

A Bayesian parameter estimation approach to pulsar time-of-arrival analysis

This article has been downloaded from IOPscience. Please scroll down to see the full text article.

2011 Class. Quantum Grav. 28 055001

(<http://iopscience.iop.org/0264-9381/28/5/055001>)

View [the table of contents for this issue](#), or go to the [journal homepage](#) for more

Download details:

IP Address: 194.94.224.254

The article was downloaded on 17/01/2012 at 12:41

Please note that [terms and conditions apply](#).

A Bayesian parameter estimation approach to pulsar time-of-arrival analysis

C Messenger^{1,4}, A Lommen², P Demorest³ and S Ransom³

¹ Albert Einstein Institut, 38 Callinstraße, Hannover 30167, Germany

² Franklin & Marshall College, Lancaster, PA, USA

³ National Radio Astronomy Observatory, Charlottesville, VA, USA

⁴ Cardiff School of Physics and Astronomy, Cardiff University, Queens Buildings, The Parade, Cardiff, CF24 3AA, UK

E-mail: chris.messenger@astro.cf.ac.uk

Received 5 October 2010, in final form 17 January 2011

Published 4 February 2011

Online at stacks.iop.org/CQG/28/055001

Abstract

The increasing sensitivities of pulsar timing arrays to ultra-low frequency (nHz) gravitational waves promise to achieve direct gravitational wave (GW) detection within the next 5–10 years. While there are many parallel efforts being made in the improvement of telescope sensitivity, the detection of stable millisecond pulsars and the improvement of the timing software, there are reasons to believe that the methods used to accurately determine the time-of-arrival (TOA) of pulses from radio pulsars can be improved upon. More specifically, the determination of the uncertainties on these TOAs, which strongly affect the ability to detect GWs through pulsar timing, may be unreliable. We propose two Bayesian methods for the generation of pulsar TOAs starting from pulsar ‘search-mode’ data and pre-folded data. These methods are applied to simulated toy-model examples and in this initial work we focus on the issue of uncertainties in the folding period. The final results of our analysis are expressed in the form of posterior probability distributions on the signal parameters (including the TOA) from a single observation.

PACS numbers: 04.30.Tv, 95.30.Sf, 95.55.Ym, 97.60.Gb

1. Introduction

Pulsar timing arrays could well be used to detect ultra-low frequency gravitational waves (GWs) within the next 5–10 years [1]. This is an especially exciting prospect given the concurrent efforts of the LIGO–Virgo Scientific collaboration (LVC) whose aim is to make direct detection of GWs (in the ~ 10 –1000 Hz regime) using the second generation of ground-based interferometric detectors within the same timescale [2].

In this work we outline the beginnings of a Bayesian approach to the detection of GWs with pulsar timing using simplistic signal and noise models onto which can be built further

levels of sophistication in the future. A key long-term aim of our analysis is to improve our ability to time-existing millisecond pulsars by a factor of 3–10 [3, 4]. One of the main problems to be overcome is to be able to sensibly account for the excess low-frequency noise seen in many stable millisecond pulsars [5]. We focus on a *single piece* of the complete pulsar timing analysis, the generation of time-of-arrival (TOA) measurements. Given a single pulsar observation⁵, this is the arrival time of the average pulse at the telescope where in this context ‘average’ means the sum of pulses produced by ‘folding’ the data with a periodicity equal to the assumed pulse period. It is from these TOAs that pulsar astronomers then model the spin evolution of pulsars taking into account the motion of the radio telescope relative to the pulsar [6]. The presence of GWs in the field between the telescope and the pulsar will result in small shifts in the arrival times of pulses [7, 8].

We choose to limit our investigation to single pulsar observations (typically 100–1000 s in duration) and since TOAs are defined in the reference frame of the telescope and the GW timescale is much greater than the timescale of a single observation, we are able to neglect any GW effect in our analysis. We will discuss two different strategies for the estimation of parameters (including the TOA) from two separate starting points, what we will call ‘search-mode’ data and ‘pre-folded’ data. In both cases we perform the analysis using a commonly used Bayesian integration algorithm in order to obtain posterior probability distributions on the signal parameters.

We note that our approach is aimed as a starting point for future more realistic scenarios and that it can be viewed as an approach being built from the bottom-up. We mean this in the sense that we try to start from the most basic datasets available (see sections 3 and 4) and attempt to build a data-analysis framework in which the multitude of physical processes affecting pulsar signals can be included and accounted for. In contrast, other work on the specifics of GW detection using pulsar timing arrays has taken a more top-down approach. These analyses have started with timing residuals, the result of a fit to the data assuming non-GW effects (effectively the end of the pulsar data processing chain), and either neglected this potential inconsistency [9, 10] or made attempts to account for it [11].

The paper is organized as follows. In section 2 we describe our simplistic signal and noise model. In sections 3 and 4 we then go on to describe the form of this signal model in two different representations of the original dataset. The basic concepts concerning our Bayesian approach to the parameter estimation problem can be found in section 5 and finally we discuss our conclusions and potential future developments in section 6.

2. The signal: a toy model

We begin with a dataset defined on a discrete two-dimensional grid of time t_j versus radio frequency f_k an example of which is shown in figure 1. Data of this kind are often referred to as ‘search-mode’ data since this data format is the kind used when performing searches for unknown pulsars. Each of the M rows of the two-dimensional grid is a time series of radio-frequency power measured within a radio-frequency band with central frequency f_k . Typical sampling times and observation durations are ~ 10 s of μ s and 100–1000s of seconds respectively. Typical frequency channel widths and total detector bandwidths are ~ 1 MHz

⁵ We discuss in section 6 that while TOAs are associated with individual pulsar observations (or subsets of an observation), in general a given TOA will depend on parameters ‘fit’ to previous observations.




