Bayesian pulsar timing and noise analysis with Vela.jl: The wideband paradigm

Abhimanyu Susobhanan

Max-Planck-Institut für Gravitationsphysik (Albert-Einstein-Institut), Leibniz Universität Hannover, Callinstraße 38, 30167 Hannover, Deutschland

(Received 23 May 2025; accepted 18 August 2025; published 12 September 2025)

Vela.jl is a package for performing Bayesian pulsar timing and noise analysis written in Julia and Python. In the wideband paradigm of pulsar timing, simultaneous time of arrival and dispersion measure measurements are derived from a radio observation using frequency-resolved integrated pulse profiles and templates without splitting the observation into multiple frequency sub-bands. We describe the implementation of the wideband timing paradigm in Vela.jl, and demonstrate its usage using the NANOGrav 12.5-year wideband data of PSR J1923 + 2515. Vela.jl is the first software package to provide this functionality.

DOI: 10.1103/n3ck-lfdy

I. INTRODUCTION

Pulsars are rotating neutron stars whose broad-band electromagnetic radiation is received as periodic pulses by terrestrial observers [1]. In addition to being fascinating objects in their own right with rich phenomenology, their high rotational stability allows us to use them as celestial clocks to study a wide range of time-domain astrophysical effects ranging from solar wind variability [e.g., [2]] to gravitational waves [e.g., [3]]. This is possible through pulsar timing, the technique of accurately tracking the rotation of a pulsar by measuring the times of arrival (TOAs) of their pulses [4].

In addition to the pulsar rotation, several astrophysical and instrumental effects influence the measured TOAs, including the binary motion of the pulsar, interstellar dispersion, solar system motion, the observatory clock, instrumental noise, etc [5]. Interstellar dispersion, for example, introduces a delay $\Delta_{\rm DM} = \mathcal{K}\mathcal{D}/\nu^2$ to the TOA, where $\mathcal{K}$ is known as the dispersion constant, $\mathcal{D}$ is the electron column density along the line of sight known as the dispersion measure (DM), and ν is the observing frequency in the solar system barycenter (SSB) frame. Pulsar timing (or more generally, pulsar timing and noise analysis) involves accurately modeling all of these contributions, and a mathematical description of these effects is called a pulsar timing and noise model.

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Open access publication funded by the Max Planck Society.

In conventional pulsar timing, also known as narrowband timing, the fundamental measurable quantity is the narrowband TOA. A frequency-resolved radio lightcurve obtained from a pulsar observation is folded using the known pulse period to obtain the integrated pulse profile, which gives the pulsar intensity as a function of observing frequency and pulse phase [6]. Narrowband TOAs are measured by dividing an integrated pulse profile into multiple frequency sub-bands and cross-correlating the profile against a noise-free template in each sub-band [7]. Template profiles are usually derived from data, and both frequency-resolved [e.g., [8]] and frequency-averaged [e.g., [9]] templates have been used in the literature. See, e.g., Wang *et al.* [10] for comparisons of different narrowband TOA generation methods.

Recent years have seen significant improvements in telescope sensitivity both through upgrades to existing observatories and the commissioning of new ones, including the advent of wide-band receivers and backends, reduction in the system noise, and large collecting areas [e.g., [11,12]]. Such improvements, along with long-running pulsar timing experiments like pulsar timing arrays [PTAs: [13,14]], have posed significant challenges to the narrowband timing paradigm. This includes inadequate modeling of frequency-dependent profile evolution [15] and interstellar scattering [16], as well as large data volumes [e.g., [17]]. The wideband timing method addresses some of these challenges by treating the integrated pulse profile as a single two-dimensional entity rather than as a collection of independent sub-bands [18,19]. In this paradigm, a single TOA and a single DM are measured from an integrated pulse profile using the full bandwidth of the observation, and the wideband DM measurement is treated as an additional data point together

with the wideband TOA. This can reduce the data volume by up to two orders of magnitude compared to the narrowband paradigm since TOAs are no longer measured per sub-band. Wideband TOAs and DMs are measured with the help of two-dimensional templates usually derived from the data using principal component analysis [20]. This method has been shown to significantly improve the representation of the frequency-dependent profile evolution in the template. It is also possible to derive a wideband TOA-DM pair by combining simultaneous multiband observations [21]. It should be noted that pulse broadening due to interstellar scattering [e.g., [22]] is not yet adequately handled in the wideband paradigm.[1]

In practice, the wideband TOAs and DMs are usually measured using the PulsePortraiture package [18,20]. Wideband timing methods, described in Sec. II, are available in the TEMPO [25] and PINT [26,27] packages, whereas the TEMPO2 package [4,5] does not support wideband timing. Bayesian noise analysis for wideband datasets, using a linear approximation of the timing model, is available in the ENTERPRISE package [28,29]. The TEMPONEST package [30], which provides Bayesian timing and noise analysis functionality for narrowband datasets without the linear approximation, does not support the wideband paradigm. Further, all the wideband timing and noise analysis efforts so far [e.g., [8,31,32]] have exclusively used a piecewise-constant model [DMX: [33]] to describe the DM variations. Gaussian process models of DM variability are not yet supported in ENTERPRISE, although a method for their implementation was recently presented in Susobhanan and van Haasteren [34].

Vela.jl[2] is a new Bayesian pulsar timing and noise analysis package written in Julia and Python, first presented in Susobhanan [35]. In this paper, we present, for the very first time, a public implementation of the full nonlinear wideband timing and noise model geared towards Bayesian inference in Vela.jl. We aim to provide a fully general implementation of the wideband timing and noise model that can support arbitrary DM models, including interstellar and solar wind dispersion models, both deterministic and stochastic, incorporating the method developed in Susobhanan and van Haasteren [34]. We also summarize the new developments in Vela.jl after the publication of Susobhanan [35] corresponding to both narrowband and wideband paradigms.

This paper is arranged as follows. Section II provides a brief overview of wideband timing and noise analysis, and Sec. III describes its implementation in Vela.jl. We demonstrate this implementation using a real dataset in Sec. IV, and we summarize our work in Sec. V. The Appendix describes a command-line interface for Vela.jl. A glossary of symbols used in this paper is provided as Supplementary Material [36].

II. A BRIEF OVERVIEW OF BAYESIAN WIDEBAND TIMING AND NOISE ANALYSIS

A wideband measurement includes the measured TOA and DM values, the corresponding uncertainties, the observing frequency, and metadata about the observing system stored as flags. It is important to note that a wideband TOA without its DM is not the same as a narrowband TOA and does not contain enough information to do timing accurately. Wideband TOAs are usually stored and distributed as `tim` files in a format that is an extension of the narrowband TOA format.

A wideband TOA $t_{\rm arr}$ measured at a terrestrial observatory is related to the pulse emission time $t_{\rm em}$ as

$$t_{\rm arr} = t_{\rm em} + \Delta_{\rm B} + \Delta_{\rm DM} + \Delta_{\rm GW} + \Delta_{\odot} + \Delta_{\rm clock} + \Delta_{\rm jump} + \mathcal{N}_{\rm R} + \cdots, \quad (1)$$

where $\Delta_{\rm B}$ contains delays arising from the binary motion of the pulsar [37], $\Delta_{\rm DM}$ is the total dispersion delay including interstellar dispersion and solar wind dispersion [38], $\Delta_{\rm GW}$ represents the delays caused by gravitational waves propagating across the line of sight to the pulsar [39], $\Delta_{\odot}$ contains the delays arising from solar system motion [5], $\Delta_{\rm clock}$ is a clock correction that converts the TOA measured against an observatory clock into an international time standard [4], and $\Delta_{\rm jump}$ contains the various instrumental delays. $\mathcal{N}_{\rm R}$ is a time-uncorrelated stochastic term that represents the radiometer noise [1]. Note that we have ignored an unmeasurable constant term representing the vacuum light travel time from the pulsar to the observatory at a fiducial epoch in the above equation.

