

Fast wavelet basis search for generic gravitational wave bursts in pulsar timing array data

Jacob A. Taylor[✉], Rand Burnette[✉], and Bence Bécsy[✉]

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

Neil J. Cornish[✉]

eXtreme Gravity Institute, Department of Physics, Montana State University, Bozeman, Montana 59717, USA



(Received 14 August 2024; accepted 22 November 2024; published 31 January 2025)

As we move into an era of more sensitive pulsar timing array datasets, we may be able to resolve individual gravitational wave sources from the stochastic gravitational wave background. While some of these sources, like orbiting massive black hole binaries, have well-defined waveform models, there could also be signals present with unknown morphology. This motivates the search for generic gravitational-wave bursts in a signal-agnostic way. However, these searches are computationally prohibitive due to the expansive parameter space. In this paper we present QuickBurst, an algorithm with a redefined likelihood that lets us expedite Markov chain Monte Carlo sampling for a subset of the signal parameters by avoiding repeated calculations of costly inner products. This results in an overall speedup factor of ~ 200 on realistic simulated datasets, which is sufficient to make generic gravitational-wave burst searches feasible on current and growing pulsar timing array datasets.

DOI: 10.1103/PhysRevD.111.022006

I. INTRODUCTION

The North American Nanohertz Observatory for Gravitational Waves (NANOGrav), the Parkes Pulsar Timing Array (PPTA), the European Pulsar Timing Array (EPTA), and the Chinese Pulsar Timing Array (CPTA) have found evidence for a stochastic gravitational-wave (GW) background (GWB) with varying level of significance [1–4]. This is exciting for both multimessenger astrophysics and future prospects for detecting individual binary systems [5–7]. In addition to the persistent stochastic signal of the GWB, there are many other possible sources of short-duration GW bursts. Examples are transient deterministic signals whose persistence is less than the dataset time span, and these signals have the potential to be observed via pulsar timing arrays (PTAs) [8,9]. In order to agnostically search for these potential sources, we can use a generic GW burst model to find transient GWs with a variety of waveforms.

Previous searches have been done for gravitational wave burst signals in LIGO data (see e.g. [10,11]) using various techniques [12–16]. Various techniques have also been proposed in other GW frequency bands, including LISA [17] and PTA data [7,18,19]. One prominent approach is to model the waveform with a collection of wavelets, where the number of these wavelets used is also a free model parameter, allowing for flexible signal modeling. This previously proposed wavelet-based technique would take

months to perform on PTA datasets in comparison to \sim week time-frame it takes to do a single analysis run for GWB searches on cutting edge hardware. This is because of the large parameter space needed to fully model a GW burst in PTA data. To feasibly do a generic GW burst search on current and future data sets, we need to improve the efficiency of our generic GW burst search algorithms.

In this paper, we present QuickBurst [20], a Bayesian generic GW burst search algorithm that consists of a trans-dimensional reversible-jump Markov chain Monte Carlo (RJMCMC) sampler [21,22] and a factorized likelihood template that utilizes the properties of a deterministic signal model for improved computational efficiency. Recent work has shown that for some deterministic signals, a faster search algorithm can be achieved by precalculating and storing components of the likelihood calculation that do not change [23]. The work presented here is built upon BayesHopperBurst and QuickCW, the algorithms presented in [7,23] respectively. BayesHopperBurst uses an RJMCMC sampler to search for generic GW burst signals, but does not utilize more efficient sampling methods, of which their usefulness has been demonstrated in QuickCW.

In Sec. II, we derive the Bayesian likelihood for a generic GW burst signal modeled as a superposition of Morlet-Gabor wavelets. We also describe the techniques employed in our Markov chain Monte Carlo (MCMC) sampling algorithm that take advantage of our likelihood

formulation. Then, we present the expected computational efficiency improvements that these techniques produce, as well as our verification tests to ensure the operational quality of the algorithm. In Sec. III, we demonstrate the capabilities of QuickBurst by performing searches on simulated datasets that resemble the NANOGrav 15-year dataset [24]. Furthermore, we test the effects of a common uncorrelated red noise (CURN) process on QuickBurst’s ability to recover a GW burst signal. Finally, in Sec. IV we discuss future applications involving QuickBurst to search for gravitational wave burst signals in PTA data.

II. QUICKBURST METHODS

A. Models

We can model the timing residuals for the PTA as a vector

$$\delta\mathbf{t} = \begin{bmatrix} \delta\mathbf{t}_1 \\ \delta\mathbf{t}_2 \\ \vdots \\ \delta\mathbf{t}_{N_{\text{psr}}} \end{bmatrix}, \quad (1)$$

where N_{psr} is the number of pulsars in the PTA. The timing residuals in the i th pulsar is then modeled as

$$\delta\mathbf{t}_i = M_i \delta\xi_i + \mathbf{n}_i + \mathbf{g}_i + \boldsymbol{\omega}_i(\boldsymbol{\theta}_s) + \mathbf{V}_i(\boldsymbol{\theta}_n). \quad (2)$$

In this formalism, M_i is our design matrix with dimensions $(N_{\text{TOA}} \times N_{\text{par}})$, where N_{TOA} represents the number of times-of-arrival (TOAs) in the i th pulsar and N_{par} represents the number of timing model parameters in this pulsar, $\delta\xi_i$ is a vector of shape N_{par} that accounts for a small error in this pulsar’s timing model, \mathbf{n}_i represents intrinsic noise in this pulsar, \mathbf{g}_i represents the contribution from spatially uncorrelated noise common to all pulsars, $\boldsymbol{\omega}_i$ is the contribution from a generic gravitational wave burst signal, and \mathbf{V}_i represents the contribution to this pulsar’s timing residuals of any incoherent noise transients that are unique to this pulsar. $\boldsymbol{\theta}_s$ and $\boldsymbol{\theta}_n$ are the parameters that describe the GW burst signal and transient noise respectively. $\delta\mathbf{t}_i$, \mathbf{n}_i , \mathbf{g}_i , $\boldsymbol{\omega}_i$, and \mathbf{V}_i are all vectors with dimensions of N_{TOA} .

Morlet-Gabor wavelets (sine-Gaussian wavelets) are a robust basis for generically modeling transient signals. First, the transient noise \mathbf{V}_i may be written as a linear combination of these wavelets:

$$\mathbf{V}_i(\boldsymbol{\theta}_n) = \sum_{j=1}^{N_i} \Psi(t_i; \lambda_{ij}), \quad (3)$$

where N_i is the number of noise transient wavelets in the i th pulsar, t_i is the vector of timing observations for the

i th pulsar, and λ_{ij} are the parameters for the j th wavelet for the i th pulsar. Each wavelet can be expressed as a vector

$$\Psi(t_i; \lambda_{ij}) = \begin{bmatrix} \Psi(t_{i1}; \lambda_{ij}) \\ \Psi(t_{i2}; \lambda_{ij}) \\ \vdots \\ \Psi(t_{iN_{\text{TOA}}}; \lambda_{ij}) \end{bmatrix}, \quad (4)$$

where the elements in this vector are described by

$$\Psi(t_i; \lambda_{ij}) = A e^{(t_{\text{in}} - t_0)^2 / \tau^2} \cos(2\pi f_0(t_{\text{in}} - t_0) + \phi_0). \quad (5)$$

In this wavelet formulation, A is an overall amplitude, t_{in} is the n th observation time in the i th pulsar, t_0 is the central time for the wavelet, τ is the characteristic duration, f_0 is the central frequency, and ϕ_0 is the initial phase.

For the signal model, the contributions from a GW burst may also be written as a sum of these Morlet-Gabor wavelets further modulated by the antenna response of the PTA based on GW polarization and source sky location. Because GWs perturb spacetime based on both GW polarization and propagation direction, we sum together our sine-Gaussian wavelets for the signal, and project the signals onto the lines of sight of the pulsars in our array, which are dependent on the corresponding antenna responses and ensures coherence between the pulsars in our array. Quantitatively, we can write the photon time integral of the metric perturbation as a linear combination of wavelets

$$\mathbf{H}_+ = \sum_{j=1}^N \Psi(t; t_{0,j}, f_{0,j}, \tau_j, A_{+,j}, \phi_{0,+,j}), \quad (6)$$

$$\mathbf{H}_\times = \sum_{j=1}^N \Psi(t; t_{0,j}, f_{0,j}, \tau_j, A_{\times,j}, \phi_{0,\times,j}), \quad (7)$$

where \mathbf{H}_+ and \mathbf{H}_\times are the plus and cross GW polarizations respectively, and have independent amplitudes A_+ and A_\times , with independent phases $\phi_{0,+}$ and $\phi_{0,\times}$. This leaves t_0 , f_0 , and τ as common parameters between both polarizations. Furthermore, we model the phases independently, as not all GWs have the same 90° phase shift relationship between the plus and cross polarizations.