Figure 1. An example of time versus frequency channel ‘search-mode’ data showing a portion of a simulated dataset consisting of a strong signal in Gaussian noise. Here we show only the first 0.05 s of data ($\Delta t = 64 \mu\text{s}$) for eight 1 MHz wide frequency channels. The signal has an amplitude $A = 5$, pulse width $w = 0.25 \text{ ms}$, period $P = 5 \text{ ms}$, phase $\Phi_0 = 0.2$ and dispersion measure $DM = 100 \text{ cm}^{-3} \text{ pc}$. The noise has unit variance.

and 100–1000s MHz, respectively. For ‘search-mode’ data we assume the following signal model:

$$x(t_j, f_k) = s(t_j, f_k) + n(t_j, f_k), \quad (1)$$

where $x(t_j, f_k)$ represents the discretely sampled dataset, $s(t_j, f_k)$ is the signal and $n(t_j, f_k)$ is the noise, which for simplicity we assume as independent Gaussian distributed random variables with zero mean. The signal itself we define as

$$s(t_j, f_k) = \sum_{\alpha=0}^{n'-1} A \exp \left[-\frac{(t_j - \mu_{\alpha k})^2}{2w^2} \right], \quad (2)$$

where α sums over all n' pulses that intersect with the observation⁶. We use A as the pulse peak amplitude, w as the pulse width and $\mu_{\alpha k}$ as the centre of the α th pulse in the k th frequency channel. Note that we are modelling each pulse as having a single Gaussian profile component and that the amplitude and width remain constant both in time and with frequency channel. The inclusion of additional Gaussian pulse components requires only a trivial modification to the model. In section 6 we discuss numerous potential additions and modifications required to make this toy model a more accurate representation of a real pulsar signal.

⁶ Due to the effects of dispersion, whilst the pulse period is equal in all frequency channels, a particular pulse near the end of the time series for a high-frequency channel can be delayed by dispersion such that it does not intersect with the observation at a lower frequency channel. The same applies to pulses near to the start of the time series in a lower frequency channel since they may arrive before the observation in a higher frequency channel.

The time at the centre of each pulse is defined as

$$\mu_{\alpha k} = (\phi_k + \alpha) P + \xi_{\alpha}, \quad (3)$$

where P is the constant pulse period and ϕ_k is the phase (defined on the range $[0, 1)$) of the first pulse in the observation for the k th frequency channel. We have also included a random pulse ‘jitter’ term where for each pulse we apply a random shift to the pulse arrival time where each shift ξ_{α} is drawn from a Gaussian distribution with zero mean and variance σ_{ξ}^2 . Such effects have been observed in several pulsars and can be attributed to unknown processes in the pulse emission mechanism and possibly related to giant pulses [12–16]. We show in section 3 that for our purposes, the effect of this particular pulse ‘jitter’ modelling can be absorbed into a subset of the other signal parameters.

The phase of the first pulse in each frequency channel ϕ_k can be related to the phase Φ_0 , defined as the phase of the pulse at the midpoint frequency channel $f_{\text{mid}} = (f_M - f_1)/2$ and with reference to the midpoint of the observation $t = T/2$ by

$$\phi_k = \text{mod} \left(\frac{T}{2P} + \frac{\Delta t(f_k)}{P} + \Phi_0, 1 \right). \quad (4)$$

The relative delay due to dispersion in the k th frequency channel $\Delta t(f_k)$ is given by

$$\Delta t(f_k) = 4.148\,808 \times 10^3 (f_k^{-2} - f_{\text{mid}}^{-2}) DM \text{ seconds}, \quad (5)$$

where DM is the dispersion measure in $\text{cm}^{-3} \text{pc}$ and the units of the frequencies are MHz. Note that in our simplistic model we do not account for dispersion smearing within individual channels.

3. Using ‘search-mode’ data in the Fourier domain

The signals received from pulsars are periodic and their frequency evolution is slow, i.e. the timescale of frequency variation is much greater than the pulse period. By Fourier transforming each channel’s time series we find that a pulsar signal can be represented as a series of narrow-band harmonics as shown in figure 2. In a realistic situation we would expect to have some prior knowledge of the pulsar frequency before performing an analysis and therefore transforming the data in such a way allows us to isolate the regions in the dataset where the signal is concentrated (at the harmonics). This in turn allows us to be economic with the data samples that we are interested in and will make any numerical likelihood computation more efficient. Let us define the discrete Fourier transform as

$$\tilde{x}(\nu_l) = \sum_{j=0}^{N-1} x(t_j) e^{-2\pi j l / N} \Delta t, \quad (6)$$

where ν_l represents the elements of a vector containing the positive discrete Fourier frequencies⁷ with frequency spacing $1/T$ and where N is the number of time samples. When applied to the time series from each frequency channel of a noise-free signal (defined by equations (2) and (3)) we obtain

$$\begin{aligned} \tilde{s}(\nu_l, f_k) &= \sum_{j=0}^{N-1} \sum_{\alpha=0}^{n'-1} A \exp \left[-\frac{(t_j - \mu_{\alpha k})^2}{2w^2} \right] e^{-2\pi j l / N} \Delta t, \\ &= \sum_{\alpha=0}^{n'-1} \tilde{s}_{\alpha}(\nu_l, f_k), \end{aligned} \quad (7)$$

⁷ Note that there is a clear distinction between the radio frequencies (or frequency channels) and the Fourier frequencies.