The wideband DM measurement d is given by

$$d = \mathcal{D} + \mathcal{D}'_{\rm jump} + \mathcal{M}_{\rm R} + \mathcal{M}_{\rm jitter}, \quad (2)$$

where $\mathcal{D}$ is the total astrophysical DM, $\mathcal{D}'_{\rm jump}$ represents the system-dependent wideband DM offsets [31], and $\mathcal{M}_{\rm R}$ and $\mathcal{M}_{\rm jitter}$ are time-uncorrelated stochastic terms arising due to radiometer noise and pulse jitter noise (random pulse shape variations) [40]. $\mathcal{D}$ is given by

$$\mathcal{D} = \sum_{n=0}^{n_D} \frac{1}{n!} D_n (t_{\rm em} - t_0)^n + D_{\rm DMN} + D_{\rm SW} + \cdots \quad (3)$$

[1]Rigorous treatment of interstellar scattering is an open problem in pulsar timing, and it is only modeled approximately even in the narrowband paradigm [16]. Promising avenues for handling interstellar scattering under active investigation include cyclic spectroscopy [e.g., [23]] and CLEAN algorithm-based deconvolution [e.g., [24]].

[2]The source code is available at https://github.com/abhisrkckl/Vela.jl. The documentation is available at https://abhisrkckl.github.io/Vela.jl/.

where D_n are Taylor coefficients that determine the long-term evolution of the astrophysical DM, $D_{\rm DMN}$ represents the stochastic DM variations (DM noise), and $D_{\rm SW}$ represents the solar wind DM contribution. $\mathcal{D}$ is related to $\Delta_{\rm DM}$ as $\Delta_{\rm DM} = \mathcal{K}\mathcal{D}/\nu^2$ as mentioned in the introduction.

The rotational phase ϕ can be written as [4]

$$\phi = \phi_0 + \sum_{n=0}^{n_F} \frac{1}{(n+1)!} F_n (t_{\rm em} - t_0)^{n+1} + \phi_{\rm glitch} + \phi_{\rm SN} + \mathcal{N}_{\rm jitter} + \cdots \tag{4}$$

where $\phi = 1$ represents a full rotation, ϕ_0 is an initial phase, F_n represent the rotational frequency and its derivatives, $\phi_{\rm glitch}$ represents a phase correction due to glitches, $\phi_{\rm SN}$ represents the slow stochastic wandering of the rotational phase known as spin noise, and $\mathcal{N}_{\rm jitter}$ is an uncorrelated noise term due to pulse jitter. The timing residual r and the DM residual δ can be written as

$$r = \frac{\phi - \mathcal{N}[\phi]}{\bar{F}}, \tag{5a}$$

$$\delta = d - \mathcal{D} - \mathcal{D}'_{\rm jump}, \tag{5b}$$

where $\mathcal{N}[\phi]$ is the closest integer to ϕ and $\bar{F} = d\phi/dt_{\rm arr}$ is the topocentric frequency.

Finally, the likelihood function for wideband timing and noise analysis is [31,34]

$$\ln L = -\frac{1}{2}\mathbf{y}^T \mathbf{N}^{-1}\mathbf{y} - \frac{1}{2}\ln\det[2\pi\mathbf{N}], \tag{6}$$

where $\mathbf{y}$ is a $2n_{\rm toa}$-dimensional column vector containing the timing residuals and the DM residuals (r_i, δ_i) corresponding to each wideband measurement computed using Eqs. (1)–(5b), and $\mathbf{N}$ is a $2n_{\rm toa} \times 2n_{\rm toa}$-dimensional TOA and DM covariance matrix, and $n_{\rm toa}$ is the number of TOAs.

Some comments are in order regarding the structure of $\mathbf{N}$. In general, $\mathbf{N}$ is a block-diagonal matrix containing the TOA and DM variances along the diagonal and off-diagonal elements representing the TOA-DM covariance at each epoch. This TOA-DM covariance is a function of the frequency within the observing band at which the measurement is made, and can be removed in principle by adjusting the observing frequency. Pennucci *et al.* [18] estimates this zero-covariance frequency using a Hessian-based point estimate of the covariance matrix assuming a Gaussian distribution. The PulsePortraiture package implements this frequency adjustment procedure, and assuming its accuracy, does not provide covariance estimates in its measurements. However, Pennucci *et al.* [18] themselves demonstrated that this frequency estimate is not always accurate, resulting in small nonzero covariances, especially when the signal-to-noise ratio is low. Further, this procedure implicitly assumes that the noise present in the profile is uncorrelated across frequency channels. This assumption is not valid when significant pulse jitter is present, since it is correlated across frequency channels [41]. Unfortunately, the currently available public wideband datasets, prepared using PulsePortraiture, do not provide any estimate of this covariance. Therefore, despite the caveat above, we treat $\mathbf{N}$ as a diagonal matrix in this paper. This covariance will be investigated in detail in a future publication.

The first $n_{\rm toa}$ diagonal elements of $\mathbf{N}$ are ς_i^2, and the rest are ε_i^2, where

$$\varsigma_i^2 = E^2(\sigma_i^2 + Q^2), \tag{7a}$$

$$\varepsilon_i^2 = \mathcal{E}^2(\epsilon_i^2 + \mathcal{Q}^2). \tag{7b}$$

σ_i and ϵ_i are the measurement uncertainties in TOAs and DMs, respectively, and ς_i and ε_i are known as the scaled TOA and DM uncertainties. The quantities E and Q that modify the TOA uncertainties are known as EFACs (error factors) and EQUADs (errors added in quadrature), respectively. Similarly, $\mathcal{E}$ and $\mathcal{Q}$ are known as DMEFACs and DMEQUADs. The EFACS, EQUADs, DMEFACs, and DMEQUADs characterize the uncorrelated part of the pulse jitter noise and imperfections in the TOA and DM measurement algorithms, leading to misestimation of uncertainties, and these depend on the observing system.

Let $\boldsymbol{a}$ represent the timing model parameters appearing in Eqs. (1)–(5b) whose prior distributions are determined by the hyperparameters $\boldsymbol{A}$, and let $\boldsymbol{b}$ be the time-uncorrelated noise parameters appearing in Eqs. (7a) and (7b). The log-posterior distribution of these parameters given a dataset $\mathfrak{D}$ and a timing and noise model $\mathfrak{M}$ can be written using Bayes theorem as

$$\begin{aligned}\ln P[\boldsymbol{a}, \boldsymbol{b}, \boldsymbol{A}|\mathfrak{D}, \mathfrak{M}] = {} & \ln L[\mathfrak{D}|\boldsymbol{a}, \boldsymbol{b}, \mathfrak{M}] + \ln \Pi[\boldsymbol{a}|\boldsymbol{A}, \mathfrak{M}] \\ & + \ln \Pi[\boldsymbol{b}|\mathfrak{M}] + \ln \Pi[\boldsymbol{A}|\mathfrak{M}] \\ & - \ln Z[\mathfrak{D}|\mathfrak{M}],\end{aligned} \tag{8}$$

where $\ln L[\mathfrak{D}|\boldsymbol{a}, \boldsymbol{b}, \mathfrak{M}]$ is given by Eq. (6), Π represents the prior distributions, and the normalizing constant Z is known as the Bayesian evidence.

