regular glitches, nulling and mode changes, and in strong stochastic irregularities known as timing noise, also referred to as red noise (RN) and spin noise, which reduce the sensitivity to GW signals (see, e.g., Caballero et al. 2016). On the other hand, MSPs have very smooth and regular rotations which result in very small levels of stochastic timing noise (e.g., Verbiest et al. 2009; Caballero et al. 2016; Lentati et al. 2016), and only a couple of glitches have been

previously observed in MSPs (Cognard & Backer 2004; McKee et al. 2016). Profile changes in MSPs are generally subtle, in the form of profile instabilities, and can be addressed in standard timing methods (e.g., Kramer et al. 1999; Liu et al. 2015; Shannon et al. 2016; Kerr et al. 2020). Profile changes that are severe and affect significantly the long-term timing of MSPs are rare, but PSR J1713+0747 has recently exhibited such a change which led to the limitation in the total data set used in this study (see Section 2).

PTAs are mostly sensitive in the nHz part of the GW spectrum, where the primary goal is to observe GWs emitted by supermassive black-hole binaries (SMBHBs). Apart from GWs from resolvable SMBHBs (see, e.g., Thorne 1989; Jenet et al. 2004; Sesana & Vecchio 2010; Lee et al. 2011), PTAs are also sensitive to stochastic GW background (GWB) signals. Such stochastic signals may be the result of the superposition of GWs from the cosmic population of SMBHBs (Rajagopal & Romani 1995; Jaffe & Backer 2003), cosmic strings (Kibble 1976; Sanidas et al. 2012), and relic GWB from quantum fluctuations in the early Universe, especially from the inflationary era (Kibble 1976). Recently, the Chinese Pulsar Timing Array (CPTA, Lee 2016) and other PTA groups have reported statistical evidence for the existence of a nHz GWB

(Agazie et al. 2023; EPTA Collaboration et al. 2023; Reardon et al. 2023; Xu et al. 2023), though each group’s PTAs have not yet detected a single GW source.

The observational GW frequency window for PTAs is determined by the total timespan of the data set, T , and the observational cadence, Δt . The upper frequency limit is determined by the Nyquist theorem and is $f_{\max} \approx 1/(2\Delta t)$ for regular sampling. For irregular sampling, which is the case for real pulsar timing data, the upper frequency limit can, in principle, be much higher than the “traditional” Nyquist frequency (based on arguments similar to that of Koen 2006). In practical PTA GW detection, on the other hand, the interesting upper frequency limit is mainly determined by the data sensitivity. The lower recoverable frequency is $f_{\min} \approx 1/T$, as any power in the pulsar ToAs below this frequency is effectively absorbed by the quadratic term of the pulsar’s rotational-frequency derivative (see Lee et al. 2012). This term is always present in timing models of radio pulsars, and accounts for the slowing-down of the pulsar rotation as it loses rotational energy through the emitted radiation.

Typical pulsar data sets have decadal timespans for a large number of pulsars, and have the highest GW sensitivity around the lowest frequencies of the spectrum. Single-telescope data seldom can have cadence below ~ 1 week, but multi-telescope data combination can achieve daily cadence. In this way, PTAs can probe single source GWs (SSGWs) up to the μHz band. So far, the μHz regime for SSGW from SMBHBs with pulsar timing has been explored using single-pulsar data (Dolch et al. 2014; Yi et al. 2014; Perera et al. 2018). As a consequence, these studies can only provide GW amplitude strain upper limits, as GW detection requires observations by multiple pulsars. The current paper also places an upper limit using single-pulsar data.

PSR J1713+0747 is one of the most observed and thoroughly studied MSPs. Due to its sky location, that permits observations from both the Northern and Southern Hemispheres, its long-term stability and high brightness that results in high ToA measurement precision, PSR J1713+0747 has always been a primary PTA target for GW search efforts. It has been continuously observed by the three founding PTAs and has also been the subject of the International Pulsar Timing Array (IPTA) 24 hr global observing campaign (Dolch et al. 2014), with nine participating telescopes observing the source as it became visible in their local sky, allowing rare investigations of noise properties on intermediate timescales. In this paper, we use PSR J1713+0747 data from the Five-hundred-meter Aperture Spherical radio Telescope (FAST; Jiang et al. 2019) in China to place new and independent upper limits for the strain amplitude of SSGWs from circular SMBHBs, in the μHz regime.

The rest of the paper is organized as follows: In Section 2, we describe the data used for the study. In Section 3, we present the pulsar timing model and the data’s noise properties,

while in Section 4 we derive the SSGW strain upper limits. Finally, we discuss our conclusions in Section 5.

2. Data

PSR J1713+0747 was observed using the 19-beam receiver of the FAST radio telescope from 2019 August to 2021 April with a cadence of about 1 week, and we conducted a dense observation campaign with nearly 2 day cadence in September of 2020. Since this pulsar was found to show a profile changing event (Xu et al. 2021), the data after 2021 April 16th are not included in this work. The receiver is centered at 1.25 GHz with a bandwidth of 500 MHz. We used the ROACH2 system to record the data in search mode with a $49.152 \mu\text{s}$ sampling time and 4094 channel mode. Each observation was typically 20–30 minutes, although some were longer, up to 3 hr. In total, we include 56.2 hr of observations. The data were off-line folded every 30 s using the standard pulsar data reduction software DSPSR (van Straten & Bailes 2011), and the polarization calibration was performed with the pulsar analysis software package PSRCHIVE (Hotan et al. 2004) using the single-axis model. Following polarization calibration, any data corrupted by radio frequency interference (RFI) were removed manually. The clean data were finally time-integrated every 15 or 20 minutes and the whole band was split into 16 subbands to allow for the measurement and corrections of frequency dependent noise, e.g., dispersion measure (DM) variation. The subband ToAs were generated following the Fourier-domain algorithm (Taylor 1992) with the PAT routine in PSRCHIVE. ToAs with signal-to-noise ratio (S/N) lower than 8 were removed from the subsequent analysis. In total, 3436 ToAs from 65 observation session were used in this paper.

3. Pulsar Timing Model and Noise Properties

The initial timing analysis of the data was performed using TEMPO2 (Hobbs et al. 2006), providing a phase-coherent timing solution for the data. As TEMPO2 uses a linearized model and the solution is calculated via least-squares fit, the presence of any RN can bias the estimated values and uncertainties of the timing parameters (Coles et al. 2011; van Haasteren & Levin 2013). We followed a Bayesian approach using TEMPONEST (Lentati et al. 2014) for the noise analysis, searching for noise components that are standard in pulsar timing (see for example Lentati et al. 2016, and see further details in the next section). We used Bayesian model selection using Bayes factors to determine how many components are supported by the data. Detailed timing and noise analysis are presented for PSR J1713+0747 and all MSPs timed by the CPTA in separate papers (Chen, S. et al. 2025, in preparation; Xu, H. et al. 2025, in preparation). Figure 1 shows the timing residuals from a generalized-least-squares fit of the timing model including the noise model, as well as the same residuals after subtracting the stochastic effects of the interstellar