Next, we perform a rotation around the propagation direction of the GW, which is given by

$$\bar{\mathbf{H}}_\times = \mathbf{H}_\times \cos(2\psi_{\text{gw}}) - \mathbf{H}_+ \sin(2\psi_{\text{gw}}), \quad (8)$$

$$\bar{\mathbf{H}}_+ = \mathbf{H}_+ \sin(2\psi_{\text{gw}}) + \mathbf{H}_\times \cos(2\psi_{\text{gw}}), \quad (9)$$

where ψ_{gw} is the polarization angle of the GW. Finally, the contribution of the signal to the timing residual can be expressed as

$$\omega_i(\theta_s) = -F_+(\Omega_i, \Omega_{\text{gw}})\bar{\mathbf{H}}_+(t_i; \lambda') - F_\times(\Omega_i, \Omega_{\text{gw}})\bar{\mathbf{H}}_\times(t_i; \lambda'), \quad (10)$$

where $F_+(\Omega_k, \Omega_{\text{gw}})$ and $F_\times(\Omega_k, \Omega_{\text{gw}})$ are the plus and cross polarization mode antenna responses, respectively. For further details on the geometric response, see [25]. Equation (10) assumes that the only contribution to the signal is from the Earth term (ET) and not the pulsar term (PT). This is a result of our current observing baseline (~ 20 years) being significantly shorter than the light travel time between the Earth and any pulsar in the PTA ($\sim 10^3$ years). Thus, any GW bursts which appear in only one pulsar will not affect the residuals in any other pulsar (unless two or more pulsars are within a few degrees of each other), nor will it reach the Earth within the time frame of the experiment. As such, this burst will appear as a single noise transient in that particular pulsar. For more details on the geometric time delay between the ET and the PT, see Section 2.1 in Ref. [7].

Then, for N_p pulsars, the ET signal-to-noise ratio (SNR) is larger than the PT SNR by a factor of $\sqrt{N_p}$. The SNR scaling for the ET also becomes larger when considering that an ET search will allow us to see a common signal in all pulsars in our PTA, while the PT search will only show a signal in one pulsar and can be modeled as a noise transient. For additional details on these arguments, see [7,26]. Moreover, it is expected that the number of events that we can possibly observe in the ET only search scales with the number of pulsars $\sim N_p^{3/2}$. However, the PT only search scales as $\sim N_p$ due to the extended timing baseline when combining individual pulsar data. As a result, the number of possible observed events in the ET only search exceeds the number of possible observed events in the PT only search by $\sim N_p^{1/2}$. Considering current PTAs contain of order 100 pulsars, there is potential to observe $\sim 10\times$ more events in the ET only search.

B. Likelihood

We use a factorized likelihood, which allows us to simplify some of our likelihood calculations depending on which parameters in our model we want to vary. We can demonstrate how to generically rewrite a standard likelihood given a specific generic GW burst model template. For any signal, the likelihood can be written as

$$\log L = \sum_{i=1}^{N_{\text{psr}}} \left[-\frac{1}{2}(\delta t_i - s_i | \delta t_i - s_i) - \frac{1}{2} \log \det(2\pi C_i) \right], \quad (11)$$

where s_i is our GW burst signal and C_i is the noise covariance matrix for the i th pulsar. Because we assume that there are no spatially correlated signals, we can guarantee that the full PTA covariance matrix is block-diagonal. Thus, we can separate the individual pulsar covariance matrices. We define a noise-weighted inner product between any two vectors \mathbf{a} and \mathbf{b} as

$$(\mathbf{a} | \mathbf{b}) = \mathbf{a}^T \mathbf{C}^{-1} \mathbf{b}, \quad (12)$$

where

$$\mathbf{C}^{-1} = \mathbf{N}^{-1} - \mathbf{N}^{-1} \mathbf{T} \Sigma^{-1} \mathbf{T}^T \mathbf{N}^{-1} \quad (13)$$

is the inverse of the covariance matrix for a particular pulsar given by the Woodbury identity [27], and

$$\Sigma = B^{-1} + T^T N^{-1} T. \quad (14)$$

B is the prior matrix for the hyperparameters, and T is the design matrix for the timing model, red noise (RN) and jitter (see [28]). The design matrix includes additional models for dispersion measure (DM), specifically modeling DM as an epoch-by-epoch DM offset. This model is referred to as DMX [28].

For the PTA, we can write the vector of GW burst contributions to the timing residuals in Eq. (1) as

$$\mathbf{s} = \begin{bmatrix} s_1 \\ s_2 \\ \vdots \\ s_{N_{\text{psr}}} \end{bmatrix}, \quad (15)$$

where we decompose the signal s_i in the i th pulsar as a sum of filter functions $\mathbf{S}^{ki}(\theta_s)$ multiplied by coefficients $\sigma_{ki}(\theta_p)$,

$$s_i = \sum_{k=1,2} \sigma_{ki}(\theta_p) \mathbf{S}^{ki}(\theta_s). \quad (16)$$

Direct substitution of Eq. (16) into Eq. (11) yields

$$\begin{aligned} \log L = & \sum_{i=1}^{N_{\text{psr}}} \left[-\frac{1}{2}(\delta t_i - s_i | \delta t_i - s_i) - \frac{1}{2} \log \det(2\pi C_i) \right. \\ & \left. + \sum_{v=1}^{N_F} \left[\sum_{k=1,2} \sigma_{ki} \mathbf{N}_i^k - \frac{1}{2} \sum_{k=1,2} \sum_{l=1,2} \sigma_{ki} \sigma_{li} M_i^{kl} \right] \right], \end{aligned} \quad (17)$$

where the last two terms are inner products between the i th pulsar residuals and the filter functions for that pulsar,

as well as the inner product between the pairs of filter functions respectively. These inner products are defined as

$$N_i^k = (\delta t_i | \mathbf{S}^{ki}), \quad (18)$$

and

$$M_i^{kl} = (\mathbf{S}^{ki} | \mathbf{S}^{li}), \quad (19)$$

where a filter function \mathbf{S}^{ki} is a vector of size N_{TOA} (which has the same structure as δt_i). Lastly, N_F is the number of wavelets present in the signal model.

Now, we combine Eqs. (6)–(10) to get the filter functions and filter coefficients for the noise transients and the GW burst signal based on Eq. (16). In doing this, we get

$$\mathbf{S}_n^{li} = e^{(t_{in}-t_0)^2/\tau^2} \cos(2\pi f_0(t_{in} - t_0)), \quad (20)$$

$$\mathbf{S}_n^{2i} = e^{(t_{in}-t_0)^2/\tau^2} \sin(2\pi f_0(t_{in} - t_0)), \quad (21)$$

which are the elements of our filter functions $\mathbf{S}^{ki}(\theta_s)$ from Eq. (16) for the i th pulsar and the n th TOA, and the corresponding signal coefficients for a single wavelet are given by

$$\begin{aligned} \sigma_{1i} &= -\cos(2\psi_{\text{gw}})[F_{i,+}A_+ \cos(\phi_{0,+}) + F_{i,\times}A_\times \cos(\phi_{0,\times})] \\ &\quad + \sin(2\psi_{\text{gw}})[F_{i,+}A_\times \cos(\phi_{0,\times}) - F_{i,\times}A_+ \cos(\phi_{0,+})], \\ \sigma_{2i} &= \cos(2\psi_{\text{gw}})[F_{i,+}A_+ \sin(\phi_{0,+}) + F_{i,\times}A_\times \sin(\phi_{0,\times})] \\ &\quad + \sin(2\psi_{\text{gw}})[F_{i,\times}A_+ \sin(\phi_{0,+}) - F_{i,+}A_\times \sin(\phi_{0,\times})]. \end{aligned} \quad (22)$$

The corresponding noise transient coefficients for a single wavelet are given by

$$\begin{aligned} \sigma_{1i} &= A \cos(\phi_0), \\ \sigma_{2i} &= -A \sin(\phi_0). \end{aligned} \quad (23)$$

We can now see why writing this factorized likelihood opens the door for more efficient calculations. For the signal model, we have separated all of the terms that are inherently tied to the pulsar TOA data into the N_i^k and M_i^{kl} terms, so the coefficients can be recalculated without doing any of the costly inner products in these matrices. On the other hand, the coefficients, which are cheap to recalculate, can be independently calculated during runtime.

Based on Eq. (22), we can now define our sets of shape and projection parameters, which are parameters that are either non-separable or separable from pulsar TOAs, respectively. For noise transients, based on Eq. (5), our shape parameters (θ_s) and projection parameters (θ_p) are

$$\begin{aligned} \theta_s &\rightarrow (\tau, t_0, f_0), \\ \theta_p &\rightarrow (A, \phi_0). \end{aligned} \quad (24)$$

For our generic GW burst signal, our shape and projection parameters are

$$\begin{aligned} \theta_s &\rightarrow (\tau, t_0, f_0), \\ \theta_p &\rightarrow (\theta_{\text{gw}}, \phi_{\text{gw}}, \psi, A_+, A_\times, \phi_{0,+}, \phi_{0,\times}). \end{aligned} \quad (25)$$

We can see from Eq. (24), the number of shape and projection parameters for noise transient wavelets in our model go as $3 \times N_n$ and $2 \times N_n$, respectively for N_n noise transient wavelets. From Eq. (25), the number of shape and projection parameters in our model will scale as $3 \times N_s$ and $3 + 4 \times N_s$ respectively for N_s GW signal wavelets, since all wavelets share the same sky location and polarization angle. Notably, in previous work using the signal template defined by [23] for continuous waves, there are always the same number of shape and projection parameters. Our model has more projection than shape parameters, which will mean a larger percent of our model will be calculated with the more efficient methods in our algorithm.

Conveniently, the shape parameters are the same for noise transients and our signal, which necessitates fewer recalculations during sampling. Due to GW information, however, the coefficients σ_{ki} depend on the antenna patterns and GW polarization, making them more tedious to recalculate.

C. Sampling

The likelihood factorization that results from utilizing our signal template, as shown in Eq. (17), allows us to avoid recalculating components of the likelihood depending on what parameters we are varying in any given step. These correspond to the two aforementioned types of parameters, shape parameters and projection parameters, as described in Eqs. (24) and (25). Sampler steps that only change projection parameters do not require recalculating any components of N_i^k and M_i^{kl} . We call these projection parameter updates “fast jumps” because the coefficients are several orders of magnitude faster to recalculate than the filter functions. This results in these steps being faster than all other parameter jumps we can do, which vary over various subsets of shape parameters (see Table I).