The time-correlated stochastic processes appearing in the timing model, such as $\phi_{\rm SN}$ and $\mathcal{D}_{\rm DMN}$, are usually modeled as Gaussian processes in the Fourier domain [30,42]. Assuming their effect on the timing residuals to be small compared to the pulsar rotational period F_0^{-1}, we can express the timing residuals as a linear function in the noise amplitudes $\boldsymbol{\alpha}$ [34]

$$\mathbf{y} = \mathbf{y}' + \mathbf{U}\boldsymbol{\alpha}, \tag{9}$$

where $\mathbf{U}$ is a noise design matrix containing derivatives of r_i and δ_i with respect to the amplitudes $\boldsymbol{\alpha}$. Equation (6) can be written in this case as

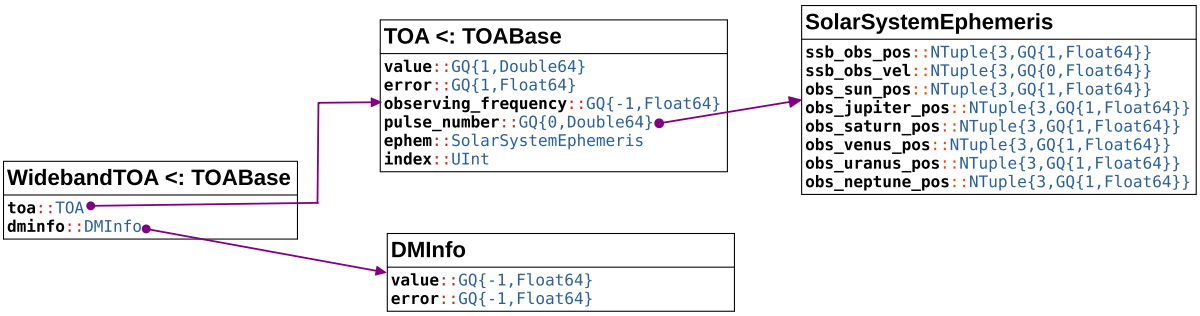


FIG. 1. The structure of the `WidebandTOA` type. <: represents subtype of. The `GQ` type represents a quantity with dimensions $[\mathrm{T}^d]$, and is described in Susobhanan [35].

$$\ln L[\boldsymbol{\theta},\boldsymbol{\alpha},\boldsymbol{b},\boldsymbol{A}|\mathfrak{D},\mathfrak{M}] = -\frac{1}{2}(\mathbf{y}-\mathbf{U}\boldsymbol{\alpha})^T\mathbf{N}^{-1}(\mathbf{y}-\mathbf{U}\boldsymbol{\alpha}) - \frac{1}{2}\ln\det[2\pi\mathbf{N}], \tag{10}$$

where we have separated $\boldsymbol{a}$ into the noise amplitudes $\boldsymbol{\alpha}$ and the remaining timing parameters $\boldsymbol{\theta}$. Substituting this expression into Eq. (8) and imposing Gaussian priors $\Pi[\boldsymbol{\alpha}|\boldsymbol{A},\mathfrak{M}] = \mathcal{N}[0,\boldsymbol{\Phi}(\boldsymbol{A})]$, we can analytically marginalize the posterior distribution over $\boldsymbol{\alpha}$ [34]. This yields

$$\ln P[\boldsymbol{\theta},\boldsymbol{b},\boldsymbol{A}|\mathfrak{D},\mathfrak{M}] = \ln\Lambda[\mathfrak{D}|\boldsymbol{\theta},\boldsymbol{b},\mathfrak{M}] + \ln\Pi[\boldsymbol{\theta}|\mathfrak{M}] + \ln\Pi[\boldsymbol{b}|\mathfrak{M}] + \ln\Pi[\boldsymbol{A}|\mathfrak{M}] - \ln Z[\mathfrak{D}|\mathfrak{M}], \tag{11}$$

with

$$\ln\Lambda[\boldsymbol{\theta},\boldsymbol{b},\boldsymbol{A}|\mathfrak{D},\mathfrak{M}] = -\frac{1}{2}\mathbf{y}^T\mathbf{C}^{-1}\mathbf{y} - \frac{1}{2}\ln\det[2\pi\mathbf{C}], \tag{12}$$

where $\mathbf{C} = \mathbf{N} + \mathbf{U}\boldsymbol{\Phi}\mathbf{U}^T$, $\boldsymbol{\Phi}$ being the prior covariance matrix imposed upon $\boldsymbol{\alpha}$. The covariance matrix $\mathbf{C}$ can also include contributions corresponding to analytically marginalized timing parameters following Haasteren and Levin [43] and Susobhanan and van Haasteren [34], and this is useful for improving the efficiency of the sampler in exploring the parameter space by reducing its dimensionality.

The implementation of the wideband timing and noise model described above in Vela.jl is summarized in the next section.

III. WIDEBAND TIMING IN Vela.jl

A wideband TOA and DM measurement is represented in Vela.jl using the `WidebandTOA` type, containing a TOA object and a `DMInfo` object. The TOA type contains the clock-corrected TOA measurement, its uncertainty, observing frequency, solar system ephemeris, etc, and is also used for representing narrowband TOAs [35]. The `DMInfo` type contains the DM measurement and uncertainty. The structure of the `WidebandTOA` type is shown in Fig. 1.

Similar to the narrowband case, the timing and noise model is represented in Vela.jl using the `TimingModel` type. A `TimingModel` contains an ordered collection of `Components` which compute the delay, phase, DM, and/or uncertainty corrections applicable to each `WidebandTOA`. In particular, the DM corrections appearing in Eq. (2) are represented by subtypes of `DispersionComponent`, which in turn is a subtype of `Component`. These types are also used for representing the timing and noise model in the narrowband case [35], and we have extended the relevant functions such as `correct_toa()` to handle wideband data. See Table 2 and Fig. 9 of Susobhanan [35] for a list of available `Components` and their type hierarchy.

The `TimingModel` also contains a `Kernel` which implements the log-likelihood computation given in Eqs. (6) and (12). In the wideband case, two types of `Kernel` are available. The `WhiteNoiseKernel` implements Eq. (6) and is applicable when the TOA and DM covariance matrix is diagonal. The `WoorburyKernel` implements Eq. (12) and is applicable to a more general class of covariance matrices which can be expressed in the reduced-rank form $\mathbf{C} = \mathbf{N} + \mathbf{U}\boldsymbol{\Phi}\mathbf{U}^T$ with a diagonal or block-diagonal $\mathbf{N}$ matrix.[3] The log-posterior evaluation for wideband datasets is parallelized using multithreading similar to the narrowband case. Further, the implementation of the prior distributions and the handling of model parameters is identical to the narrowband case described in Susobhanan [35].

Partial analytic marginalization over a subset of timing parameters [43] for both narrowband and wideband datasets is implemented with the help of a special `Component` type named `MarginalizedTimingModel` that can only be used as part of a `WoodburyKernel`. Only

[3]The `WoodburyKernel` was implemented after the publication of Susobhanan [35]. It also handles narrowband datasets following, e.g., Lentati *et al.* [30]. The name reflects the fact that the matrix inverse ($\mathbf{C}^{-1}$) appearing in Eq. (12) is evaluated using the Woodbury identity.