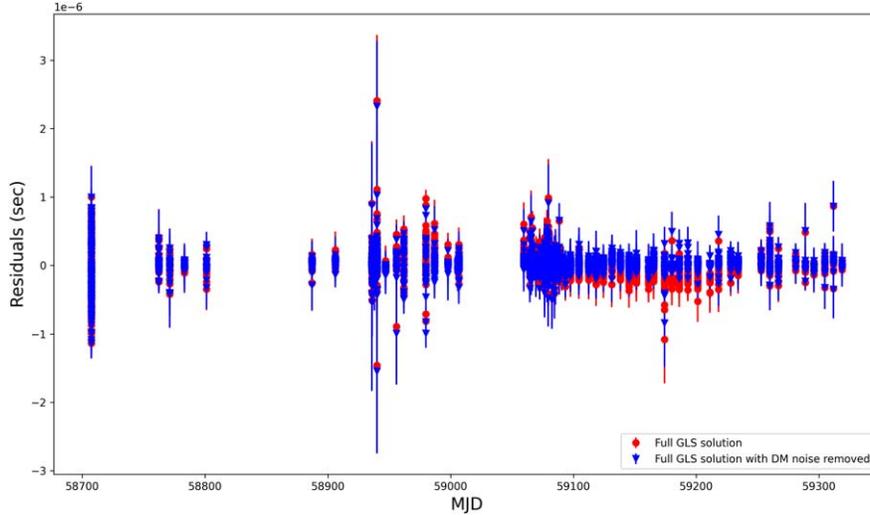


Figure 1. Timing residuals (red circles) of PSR J1713+0747 after a generalized-least-squares fit of the timing model with TEMPO2. The covariance matrix is constructed based on the maximum likelihood values of the inferred noise parameters. We also overplot (blue triangles) the same residuals after subtracting the stochastic noise due to DM variations. The weighted rms values of the two timing residuals are 115 and 93 ns, respectively.

medium (see next section). The latter residuals have a root mean square (rms) of only 93 ns.

3.1. Noise Model

In general, for a one-dimensional (1D) time series, noise can be divided into time-uncorrelated and time-correlated noise. The former are referred to as white noise parameters and they include: (i) a corrective multiplying factor on the ToA uncertainties calculated during the ToA estimation (Taylor & Weisberg 1982) named EFAC, (ii) a factor added in quadrature to the ToA uncertainties (Ekers & Moffet 1968; Shannon & Cordes 2010), named EQUAD, which tracks effects of pulse phase jitter, and (iii) a parameter named ECORR (Arzoumanian et al. 2016), which models the pulse jitter effects that are specifically correlated in subband ToAs across the observing bandwidth. The interested reader can read more details in the CPTA pulsar-noise paper (S. Chen et al. 2025, in preparation).

The time-correlated noise components included are: (i) the red (achromatic) noise, stochastic noise presumably associated with pulsar spin irregularities (Shannon & Cordes 2010) and (ii) stochastic DM variations (see You et al. 2007). DM is the column density of free electrons between the pulsar and Earth, causing a delay $\delta\tau_{\text{dm}}$ in the radio signal proportional to the inverse square of the observing frequencies, ν (i.e., $\delta\tau_{\text{dm}} \propto \nu^{-2}$), following the dispersion law of cold plasma (e.g., Landau & Lifshitz 1960). These types of stochastic, correlated noise are typically assumed to be wide-sense stationary signals, and are modeled in the time-frequency domain (Lentati et al. 2013). Following the literature, we use a

power-law spectrum for RN of the form,

$$S(f) \propto A^2 f^{-\gamma}, \quad (1)$$

where $S(f)$ is the power spectral density, A is the spectrum's amplitude and γ is the spectral index.

3.2. Bayesian Parameter Estimation and Model Selection

We use Bayesian inference to perform parameter estimation, where the model selection is based on the Bayes evidence comparison as previously described (Lentati et al. 2016; Chalumeau et al. 2022). We employed the software package TEMPO2 to perform the Bayesian noise modeling. Here, we give a brief mathematical description of the process, which is also (partially) applied in the calculations of the SSGW amplitude limits discussed in the next section.

Bayes' theorem relates the posterior distribution (i.e., the “distribution” of inferred parameters, given the data) and the likelihood (the “distribution” of data, given the model), as

$$P(\lambda|\mathcal{D}) = \frac{L(\mathcal{D}|\lambda)p(\lambda)}{Z}, \quad (2)$$

where \mathcal{D} denotes our data, H denotes the hypothesis (i.e., the model), and λ denotes the model parameters. P is the posterior distribution, L is the likelihood function, p is the prior distribution, and Z is the Bayesian evidence.

The likelihood function of pulsar timing was assumed to be Gaussian (van Haasteren et al. 2009), i.e.,

$$L = \frac{e^{-\frac{1}{2}(t-t_{\text{m}})^T C^{-1}(t-t_{\text{m}})}}{\sqrt{(2\pi)^n |\mathbf{C}|}}, \quad (3)$$

where t is the recorded ToA, t_{tm} is the ToA predicted by the timing model, ($t - t_{\text{tm}}$ is, therefore, equal to the timing residuals) C is the total noise covariance matrix, and n is the number of ToAs in the data.

The prior distribution represents our prior knowledge or belief for the distribution of the unknown parameter. This choice therefore can add information to the parameter estimation and can lead to erroneous results if our belief regarding the parameter distribution is informative but wrong. Theoretically, Jeffreys prior represents a choice of prior being the least informative. However its formal calculation can be difficult and approximations are used. For example, for scale invariant parameters, such as the amplitude of the power-law RN spectrum, we use a uniform prior (flat) in the log space as an uninformative prior (Gregory 2005), while one would use a uniform prior in linear space to estimate conservative upper limits.⁷

The evidence Z is accurately estimated when we want to perform model comparison and selection in order to decide which model better fits the data. The evidence is the integrated likelihood over the prior, and for N parameters is defined as,

$$Z = \int L(\mathcal{D}|\lambda)p(\lambda) d^N\lambda. \quad (4)$$

Larger Z values signify more favorable models. As a metric to choose whether a more complicated model is required, Z should be larger than a threshold. We can measure this using the posterior odds ratio for two models, e.g., 0 and 1,

$$R = \frac{Z_1 p_1}{Z_0 p_0} = \mathcal{B}_{10} \frac{p_1}{p_0}, \quad (5)$$

where the ratio of the evidences \mathcal{B}_{10} is called the Bayes factor. As the prior distributions for all models we discuss in model selection analysis in this paper are the same, the Bayes factor is equal to the posterior odds ratio. The distribution function of Bayes factor depends on the prior and likelihood. As a rule of thumb, a threshold of $\log_{10}(\mathcal{B}_{10}) > 2$ is taken in the current paper to choose the preferred model, following Kass & Raftery (1995).

3.3. Noise Analysis

In the noise analysis, we used the priors shown in Table 1, and employed TEMPONEST's importance nested sampling algorithm to perform posterior and evidence computation (see Skilling 2004; Feroz & Hobson 2008). Our model selection process uses Bayes factors to select our preferred noise model from a set of nested models, where the simplest models only include EFAC. We proceed trying the models with white-noise

⁷ By taking the "log," the physical units (e.g., Volts) of signal amplitude becomes dimensionless, and one removes the information inherited in the units system, i.e., the scale. Taking the uniform prior, one "believes" in higher value and assigns a higher probability to it, which aids in calculating conservative upper bounds.