We use similar jump techniques as BayesHopper-Burst [7] when exploring intrinsic pulsar RN, white noise (WN), or a CURN process. It is worth noting that these do require recalculating more pieces of the likelihood than signal parameter jumps. Based on Eq. (17), our options for parameter types to vary are (in order of decreasing computational expense):

- (1) Intrinsic pulsar noise and CURN: All terms in the likelihood are recalculated as a result of the noise covariance matrix C changing.
- (2) Shape parameters: Parts of the terms which include our signal model need to be recalculated when varying shape parameters. We only recalculate the

elements of M_i and N_i that are changing, while all other elements are stored for the next likelihood calculation.

(3) Projection parameters: Only the filter function coefficients are recalculated.

We also need to allow the dimensions of our parameter space to change during sampling depending on how many GW signal wavelets or transient wavelets the model prefers. To properly explore these models, we apply transdimensional proposals to probe more or less signal or noise transient parameters [21,22] as the sampler evolves. These transdimensional proposal steps are labeled as TD in Table I. Additionally, we apply Fisher matrix proposals and τ -scan proposals to increase sampling efficiency.

For Fisher matrix proposals, we can numerically construct a matrix from the covariance between our model parameters by computing second derivatives of the likelihood via finite differences. Then, the eigenvectors of the Fisher matrix can be used to change model parameters between MCMC steps based on the covariance between these model parameters. These Fisher matrices are constructed for various sets of model parameters separately from each other (i.e. for intrinsic pulsar noise, signal and noise transient wavelets, CURN, etc.), which allows us to make more informed MCMC jump proposals based on the type of step taken during sampling. Additionally, these Fisher matrices can be recalculated during runtime based on the location of the MCMC sampler in parameter space to better inform our Fisher steps during sampling. For more details on how these matrices are computed, see [29].

Additionally, we can compute a global proposal for all wavelets to draw from in what we call τ -scan proposals.

TABLE I. Timing of different parameter jump types in seconds, from runs on the SIM-MID dataset as described in Table II. The total speed up is calculated with a ratio of 10,000 Fast samples to any one other jump sample. All results were obtained using single personal computers (PCs) utilizing Intel I9-10900k CPUs with 20 cores at 3.7 GHz and 64 GBs of memory at speeds of 3200 MHz. The steps labeled as “TD” are averages over both the whole step, while all other timings come from the likelihood calls. This is done because the time per likelihood call for TD steps can vary greatly depending on if the model is adding or removing parameters.

Jump types	QuickBurst (s)	BayesHopperBurst (s)
TD signal	0.50	0.20
TD noise transient	0.41	0.18
Signal τ scan	0.30	0.21
Noise transient τ scan	0.24	0.18
Regular Fisher	0.23	0.13
Noise	0.66	0.08
Fast	1.5×10^{-5}	N/A
Total speed up	$\times 186$ for 10,000 Fast/Slow	

These proposals are informed by a 3D map over our shape parameters [see Eqs. (24) and (25)] where we find some excess power in our data. τ -scan proposals have been used previously in BayesHopperBurst, and can be utilized in our algorithm to improve sampling efficiency specifically in our shape parameters. More details on τ -scan proposals can be found in Sec. 2 of [7].

D. Verification and timing

Before analyzing efficiency improvements, we first present comparisons between the Bayesian likelihood implemented in QuickBurst and BayesHopperBurst. As BayesHopperBurst utilizes the software package ENTERPRISE [30] to evaluate the Bayesian likelihood, we use these values as a benchmark for our algorithm. The likelihood comparisons are evaluated using a simple data set consisting of two sine-Gaussian wavelets. Our comparisons show that the two signal models yield likelihoods that are equivalent up to 10^{-10} precision. When including either intrinsic pulsar red noise or CURN, compounding numerical errors in likelihood evaluations increase differences to $\sim 10^{-5}$. However, the spread in likelihood differences is still centered around a difference of zero. Additionally, the signal reconstruction for a single pulsar from this run can be seen in Fig. 1.

With consistency between ENTERPRISE and QuickBurst established, we can now proceed to demonstrating the improved computational efficiency of QuickBurst. To do so, we simulate datasets as close to the 15 year NANOGrav dataset as possible (see Sec. III for simulation details) and time our various types of steps. The timing for these steps correspond to runs on the SIM-MID dataset detailed in Table II, which only vary our GW signal and noise transient parameters. However, in order to time noise jumps, we allow intrinsic pulsar RN to vary. It is worth noting that the timing runs do not include CURN, which is handled in the “regular Fisher” jump proposals. The average timing for these jumps is around 0.51 s if CURN is included. The results can be seen in Table I.

There are some notable takeaways from this timing data. First, we see QuickBurst yields a $\sim 1.5 \times - 2 \times$ increase in time per step for jumps that vary shape parameters compared to ENTERPRISE. These include τ -Scan, transdimensional (TD), and Regular Fisher jumps. Second, we see QuickBurst has an $8 \times$ increase in time per step compared to ENTERPRISE for Noise jumps, which vary intrinsic pulsar noise. This is due to an increase in the number of inner products required for deterministic signal models. In expanding the likelihood to account for this deterministic model [as shown in Eq. (17)], we have a factor of $\frac{N_F^2 + 5N_F}{2}$ increase in the number of inner products we need to calculate during these Noise jumps in the M and N matrices seen in both Eqs. (18) and (19), which accounts for the time increase shown.

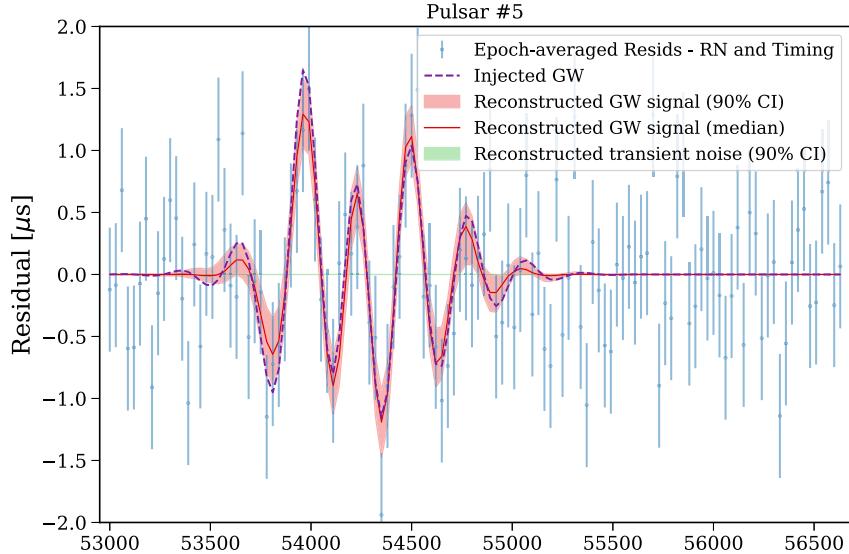


FIG. 1. Waveform reconstructions for a simulated dataset containing two sine-Gaussian signal wavelets with an SNR of 17.9. The recovered waveform has an SNR of 16.3, and a match of $M = 0.97$ with the injected signal [see Eq. (27)].

Despite these increases, we still see a significant speed up thanks to our parameter separation as described in Sec. II B. QuickBurst samples our “fast” jumps—which only recalculate $\sigma_{ki}(\theta_p)$ —up to 10,000× faster than ENTERPRISE. With this advantage over ENTERPRISE, we see a factor of 186× speed up overall for generic GW burst searches in comparison to BayesHopperBurst, even when including RN and CURN timing. If we exclude both CURN and RN, this factor improves up to approximately 204×.

These times were calculated by averaging over the timing table jump types for QuickBurst and BayesHopperBurst, and computing the average time over 100 samples for each algorithm. Note that the timing for an ENTERPRISE likelihood calculation for an equivalent fast jump is on average 100 μ s, which is omitted in the Table I given that BayesHopperBurst does not have the ability to perform this type of parameter jump.

III. ANALYSIS OF SIMULATED DATA SETS

A. Datasets

To test QuickBurst, we simulate four datasets that include a burst signal of varying amplitude. These

TABLE II. Table consisting of luminosity distance, and signal SNR in the cases of including only intrinsic pulsar WN or both intrinsic pulsar RN and WN. The quoted SNR values were computed with only 24 of 67 NANOGrav pulsars, which correspond to pulsars with more than 10 years of observations.

Parameters	SIM-LOW	SIM-MID	SIM-HIGH
luminosity distance	120 Mpc	60 Mpc	30 Mpc
WN SNR	4.7	9.5	19.0
WN + RN SNR	3.9	7.8	15.7

simulations are based on the Astro4Cast simulations presented in [31] and realistically emulate the 15 year NANOGrav dataset. This means our observing epochs (including gaps) as well as the injected RN and WN in every pulsar are based on measured values. However, to ensure a realistic-sized data set, we do not average the residuals over each observing epoch, and instead retain and analyze all of them. One dataset contains no GW burst signal injection, while the other three datasets include an injected GW burst signal in the form of a parabolic encounter of two SMBHs [32] each with a mass of $10^9 M_\odot$, an impact parameter of $120 M_\odot$, a sky location of ($\theta_{gw} = \pi/2$, $\phi_{gw} = 4.0$) and while varying the source luminosity distance. These four datasets form four signal regimes: a high amplitude (SIM-HIGH), medium amplitude (SIM-MID), low amplitude (SIM-LOW), and no signal case. All datasets contain pulsar noise properties as found in Sec. V in [33]. For these four datasets, we are disregarding the inclusion of a CURN process. See Sec. III B 3 for generic GW burst analysis including CURN.