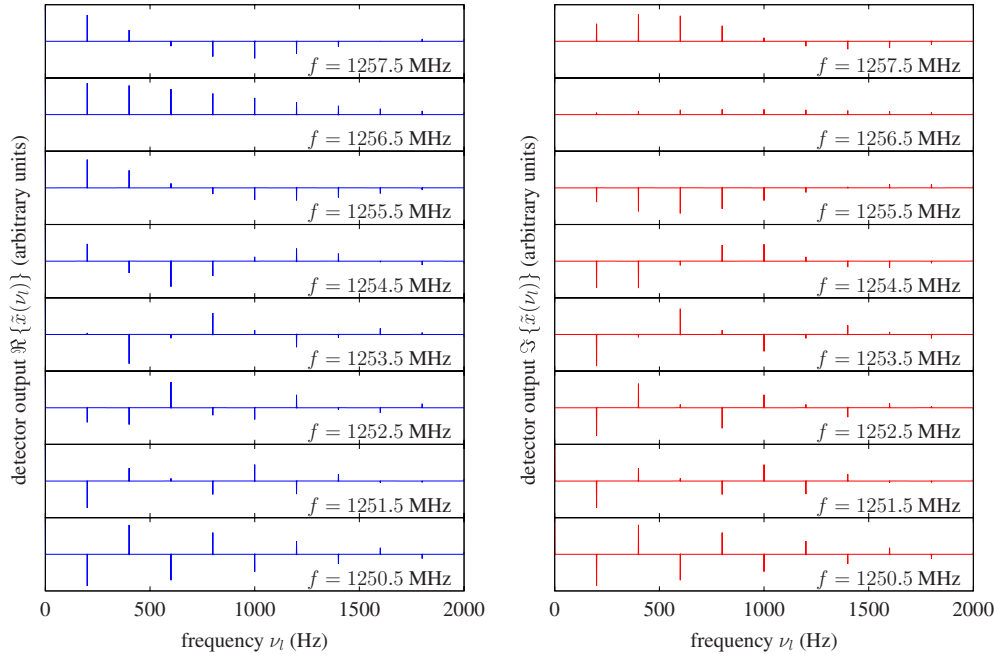


Figure 2. A example of Fourier-transformed ‘search-mode’ data showing a portion of a simulated dataset consisting of a strong signal in Gaussian noise. The panels on the left show (in blue) the real part of the complex Fourier transform of the data as a function of Fourier frequency for eight 1 MHz wide frequency channels. The imaginary parts are shown (in red) on the right. The dataset used to generate this plot is identical to that shown in figure 1 and we have truncated the frequency range at 2 kHz since there the harmonic content of the signal is significantly reduced beyond this frequency.

where we have decomposed the complete Fourier transform into the Fourier transform of each pulse. We then have

$$\begin{aligned}
 \tilde{s}_\alpha(\nu_l, f_k) &= A \sum_{j=0}^{N-1} \exp \left[-\frac{(t_j - \mu_{\alpha k})^2}{2w^2} - 2\pi j l / N \right] \Delta t, \\
 &\approx A \exp\{-2\pi i \nu_l \mu_{\alpha k}\} \int_{-\infty}^{\infty} \exp \left\{ \frac{y^2}{2w^2} - 2\pi i \nu_l y \right\} dy, \\
 &= A w \sqrt{2\pi} \exp\{-2\pi i \nu_l \mu_{\alpha k} - 2\pi^2 w^2 \nu_l^2\}, \tag{8}
 \end{aligned}$$

where we have approximated the discrete sum over time samples with the continuous integral over the dummy variable $y = t - \mu_{\alpha k}$ assuming that each pulse itself spans much greater than 1 time bin and is not truncated by the edges of the time series⁸. We can now perform the sum over α (the individual pulses) to obtain the complete Fourier transform. However, note that $\mu_{\alpha k}$ is a function of ξ_α , the random individual pulse arrival time jitter. We choose to average over this random variable under the assumption that there are a large number of pulses within

⁸ Clearly, of the n' pulses that intersect the time–frequency plane there will be some frequency channels for which a pulse does not appear in the time series due to dispersion. Equation (8) is therefore only applicable to those pulses in a particular frequency channel that are found to intersect the time series.

the observation time. This averaging procedure leads to the following replacement:

$$\begin{aligned} e^{-2\pi i v_l \xi} &\rightarrow \langle e^{-2\pi i v_l \xi} \rangle = \int_{-\infty}^{\infty} \frac{e^{-\xi^2/2\sigma_\xi^2}}{\sqrt{2\pi\sigma_\xi^2}} e^{-2\pi i v_l \xi} d\xi, \\ &= e^{-2\pi^2 v_l^2 \sigma_\xi^2}, \end{aligned} \quad (9)$$

where we have replaced the pulse jitter term with its expectation value and used a Gaussian distribution for the pulse jitter with a zero mean and variance of σ_ξ^2 . Finally we obtain the following expression for the Fourier transform of the signal-only time series:

$$\tilde{s}(v_l, f_k) = \frac{A_\xi w_\xi T}{P} \sqrt{2\pi} \exp\{-2\pi^2 v_l^2 w_\xi^2\} \exp\{-2\pi i \phi_k v_l\} \tilde{W}_l. \quad (10)$$