Table 1
Ranges and Types for the Priors Used in the Single-pulsar Noise Analysis of PSR J1713+0747

Parameter	Range	Type
EFAC	[0.1, 5]	Uni
$\log_{10}(\text{EQUAD})$	[-9, -5]	log-Uni
$\log_{10}(\text{ECORR})$	[-9, -5]	log-Uni
$\log_{10}(A_{\text{DM}})$	[-20, -8]	log-Uni
γ_{DM}	[0, 7]	Uni
$\log_{10}(A_{\text{RN}})$	[-20, -8]	log-Uni
γ_{RN}	[0, 7]	Uni

Note. The prior types Uni and log-Uni correspond to uniform priors in linear and logarithmic space, respectively.

parameters only as EFAC+EQUAD, EFAC+ECORR, EFAC+EQUAD+ECORR. The analysis results suggested that all three terms were required. To test whether the data support red and/or DM stochastic noise, we first assume EFAC+RN and EFAC+DM models with increasing number of frequency bins, to find the optimal number of frequency bins for the RN and DM power-law spectrum models using Bayes factors evaluations. Once that number was defined we tested the EFAC+RN, EFAC+DM and EFAC+RN+DM models with the respective optimal numbers of frequency bins. The analysis suggested that the data are sufficiently modeled using white noise+DM noise, and that there is not sufficient support to add the RN noise component. We note that this does not mean that an RN component, including a GWB, is not present. It means that if present, the current data set is not sufficient to detect it with such high statistical significance, and therefore model selection prefers the simpler model that does not include it. This is expected for a GWB from GW-driven SMBHBs whose stochastic signal has a spectral index of 4.33, as a longer data span would be required to detect it: see for example Jenet et al. (2005) where (simulated) data similar to CPTA data require $\gtrsim 3$ yr for such a GWB detection, depending on the exact pulsar number and noise properties, and Xu et al. (2023) for the case of CPTA data. The detected DM noise in the current data is a shallow power-law with spectral index of ≈ 1.7 (see Table 2), thus allowing its detection at the higher frequencies where our data are more sensitive. We confirmed that from all the white-noise models with DM noise added, the most supported model is the EFAC+EQUAD+ECORR+DM model. We then confirmed that changing the number of frequency bins for the RN does not improve the support to add the RN in the EFAC+EQUAD+ECORR+DM model, and that the optimal number of DM noise frequency bins (43) remains the same with the optimal model as in the case of the simpler EFAC+DM model. Thus, the DM power-law signal spans the frequency spectrum from $1/T$ (1.9×10^{-8} Hz) to the highest frequency of $1/43 \text{ day}^{-1}$ (2.7×10^{-7} Hz).

Table 2

Overview of the Results from the Single-pulsar Noise Analysis of PSR J1713+0747, for the Optimal Noise Model (See Main Text for Details)

Parameter	Median	ML
EFAC	$1.037^{+0.017}_{-0.016}$	1.039
$\log_{10}(\text{EQUAD})$	$-7.531^{+0.021}_{-0.022}$	-7.533
$\log_{10}(\text{ECORR})$	$-7.159^{+0.0255}_{-0.0246}$	-7.162
$\log_{10}(A_{\text{DM}})$	$-12.176^{+0.163}_{-0.150}$	-12.196
γ_{DM}	$1.804^{+0.472}_{-0.395}$	1.678

Note. The table lists the median values of the 1D marginalized posterior distribution for each parameter, with the (asymmetric) 68% uncertainties, as well as the maximum likelihood (ML) values.

Here, we present the results from the optimal model’s noise analysis. Figure 2 is the corner plot showing the two-dimensional (2D) and marginalized 1D histograms of the noise parameters. We note that while EFAC and EQUAD appear correlated, ECORR does not appear to correlate with either EFAC or EQUAD. On the other hand, all three of these parameter show correlation with the spectral index of DM. This happens due to the very shallow DM power-law, i.e., because the DM signal has significant power in higher frequencies. Table 2 gives the median and maximum likelihood values of the noise-parameter posterior distributions.

4. Continuous Gravitational-wave Strain Upper Limits

In the context of PTAs, the problem of SMBHB GW detection has been discussed extensively (e.g., Jenet et al. 2004; Lee et al. 2011; Sesana 2013). In general, GWs from a SMBHB have two terms, which we refer to as the “Earth term” and the “pulsar term,” which quantify the spacetime distortion of the passing GW at the Earth’s and at the pulsar’s vicinities, respectively. The two-component SSGW signal is a direct result of the large (thousands of years) time-delay between the pulsar and the Earth. In this work, we focus on the case of monochromatic GWs and neglect the evolution of the GW source. We adopt the mathematical framework described in Lee et al. (2011). The monochromatic GW is then described by seven parameters, namely the GW strain amplitude (A), the two GW-source position sky coordinates, namely the right ascension (RA) and declination (DEC), the GW angular frequency (ω_{gw}), the orbital inclination angle (i), the orientation angle of ascending node (ψ), and the GW initial phase (ϕ_0) at reference epoch t_0 . The term “monochromatic” refers to signals where we have in effect the condition $\delta f \ll \omega_{\text{gw}}/2\pi$, where δf is the difference between the frequency of the Earth and pulsar terms.

As only one pulsar is considered, our SSGW analysis is limited to the derivation of robust upper limits for the GW strain amplitude. We have employed two methods to estimate the upper limits. We use (1) analytic calculations based on the

Cramér–Rao lower bound (CRLB) and (2) Bayesian inference analysis. In contrast to the Bayesian inference, the CRLB calculations require less computational resources and are used to cross-check the Bayesian result, as the CRLB provides the theoretically best possible result (Fisz 1963) and the Bayesian-inference result must be “worse” than the CRLB in any case (see discussion in Lee et al. 2011; Caballero et al. 2016). We are interested in two results: (i) the sky-averaged and (ii) the best-sky. Because we only use one pulsar, our analysis is insensitive to the sky-direction opposite to the pulsar position (henceforth anti-pulsar direction), while it is optimal around the direction of the pulsar. For this reason, when computing the all-sky upper limit, we exclude an area around the anti-pulsar direction (see Section 4.2). All of the SSGW analysis was performed using the PTA data analysis package FORTYTWO⁸ (Caballero et al. 2016) which performs pulsar-timing noise and GW analysis.