We quantify the signal strength by calculating the SNR of our GW burst injection. The SNR is given by

$$\text{SNR}(h) = \sqrt{(h|h)} \quad (26)$$

for any signal h . For all three nonzero amplitude datasets, the signal SNR and luminosity distances can be found in Table II. We calculate the SNR in two cases for each dataset: one case where we only include intrinsic pulsar WN in our covariance matrix, and another case where we include intrinsic pulsar RN and WN in our model. The difference between these SNR values indicates how much the detectability of the burst is affected by RN.

TABLE III. Table of prior ranges used for signal parameters in analysis of the datasets described in Table II. T_{\max} refers to the maximum observation time span in a given dataset with respect to the starting MJD, which is an MJD of ~ 5855 for these datasets. All angles are in units of radians.

Parameters	Shape parameters				Projection parameters					
	τ (yrs)	t_0 (yrs)	$\log_{10}(f_0)$	$\cos \theta_{\text{gw}}$	ϕ_{gw}	ψ_{gw}	A_+	A_x	$\phi_{0,+}$	$\phi_{0,x}$
Range	0.2, 5.0	0.0, T_{\max}	-8.46, -7	-1, 1	$0, 2\pi$	$0, \pi$	$10^{-10}, 10^{-5}$	$10^{-10}, 10^{-5}$	$0, 2\pi$	$0, 2\pi$

The subsequent analysis performed on these four datasets in Sec. III B includes both intrinsic pulsar RN and WN.

Another important note is that the burst epoch in all datasets is at 55,000 MJD, which is in the first half of the dataset. This is an arbitrary choice. Many pulsars do not have TOAs around this burst epoch, and will not provide additional significance on GW burst recovery. As such, we truncate the list of pulsars involved in our searches on the datasets described in Table II to only include pulsars with more than 10 years of data, which is 24 of the 68 pulsars in the NANOGrav 15 year PTA. Note that we only use 67 of these pulsars (see [1,24] for details). Doing so improves the efficiency of our generic GW burst search, as the model has fewer pulsars to fit a GW burst to. Such a PTA reduction may be possible for real datasets, even without knowing the burst epoch, by investigating the SNR contributions from individual pulsars. See Appendix A for further discussion. However, this does reduce our sky coverage, which is discussed in Sec. III B 2 in further detail.

B. Analysis

Now we present the results of the analysis of datasets of varying signal strengths. We first analyze the no signal regime, and then analyze the datasets that have a signal present. Since the parameters of our generic model have no physical meaning, the pipeline's performance is best characterized by how well the waveform of the GW signal is reconstructed. To quantify our signal recovery, we can see how close the injected (h) and recovered (h') waveforms are to each other by computing a normalized noise weighted inner product between h and h' over the entire PTA. This match M is defined as

$$M = \frac{(h|h')}{\text{SNR}(h) \times \text{SNR}(h')} \quad (27)$$

where a value of $M = 1$ is a perfect overlap between h and h' , while $M = 0$ is no overlap between the two waveforms. More details on this statistic can be found in Sec. 2 of [7].

Described in Table III are the priors assigned to both shape and projection parameters. Given the exploratory nature of our testing, we use wide ranges on many of our priors. We use a uniform prior for all of our signal parameters.

1. Zero amplitude dataset

To establish a baseline, we have tested QuickBurst on a dataset with no GW burst injection, while keeping all intrinsic pulsar noise as realistic as possible. As expected, our analysis shows preference for no signal wavelets being included in our signal model, in addition to no noise transient wavelets, as shown by the histogram in Fig. 2.

2. Nonzero signal amplitude datasets

Here we show our analysis of the three non-zero SNR signal simulated datasets described in Table II. Given that the waveform shape can vary widely based on the orbit parameters of the two SMBHs, a good indicator for QuickBurst's ability to recover our signal is to accurately reconstruct the signal waveform after performing searches in our datasets.

In Fig. 3, we show the results of our GW signal and noise transient reconstructions in PSR B1937 + 21. This pulsar has a significant burst amplitude present in all three datasets, with single pulsar GW burst SNR values of 2.8, 1.4, and 0.7 for the SIM-HIGH, SIM-MID, and SIM-LOW datasets respectively. The plots on the left show the epoch averaged residuals for B1937 + 21 with the RN and best fit timing model parameters [$\mathbf{n}_i + \mathbf{g}_i$ and $M_i \delta \xi_i$ in Eq. (2), respectively] subtracted out. Additionally, the solid red line is the median waveform reconstruction for the GW burst, and the red shaded region is the 90% credible region.

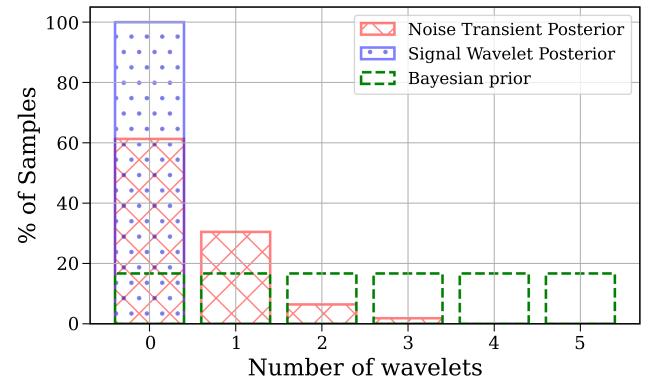


FIG. 2. Histogram for recovered noise transients and signal wavelets for our NANOGrav 15-year like dataset with no GW burst injection. We note that the model prefers no signal wavelets to be present in the model at all times in sampling.