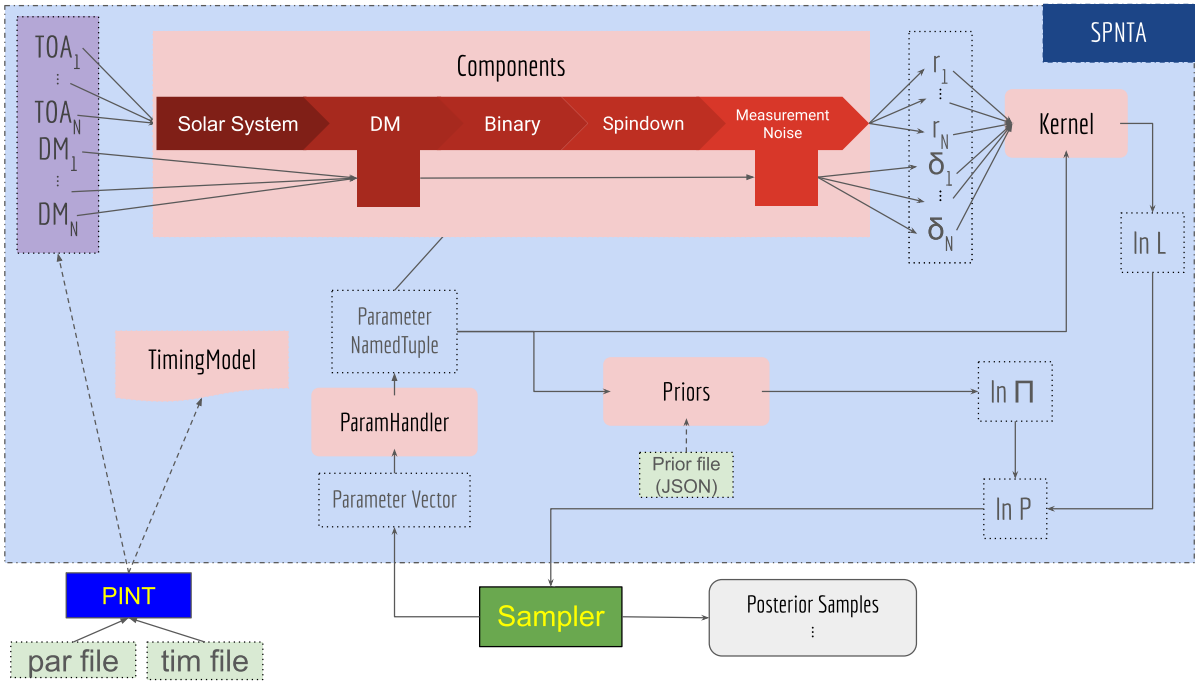


FIG. 2. Schematic diagram summarizing Vela.jl's working in the wideband paradigm. The `Components`, the `Kernel`, the `ParamHandler`, and the `Priors` are part of the `TimingModel` although they are shown separately for clarity. The Measurement Noise block includes EFACs, EQUADs, DMEFACs, and DMEQUADs, and acts on TOA and DM measurement uncertainties, which are not shown explicitly. Except for the `Components` and `Kernel` that act on the DM measurements, all components of Vela.jl behave identically to the narrowband case.

```python
from pyvela import SPNTA
import emcee
import numpy as np

spnta = SPNTA(
    parfile="J1234+5678.par",
    timfile="J1234+5678.tim",
    custom_priors="J1234+5678_priors.json",
    cheat_prior_scale=100,
    analytic_marginalized_params=["PHOFF", "F0", "F1"],
)

nwalkers = spnta.ndim * 5
p0 = np.array([
    spnta.prior_transform(cube)
    for cube in np.random.rand(nwalkers, spnta.ndim)
])

sampler = emcee.EnsembleSampler(
    nwalkers,
    spnta.ndim,
    spnta.lnpost,
)

sampler.run_mcmc(p0, 6000, progress=True)
samples_raw = sampler.get_chain(flat=True, discard=1000, thin=50)

samples = spnta.rescale_samples(samples_raw)
```

FIG. 3. An example script demonstrating the usage of the `pyvela` interface with EMCEE [44]. User-defined priors are read from a `JSON` file, and cheat priors are set based on the frequentist uncertainties scaled by `cheat_prior_scale`; see Appendix B of Susobhanan [35]. The `analytic_marginalized_params` option specifies which parameters should be analytically marginalized. The `emcee.EnsembleSampler` object is initialized with prior samples drawn with the help of `SPNTA.prior_transform()`. The posterior samples are converted from Vela.jl's internal units to their common units using `SPNTA.rescale_samples()`. The high-level syntax shown here is identical for both narrowband and wideband datasets.

parameters whose effects on the timing residuals are exactly or approximately linear are allowed to be analytically marginalized. This currently includes the overall phase offset, rotational frequency, its derivatives, instrumental time jumps, DMX parameters, and wideband DM jumps.

Similar to the narrowband case, wideband timing and noise analysis is also accessible through the `pyvela` Python interface with identical syntax; see Fig. 3. In particular, the `SPNTA` type in `pyvela` also supports wideband data, and the `SPNTA.is_wideband()` method can be used to check if a dataset is wideband. Additionally, we have also implemented a new command-line interface for Vela.jl, and this is described in the Appendix.

The working of Vela.jl in the wideband paradigm is illustrated in Fig. 2. A Python script demonstrating the use of Vela.jl is shown in Fig. 3. Figure 4 shows the execution time of Vela.jl's log-likelihood computation as a function of the number of TOAs using simulated data.

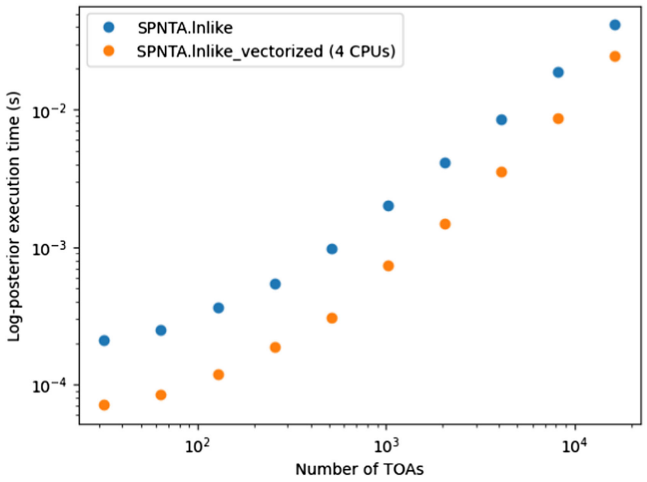


FIG. 4. Execution time of Vela.jl's log-likelihood function as a function of number of TOAs for a simulated narrowband dataset. The execution times for wideband datasets are similar. The timing model includes quadratic spin-down of the pulsar, spin noise modeled as a reduced-rank Gaussian process with 30 harmonics, the relativistic binary model of Damour and Deruelle [37], dispersion measure variations modeled using a quadratic polynomial and a reduced-rank Gaussian process with 30 harmonics, solar system effects (including proper motion and parallax), and instrumental noise modeled using an EFAC, EQUAD, and a DMEFAC, and frequency-dependent profile evolution modeled using FD parameters [33]. The DM and spin noise amplitudes were analytically marginalized. The execution times were estimated on an AMD Ryzen 7 CPU with 16 GB of RAM. The difference between `SPNTA.lnlike` and `SPNTA.lnlike_vectorized` is that the former computes the likelihood function serially, whereas the latter computes the likelihood function in parallel using multiple threads for multiple points in the parameter space (100 points in this case). These numbers depend on the structure of the dataset, the complexity of the timing and noise model, and the executing machine, and therefore must be treated only as order-of-magnitude estimates.

IV. APPLICATION TO PSR J1923+2515

We now demonstrate the application of Vela.jl using the NANOGrav 12.5-year wideband dataset of the isolated pulsar J1923 + 2515 [31]. This dataset contains 119 wideband measurements (TOAs, DMs, and the corresponding uncertainties) taken at the Arecibo observatory between MJDs 55791–57927. The TOA-DM covariances are assumed to be zero. Two receivers, one centered around 1.4 GHz (L-wide) and one centered around 430 MHz (430), and two backends (ASP and PUPPI) were used to record the data.