4.1. Cramér–Rao Lower-bound

The calculation of GW sensitivity curves for PTA data using the CRLB was previously described in Lee (2016), Caballero et al. (2016). The CRLB is defined as the inversion of the Fisher information matrix, \mathcal{I} . Given the likelihood function, $L(\boldsymbol{\lambda}, \mathbf{x})$, where \mathbf{x} is the data and $\boldsymbol{\lambda}$ are the model parameters, the CRLB is described by

$$\text{Cov}(\boldsymbol{\lambda}) = \langle \sigma_{\lambda_i} \sigma_{\lambda_j} \rangle \geq \mathcal{I}_{ij}^{-1}, \quad (6)$$

where the indices i and j denote the different parameters, $\text{Cov}(\boldsymbol{\lambda})$ is the covariance of the parameters $\boldsymbol{\lambda}$, and \mathcal{I}_{ij} is defined as

$$\mathcal{I}_{ij} = \left\langle \frac{\partial \ln L(\mathbf{x}, \boldsymbol{\lambda})}{\partial \lambda_i} \frac{\partial \ln L(\mathbf{x}, \boldsymbol{\lambda})}{\partial \lambda_j} \right\rangle \equiv - \left\langle \frac{\partial^2 \ln L(\mathbf{x}, \boldsymbol{\lambda})}{\partial \lambda_i \partial \lambda_j} \right\rangle. \quad (7)$$

For Gaussian likelihood functions, like the pulsar timing likelihood function, \mathcal{I} can be analytically calculated and the result is known as the Slepian–Bangs formula (Slepian 1954; Bangs 1971), that is,

$$\mathcal{I}_{ij} = \frac{\partial \mathbf{C}}{\partial \beta_i} \mathbf{C}^{-1} \frac{\partial \mathbf{C}}{\partial \beta_j} + \frac{1}{2} \text{tr} \left[\mathbf{C}^{-1} \frac{\partial \mathbf{S}(\boldsymbol{\lambda})^T}{\partial \lambda_i} \mathbf{C}^{-1} \frac{\partial \mathbf{S}(\boldsymbol{\lambda})}{\partial \lambda_j} \right]. \quad (8)$$

Here, β_i are the model parameters describing the covariance matrix, λ_i are the parameters describing the unknown waveform \mathbf{S} and tr is the matrix trace. In the context of SSGW in this paper, we use the reduced likelihood function with the timing parameters marginalized (van Haasteren et al. 2009), where we now separate the timing residuals in the residuals induced by stochastic processes, $\delta \mathbf{t}_s$ (for example RN or a stochastic GWB) and the deterministic residuals induced by the SSGW, $S(\boldsymbol{\lambda})$, thereby $\delta \mathbf{t} = \delta \mathbf{t}_s - S(\boldsymbol{\lambda})$. In our specific problem,

⁸ <https://github.com/caballero-astro/fortytwo>

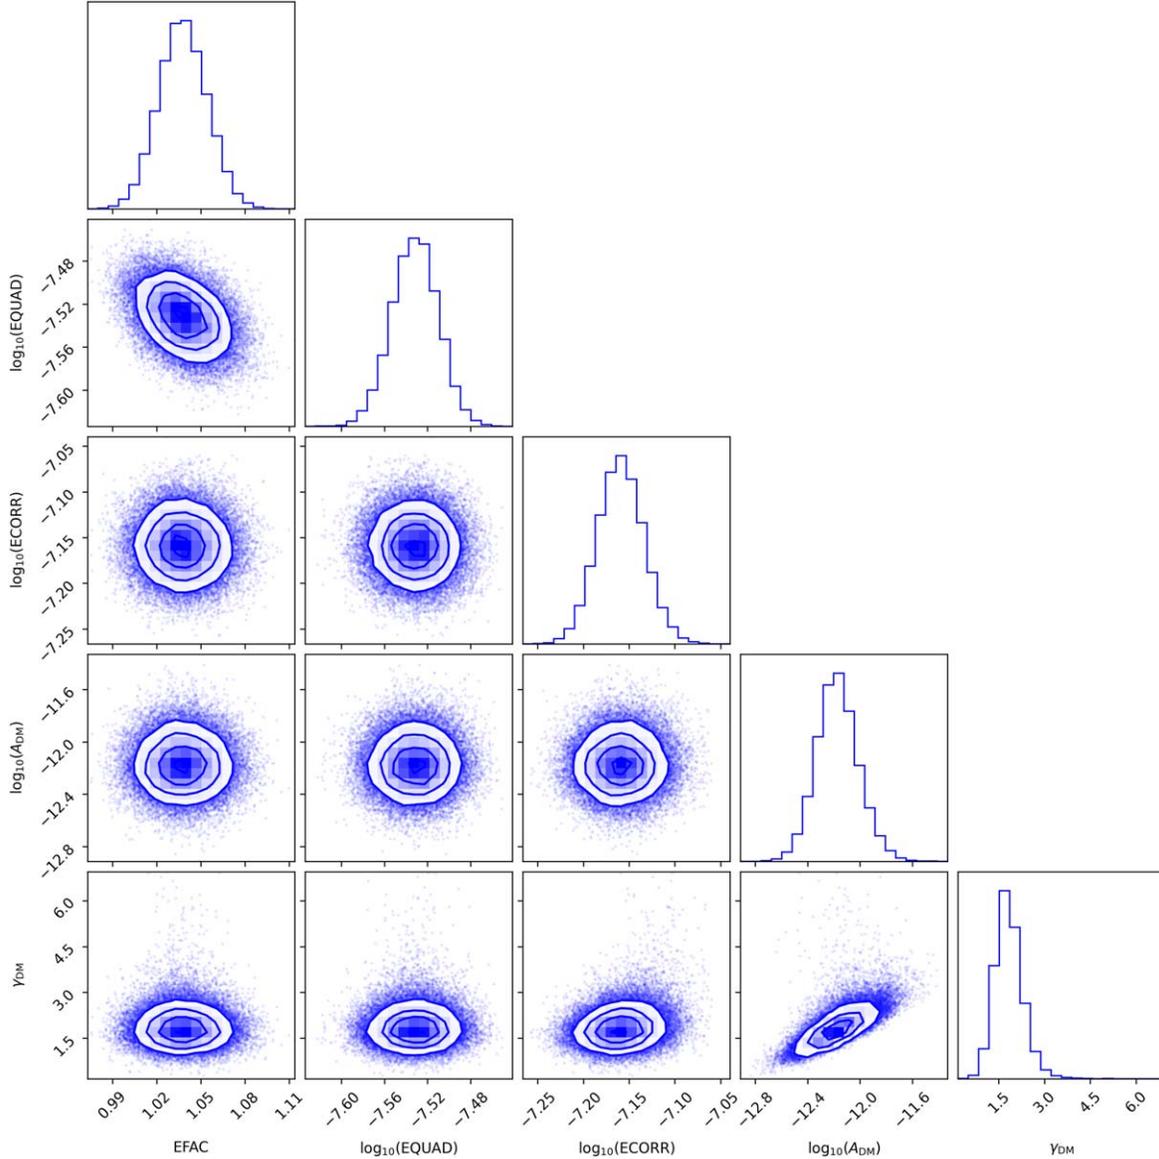


Figure 2. Corner plot for the posterior distributions of the single-pulsar noise analysis for the optimal noise model (see main text for details), displaying 2D and 1D marginalized distributions.