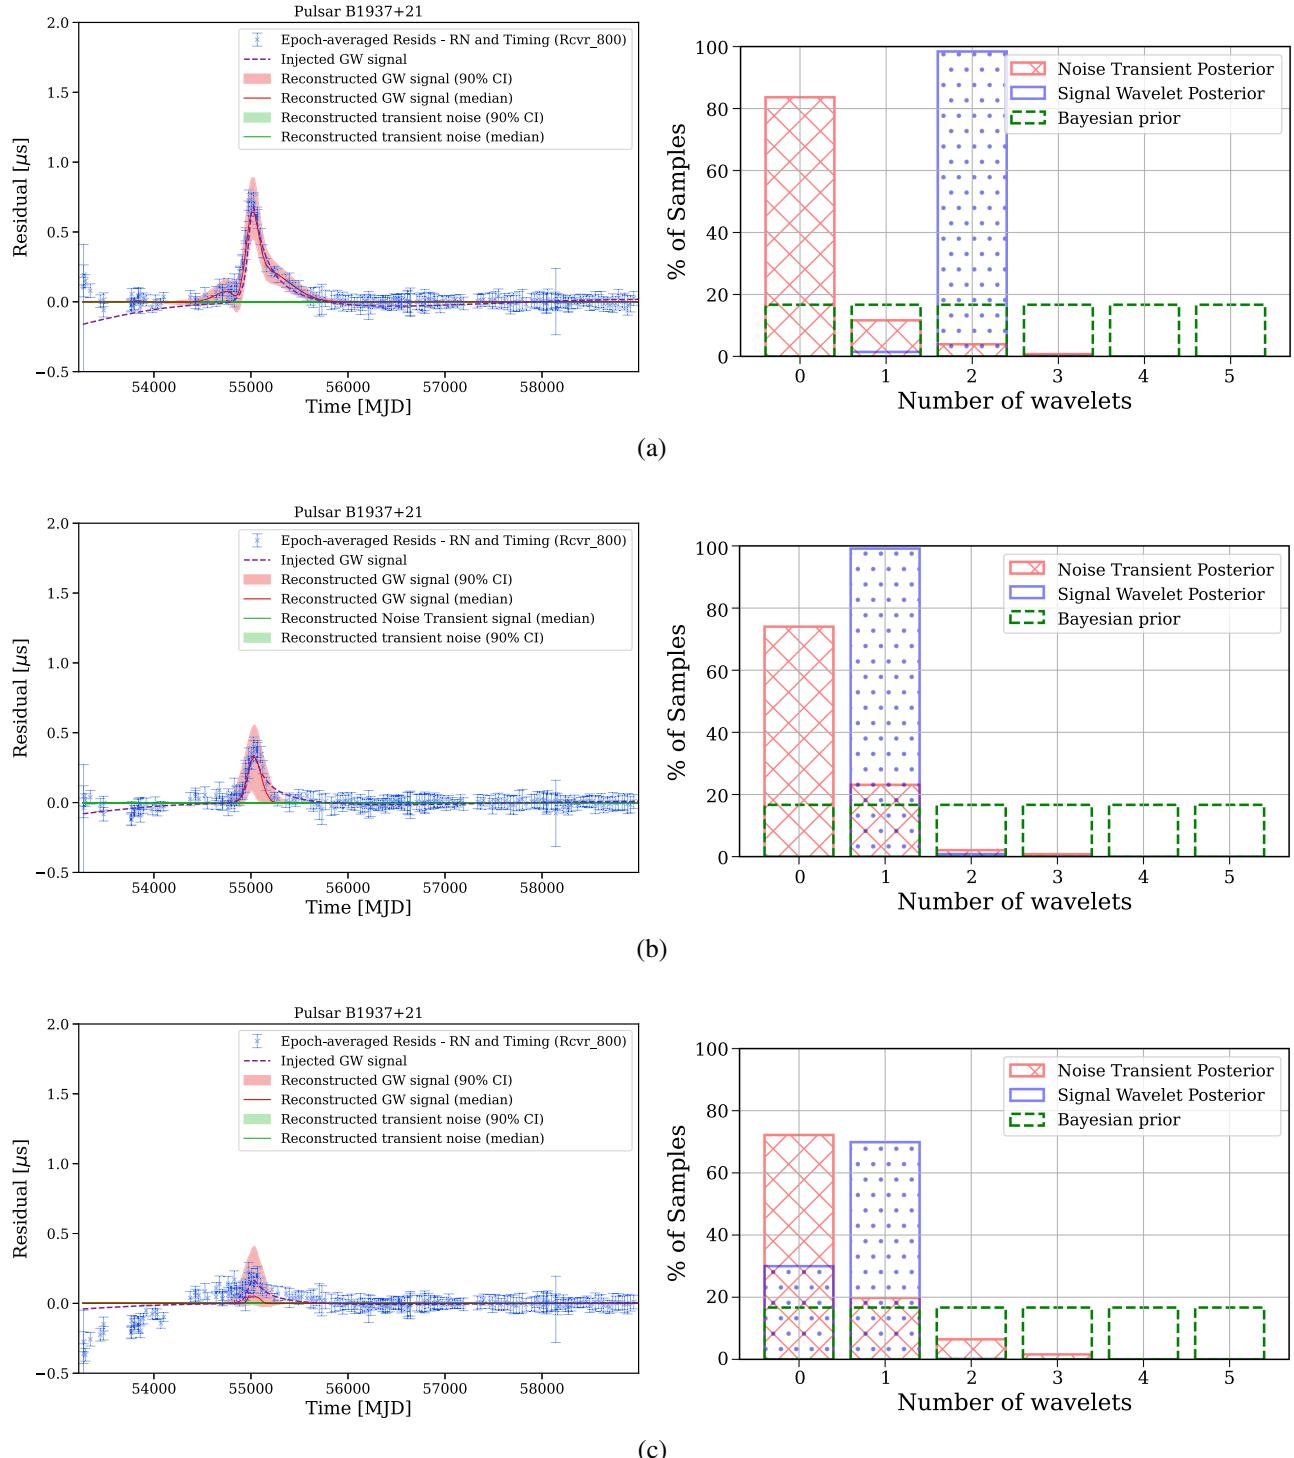


FIG. 3. (a) 30 Mpc SMBHB parabolic encounter dataset with an injected signal SNR of 15.7 and $M = 0.96$. (b) 60 Mpc SMBHB parabolic encounter dataset with an injected signal SNR of 7.8 and $M = 0.86$. (c) 120 Mpc SMBHB parabolic encounter dataset with an injected signal SNR of 3.9 and $M = 0.84$. For each dataset, (left) Signal reconstruction in B1937 + 21. (right) Number of recovered signal and noise transient wavelets. The waveform reconstructions are created using the La Forge software package [34].

The solid green line is the median noise transient reconstruction, and the shaded green region is the 90% credible region for the noise transient (if any are present). The plots on the right show the posteriors for the number of

signal wavelets and noise transient wavelets for these analyses. The three columns from top to bottom correspond to the three analyzed datasets, SIM-HIGH, SIM-MID, and SIM-LOW, respectively (see Table II).

First, we see that the waveform for the GW bursts in these three datasets are accurately reconstructed. We find that the median reconstructions match the injected waveforms with values $M = 0.96$, $M = 0.86$, $M = 0.84$ for the SIM-HIGH, SIM-MID, and SIM-LOW dataset respectively. We also see the match monotonically increase with GW burst SNR as expected [35].

Next, for the recovered signal wavelets, it has been seen in other generic GW burst algorithms that utilize a sine-Gaussian basis (see [36]), where the number of recovered wavelets scales with GW burst SNR, as the model can fit for more features in a higher SNR signal. This is consistent with the signal wavelet posteriors in Fig. 3. Regarding the SIM-LOW dataset, we find that due to the significantly lower signal SNR, convergence took $\sim 3\times$ longer than the SIM-MID and SIM-HIGH analyses. This is a result of marginal differences in the Bayesian likelihood when proposing GW signal wavelets at this SNR. Furthermore, we see there is a significant preference for no noise transient wavelets in all three cases, and that our reconstructions are always consistent with zero. Despite this, there is still a non-negligible percentage of samples at greater than zero noise transients. Having no noise transients in our model is preferred, but having a few noise transients cannot be completely ruled out. This is expected, as white noise fluctuations can be fitted with a small amplitude wavelet.

We also test that QuickBurst can properly recover the sky location of a GW burst. In Fig. 4 we show the sky recovery results of three different runs. In the top two plots, we plot the sky location recovery for a search on the SIM-MID dataset, both where we include only our longest timed pulsars (top figure), and where we include all the pulsars in the NANOGrav 15 year PTA (middle figure). We then display the sky location recovery for a run on the SIM-MID dataset, but with a shifted GW burst sky location (bottom figure, referred to as the “SIM-MID sky shifted” dataset).

For the original SIM-MID search, the burst is only constrained to around a quadrant of the sky. We see, however, that including all NANOGrav 15 year pulsars (as shown in the middle figure) does not provide additional constraints on our sky location recovery for this particular GW burst signal. The pulsars labeled by black dots are the additional pulsars we are adding, but there is limited data for those pulsars around the burst epoch, despite lying in a more sensitive sky region. It is also important to note that GW burst signals of this type have a finite width. While there are small changes in the posterior between the top two panels, the overall shape is the same between the 24 pulsar case and the 67 pulsar case. This is expected behavior given that these additional pulsars contribute little new GW burst information.

To explore whether this is a feature of the SIM-MID dataset or our pipeline, we can simply change our sky

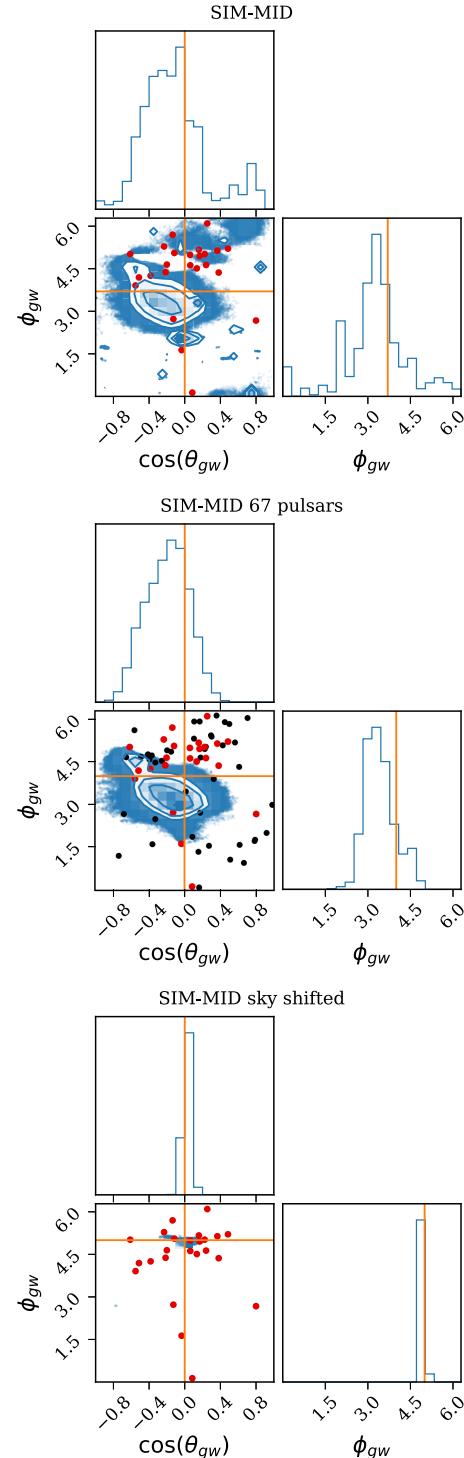


FIG. 4. Sky location recovery for a search on the SIM-MID dataset that includes either 24 NANOGrav pulsars, (top) and (bottom), or 67 NANOGrav pulsars (middle). The 10 year and higher baseline pulsars are displayed as red dots, while the pulsars with less than a 10 year baseline are displayed as black dots. The lines indicate injected sky location parameters, which are at $\phi_{gw} = 4$ and $\theta_{gw} = \pi/2$ for (top) and (middle), but are shifted to $\phi_{gw} = 5.0$, $\theta_{gw} = \pi/2$, and $\psi_{gw} = -0.3$ for (bottom). All angles are in units of radians.

location of the burst to have overlap with more of our longer timed pulsars. Moreover, we ensure that the total GW burst SNR is close to the signal SNR in the SIM-MID dataset. In doing so, we create a new dataset that is almost identical to the SIM-MID dataset, except the burst location is now at $\phi_{\text{gw}} = 5.0$ rad, $\theta_{\text{gw}} = \pi/2$, and $\psi_{\text{gw}} = -0.3$ rad. Based on the bottom figure in Fig. 4, we clearly recover the injected burst location with a shifted burst location. In comparison with the sky location recovery efforts (as shown in the top and middle panels of Fig. 4), we see that having better sky coverage allows us to better recover both θ_{gw} and ϕ_{gw} , so these sky recovery differences are features of the datasets, and not a limitation of the sampler. We will discuss the issue of sky coverage and GW burst sensitivity more in Appendix A.

3. Common red noise effects

The analysis we have discussed so far does not include a common uncorrelated red noise, or CURN. This shows the effectiveness of QuickBurst in a simpler context, but as we have seen from the NANOGrav 15 year GWB results [1], there is strong evidence for a correlated GWB. While it is beneficial to analyze the impacts of a correlated RN process on GW burst analysis, it is not expected that adding in Hellings-Downs (HD) correlations produces significant additional covariance with our wavelet model (see Appendix B for more details). Therefore, we will show QuickBurst can disentangle a GW burst from a CURN process.