The timing and noise model includes pulsar rotational frequency and its derivative, solar system delays, interstellar DM and its derivative, a time jump between the two receivers, wideband DM jumps corresponding to each receiver, EFACs, EQUADs, and DMEFACs corresponding to each receiver-backend combination, and reduced-rank Gaussian process models for DM noise and achromatic red noise with power law spectra. Achromatic red noise includes the pulsar spin noise and possibly the effect of a gravitational wave background; these stochastic processes cannot be separated in a single-pulsar analysis. We use 16 linearly spaced harmonics for the DM noise and achromatic red noise with a fundamental frequency of 5.42 nHz corresponding to the total observing time span. Following Haasteren and Vallisneri [42], we also include 4 logarithmically spaced harmonics below the fundamental frequency to better model the lower frequency components of the DM and spin noise. The solar system delays are computed using the DE440 solar system ephemeris [45], and the clock corrections are computed using the BIPM2021 realization of the TT timescale.

Following Susobhanan and van Haasteren [34], we analytically marginalize the DM noise and achromatic red noise amplitudes as well as the overall phase offset, rotational frequency, its derivative, instrumental time jumps, and wideband DM jumps assuming improper Gaussian priors with infinite width to reduce the dimensionality of the parameter space. We draw samples from the partially marginalized posterior distribution using the EMCEE package, which implements the affine-invariant ensemble sampler algorithm [44]. The Python script used to perform this analysis is similar to the one shown in Fig. 3.

The prior distributions used in this analysis are listed in Table I. The prior distribution on parallax is based on a distance estimate obtained from the DM using the YMW16 electron density model [46]. This pulsar is at a DM of around 18.86 pc/cm^3, and this corresponds to a distance of ∼1200 kpc (parallax of ∼0.83 mas) at its sky location. Unfortunately, DM-based distance estimates are known to suffer from large systematic errors. Therefore, we set the parallax prior distribution to be a Gaussian with a mean of 0.83 mas and a 50% standard deviation, truncated at 0. We use log-normal priors for EFACs and DMEFACs so that

TABLE I. The timing and noise model parameters estimated in Sec. IV. The prior distribution and the measurement/constraint for each parameter derived from the posterior samples are listed. Analytically marginalized parameters (overall phase offset, rotational frequency and its derivative, instrumental time jumps, and instrumental DM jumps) are not listed; irregular Gaussian priors with infinite variance are assumed for these parameters. "Trunc" indicates a truncated distribution. The DM derivative is defined at the epoch MJD 56859. Cheat priors, namely uniform priors centered at the frequentist estimate (obtained using PINT) with a width equal to 100 times the frequentist uncertainty, are used for some parameters, and such priors are not listed in the table (see Appendix B of Susobhanan [35] for a detailed explanation). We have checked that increasing the cheat prior widths does not appreciably change the posterior distribution. The marginalized posterior distribution for the spin noise spectral index (TNREDGAM) is consistent with its prior distribution, and this is indicated as Not measured in the table. The posterior distribution constrains the achromatic red noise log-amplitude from above, and its 95.45% upper limit is indicated with †. The measurements/constraints for the white noise parameters EFAC, EQUAD, and DMEFAC are not listed since they are not astrophysically relevant.

Parameter	Description and Unit	Prior distribution	Measurement/Constraint
ELONG	Ecliptic longitude (deg)		297.98094796(1)
ELAT	Ecliptic latitude (deg)		46.69619879(2)
PMELONG	Proper motion in ecliptic longitude (mas/yr)		−9.74(2)
PMELAT	Proper motion in ecliptic latitude (mas/yr)		−12.51(3)
PX	Parallax (mas)	Trunc[Normal[0.83, 0.415], 0, ∞]	0.5(2)
DM	Dispersion measure (pc/cm^3)		18.8558(3)
DM1	Dispersion measure derivative ($pc/cm^3/yr$)		0.00010(4)
EFAC	EFAC for each receiver-backend combination	LogNormal[0.0, 0.25]	
EQUAD	EQUAD for each receiver-backend combination	LogUniform[10^{-3}, 10^2]	
DMEFAC	DMEFAC form each receiver-backend combination	LogNormal[0.0, 0.25]	
TNREDAMP	Achromatic red noise log-amplitude	Uniform[−18, −9]	−13.7†
TNREDGAM	Achromatic red noise spectral index	Uniform[0, 7]	*Not measured*
TNDMAMP	DM noise log-amplitude	Uniform[−18, −9]	−13.7(2)
TNDMGAM	DM noise spectral index	Uniform[0, 7]	2.7(8)

they remain positive, but favor values close to 1. Log-uniform priors are used for EQUADs, and uniform priors are used for achromatic red noise and DM noise log-amplitudes and spectral indices. The data strongly constrain the sky coordinates, proper motion, DM, and DM derivative, and we use uniform cheat priors for these parameters centered around the frequentist estimate, with a width equal to 100 times the frequentist uncertainty (see Appendix B of Susobhanan [35] for a detailed discussion on cheat priors). We have checked that increasing the width of the cheat priors has no discernible impact on the posterior distribution.

The posterior distribution and the postfit residuals obtained from this analysis are plotted in Fig. 5, and the measurement/constraint for each parameter obtained therefrom is listed in Table I. The prior distributions for the different parameters are also plotted in Fig. 5 for comparison. Proper motion is measured, but parallax shows a marginal detection that is consistent with 0 within the 3σ credible interval. Achromatic red noise is not detected, and the posterior distribution shows a 95.45% upper limit of −13.7 for its log-amplitude. This upper limit is consistent with the amplitude of the recently found common achromatic signal, possibly due to a gravitational wave background, discussed in Agazie *et al.* [3]. DM noise is detected, and a hundred random DM time series realizations generated using parameters drawn from the posterior distribution are plotted in Fig. 6.

V. SUMMARY

We implemented support for wideband timing in the Vela.jl Bayesian pulsar timing and noise analysis package. This enables Vela.jl to compute the log-likelihood, log-prior, prior transform, and log-posterior functions for wideband timing datasets using the full nonlinear timing model, optionally marginalizing over the noise amplitudes as well as a subset of timing parameters whose effect on the residuals is approximately linear. Vela.jl is the first package to have this functionality. The log-likelihood and log-posterior computations can be parallelized across multiple CPU cores using multithreading.

We extended the `pyvela` Python interface of Vela.jl such that it provides an interface for wideband datasets that is *identical* to the one for narrowband datasets. This can be used alongside any Markov Chain Monte Carlo [47] or nested sampling [48] library to perform Bayesian inference. We have also developed the `pyvela` command-line interface, which can be used to run quick analyses when fine-tuning of the timing and noise model and the sampler is not necessary, combining Vela.jl with the EMCEE sampling package. We demonstrated the application of Vela.jl on

the NANOGrav 12.5-year wideband dataset of PSR J1923 + 2515. This development should enable, for the first time, Bayesian single-pulsar noise and timing analysis to be performed on wideband datasets, including PTA datasets.

The implementation of wideband timing in Vela.jl is summarized below.

(1) A clock-corrected wideband TOA and DM measurement is represented using the `WidebandTOA` type.
(2) A timing and noise model is represented using the `TimingModel` type. Its structure remains the same as in the narrowband case, but we have extended the various functions operating on `TimingModel` and its constituents to handle `WidebandTOAs`.
(3) The `pyvela` Python interface retains the same interface as in narrowband timing, but now also handles wideband datasets.