λ is the set of seven parameters required to describe the monochromatic SSGW. When focusing on the SSGWs, the terms with partial derivatives of \mathbf{C} are zero and Equation (8) reduces to

$$\mathcal{I}_{ij} = \frac{1}{2} \text{tr} \left[\mathbf{C}^{-1} \frac{\partial \mathcal{S}(\lambda)^T}{\partial \lambda_i} \mathbf{C}^{-1} \frac{\partial \mathcal{S}(\lambda)}{\partial \lambda_j} \right]. \quad (9)$$

For the CRLB calculations, the noise covariance matrix \mathbf{C} is fixed using the maximum likelihood values of the noise components, as derived from the pulsar noise analysis. We run 1000 analyses by selecting random values from uniform

distributions for the six SSGW parameters (excluding the amplitude), and calculate the covariance of the amplitude via Equation (6). The sky-position-parameter range is set accordingly depending on whether we are interested in the full sky-averaged amplitude limits or the “best-sky” limits. Figure 3 shows both results. Note that by default the estimated CRLB value corresponds to the 68% confidence level (which corresponds to the 1σ uncertainty, as we are dealing with a Gaussian likelihood). From this result, we may expect the sky-averaged and best-sky upper limits to have a factor of a few difference. More specifically, in the frequency range 10^{-7} – 10^{-6} Hz, the average difference is a factor of 3.35, while

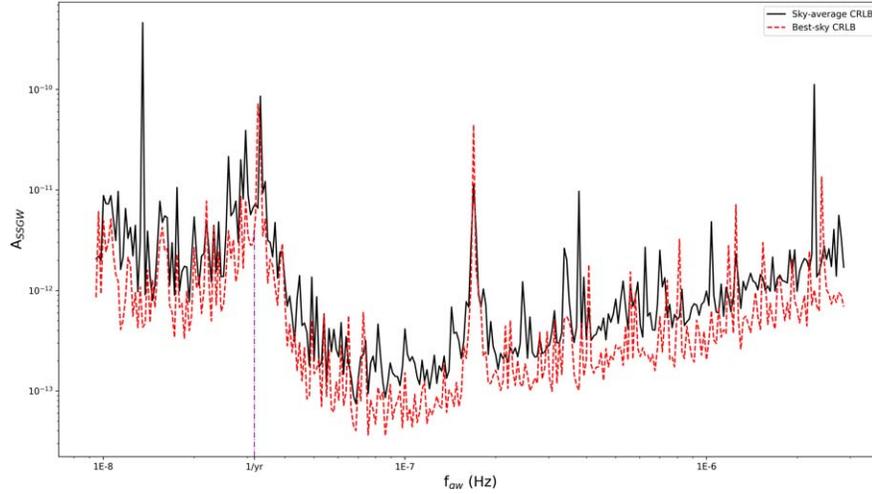


Figure 3. The sky-averaged (solid, black line) and best-sky (dashed, red line) CRLBs for the SSGW strain amplitude from GW-driven, circular SMBHBs. Note that the CRLB corresponds to the 1σ error bar. The plot shows the 1 yr^{-1} frequency with the dashed-dot, purple line. Note that due to the limited sensitivity below 10^{-7} Hz, the 1 yr^{-1} spike due to the Earth’s orbital frequency is not very narrow and distinct.

at frequency $1\ \mu\text{Hz}$, the difference is a factor of 2.2. We also note that for frequencies below 10^{-7} , due to the limited timespan, the sensitivity drops sharply and becomes non-comparable to previously published results. For this reason, we opted to only include frequencies above 10^{-7} in the Bayesian analysis, saving computational resources and time.

4.2. Bayesian Inference

Our main SSGW analysis follows the Bayesian inference principles outlined in Section 3.2. As we are only interested in amplitude upper limits, we do not need to perform any model selection. We perform two analyses. First, as in the derivation of the CRLB, we perform fixed-noise analysis using the maximum likelihood noise-parameter values. This allows a direct comparison of the methods. In addition, we perform a complementary analysis with varying noise for five key GW frequencies, namely 0.1, 0.3, 0.5, 1 and $2\ \mu\text{Hz}$. This allows for the derivation of the most conservative upper limit at our reference frequency of $1\ \mu\text{Hz}$. We use uniform priors for the amplitude with wide range in order to derive robust upper limits. The prior ranges are presented in Table 3.

The sky-averaged results are derived by averaging the upper limits derived in 400 sky cells, formed by having a 20×20 grid which is equally large in RA and $\cos(\text{DEC})$. As explained earlier, motivated by the CRLB results, we limit our analysis in the GW frequency range of 10^{-7} – 3×10^{-6} Hz. The amplitude upper limit is evaluated at each GW frequency with a fixed frequency analysis, using a frequency grid uniform in \log_{10} for a total of 180 frequencies per sky cell. The best-sky limit was calculated with one single analysis where the RA and DEC priors were set at $\pm 10^\circ$ around the pulsar position coordinates.

Table 3
Ranges and Types for the Priors Used in the Bayesian SSGW Upper-limit Analysis

Parameter	Priors	Type
A_{SSGW}	$[10^{-10}, 10^{-20}]$	Uni
RA^*	$[0, 2\pi]$	Uni
$\cos(\text{DEC})^*$	$[1, -1]$	Uni
ω_{gw}	$[10^{-7}, 3 \times 10^{-6}]$	Uni/fixd from grid
ι	$[0, \pi]$	Uni
ψ	$[0, 2\pi]$	Uni
ϕ_0	$[0, 2\pi]$	Uni

Note. The prior types Uni and log-Uni correspond to uniform priors in linear and logarithmic space, respectively. The sky coordinates RA and $\cos(\text{DEC})$ are marked with an asterisk (*) to signify that the quoted range is the total covered, but each analysis used only the coordinates range corresponding to each sky cell. The GW angular frequency ω_{gw} is fixed for each analysis, each time using a value from the pre-calculated grid.

We first make a comparison between the Bayesian and the CRLB estimates. Figure 4 shows the two results. As expected, the results are compatible and, in fact, almost identical, which suggests that the Bayesian result approaches that of a fully efficient estimator (Fisz 1963).

Figure 5 presents the results from the Bayesian analyses. It compares the sky-average with the best-sky 95% upper limits of the amplitude for the full 180-frequency spectrum using fixed-noise analysis. In addition, we plot the results for the varying-noise, full-sky analysis for the five key frequencies, noting the corresponding 95% amplitude upper limits. At the reference GW frequency of $1\ \mu\text{Hz}$, the most conservative sky-averaged upper limit, using the varying pulsar-noise analysis, is