We can see from this equation that in the Fourier domain the signal can be decomposed into four parts. There is a real positive amplitude term proportional to the pulse amplitude, width and number of pulses ($n \approx T/P$) which is multiplied by a frequency-dependent envelope function that decays with increasing frequency at a rate proportional to the pulse width. There is also a unit amplitude complex phase term dependent upon the initial phase of the pulse in the given frequency channel multiplied by a second complex phase term \tilde{W}_l given by

$$\tilde{W}_l = \frac{P}{T} \exp\{2\pi i v_l P\} \left[\frac{1 - \exp\{-2\pi i v_l T\}}{\exp\{2\pi i v_l P\} - 1} \right], \quad (11)$$

which, in the limit of $T \gg P$, can be written as

$$\tilde{W}_l = \sum_{\beta=0}^n \left\{ \frac{\sin(2\pi \Delta v_{l\beta} T)}{2\pi \Delta v_{l\beta} T} + i \left[\frac{\cos(2\pi \Delta v_{l\beta} T) - 1}{2\pi \Delta v_{l\beta} T} \right] \right\}, \quad (12)$$

where $\Delta v_{l\beta} = v_l - \beta/P$ and β labels the individual signal harmonics of which there are n . This final complex phase term contains the information regarding the location and phase of the signal harmonics. We can now see that each signal harmonic is identical in shape but will each have a different phase and amplitude. In addition, as one moves to different frequency channels the phase of a given harmonic will be rotated by a quantity dependent upon the dispersion measure.

Note that we have also re-parameterized the pulse amplitude and width using

$$A_\xi = \frac{Aw}{\sqrt{w^2 + \sigma_\xi^2}}, \quad (13)$$

$$w_\xi = \sqrt{w^2 + \sigma_\xi^2}, \quad (14)$$

since with the addition of pulse jitter there exists a degeneracy between the original pulse amplitude and width. The product of the amplitude and width determines the overall amplitude of the Fourier transform of the signal and the sum of the squares of the width and the pulse jitter parameter determines the rate of the reduction in the amplitude of the harmonics with increasing frequency. Using the data to measure this amplitude and its attenuation with increasing frequency will therefore not allow us to constrain all three parameters⁹.

⁹ We note that strictly speaking it would be possible to identify the values of all three parameters for a very strong signal. Pulse arrival time jitter acts to remove a small fraction of power from the harmonics and distribute it amongst the inter-harmonic frequency bins. Our analysis is restricted to localized regions at each harmonic and so we treat this information as lost.

4. Using folded data

The majority of pulsar timing data are pre-processed and reduced in volume by the process of folding. In this process, sections of the time series from each frequency channel of an observation will be folded with an assumed pulse period¹⁰. At the time of folding this pulse period will not necessarily be the most accurate value. The pulse period itself is updated and refined with each subsequent observation. However, once data have been folded, most notably for older observations, the original search mode data may be lost, meaning that re-folding with the more refined period is not possible.

We will focus on the effect of folding with an inaccurate pulse period. One can argue that since the most basic initial pulse period estimates will require a coherent measurement over some prior observation spanning many pulses, we should expect an initial worst case fractional uncertainty in the pulse period of $\sim P/T$ which for a 10 ms pulsar period and a 100 s coherent observation equates to a period error of $\sim 1 \mu\text{s}$. In addition to the pulse period, for realistic analyses a number of other parameters are used in the folding procedure such as the sky position coordinates, the intrinsic pulsar spin-frequency derivatives, the dispersion measure plus orbital parameters if the source is in a binary system. In our toy model we ignore these complications.

We choose to define the result of the folding process for a single observation as a two-dimensional grid of pulse profiles labelled by time and channel frequency, an example of which is shown in figure 3. To perform a consistent analysis of such a dataset we take into account the fact that profiles have been obtained using a non-precise value of the pulse period. If we consider a dataset that has already been folded at a specific (non-exact) pulse period $P' = P + \Delta P$, then we can define a new folded dataset as

$$X(\phi', P', f_k) = \sum_{\beta=0}^{n-1} x((\beta + \phi')P', f_k), \quad (15)$$

where β indexes each fold up to $n = \text{floor}(T/P')$. Substituting in our signal model (equations (2) and (3)) we can accurately approximate the discretely summed noise-free pulse profile as

$$S(\phi', P', f_k) \approx \frac{A_\xi w_\xi}{|\Delta P|} \sqrt{\frac{\pi}{2}} \sum_{z=-1}^1 [\text{erf}(a_z + b) - \text{erf}(a_z)], \quad (16)$$

where we have used

$$a_z = \frac{|\Delta P|}{\Delta P} \left[\frac{(P + \Delta P)(\phi' + z) - \phi_k P}{w_\xi \sqrt{2}} \right], \quad (17)$$

$$b = \frac{|\Delta P|(n-1)}{w_\xi \sqrt{2}}. \quad (18)$$

In the calculation of equation (16) we have again replaced the pulse arrival time jitter term with its expectation value (as done in section 3) and approximated the sum over pulses with a continuous integral. We have also been forced to reparameterize the pulse amplitude and width parameters for the same reasons as described in the previous section and have chosen to use an identical reparameterization (defined in equations (13) and (14)). The summation over

¹⁰ The folding procedure can also include de-dispersion over a limited range of frequency channels where, just as with folding, an assumed value of the dispersion is used. Hence a large number of frequency channels can be grouped together into a single pulse profile measurement. We do not consider this potential feature of the folding procedure in this work.