To isolate the effects a CURN process would have on our dataset, we test on a simple simulated data set of 20 pulsars with evenly sampled data that includes simulated pulsar WN, CURN, and a GW burst signal from a parabolic SMBH encounter at a sky location of $\phi_{\text{gw}} = 3.7$ rad, $\theta_{\text{gw}} = \pi/2$, and $\psi_{\text{gw}} = 0.2$ rad (unlike the results in Fig. 3 which used a more realistic dataset made to resemble the NANOGrav 15 yr dataset). This dataset consists of a CURN injection with an amplitude of $A = 6.0 \times 10^{-15}$, a spectral index of $\gamma = 13/3$, and 1 μs of intrinsic pulsar WN (which sets the CURN process to be larger than the WN at the lowest frequencies).

First, we perform a search with QuickBurst where we fix our CURN process to the injected CURN values. In doing so, we recover a significant portion of our injected signal, boasting a recovered SNR of 6.4 (injected SNR of 6.8) and a match of $M = 0.82$. We then perform a run on the same dataset where we simultaneously vary both wavelet parameters and CURN spectral parameters, which is displayed in Fig. 5(a). Results indicate a match of $M = 0.79$ and a recovered SNR of 10.2 (injected SNR of 9.0). Furthermore, in comparison to histograms shown in Fig. 3, the distribution of signal wavelets and noise transient wavelets are wider. This is likely due to covariance between the CURN process at overlapping frequencies, which can result in some of the GW burst power being

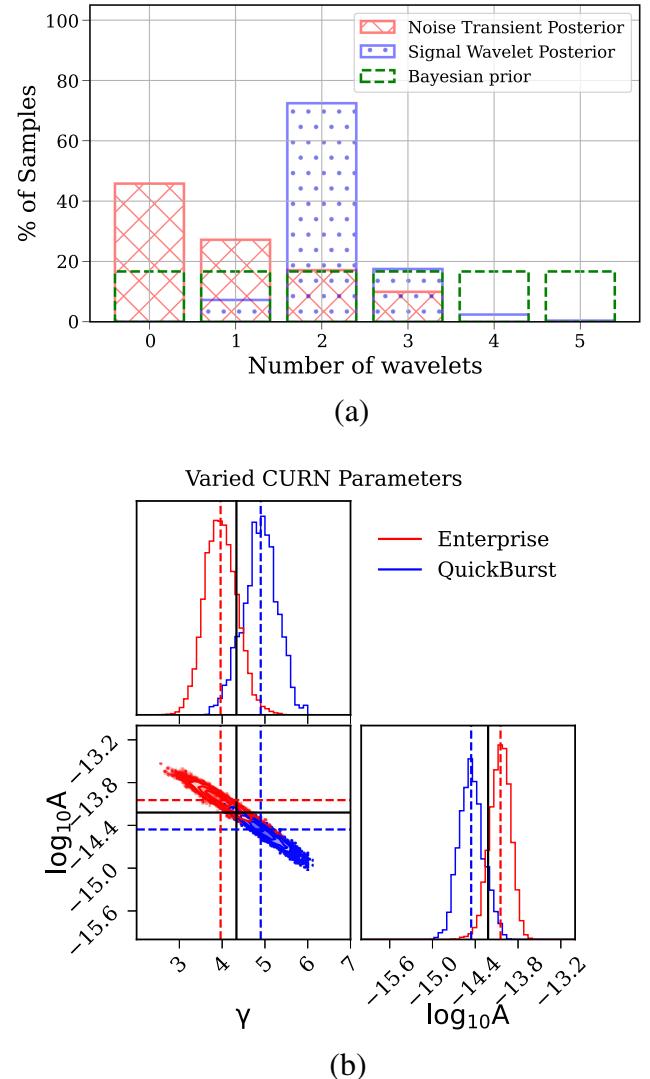


FIG. 5. Plots using a dataset with 20 pulsars and a CURN of amplitude $A = 6.0 \times 10^{-15}$. CURN parameters were allowed to vary. (a) Histogram of recovered signal and noise transient wavelets as found in QuickBurst while simultaneously searching over CURN spectral parameters. (b) Corner plot of recovered CURN power law parameters from both the QuickBurst run that generated the top plot (blue), and a search using enterprise for only CURN and other noise intrinsic pulsar (red). The injected values were $\gamma = 4.33$ and $\log_{10} A = -14.22$ overlay-ed as black lines, while the dotted lines correspond to the median recovered values for each run. QuickBurst recovered median values of $\gamma = 4.90$ and $\log_{10} A = -14.45$, while the search with ENTERPRISE yielded median values of $\gamma = 3.96$ and $\log_{10} A = -14.05$. Both of these recoveries are consistent with the injected signal within a 65% credible region.

absorbed into the CURN model. This suggests that even in strongly red noise dominated data, it is possible to disentangle the signal. There may be additional complications on real data in the future, where additional methods may help to better disentangle these two signal types, if a GW burst is present.

It is also informative for future searches to see the effects of excluding a GW burst from our overall signal model. To analyze these effects, we perform a search using ENTERPRISE on our simple dataset where only CURN is being modeled (and varied), and compare to a search using QuickBurst where both CURN and our GW signal wavelet parameters are varied. Figure 5(b) shows the posterior distribution of the CURN amplitude (A) and the spectral index (γ) from these two runs. We can see that the difference in posteriors illuminates a bias in the recovery of a CURN process.

In the QuickBurst search, given that there is overlapping frequency content between a CURN signal and a GW burst, the GW burst signal model can absorb some of the CURN power at the overlapping frequencies. This bias sharpens the CURN spectrum, raising the spectral index and lowering its amplitude at higher frequencies. Inversely, in the ENTERPRISE search, our recovered CURN parameters are biased toward a lower spectral index and a higher amplitude. This is a result of the GW burst having higher frequency content, but still having a larger SNR than the WN present. This high frequency GW burst content flattens the CURN spectrum while increasing its overall amplitude. This effect means that we will have to balance using fixed vs varied runs on real datasets to try and most accurately recover any signals that may be present.

IV. CONCLUSIONS

We have developed a pipeline that can search for generic GW bursts in PTA datasets using a basis of Morlet-Gabor wavelets. This method has been demonstrated to improve search efficiency by $\sim 200\times$. QuickBurst is capable of recovering detectable signals in a variety of signal strength regimes, while buried in a realistic simulated PTA with intrinsic pulsar RN, intrinsic pulsar WN, and non-uniform sampling/sky coverage. QuickBurst can also be used to efficiently find noise transients in PTA data. Mitigating such noise transients can have a positive effect on other GW analyses, as it can help make sure the presence of such transients do not bias the GW parameter recovery.

With this pipeline developed, it is now computationally feasible to search agnostically for any kind of gravitational wave burst signal in the newest PTA datasets, like the NANOGrav 15 yr dataset, the EPTA DR2 dataset, or the PPTA DR3 dataset [24,37,38]. QuickBurst will also allow us to scale this analysis to future combined datasets, like the upcoming IPTA DR3 [39]. Future work includes implementing the ability to simultaneously searching for a generic GW burst signal alongside a common correlated RN process, which would take into account the most recent results found in [1–4]. Future work also includes expanding this approach to include radio-frequency-dependent wavelets, which can be useful for modeling transient changes in the dispersion measure or other chromatic effect (for a

review, see [33]). Lastly, DMX is a flexible model that can easily fit features that are not in fact DM variations. This means that DMX can make the covariance between the DM and burst models worse than it intrinsically is, especially on short timescales. Future work may include investigating this covariance through modeling DM as a Gaussian process rather than epoch-by-epoch offsets [40,41].

ACKNOWLEDGMENTS

We would like to thank Xavier Siemens and Jerry Sun for their insightful discussions throughout the project. This NANOGrav project receives support from National Science Foundation (NSF) Physics Frontiers Center Awards No. 1430284 and No. 2020265. This work was also partly supported by the George and Hannah Bolinger Memorial Fund in the College of Science at Oregon State University.

APPENDIX A: PER PULSAR GENERIC GW BURST SNR PREDICTIONS

In the search for a stochastic background, the SNR of the HD correlations scales with the number of pairs in our PTA [42], which motivates using all the available pulsars in the analysis. However, for a generic GW burst, we need a large enough observing time span for any given pulsar to resolve burst features, and pulsars whose noise properties are understood well enough to be subtracted out as accurately as possible. In addition, removing pulsars decreases the computational cost of these searches, giving us more time to resolve the features in the most significant pulsars. In summary, it may be more efficient to remove pulsars with an observation time span less than the maximum width (given by τ) of a GW burst allowed in a GW burst search. These short timed pulsars will therefore be uninformative regarding the morphology of a GW burst, and can bias SNR values. These pulsars will then contribute minimally to the total GW burst SNR [43]. For this reason, the maximum width on the GW signal wavelets is set to $\tau = 5$ years for analysis displayed in Sec. III B 2.