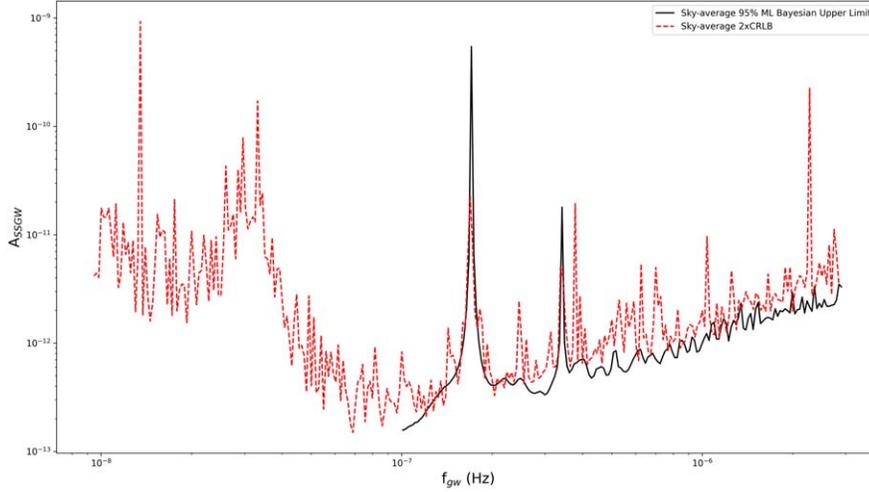


Figure 4. Comparison between the Bayesian-derived upper limit (solid, black line) and the CRLB (dashed, red line) for the SSGW strain amplitude from GW-driven, circular SMBHBs. Both analyses use fixed pulsar noise parameters, corresponding to the maximum likelihood (ML) from the single-pulsar noise analysis. The Bayesian result corresponds to the 95% confidence level. We plot the double of the CRLB to approximate the 2σ error bar, under the assumption of a Gaussian parent distribution.

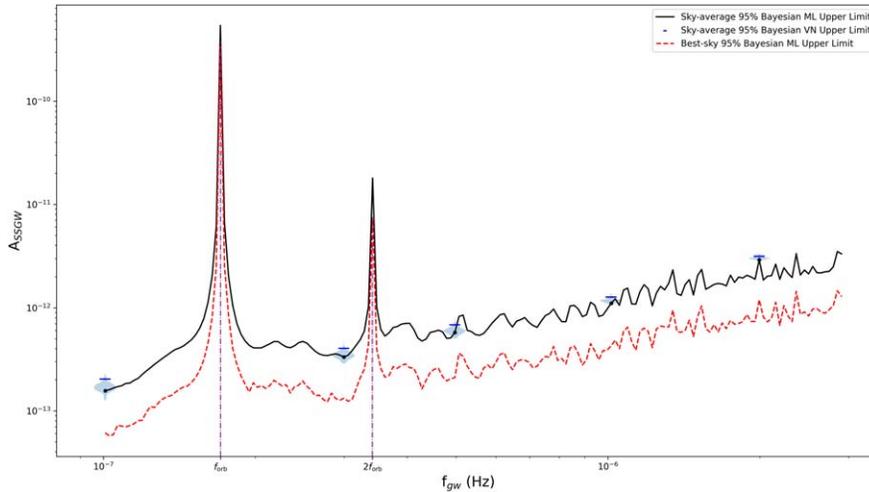


Figure 5. Results from the Bayesian analyses: The figure shows the sky-averaged fixed-pulsar noise analysis using the maximum likelihood (ML) noise parameters (solid, black line), and the same limit in the pulsar direction (best-sky position; dashed, red line). Results correspond to the 95% confidence level. The blue violin plots show the posterior distribution for the varying pulsar noise (VN) analysis performed for five GW frequencies. The 95% upper limits are marked with (blue) horizontal lines. For better comparison, we also signify the corresponding upper limits on these frequencies on the ML line, using filled circles. Finally, note that the two large spikes are associated with the pulsar's orbital frequency and its first harmonic (marked with dashed-dot purple lines), and neither should be confused with the 1 yr^{-1} spike, which is not in this figure's frequency range (see also note in Figure 3).

1.26×10^{-12} . The corresponding best-sky limit is 4.77×10^{-13} . For the fixed-noise analysis, the same sky-average and best-sky limits are 1.10×10^{-12} and 1.01×10^{-13} , respectively. In order to further highlight the effect of using a single pulsar, in which case the sensitivity to GWs tends to zero in the anti-pulsar direction, we note that the sky-averaged amplitude limit in half the sky of the pulsar

direction is 7.01×10^{-13} , in contrast to 1.41×10^{-10} in half the sky of the anti-pulsar direction.

4.2.1. High-resolution Sky Map at $1 \mu\text{Hz}$

We repeat the Bayesian analysis on a finer, 40×40 grid of 1600 sky cells at our reference GW frequency of $1 \mu\text{Hz}$ in order

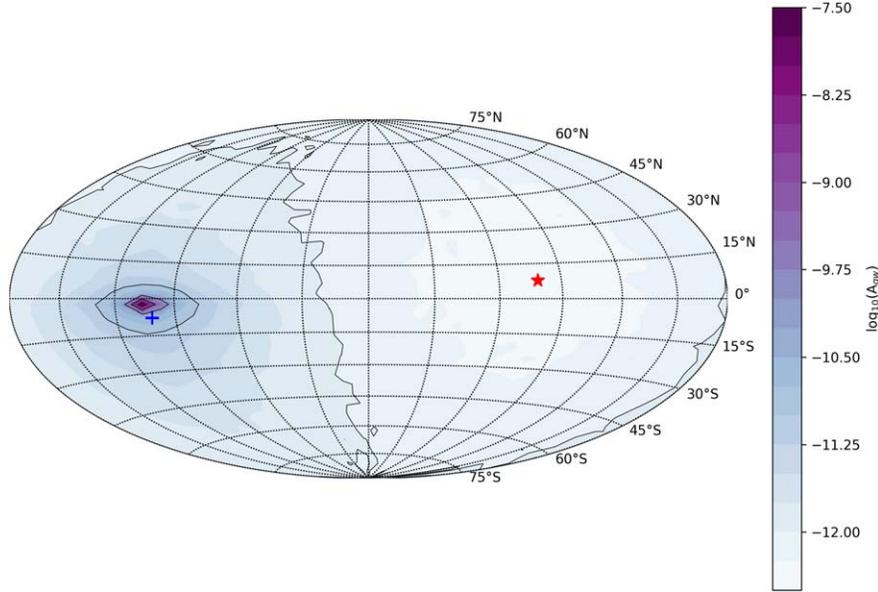


Figure 6. Sky map of the logarithmic amplitude's upper limit (95% confidence level) for SSGWs from circular SMBHBs, using the CPTA PSR J1713+0747 data, at $1 \mu\text{Hz}$ GW frequency. The sky map is produced on the 40×40 grid in right ascension and declination. The red star and cross indicate the location of PSR J1713+0747 and its opposite direction in the sky (anti-pulsar direction), respectively. The slight offset between the anti-pulsar direction and worst sensitivity is an artifact due to finite grid size.

to create a high-resolution sky map of the SSGW strain amplitude upper limit. We use the fixed noise approach to reduce computational cost, as the numerical results are identical to those using the 20×20 case, and thus the sky map serves primarily as a visual representation of the upper limit as a function of the sky location. The sky map is presented in Figure 6 and shows very clearly the dipolar sensitivity of the single-pulsar GW detector: the sensitivity is maximized in the pulsar direction and approaches zero in the opposite direction, which leads to the measured multiple orders of magnitude difference in the sensitivity in their two respective sky hemispheres.