Figure 3. An example-simulated folded dataset showing folded pulse profiles for four sub-intervals each spanning 25 s and for eight frequency channels each spanning 1 MHz. The simulated signal parameters are equal to those defined and used in figures 1 and 2 with the exception that here the signal amplitude $A = 0.1$ is significantly lower. Two curves are plotted in each panel. The blue curves are profiles obtained through folding with the true pulse period P . The red curves are the profiles obtained through folding with a pulse period error $\Delta P' = 10$ ns. Note that this size period error is equivalent to a phase error of ~ 0.01 cycles over the course of a sub-integration.

the index z is simply to account for the fact that folding a signal with an arbitrary initial phase may separate the pulse profile into significant contributions spanning the $\phi' = 0 = 1$ point. This also acts to account for the fact that if folding with an incorrect pulse period the true pulse will slowly drift across the ϕ' space. In this scenario the tails of neighbouring pulses begin to contribute to the sum and by including the $z = \pm 1$ terms we are accurately modelling this effect.

5. A Bayesian analysis

The Bayesian component to our approach can be viewed as standard in the sense that we aim to simply apply Bayes probability theorem to the TOA problem with the intention of computing marginalized posterior probability distributions on the signal parameters.

The Bayes theorem can be expressed as

$$p(\boldsymbol{\theta}|\boldsymbol{x}, \mathcal{M}) = \frac{L(\boldsymbol{x}|\boldsymbol{\theta}, \mathcal{M})\pi(\boldsymbol{\theta}|\mathcal{M})}{E(\mathcal{M}|\boldsymbol{x})}, \quad (19)$$

where the term on the left-hand side is the joint posterior probability distribution on the parameter set $\boldsymbol{\theta}$ given a dataset defined by the vector \boldsymbol{x} and a chosen model represented by \mathcal{M} . The function $L(\boldsymbol{x}|\boldsymbol{\theta}, \mathcal{M})$ is the likelihood function describing the dataset \boldsymbol{x} given

the parameter set θ and the model \mathcal{M} . The function $\pi(\theta|\mathcal{M})$ is the joint prior probability distribution on the parameter set θ given the model \mathcal{M} . Finally we have the Bayesian evidence $E(\mathcal{M}|\mathbf{x}, \theta)$ representing the probability of the model \mathcal{M} given the dataset \mathbf{x} .

To obtain marginalized posterior distributions on a particular signal parameter we are required to perform a multidimensional integration of the joint posterior distribution over the remaining parameters. Formally this can be written as

$$p(\theta_m|\mathbf{x}, \mathcal{M}) \propto \int_{\mathcal{S}} d^n \theta' L(\mathbf{x}|\theta, \mathcal{M}) \pi(\theta|\mathcal{M}), \quad (20)$$

where the parameter vector θ' consists of the subset of parameters in the vector θ excluding the parameter θ_m and where \mathcal{S} defines the volume of integration on that space. Note that there is no dependence upon the Bayesian evidence in the calculation of the marginalized posterior distribution since it is independent of the parameter values themselves and can be absorbed into the normalization of the posterior distribution.

In practice the calculation of posterior distribution functions can be a difficult and computationally intensive procedure. Over the last decade much work has been dedicated to the efficient numerical computation of posterior probability distributions and more recently to the evaluation of the Bayesian evidence. One of the now standard tools available for Bayesian data analysis is the Markov chain Monte Carlo (MCMC) [17, 18]¹¹, an efficient method for obtaining random samples drawn from a posterior probability distribution of which there are a number of variations [19–24]. More recently the strategy known as ‘nested sampling’ [25, 26]¹² has given the data analyst the ability to accurately estimate the Bayesian evidence, a model-dependent quantity used to perform model selection. The first direct application of this strategy was to perform cosmological model selection using *WMAP* data [27]. For this work we chose to perform our analysis using the freely available nested sampling algorithm *MultiNest* [28]. Note that this algorithm has been specifically designed to be robust with respect to multi-modal posterior distributions and to compute the Bayesian evidence. For this work we use it purely to obtain posterior probability distributions on the pulsar parameters.

Let us now define the likelihood functions specific to the two approaches described in sections 3 and 4. The likelihood function for the Fourier domain approach to the ‘search-mode’ data is defined as

$$L^{\text{sm}}(\tilde{\mathbf{x}}|\theta) = (2\pi\sigma_f^2)^{-NM/4} \exp \left\{ -\frac{1}{2\sigma_f^2} \sum_{j=0}^{N/2-1} \sum_{k=0}^{M-1} |\tilde{x}_{jk} - \tilde{s}_{jk}(\theta)|^2 \right\}, \quad (21)$$

where $N/2$ and M are the total number of Fourier-frequency and radio-frequency bins, respectively, and we define $\theta = \{A_\xi, w_\xi, DM, P, \Phi_0\}$ as the vector of signal parameters. We have used σ_f^2 to represent the frequency domain noise variance which we assume to be Gaussian, white and stationary and therefore constant for all Fourier and radio-frequency bins. In this ideal scenario the frequency domain noise variance is related to the time domain noise variance σ_t^2 by $\sigma_f^2 = N(\Delta t)^2\sigma_t^2$.