To test for uninformative pulsar SNR contributions, we can calculate the SNR^2 for each pulsar in the SIM-LOW dataset, but not limited to pulsars with the longest observation time-span. We compute these values over 10,000 realizations and average them in two different cases: (1) the burst epoch is fixed to MJD 55,000 (to match the analysis described in Sec. III B 2), while the GW burst sky location is randomized between each realization, and (2) where both the GW burst epoch and sky location are randomized between each realization. Assuming there is a GW burst present, the second case is the most agnostic, as in real datasets, we will not know when or where a GW burst will present itself. We then take the mean SNR^2 value over all 10,000 realizations in each pulsar along with the 65%

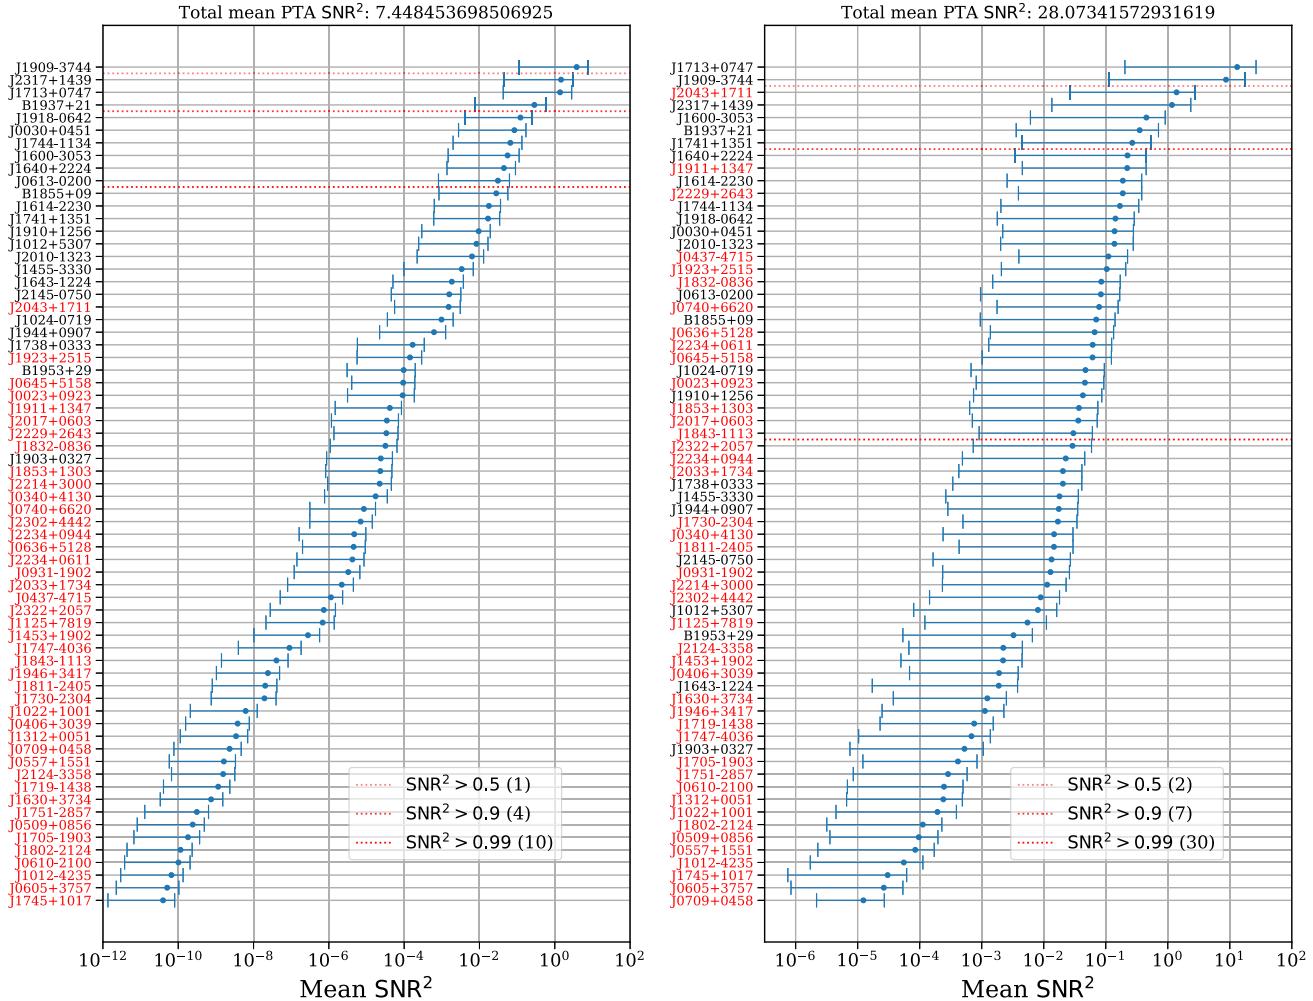


FIG. 6. Sky averaged burst SNR^2 of pulsars in the SIM-LOW dataset. The mean SNR^2 for the full PTA is listed at the top of each plot for each set of 10,000 realizations. Left: mean SNR^2 for all pulsars over 10,000 realizations of randomized burst locations in the sky at a fixed burst epoch of MJD = 55,000. Right: mean SNR^2 for all pulsars over 10,000 realizations for both randomized burst locations on the sky and randomized burst epochs. Pulsars in red indicate the pulsars with less than 10 years of data, which are pulsars that were dropped from our analyzed datasets described in Sec. III. The dashed red lines indicate the cumulative SNR^2 at 50%, 90% and, 99%, and in parenthesis are the number of pulsars that are included in the cumulative sum.

credible regions for both cases, and these are plotted in Fig. 6.

In the fixed epoch case (left figure), the pulsars which have an observing time less than 10 years of data (highlighted in red) show a clear trend to lower SNR^2 values. This directly results from the GW burst epoch occurring at a time when those pulsars have no observational data and miss most GW burst features. In the random epoch case (right figure), the trend is less clear. Generally, there is a higher density of longer observed pulsars having larger SNR contributions. However, there are several exceptions to this trend (namely pulsar J2043 + 1711) that do not allow us to simply cut pulsars from the analysis based on observation time span. In this case, there are 30 pulsars whom contribute to 99% of the total GW burst SNR , but

not all of these 30 pulsars are ones with a greater than 10 year baseline. Therefore, simply dropping pulsars from the analysis is not an advisable solution for additional computational efficiency.

However, the expected SNR calculation presented here provides a compelling alternative, which can be used to at least halve the number of pulsars to analyze without losing any significant sensitivity. Such an approach will make future analyses of real data even more efficient. This may also suggest performing future searches with a tiered approach. One could start by analyzing a dataset with the best ~ 10 pulsars (which would encapsulate $> 90\%$ of the SNR), and then perform follow up searches where more pulsars are added. This tiered approach will increase search efficiency by reducing the data volume being analyzed,

but at the cost of GW burst sky location sensitivity, as shown in Fig. 4.

APPENDIX B: HD CORRELATIONS COVARIANCES WITH GENERIC WAVELET MODELING

Recently, the first evidence for the stochastic GWB has been found in PTA data [1–4]. This motivates considering how HD correlations in the common red noise could impact GW burst analysis. In this analysis, we first consider the spectral characterizations of these common signals to see how significant we expect the impact to be.

As seen in the results from various PTAs, the stochastic GWB has its highest contributions in the lower frequency bins based on free spectrum analysis, which are frequencies ranging from $\sim 1 - 10$ nHz. Because these GW burst searches are intended to search for signals that persist for less than the time span of a dataset, it is not expected that the spectral power from the GWB will significantly overlap with the peak spectral power of a GW burst. An example can be seen in Fig. 3, as the wavelet model prefers a frequency between $f_0 \sim 7 - 9$ nHz, which does not overlap with the highest power contributions of a stochastic GWB with the most recent spectral characterization of the GWB.

To test this, we create a simulated dataset that is identical in construction to the dataset used to analyze the inclusion of CURN (as shown in Fig. 5), with the only difference being including a correlated RN process with HD correlations versus CURN, both with the same spectral parameters. We then perform identical searches with QuickBurst on both datasets. In both cases, we allow the common process in both analyses vary, full knowing QuickBurst only has the ability to model both common processes as CURN.

In Fig. 7, we see the results of this analysis. In the top panel, we compute the match statistic for 10,000 randomly drawn samples from the chain returned by QuickBurst for each dataset, and plot the distribution for each. For the dataset containing CURN, the median value of the match statistic is $M = 0.84 \pm 0.05$, and for the median for the dataset containing HD is $M = 0.82 \pm 0.04$. We see that the match statistic distributions are consistent within 1σ of each other. The slight decrease in the match statistic distribution for the dataset with an HD GWB is insignificant due to significant variations between noise realizations in simulations. Therefore, this shows a negligible increase in covariance between the inclusion of HD correlations and our wavelet model.