The planned future developments in Vela.jl include the implementation of photon-domain timing for high-energy observations [49], the implementation of a sampler optimized for single-pulsar noise and timing analysis [e.g., following [50]], reconstruction of cross-pulsar correlations from `SPNTA` results [51], and the investigation of pulse jitter-induced correlations in wideband measurements. A more general framework for linear fitting of wideband datasets based on this work and Susobhanan and van Haasteren [34] is being implemented in PINT. Further, a nonlinear fitting algorithm for PINT using Vela.jl is also under development.

Software: PINT [26,27], EMCEE [44], PyGEDM [52], NumPy [53], Astropy [54], Matplotlib [55], CORNER [56], DoubleFloats.jl [57], Distributions.jl [58], JuliaCall/PythonCall.jl [59], JLSO.jl [60], and `git`.[4]

ACKNOWLEDGMENTS

I thank David Kaplan for providing valuable feedback on the manuscript and for contributing to the development and testing of Vela.jl. I thank Rutger van Haasteren for fruitful discussions, Chiara Mingarelli for valuable comments on this manuscript, and Paul Ray for valuable feedback on the Vela.jl package.

DATA AVAILABILITY

The data that support the findings of this article are openly available [61].

APPENDIX: THE `pyvela` COMMAND-LINE INTERFACE

The `pyvela` script (not to be confused with the `pyvela` Python interface) provides a convenient way to run a Bayesian pulsar timing and noise analysis from the command line using Vela.jl in combination with the EMCEE sampler. It sacrifices flexibility and fine-tuning capabilities in favor of convenience.

This script takes a pair of `par` and `tim` files, a prior file in the `JSON` format, and an output directory as input, and saves the posterior samples and related metadata such as parameter names, units, etc, into the output directory. An example invocation of the `pyvela` script is given below:

```
pyvela J1234-5678.par J1234-5678.tim
 -o J1234-5678_out
 -A PHOFF F0 F1 -C 200
 -P J1234-5678_priors.json
 -N 6000 -b 1500 -t 40
```

Here, the first two positional arguments are the `par` and `tim` files, the `-o` option specifies the output directory, the `-A` option enables partial analytic marginalization over a subset of timing parameters, the `-P` option provides the prior file, `-C` option specifies the factor with which to multiply the frequentist uncertainties to obtain the cheat priors of parameters whose priors are otherwise unspecified, `-N` option specifies the number of iterations for the sampler, `-b` option specifies the burn-in length, and `-t` option specifies the thinning factor for the posterior samples.

Other scripts included with Vela.jl include `pyvela-plot`, which creates a plot similar to Fig. 5 given `pyvela` output, `pyvela-rethin`, which modifies the burn-in length and thinning factor of the posterior samples, and `pyvela-jlso`, which saves the data, model, and prior distributions into a `JLSO` file [60] for faster reading.

[4] https://git-scm.com/

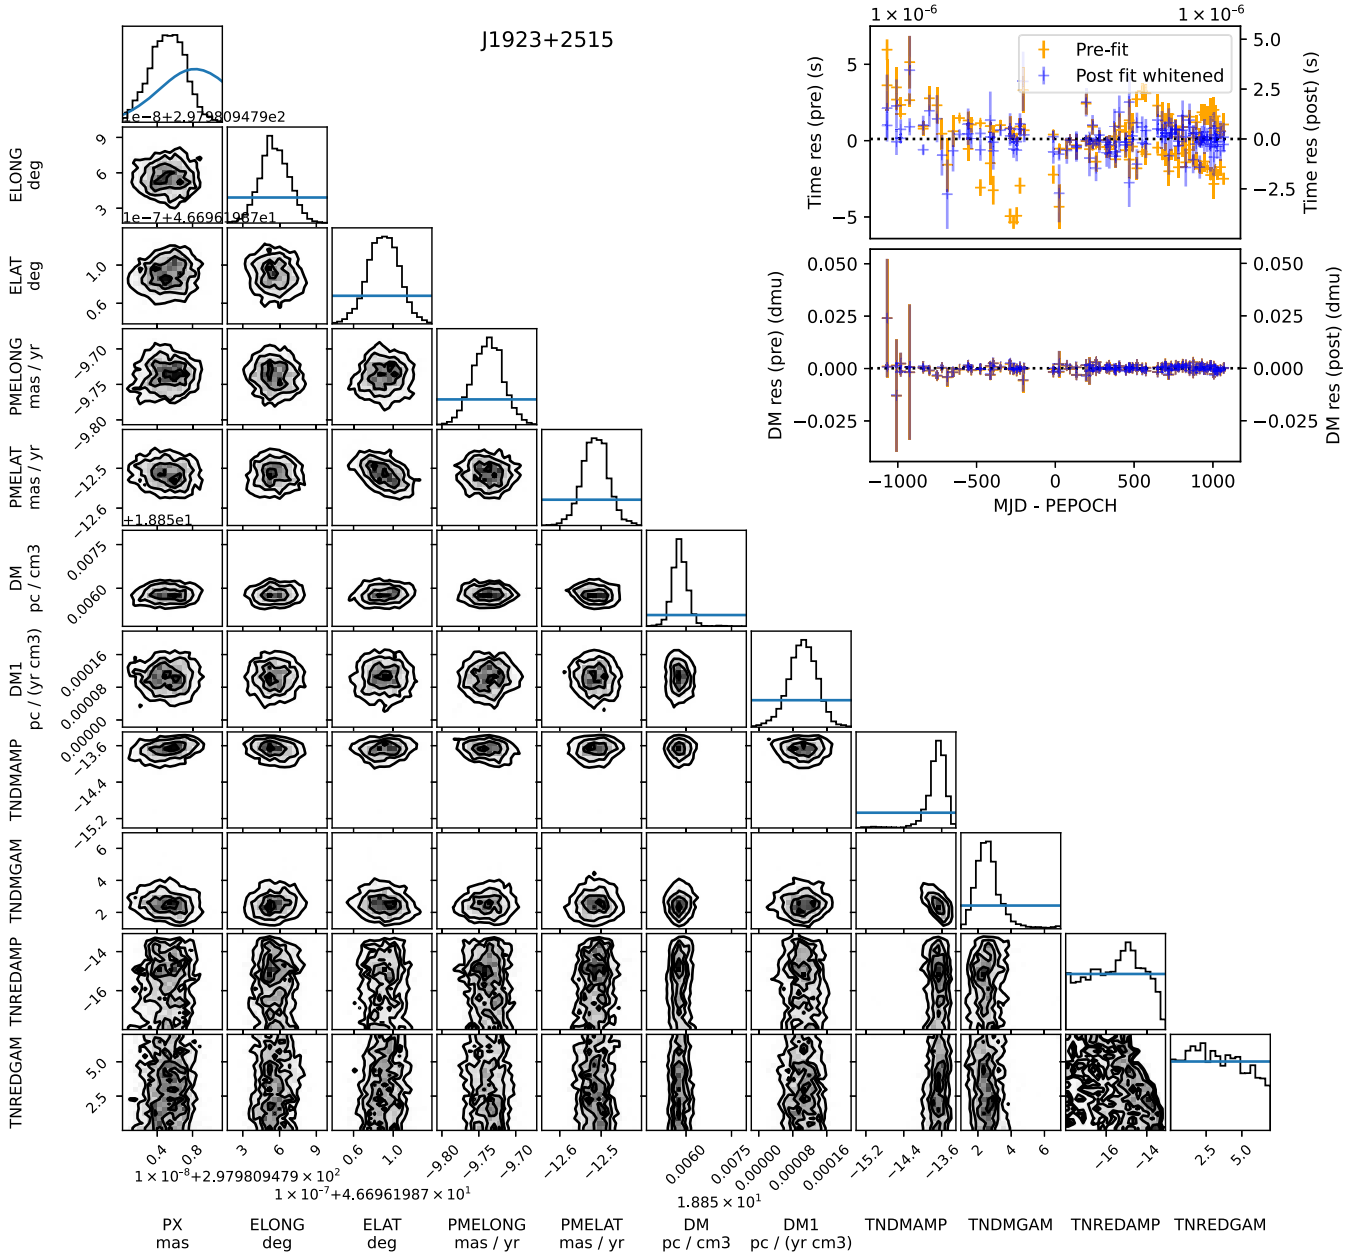


FIG. 5. The posterior distribution for timing and noise model parameters for PSR J1923 + 2515 as described in Sec. IV. Blue curves plotted with marginalized univariate posterior distributions represent the corresponding prior distribution. The parameters and prior distributions shown in this figure are described in Table I. EFACs, EQUADs, and DMEFACs, as well as the analytically marginalized parameters, are not plotted. The time and DM residuals obtained using the median sample are plotted in the inset, along with their prefit counterparts. PEPOCH is the spin frequency epoch, and its value in MJD is 56859.