The likelihood function for the folded data can similarly be written as

$$L^{\text{fold}}(\mathbf{X}|\theta) = (2\pi\sigma_x^2)^{-N_s M/2} \exp \left\{ -\frac{1}{2\sigma_x^2} \sum_{j=0}^{N_s-1} \sum_{k=0}^{M-1} (X_{jk} - S_{jk}(\theta))^2 \right\}, \quad (22)$$

¹¹ Gelman *et al* summarize the full breadth of Bayesian analysis. A strong point of this presentation is its thoughtful and informative style, and avoidance of detailed mathematical derivations.

¹² Divinder Sivia presents a very straightforward account of Bayesian analysis. Rather than going for mathematical rigor and an axiomatic approach, he introduces concepts as they are needed to solve a sequence of increasingly complex analysis problems. Very readable and informative.

where N_s is the number of equal length sub-intervals into which each frequency channel's time series has been divided. The noise contribution in a particular folded phase bin is simply the sum of $n = \text{floor}(T/P')$ Gaussian distributed variables of variance σ_i^2 and hence $\sigma_X^2 = n\sigma_i^2$. The parameter vector θ is identical to that defined for the search-mode data.

In general the choice of prior probability distribution functions on the parameters θ would be chosen according to one's prior beliefs on the values of those parameters. However, for the purpose of our toy model investigation we choose 'flat' prior distributions for all parameters with prior ranges chosen to be far greater than the expected span of the posterior distributions. In this case we do not favour any particular choice of parameter values over any others. We note that in making this choice we are disregarding a powerful feature of the Bayesian analysis, the ability to correctly incorporate prior information into the result. However, one can show that for strong signal-to-noise ratios the effect of the prior on the posterior is dominated by that of the likelihood function itself.

To conclude this section we would like to make it clear that what we have described in sections 3 and 4 do *not* constitute two separate models. We have described two separate representations of the same original dataset and have in fact used the same signal model. Model selection therefore could *not* be applied to these two methods. Our aim is to compare the effectiveness of each choice of dataset representation by contrasting the posterior distributions on the signal parameters when a single common time–radio frequency dataset is used to generate both the Fourier radio frequency and a folded dataset. Model selection using the Bayesian evidence and the computation of the Bayes factor (the ratio of model evidences) and odds ratio (the Bayes factor multiplied by the ratio of prior model probabilities) is a potentially powerful tool in future advanced implementations of our analysis strategy. Our choice of nested sampling implementation, MultiNest, has been designed specifically to compute the Bayesian evidence, making model selection between different pulsar signal models an obvious and easy to implement extension of our approach.

6. Discussion

Shown in figure 4 is an example of typical marginalized posterior probability distributions on the signal parameters $\theta = \{A_\xi, w_\xi, DM, P, \Phi_0\}$ plus the TOA parameter t_{TOA} . The latter is not independent of the other parameters and is a function of both the phase parameter and the pulse period such that $t_{\text{TOA}} = \Phi_0 P$ and is therefore defined as the arrival time of the first pulse received at the mid-point frequency channel immediately following the mid-point of the observation.

Our results show that the ability to determine the signal parameters is unaffected by the choice of data representation when comparing the Fourier domain approach and the folded data. This is apparent from the consistent widths of the posterior distributions which define the uncertainty in parameter estimation. The clear effect that we see is the discrepancy between the location of the posterior distributions for the case where the error in the folding period has been accounted for and where it has not. We see that the estimation of the dispersion, pulse amplitude, pulse width and pulse frequency is only marginally affected. However, the phase parameter, and therefore the TOA estimate, is strongly biased by the false assumption that the signal has been folded with the correct pulse period. For the results shown in figure 4 the pulse period error of 10 ns is equivalent to an accumulated phase error of only 3.6 degrees