Moreover, in the bottom panel of Fig. 7, we see the spectral recovery for the common signals of each dataset. The black crosshair signifies the injected signal parameters for each, and the dashed lines represent the median recovered spectral parameters for each run. For the CURN dataset, the median recovery for the spectral

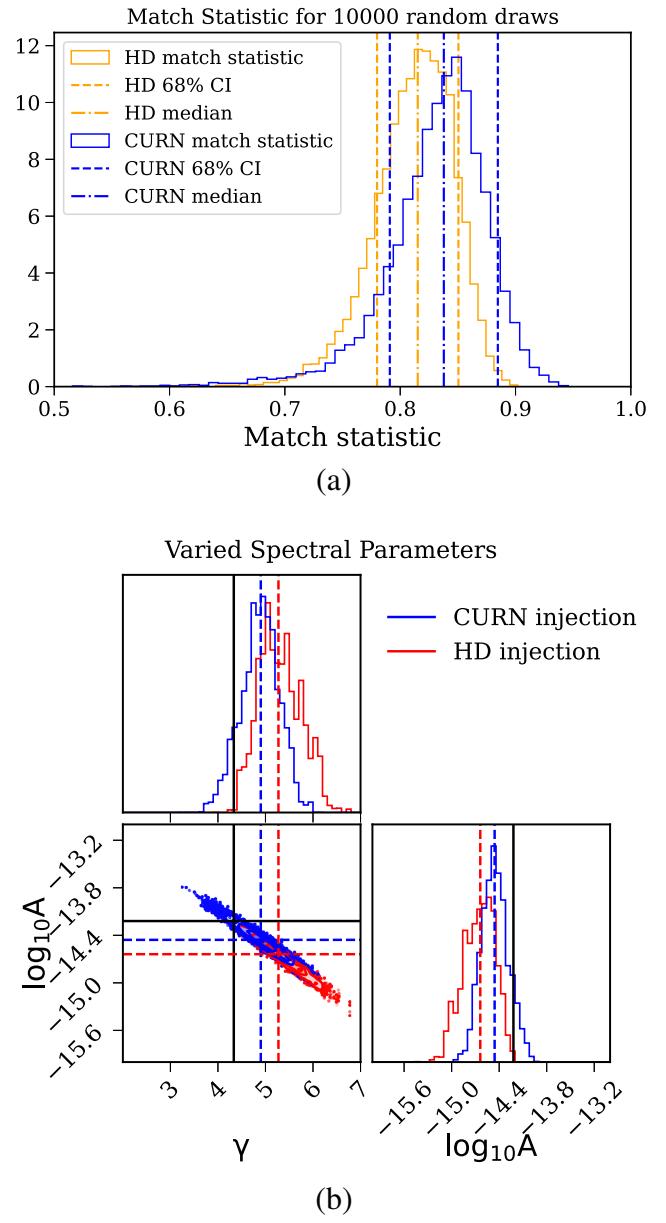


FIG. 7. Comparison between two datasets consisting of identical GW burst signals and intrinsic pulsar noise, but only differing in the common noise. One dataset contains an HD correlated signal and the other contains a CURN signal, yet both have the same spectral parameters of $A = 6 \times 10^{-15}$ and $\gamma = 13/3$. (a) Distribution of the match statistic from identical analyses on both datasets. (b) Corner plot of spectral parameter recovery between both datasets. The black crosshair indicates the injected parameter values for both datasets, while the dashed lines represent the median recovered values for each dataset.

parameters are $A = 3.5 \times 10^{-15}$ and $\gamma = 4.9$. For the HD dataset, the median recovery for the spectral parameters are $A = 2.5 \times 10^{-15}$ and $\gamma = 5.3$. Given that there are minimal changes in the posteriors for the spectral parameters, this implies that the model for HD correlations absorbs little additional power than the CURN model has shown to

absorb from the GW burst that is present. While we cannot conclusively say for sure that a GW burst of any morphology will not have significant covariance with an HD correlated GWB based on this one example, it is promising

that adding in these correlations does not produce any issues with our wavelet model. Similar simulations can easily validate this for other waveforms with different time-frequency content in the future.

[1] Gabriella Agazie, Akash Anumarlapudi, Anne M. Archibald *et al.*, *Astrophys. J. Lett.* **951**, L8 (2023).

[2] J. Antoniadis, P. Arumugam, S. Arumugam *et al.*, *Astron. Astrophys.* **690**, A118 (2024).

[3] Daniel J. Reardon, Andrew Zic, Ryan M. Shannon *et al.*, *Astrophys. J. Lett.* **951**, L6 (2023).

[4] Heng Xu, Siyuan Chen, Yanjun Guo *et al.*, *Res. Astron. Astrophys.* **23**, 075024 (2023).

[5] Zaven Arzoumanian, Paul T. Baker, Laura Blecha *et al.*, *Astrophys. J.* **951**, L28 (2023).

[6] Gabriella Agazie, Akash Anumarlapudi, Anne M. Archibald *et al.*, *Astrophys. J. Lett.* **951**, L50 (2023).

[7] Bence Bécsy and Neil J. Cornish, *Classical Quantum Gravity* **38**, 095012 (2021).

[8] Gabriella Agazie, Zaven Arzoumanian, Paul T. Baker *et al.*, *Astrophys. J.* **963**, 61 (2024).

[9] Adeela Afzal, Gabriella Agazie, Akash Anumarlapudi *et al.*, *Astrophys. J. Lett.* **951**, L11 (2023).

[10] R. Abbott, T. D. Abbott, F. Acernese *et al.*, *Phys. Rev. D* **104**, 122004 (2021).

[11] R. Abbott, T. D. Abbott, F. Acernese *et al.*, *Phys. Rev. D* **104**, 102001 (2021).

[12] Neil J. Cornish and Tyson B. Littenberg, *Classical Quantum Gravity* **32**, 135012 (2015).

[13] Tyson B. Littenberg and Neil J. Cornish, *Phys. Rev. D* **91**, 084034 (2015).

[14] Neil J. Cornish, Tyson B. Littenberg, Bence Bécsy, Katerina Chatzilioannou, James A. Clark, Sudarshan Ghonge, and Margaret Millhouse, *Phys. Rev. D* **103**, 044006 (2021).

[15] S. Klimenko, G. Vedovato, M. Drago, F. Salemi, V. Tiwari, G. A. Prodi, C. Lazzaro, K. Ackley, S. Tiwari, C. F. Da Silva, and G. Mitselmakher, *Phys. Rev. D* **93**, 042004 (2016).

[16] Ryan Lynch, Salvatore Vitale, Reed Essick, Erik Katsavounidis, and Florent Robinet, *Phys. Rev. D* **95**, 104046 (2017).

[17] Travis Robson and Neil J. Cornish, *Phys. Rev. D* **99**, 024019 (2019).

[18] X. J. Zhu, L. Wen, G. Hobbs, Y. Zhang, Y. Wang, D. R. Madison, R. N. Manchester, M. Kerr, P. A. Rosado, and J. B. Wang, *Mon. Not. R. Astron. Soc.* **449**, 1650 (2015).

[19] Heling Deng, Bence Bécsy, Xavier Siemens, Neil J. Cornish, and Dustin R. Madison, *Phys. Rev. D* **108**, 102007 (2023).

[20] Code publicly available at: <https://github.com/JacobATaylor/QuickBurst>.

[21] Peter J. Green, *Biometrika* **82**, 711 (1995).

[22] David I. Hastie and Peter J. Green, *Stat. Neerl.* **66**, 309 (2012).

[23] Bence Bécsy, Neil J. Cornish, and Matthew C. Digman, *Phys. Rev. D* **105**, 122003 (2022).

[24] Gabriella Agazie *et al.*, *Astrophys. J. Lett.* **951**, L9 (2023).

[25] J. A. Ellis, X. Siemens, and J. D. E. Creighton, *Astrophys. J.* **756**, 175 (2012).

[26] Matthew Pitkin, *Mon. Not. R. Astron. Soc.* **425**, 2688 (2012).

[27] Max A. Woodbury, Memo. Rep. No 42, Statistical Research Group, 1950, p. 4.

[28] Stephen R. Taylor, [arXiv:2105.13270](https://arxiv.org/abs/2105.13270).

[29] Bence Bécsy and Neil J. Cornish, *Classical Quantum Gravity* **37**, 135011 (2020).

[30] Aaron D. Johnson, Patrick M. Meyers, Paul T. Baker *et al.*, *Phys. Rev. D* **109**, 103012 (2024).

[31] Nihan S. Pol, Stephen R. Taylor, Luke Zoltan Kelley *et al.*, *Astrophys. J. Lett.* **911**, L34 (2021).

[32] Jonah B. Kanner, Tyson B. Littenberg, Neil Cornish, Meg Millhouse, Enia Xhakaj, Francesco Salemi, Marco Drago, Gabriele Vedovato, and Sergey Klimenko, *Phys. Rev. D* **93**, 022002 (2016).

[33] Gabriella Agazie, Akash Anumarlapudi, Anne M. Archibald *et al.*, *Astrophys. J. Lett.* **951**, L10 (2023).

[34] Jeffrey Shafiq Hazboun, La Forge, Zenodo, 10.5281/zenodo.4152550 (2020).

[35] Bence Bécsy, Peter Raffai, Neil J. Cornish, Reed Essick, Jonah Kanner, Erik Katsavounidis, Tyson B. Littenberg, Margaret Millhouse, and Salvatore Vitale, *Astrophys. J.* **839**, 15 (2017).

[36] Jonah B. Kanner, Tyson B. Littenberg, Neil Cornish, Meg Millhouse, Enia Xhakaj, Francesco Salemi, Marco Drago, Gabriele Vedovato, and Sergey Klimenko, *Phys. Rev. D* **93**, 022002 (2016).

[37] J. Antoniadis *et al.*, *Astron. Astrophys.* **678**, A48 (2023).

[38] Andrew Zic *et al.*, *Publ. Astron. Soc. Aust.* **40**, e049 (2023).

[39] G. Agazie *et al.*, *Astrophys. J.* **966**, 105 (2024).

[40] Bjorn Larsen, Chiara M. F. Mingarelli, Jeffrey S. Hazboun *et al.*, *Astrophys. J.* **972**, 49 (2024).

[41] G. Agazie, J. Antoniadis, A. Anumarlapudi *et al.*, *Astrophys. J.* **966**, 105 (2024).

[42] Xavier Siemens, Justin Ellis, Fredrick Jenet, and Joseph D. Romano, *Classical Quantum Gravity* **30**, 224015 (2013).

[43] Lorenzo Speri, Nataliya K. Porayko, Mikel Falxa, Siyuan Chen, Jonathan R. Gair, Alberto Sesana, and Stephen R. Taylor, *Mon. Not. R. Astron. Soc.* **518**, 1802 (2023).