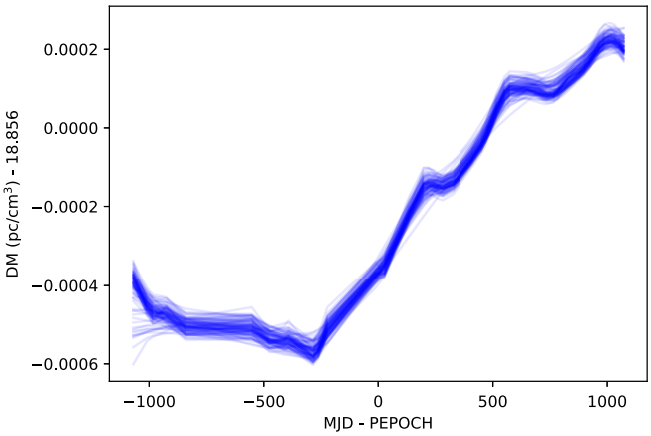


FIG. 6. Hundred random DM time series realizations of PSR J1923 + 2515. These were generated using parameters drawn from the posterior distribution shown in Fig. 5. PEPOCH is the spin frequency epoch, and its value in MJD is 56859.

[1] D. R. Lorimer and M. Kramer, *Handbook of Pulsar Astronomy* (Cambridge University Press, Cambridge, United Kingdom, 2012).

[2] S. C. Susarla, A. Chalumeau, C. Tiburzi, E. F. Keane, J. P. W. Verbiest *et al.*, Exploring the time variability of the solar wind using LOFAR pulsar data, Astron. Astrophys. **692**, A18 (2024).

[3] G. Agazie, J. Antoniadis, A. Anumarlapudi, A. M. Archibald, P. Arumugam *et al.*, Comparing recent pulsar timing array results on the nanohertz stochastic gravitational-wave background, Astrophys. J. **966**, 105 (2024).

[4] G. B. Hobbs, R. T. Edwards, and R. N. Manchester, TEMPO2, A new pulsar-timing package—I. An overview, Mon. Not. R. Astron. Soc. **369**, 655 (2006).

[5] R. T. Edwards, G. B. Hobbs, and R. N. Manchester, TEMPO2, A new pulsar timing package—II. The timing model and precision estimates, Mon. Not. R. Astron. Soc. **372**, 1549 (2006).

[6] A. W. Hotan, W. van Straten, and R. N. Manchester, PSRCHIVE and PSRFITS: An open approach to radio pulsar data storage and analysis, Publ. Astron. Soc. Aust. **21**, 302 (2004).

[7] J. H. Taylor, Pulsar timing and relativistic gravity, Phil. Trans. R. Soc. A **341**, 117 (1992).

[8] P. Tarafdar, K. Nobleson, P. Rana, J. Singha, M. A. Krishnakumar *et al.*, The Indian Pulsar Timing Array: First data release, Publ. Astron. Soc. Aust. **39**, e053 (2022).

[9] P. B. Demorest, R. D. Ferdman, M. E. Gonzalez, D. Nice, S. Ransom *et al.*, Limits on the stochastic gravitational wave background from the North American Nanohertz Observatory for gravitational waves, Astrophys. J. **762**, 94 (2012).

[10] J. Wang, G. M. Shaifullah, J. P. W. Verbiest, C. Tiburzi, D. J. Champion *et al.*, A comparative analysis of pulse time-of-arrival creation methods, Astron. Astrophys. **658**, A181 (2022).

[11] G. Hobbs, R. N. Manchester, A. Dunning, A. Jameson, P. Roberts *et al.*, An ultra-wide bandwidth (704 to 4 032 MHz) receiver for the Parkes radio telescope, Publ. Astron. Soc. Aust. **37**, e012 (2020).

[12] H.-F. Liu, P. Jiang, C. He, F. Yang, H.-J. Liu *et al.*, An ultra-wide bandwidth low-frequency radio astronomical cryogenic receiver for FAST telescope, Res. Astron. Astrophys. **22**, 115016 (2022).

[13] R. S. Foster and D. C. Backer, Constructing a pulsar timing array, Astrophys. J. **361**, 300 (1990).

[14] B. B. P. Perera, M. E. DeCesar, P. B. Demorest, M. Kerr, L. Lentati *et al.*, The International Pulsar Timing Array: Second data release, Mon. Not. R. Astron. Soc. **490**, 4666 (2019).

[15] T. H. Hankins and B. J. Rickett, Frequency dependence of pulsar profiles, Astrophys. J. **311**, 684 (1986).

[16] D. A. Hemberger and D. R. Stinebring, Time variability of interstellar scattering and improvements to pulsar timing, Astrophys. J. **674**, L37 (2008).

[17] G. Agazie, M. F. Alam, A. Anumarlapudi, A. M. Archibald, Z. Arzoumanian *et al.*, The NANOGrav 15 yr data set: Observations and timing of 68 millisecond pulsars, Astrophys. J. **951**, L9 (2023).

[18] T. T. Pennucci, P. B. Demorest, and S. M. Ransom, Elementary wideband timing of radio pulsars, Astrophys. J. **790**, 93 (2014).

[19] K. Liu, G. Desvignes, I. Cognard, B. W. Stappers, J. P. W. Verbiest, K. J. Lee, D. J. Champion, M. Kramer, P. C. C. Freire, and R. Karuppusamy, Measuring pulse times of arrival from broad-band pulsar observations, Mon. Not. R. Astron. Soc. **443**, 3752 (2014).

[20] T. T. Pennucci, Frequency-dependent template profiles for high-precision pulsar timing, Astrophys. J. **871**, 34 (2019).

[21] A. K. Paladi, C. Dwivedi, P. Rana, K. Nobleson, A. Susobhanan *et al.*, Multiband extension of the wideband timing technique, Mon. Not. R. Astron. Soc. **527**, 213 (2023).