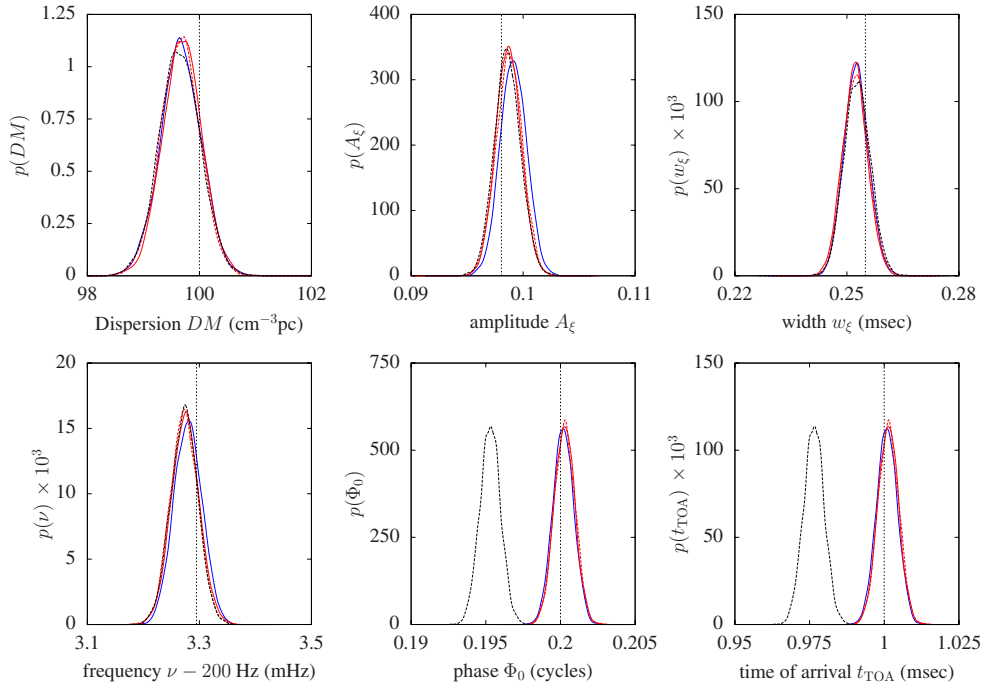


Figure 4. The marginalized posterior distributions on the signal parameters for a simulated signal in Gaussian noise. In solid blue we show the results obtained when using the Fourier domain representation of the ‘search-mode’ data. In solid red we show the results obtained when using folded data as the input for the case where the data were folded with the true pulse period. In dashed red we show the results for the folding scenario where an incorrect period, $\Delta P = 10$ ns, has been used to fold the data and we have accounted for this within the signal model. The dashed black curves show the result where this effect has not been accounted for. The vertical dotted black lines indicate the values of the true signal parameters. For all results the data were converted into the Fourier and folded representations from a single common time–frequency dataset of length 100 s with sampling time $64 \mu\text{s}$ and frequency range of 8 MHz consisting of eight channels each of 1 MHz bandwidth.

over the length of the 25 s sub-integrations. This appears as an ~ 1.8 degree ($\equiv 0.005$ cycles) error in the estimate of Φ_0 leading to an $\sim 25 \mu\text{s}$ error in the estimate of the TOA value¹³.

It is clear that the work presented here is intended only as a potential starting point for more advanced applications of Bayesian data analysis techniques to the problem of pulsar timing. A clear difference between our approach as described here and established techniques is that we have obtained our pulsar parameter estimates from a single simulated observation. The standard approach is to employ a more global strategy in which the process of producing a TOA measurement for a given observation is not just a function of the given observation but of all existing observations of the pulsar. Each TOA represents the reduction of an entire observation into a single number after having performed a global fit (over all observations) for a set of common pulsar parameters, e.g. pulse period, the period derivatives, the sky position, proper motion, pulse shape parameters, the dispersion measure, etc. When new observations are taken, the procedure is repeated and these parameters are refined. As discussed in section 4, as data are recorded they are often reduced (in terms of data volume) by folding at an assumed

¹³ This observed phase error is half of the total accumulated phase error because the phase parameter value is defined at the midpoint of the observation and therefore the phase error effectively accumulates over $T/2$ rather than T .

pulse period and in addition may be partially de-dispersed with an assumed dispersion measure. The detrimental effect of this process (as seen in our results) will rapidly diminish as more and more observations are made but further analysis is required to rule out such effects as contributors to the low-frequency timing noise seen in the ms pulsars.

The scope of this work is limited to the generation of TOAs but we would also like to briefly discuss the specific aim of gravitational wave detection using pulsar timing arrays. From a purely theoretical Bayesian data analysis perspective in an ideal scenario, firstly one would use an unreduced dataset spanning all observations of all relevant pulsars. Secondly one would construct a model including all pulsar signal parameters *and* all gravitational wave signal parameters. After applying sensible prior distributions to all of these parameters one would compute marginalized posterior distributions on both pulsar and gravitational wave parameters and perform model selection. We could then establish whether the observations coupled with our prior beliefs were consistent with the presence of gravitational waves. In practice this is a very difficult task for various reasons but most notably due to the vast computational resources required to process the vast volume of unreduced data and to explore the multidimensional parameter space describing the entire pulsar array and the intervening gravitational wave. For this reason, in terms of gravitational wave detection, constructing a reduced dataset is highly desirable. In fact, the problem of gravitational wave detection using timing residuals (the difference between the TOA values and those attained by fitting a gravitational wave free pulsar model) as the initial dataset has already been applied to the specific case of searching for the gravitational wave stochastic background [10, 11].

The apparent separation of the complete gravitational wave detection problem into a gravitational wave free component, from which a reduced dataset is produced, and then a second component in which this reduced dataset is then analysed including the effects of gravitational waves seems potentially problematic. Under the assumption that each TOA measurement is independent of all others one can argue strongly that the effect of a low-frequency gravitational wave on each measurement is negligible and that the TOA truly represents the unambiguous arrival time of an average pulse within that observation and defined at some epoch. As soon as one performs a global fit (neglecting gravitational waves) over all observations of a given pulsar a gravitational wave of sufficient amplitude will affect the best-fit pulsar parameters. Such a procedure could absorb some fraction of a gravitational wave into the pulsar parameter estimates (e.g. the pulsar period derivatives). In future work we hope to address this issue and to provide a comparison between an analysis using independent TOA measurements as a dataset for gravitational wave detection and an analysis using globally estimated TOAs.

In addition we hope to be able to include, and account for, many of the physical effects and data analysis issues that we have ignored in our toy model approach. These include a more robust treatment of the noise where we allow time and frequency variation and investigate the validity of the assumption of Gaussianity. In reference to this we hope to also include the effects of radio-frequency interference (RFI) and investigate methods in which we are able to analytically marginalize over the noise and therefore potentially avoid the need to estimate it. We also aim to include the effect of polarization into the analysis. A search-mode dataset is itself the product of two independent radio signal polarization measurements which are combined as a function of the Stokes parameters. These parameters can be incorporated into the Bayesian framework and uncertainties on these parameters can be marginalized over in parallel with the signal parameters. Less well-defined effects to consider include a time- and frequency-varying pulse profile parameterization, time-varying dispersion measure, scattering, scintillation and nulling. Finally, we hope to develop this work beyond the toy model to a point at which it can be applied to real pulsar data. In such a scenario we will also have

to incorporate barycentric routines [6] to include the obvious effects of detector motion, sky position uncertainty and, where applicable, binary orbital motion.