5. Conclusion

We have used CPTA data collected with the FAST radio telescope from the MSP PSR J1713+0747 to derive robust upper limits on the strain amplitude of monochromatic SSGW from GW-driven SMBHBs in circular orbits at the μHz regime. This is one of the most observed MSPs in the context of PTAs, due to its brightness that leads to high-precision ToA measurements and low levels of time-correlated noise. As it is observable from all PTAs both in the Northern and Southern Hemispheres, it has been a source used for various studies, including noise analysis at very high frequencies, scintillation and pulse jitter studies. So far, it has provided the best SSGW strain amplitude limits at μHz frequencies. Due to their very high measurement precision, the CPTA FAST data are especially suitable for SSGW studies, and we can probe the

μHz regime thanks to dense observation campaigns that give the data set a best cadence of 1–2 days. Due to the recent profile-change event of PSR J1713+0747 in 2021 April, the data beyond this date are not suitable for this study and are excluded. As the FAST observations began in mid-2019, our data set has a timespan of only 1.67 yr, which causes some limitations in the sensitivity to SSGW across the spectrum, and especially to frequencies below 10^{-7} Hz. However, the high ToA precision offsets this loss and allows us to derive upper limits comparable to the previously published European Pulsar Timing Array (EPTA) limits (Perera et al. 2018) that used 4.3 yr of data with roughly daily cadence. Using Bayesian inference with the analysis software package FORTYTWO, we derived a sky-averaged strain amplitude upper limit of 1.26×10^{-12} at $1 \mu\text{Hz}$ at the 95% confidence level, and a limit of 4.77×10^{-13} in the direction of PSR J1713+0747, where the sensitivity is maximum. We also produced a high-resolution sky map of the amplitude limit by running the analysis on a fine grid of 1600 sky cells.

By comparison to Perera et al. (2018), the sky-averaged sky limit is ~ 3.5 times higher, and the best-sky limit ~ 2.2 times higher. This is primarily due to the shorter timespan in CPTA data, caused by the profile change that limits the data set. Figure 7 (left panel) shows the effect of the short timespan using simulated PSR J1713+0747 data. The two data sets have properties similar to the CPTA data, i.e., ToA uncertainty of 50 ns and a 2 day cadence. White noise includes only radiometer noise to exclude jitter effects. The first data set has the timespan of the present CPTA data (1.67 yr) and the

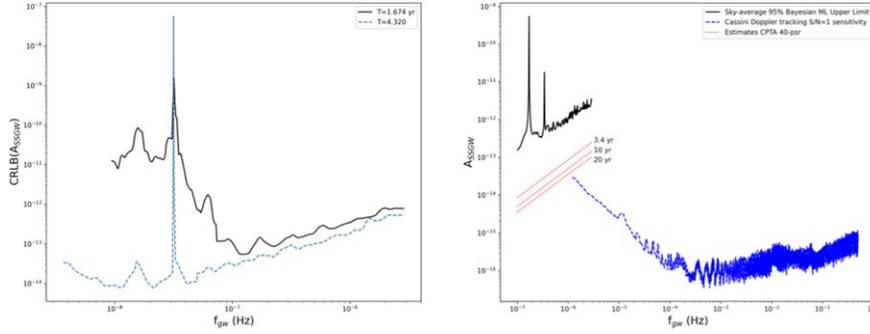


Figure 7. Left panel: CRLB sensitivity curves for PSR J1713+0747 simulated data with a 2 day cadence for timespans of 1.67 (solid, black line) and 4.32 yr (dashed, cyan line). The curves are smoothed for better visualization of their difference. The short timespan causes a sensitivity loss of 2.2 at the μHz regime. Right panel: 95% confidence level sky-averaged upper limits for GW strain of PSR J1713+0747 (solid, black line), and $S/N = 1$ sensitivity of Cassini 2001–2002 GW observations (dashed, blue line), where the data come from Armstrong (2006). The red, dotted lines show estimates for the CPTA, using 40 pulsars and timespans of 3.4, 10 and 20 yr.

second has a timespan equal to the data in Perera et al. (2018) (4.32 yr). Because of the short timespan and the $1/T$ limit near the 1 yr^{-1} frequency spike, the shape of the sensitivity curve for the first case rises at such high frequency that it affects the sensitivity in the μHz regime. By comparison to the second curve, the sky-average limit is 2.2 times higher. Note that this is the same difference with the best-sky limit when comparing the two real data sets, which satisfy the exact same condition ($\pm 10^\circ$ around the pulsar position). The exact area around the opposite direction to the pulsar excluded in the sky-average result in Perera et al. (2018) is not mentioned, and therefore we cannot have an exact comparison. Nevertheless, the simulations show that the majority of the difference between the two results is explained by the short CPTA timespan. Details on effects on sensitivity by other mechanisms, such as jitter, is referred for future work. We note however, that the majority of CPTA pulsar data are not jitter-, but radiometer-noise limited.

In the current study, we only used data of a single pulsar, i.e., PSR J1713+0747, to prove the concept and verify the validity of the analysis software. Such a GW sensitivity is ten times worse compared to the spacecraft tracking experiment (Armstrong 2006) at a similar frequency band (see Figure 7). In future studies, data from multiple pulsars can be combined. The SSGW sensitivity increases with a rough scale of $N^{-1/2}$, where N is the total number of data points. We expect a future combination of CPTA data (50+ pulsars) will lower the current upper limit by about 10 times. Since most CPTA pulsars are radiometer-noise limited, we compute a first estimate on the improvement in the SSGW upper limits we can achieve when including the rest of the CPTA data in the analysis, without being concerned about jitter effects. In Figure 7 (right panel), we show an estimate of the upper limit if we assume the use of 40 pulsars with timing precision at the average of the 40 brighter CPTA sources (Xu, H. et al. 2025, in preparation). We assume timespans of 3.4 (timespan of CPTA first data release),

10 and 20 yr. For simplicity, we assume the pulsars to be isotropically distributed in the sky. We assume that all pulsars except PSR J1713+0747 can contribute beyond the 1.67 yr limitation, and therefore counteract the corresponding sensitivity loss of ~ 2 discussed above. One can expect more improvement over time with dense observing campaigns for various pulsars, as the one made for PSR J1713+0747 and the addition of more pulsars in PTA target lists. The longer data span will also further aid us in probing the lower frequency band, i.e., $f \leq 10^{-8}$. We can thus expect an upper limit at the level of $\lesssim 10^{-14}$ in the near future for the SSGW, resulting in pulsar timing probably becoming be the most sensitive μHz GW detector.

Acknowledgments

Observation of CPTA is supported by the FAST Key project. FAST is a Chinese national mega-science facility, operated by National Astronomical Observatories, Chinese Academy of Sciences. This work is supported by the National SKA Program of China (2020SKA0120100), the National Natural Science Foundation of China (NSFC, Grant Nos. 12041303 and 12250410246), the CAS-MPG LEGACY project, and funding from the Max-Planck Partner Group. K.J.L. acknowledges support from the XPLOER PRIZE and 20 yr long-term support from Dr. Guojun Qiao. H.X. is supported by Major Science and Technology Program of Xinjiang Uygur Autonomous Region No. 2022A03013-4. The data analyses are performed with computer clusters DIRAC and C*-SYSTEM of PSR@pku and computational resources provide by the PARATERA company.