[22] M. A. Krishnakumar, D. Mitra, A. Naidu, B. C. Joshi, and P. K. Manoharan, Scatter broadening measurements of 124 pulsars at 327 MHz, Astrophys. J. **804**, 23 (2015).
[23] J. E. Turner, D. R. Stinebring, M. A. McLaughlin, A. M. Archibald, T. Dolch, and R. S. Lynch, Scattering delay mitigation in high-accuracy pulsar timing: Cyclic spectroscopy techniques, Astrophys. J. **944**, 191 (2023).
[24] O. Young and M. T. Lam, Redeveloping a clean deconvolution algorithm for scatter-broadened radio pulsar signals, Astrophys. J. **962**, 131 (2024).
[25] D. Nice, P. Demorest, I. Stairs, R. Manchester, J. Taylor, W. Peters, J. Weisberg, A. Irwin, N. Wex, and Y. Huang, TEMPO: Pulsar timing data analysis, Astrophysics Source Code Library, record ascl:1509.002 (2015), https://ui.adsabs.harvard.edu/abs/2015ascl.soft09002N.
[26] J. Luo, S. Ransom, P. Demorest, P. S. Ray, A. Archibald *et al.*, PINT: A modern software package for pulsar timing, Astrophys. J. **911**, 45 (2021).
[27] A. Susobhanan, D. L. Kaplan, A. M. Archibald, J. Luo, P. S. Ray *et al.*, PINT: Maximum-likelihood estimation of pulsar timing noise parameters, Astrophys. J. **971**, 150 (2024).
[28] J. A. Ellis, M. Vallisneri, S. R. Taylor, and P. T. Baker, ENTERPRISE: Enhanced numerical toolbox enabling a robust pulsar inference suite, Zenodo (2020).
[29] A. D. Johnson, P. M. Meyers, P. T. Baker, N. J. Cornish, J. S. Hazboun *et al.*, NANOGrav 15-year gravitational-wave background methods, Phys. Rev. D **109**, 103012 (2024).
[30] L. Lentati, P. Alexander, M. P. Hobson, F. Feroz, R. van Haasteren, K. J. Lee, and R. M. Shannon, TEMPONEST: A Bayesian approach to pulsar timing analysis, Mon. Not. R. Astron. Soc. **437**, 3004 (2014).
[31] M. F. Alam, Z. Arzoumanian, P. T. Baker, H. Blumer, K. E. Bohler *et al.*, The NANOGrav 12.5 yr data set: Wideband timing of 47 millisecond pulsars, Astrophys. J. **252**, 5 (2021).
[32] M. Curyło, T. T. Pennucci, M. Bailes, N. D. R. Bhat, A. D. Cameron *et al.*, Wide-band timing of the parkes pulsar timing array UWL data, Astrophys. J. **944**, 128 (2023).
[33] Z. Arzoumanian, A. Brazier, S. Burke-Spolaor, S. Chamberlin, S. Chatterjee *et al.*, The NANOGrav nine-year data set: Observations, arrival time measurements, and analysis of 37 millisecond pulsars, Astrophys. J. **813**, 65 (2015).
[34] A. Susobhanan and R. van Haasteren, Gaussian process representation of dispersion measure noise in pulsar wideband datasets, Mon. Not. R. Astron. Soc., staf1422 (2025).
[35] A. Susobhanan, Bayesian pulsar timing and noise analysis with Vela.jl: An overview, Astrophys. J. **980**, 165 (2025).
[36] See Supplemental Material at http://link.aps.org/supplemental/10.1103/n3ck-lfdy for a glossary of mathematical symbols used in this paper.
[37] T. Damour and N. Deruelle, General relativistic celestial mechanics of binary systems. II. The post-Newtonian timing formula, Ann. l'Inst. Henri Poincare Sect. A Phys. Theor. **44**, 263 (1986).https://www.numdam.org/item/AIHPA_1986__44_3_263_0/
[38] D. C. Backer and R. W. Hellings, Pulsar timing and general relativity, Annu. Rev. Astron. Astrophys. **24**, 537 (1986).
[39] F. B. Estabrook and H. D. Wahlquist, Response of Doppler spacecraft tracking to gravitational radiation, Gen. Relativ. Gravit. **6**, 439 (1975).
[40] A. Parthasarathy, M. Bailes, R. M. Shannon, W. van Straten, S. Osłowski *et al.*, Measurements of pulse jitter and single-pulse variability in millisecond pulsars using MeerKAT, Mon. Not. R. Astron. Soc. **502**, 407 (2021).
[41] A. D. Kulkarni, R. M. Shannon, D. J. Reardon, M. T. Miles, M. Bailes, and M. Shamohammadi, An insight into chromatic behaviour of jitter in pulsars and its modelling: A case study of PSR J0437-4715, Mon. Not. R. Astron. Soc. **528**, 3658 (2024).
[42] R. van Haasteren and M. Vallisneri, Low-rank approximations for large stationary covariance matrices, as used in the Bayesian and generalized-least-squares analysis of pulsar-timing data, Mon. Not. R. Astron. Soc. **446**, 1170 (2014).
[43] R. van Haasteren and Y. Levin, Understanding and analysing time-correlated stochastic signals in pulsar timing, Mon. Not. R. Astron. Soc. **428**, 1147 (2013).
[44] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, EMCEE: The MCMC hammer, Publ. Astron. Soc. Pac. **125**, 306 (2013).
[45] R. S. Park, W. M. Folkner, J. G. Williams, and D. H. Boggs, The JPL planetary and lunar ephemerides DE440 and DE441, Astron. J. **161**, 105 (2021).
[46] J. M. Yao, R. N. Manchester, and N. Wang, A new electron-density model for estimation of pulsar and FRB distances, Astrophys. J. **835**, 29 (2017).
[47] P. Diaconis, The Markov Chain Monte Carlo revolution, Bull. Am. Math. Soc. **46**, 179 (2009).https://www.ams.org/journals/bull/2009-46-02/S0273-0979-08-01238-X/
[48] G. Ashton *et al.*, Nested sampling for physical scientists, Nat. Rev. Methods Primers **2**, 39 (2022).
[49] H. J. Pletsch and C. J. Clark, Gamma-ray timing of redback PSR J2339-0533: Hints for gravitational quadrupole moment changes, Astrophys. J. **807**, 18 (2015).
[50] N. Laal, W. G. Lamb, J. D. Romano, X. Siemens, S. R. Taylor, and R. van Haasteren, Exploring the capabilities of Gibbs sampling in pulsar timing arrays, Phys. Rev. D **108**, 063008 (2023).
[51] S. Valtolina and R. van Haasteren, Regularizing the pulsar timing array likelihood: A path towards Fourier space, Phys. Rev. D **112**, 043046 (2025).
[52] D. C. Price, C. Flynn, and A. Deller, A comparison of Galactic electron density models using PyGEDM, Publ. Astron. Soc. Aust. **38**, e038 (2021).
[53] C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen *et al.*, Array programming with NumPy, Nature (London) **585**, 357 (2020).
[54] T. P. Robitaille, E. J. Tollerud, P. Greenfield, M. Droettboom, E. Bray *et al.*, Astropy: A community Python package for astronomy, Astron. Astrophys. **558**, A33 (2013).
[55] J. D. Hunter, Matplotlib: A 2d graphics environment, Comput. Sci. Eng. **9**, 90 (2007).
[56] D. Foreman-Mackey, CORNER.PY: Scatterplot matrices in Python, J. Open Source Software **1**, 24 (2016).
[57] J. Sarnoff *et al.*, DoubleFloats, version 1.2.2, MIT, (2022), https://github.com/JuliaMath/DoubleFloats.jl.

[58] M. Besançon, T. Papamarkou, D. Anthoff, A. Arslan, S. Byrne, D. Lin, and J. Pearson, Distributions.jl: Definition and modeling of probability distributions in the Juliastats ecosystem, J. Stat. Software **98**, 1 (2021).
[59] C. Rowley, PythonCall.jl: Python and Julia in harmony (2022), https://github.com/JuliaPy/PythonCall.jl.
[60] R. Finnegan, L. White *et al.*, INVENIA/JLSO.jl, Zenodo, 2022, https://doi.org/10.5281/zenodo.3992374.
[61] T. T. Pennucci *et al.*, The NANOGrav 12.5-year wideband data set (version 12yv4), 10.5281/zenodo.4312887 (2020).