Acknowledgments

We thank Maura McLaughlin, Benjamin Knispel, Reinhard Prix, Christian Röver and Xavier Siemens for insightful discussions and invaluable input.

References

- [1] Hobbs G 2010 Pulsar timing array projects *IAU Symp.* ed S A Klioner, P K Seidelmann and M H Soffel **261** 228–33
- [2] The LIGO Scientific Collaboration 2009 LIGO: the Laser Interferometer Gravitational-Wave Observatory *Rep. Prog. Phys.* **72** 076901
- [3] Jenet F *et al* 2009 The North American Nanohertz Observatory for Gravitational Waves arXiv:0909.1058v1
- [4] Hobbs G, Lyne A and Kramer M 2006 Pulsar timing noise *Chin. J. Astron. Astrophys.* **6** (Suppl. 2) 169–75
- [5] Hobbs G *et al* 2010 The International Pulsar Timing Array project: using pulsars as a gravitational wave detector *Class. Quantum Grav.* **27** 084013
- [6] Hobbs G B, Edwards R T and Manchester R N 2006 TEMPO2, a new pulsar-timing package: I. An overview *Mon. Not. R. Astron. Soc.* **369** 655–72
- [7] Detweiler S 1979 Pulsar timing measurements and the search for gravitational waves *Astrophys. J.* **234** 1100–4
- [8] Hellings R W and Downs G S 1983 Upper limits on the isotropic gravitational radiation background from pulsar timing analysis *Astrophys. J. Lett.* **265** L39–42
- [9] Jenet F A, Lommen A, Larson S L and Wen L 2004 Constraining the properties of supermassive black hole systems using pulsar timing: application to 3C 66B *Astrophys. J.* **606** 799–803
- [10] Anholm M, Ballmer S, Creighton J D E, Price L R and Siemens X 2009 Optimal strategies for gravitational wave stochastic background searches in pulsar timing data *Phys. Rev. D* **79** 084030
- [11] van Haasteren R, Levin Y, McDonald P and Lu T 2009 On measuring the gravitational-wave background using pulsar timing arrays *Mon. Not. R. Astron. Soc.* **395** 1005–14
- [12] Kinkhabwala A and Thorsett S E 2000 Multifrequency observations of giant radio pulses from the millisecond pulsar B1937+21 *Astrophys. J.* **535** 365–72
- [13] Kramer M, Johnston S and van Straten W 2002 High-resolution single-pulse studies of the Vela pulsar *Mon. Not. R. Astron. Soc.* **334** 523–32
- [14] Cognard I, Shrauner J A, Taylor J H and Thorsett S E 1996 Giant radio pulses from a millisecond pulsar *Astrophys. J.* **457** L81
- [15] Kondratiev V I, Popov M V, Soglasnov V A, Kovalev Y Y, Bartel N, Cannon W and Novikov A Y 2007 Probing cosmic plasma with giant radio pulses *Astron. Astrophys. Trans.* **26** 585–95
- [16] Kuzmin A D and Ershov A A 2006 Detection of giant radio pulses from the pulsar PSR B0656+14 *Astron. Lett.* **32** 583–7
- [17] Gilks W R, Richardson S and Spiegelhalter D J 1996 *Markov Chain Monte Carlo in Practice* (Boca Raton, FL: Chapman and Hall/CRC)
- [18] Gelman A, Carlin J B, Stern H S and Rubin D B 1995 *Bayesian Data Analysis* (London: Chapman and Hall)
- [19] Marinari E and Parisi G 1992 Simulated tempering: a new Monte Carlo scheme *Europhys. Lett.* **19** 451
- [20] Gramacy R, Samworth R and King R 2010 Importance tempering *Stat. Comput.* **20** 1–7
- [21] Cai B, Meyer R and Perron F 2008 Metropolis–Hastings algorithms with adaptive proposals *Stat. Comput.* **18** 421–33
- [22] Mira A 2001 On Metropolis–Hastings algorithms with delayed rejection *Metron* **59** 231–41
- [23] Hernandez-Marin S, Wallace A M and Gibson G J 2007 Bayesian analysis of lidar signals with multiple returns *IEEE Trans. Pattern Anal. Mach. Intell.* **29** 2170–80
- [24] Trias M, Vecchio A and Veitch J 2009 Delayed rejection schemes for efficient Markov chain Monte Carlo sampling of multimodal distributions arXiv:0904.2207
- [25] Skilling J 2004 *Nested Sampling AIP Conf. Proc.* **735** 395–405
- [26] Sivia D S 1996 *Data Analysis: A Bayesian Tutorial* (Oxford: Clarendon, Oxford University Press)
- [27] Mukherjee P, Parkinson D and Liddle A R 2006 A nested sampling algorithm for cosmological model selection *Astrophys. J. Lett.* **638** L51–4
- [28] Feroz F and Hobson M P 2008 Multimodal nested sampling: an efficient and robust alternative to Markov chain Monte Carlo methods for astronomical data analyses *Mon. Not. R. Astron. Soc.* **384** 449–63