H.X., S.Y.C., Y.J.G., J.J.C., B.J.W., J.W.X., Z.H.X., R.N.C., Y.H.X., and K.J.L. are the core team that performed the data analysis for the current paper, where H.X. worked on data reduction, timing and data analysis; R.N.C. and Z.H.X. performed the noise and SSGW analysis; B.J.W. carried out

single pulse studies; J.J.C. and J.W.X. calibrated data polarization; Z.H.X. and R.N.C. searched for GW single sources; Y.H.X. studied scintillation processes.

ORCID iDs

R. Nicolas Caballero  <https://orcid.org/0000-0001-9084-9427>
 Heng Xu  <https://orcid.org/0000-0002-5031-8098>
 Siyuan Chen  <https://orcid.org/0000-0002-3118-5963>
 Jinchen Jiang  <https://orcid.org/0000-0002-6465-0091>

References

- Agazie, G., Anumarlapudi, A., Archibald, A. M., et al. 2023, *ApJL*, **951**, L8
- Armstrong, J. W. 2006, *LRR*, **9**, 1
- Arzoumanian, Z., Brazier, A., Burke-Spolaor, S., et al. 2016, *ApJ*, **821**, 13
- Bangs, G. W. 1971, Array Processing with Generalized Beamformers., PhD thesis, Yale University
- Caballero, R. N., Lee, K. J., Lentati, L., et al. 2016, *MNRAS*, **457**, 4421
- Chalumeau, A., Babak, S., Petiteau, A., et al. 2022, *MNRAS*, **509**, 5538
- Cognard, I., & Backer, D. C. 2004, *ApJL*, **612**, L125
- Coles, W., Hobbs, G., Champion, D. J., Manchester, R. N., & Verbiest, J. P. W. 2011, *MNRAS*, **418**, 561
- Dolch, T., Lam, M. T., Cordes, J., et al. 2014, *ApJ*, **794**, 21
- Ekers, R. D., & Moffet, A. T. 1968, *Natur*, **220**, 756
- EPTA Collaboration, InPTA Collaboration, Antoniadis, J., et al. 2023, *A&A*, **678**, A50
- Feroz, F., & Hobson, M. P. 2008, *MNRAS*, **384**, 449
- Fisz, M. 1963, in Probability Theory and Mathematical Statistics, ed. S. Geisser (New York: Wiley)
- Foster, R. S., & Backer, D. C. 1990, *ApJ*, **361**, 300
- Gregory, P. C. 2005, Bayesian Logical Data Analysis for the Physical Sciences: A Comparative Approach with “Mathematica” Support (Cambridge: Cambridge Univ. Press)
- Hobbs, G., Lyne, A. G., & Kramer, M. 2010, *MNRAS*, **402**, 1027
- Hobbs, G. B., Edwards, R. T., & Manchester, R. N. 2006, *MNRAS*, **369**, 655
- Hotan, A. W., van Straten, W., & Manchester, R. N. 2004, *PASA*, **21**, 302
- Jaffe, A. H., & Backer, D. C. 2003, *ApJ*, **583**, 616
- Jenet, F. A., Hobbs, G. B., Lee, K. J., & Manchester, R. N. 2005, *ApJL*, **625**, L123
- Jenet, F. A., Lommen, A., Larson, S. L., & Wen, L. 2004, *ApJ*, **606**, 799
- Jiang, P., Yue, Y., Gan, H., et al. 2019, *SCPMA*, **62**, 959502
- Kass, R. E., & Raftery, A. E. 1995, *JASA*, **90**, 773
- Kerr, M., Reardon, D. J., Hobbs, G., et al. 2020, *PASA*, **37**, e020
- Kibble, T. W. B. 1976, *JPhA*, **9**, 1387
- Koen, C. 2006, *MNRAS*, **371**, 1390
- Kramer, M., Xilouris, K. M., Camilo, F., et al. 1999, *ApJ*, **520**, 324
- Landau, L. D., & Lifshitz, E. M. 1960, *Electrodynamics of Continuous Media* (Oxford: Pergamon Press)
- 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
- Lee, K. J., Bassa, C. G., Janssen, G. H., et al. 2012, *MNRAS*, **423**, 2642
- Lee, K. J., Wex, N., Kramer, M., et al. 2011, *MNRAS*, **414**, 3251
- Lentati, L., Alexander, P., Hobson, M. P., et al. 2014, *MNRAS*, **437**, 3004
- 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, K., Karuppusamy, R., Lee, K. J., et al. 2015, *MNRAS*, **449**, 1158
- McKee, J. W., Janssen, G. H., Stappers, B. W., et al. 2016, *MNRAS*, **461**, 2809
- Parthasarathy, A., Shannon, R. M., Johnston, S., et al. 2019, *MNRAS*, **489**, 3810
- Perera, B. B. P., Stappers, B. W., Babak, S., et al. 2018, *MNRAS*, **478**, 218
- Rajagopal, M., & Romani, R. W. 1995, *ApJ*, **446**, 543
- Reardon, D. J., Zic, A., Shannon, R. M., et al. 2023, *ApJL*, **951**, L6
- Sanidas, S. A., Battye, R. A., & Stappers, B. W. 2012, *PhRvD*, **85**, 122003
- Sesana, A. 2013, *MNRAS*, **433**, L1
- Sesana, A., & Vecchio, A. 2010, *PhRvD*, **81**, 104008
- Shannon, R. M., & Cordes, J. M. 2010, *ApJ*, **725**, 1607
- Shannon, R. M., Lentati, L. T., Kerr, M., et al. 2016, *ApJL*, **828**, L1
- Skilling, J. 2004, in AIP Conf. Ser. 735, *Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, ed. R. Fischer, R. Preuss, & U. von Toussaint (Melville, NY: AIP), 395
- Slepian, D. 1954, *IREPGIT*, **3**, 68
- Taylor, J. H. 1992, *RSPTA*, **341**, 117
- Taylor, J. H., & Weisberg, J. M. 1982, *ApJ*, **253**, 908
- Thorne, K. S. 1989, in *Three Hundred Years of Gravitation*, ed. S. W. Hawking & W. Israel (Cambridge: Cambridge Univ. Press), 330
- van Haasteren, R., & Levin, Y. 2013, *MNRAS*, **428**, 1147
- van Haasteren, R., Levin, Y., McDonald, P., & Lu, T. 2009, *MNRAS*, **395**, 1005
- van Straten, W., & Bailes, M. 2011, *PASA*, **28**, 1
- Verbiest, J. P. W., Bailes, M., Coles, W. A., et al. 2009, *MNRAS*, **400**, 951
- Xu, H., Chen, S., Guo, Y., et al. 2023, *RAA*, **23**, 075024
- Xu, H., Huang, Y. X., Burgay, M., et al. 2021, *ATel*, **14642**, 1
- Yi, S., Stappers, B. W., Sanidas, S. A., et al. 2014, *MNRAS*, **445**, 1245
- You, X. P., Hobbs, G., Coles, W. A., et al. 2007, *MNRAS*, **378**, 493