7-10-2012

Practical methods for continuous gravitational wave detection using pulsar timing data

J. A. Ellis

F. A. Jenet

M. A. McLaughlin

PRACTICAL METHODS FOR CONTINUOUS GRAVITATIONAL WAVE DETECTION USING PULSAR TIMING DATA

J. A. ELLIS[1,2], F. A. JENET[3], AND M. A. MCLAUGHLIN[2,4,5]

[1] Center for Gravitation, Cosmology, and Astrophysics, University of Wisconsin-Milwaukee, Milwaukee, WI 53211, USA
[2] Physics Department, West Virginia University, Morgantown, WV 26505, USA
[3] Center for Gravitational Wave Astronomy, University of Texas, Brownsville, TX 78520, USA
[4] Adjunct at the National Radio Astronomy Observatory, Green Bank, WV 24944, USA

ABSTRACT

Gravitational waves (GWs) are tiny ripples in the fabric of space time predicted by Einstein's general relativity. Pulsar timing arrays (PTAs) are well poised to detect low-frequency (10^{-9}–10^{-7} Hz) GWs in the near future. There has been a significant amount of research into the detection of a stochastic background of GWs from supermassive black hole binaries (SMBHBs). Recent work has shown that single continuous sources standing out above the background may be detectable by PTAs operating at a sensitivity sufficient to detect the stochastic background. The most likely sources of continuous GWs in the pulsar timing frequency band are extremely massive and/or nearby SMBHBs. In this paper we present detection strategies including various forms of matched filtering and power spectral summing. We determine the efficacy and computational cost of such strategies. It is shown that using an optimal matched filter explicitly including the poorly constrained pulsar distances with a grid-based method is computationally infeasible. We show that an Earth-term-matched filter constructed using only the correlated signal terms is robust, computationally viable and highly sensitive to GW signals. We further show that a simple power spectral summing technique is nearly equivalent to the Earth-term-matched filter in terms of the minimum detectable amplitude. Both of these techniques are only a factor of two less sensitive than the computationally unrealizable optimal matched filter. We also show that a pairwise matched filter, taking the pulsar distances into account, is comparable to the optimal matched filter for the single template case and comparable to the Earth-term-matched filter for many search templates. Finally, using simulated data optimal quality, we place a theoretical minimum detectable strain amplitude of $h > 2 \times 10^{-15}$ from continuous GWs at frequencies on the order $\sim 1/T_{\rm obs}$.

Key words: gravitational waves – methods: data analysis – pulsars: general

Online-only material: color figures

1. INTRODUCTION

Low-frequency (10^{-9}–10^{-7} Hz) gravitational waves (GWs) are expected from supermassive black hole binary (SMBHB) systems, cosmic strings, the big bang, and inflationary era of the early universe. GWs from these sources can manifest themselves in different ways. Single nearby SMBHBs can produce resolvable waves with periods on the order of years (Wyithe & Loeb 2003; Sesana et al. 2009; Sesana & Vecchio 2010). SMBHBs and cosmic strings can also produce GW bursts (Damour & Vilenkin 2001; Siemens et al. 2007; Leblond et al. 2009) in which the duration of the GW signal is much less than the observation time. We also expect pulsar timing arrays (PTAs) to be sensitive to a stochastic background of unresolvable sources. PTAs offer an opportunity to detect low-frequency GWs from all of these sources. The concept of a PTA composed of the best timed Millisecond pulsars (MSPs) was first developed over two decades ago (Romani 1989; Foster & Backer 1990). Today there are three main PTAs in existence with the goal of GW detection using pulsars: the European Pulsar Timing Array (EPTA; Janssen et al. 2008), the North American Nanohertz Observatory for Gravitational waves (NANOGrav; Jenet et al. 2009), and the Parkes Pulsar Timing Array (PPTA; Manchester 2008), all of which are in collaboration to form the International Pulsar Timing Array (IPTA; Hobbs et al. 2010). There is also a large international effort for the construction of future generation radio telescope arrays such as the Square Kilometer Array

(SKA; Lazio 2009), for which a primary science goal will be GW astrophysics.

Previous authors have developed statistical data-analysis methods for searches of the PTA data sets for stochastic backgrounds (Jenet et al. 2005; Anholm et al. 2009; van Haasteren et al. 2009; Yardley et al. 2011; Demorest et al. 2012) and burst sources (Finn & Lommen 2010). However, studies into continuous GW detection have been more theoretical or "proof of principle" in nature, as opposed to a more rigorous detection method aimed at real data-analysis pipeline implementation. Prior to the establishment of PTAs, Jenet et al. (2004) used existing pulsar data to rule out the proposed SMBHB system 3C66B, a possible source of continuous GWs. This work looked for the signature of a continuous GW in real pulsar data through the use of Lomb–Scargle periodograms and suggested a method for directed searches of known sources. Yardley et al. (2010) also relied on the Lomb–Scargle periodogram to determine the sensitivity of a PTA to continuous GW sources as a function of GW frequency. Sesana & Vecchio (2010) developed a Bayesian framework for the detection of continuous GWs from monochromatic SMBHBs in circular orbits. This work only included the Earth term in the GW signal model and estimated the uncertainties one would expect on search parameters via the Fisher Information matrix, which is known to perform well in the high signal-to-noise ratio (S/N) regime (Vallisneri 2008). Corbin & Cornish (2010) have developed a Bayesian Markov Chain Monte Carlo (MCMC) technique for parameter estimation of an evolving SMBHB system in which

---

[5] Alfred P. Sloan Research Fellow.

the pulsar term is taken into account in the detection scheme. This work took advantage of a signal model in which the GW frequency evolves significantly in time to determine the pulsar distance. Lee et al. (2011) have developed parameter estimation techniques based on vector Ziv-Zakai bounds incorporating the pulsar term and have placed limits on the detectable amplitude of a continuous GW. In their work, a future PTA with SKA sensitivity is assumed, resulting in high S/N signals for their parameter estimation studies. All of these methods are promising for parameter estimation in a relatively high S/N (e.g., S/N = 20 in Corbin & Cornish 2010) limit and a favorable signal model. Very recently Babak & Sesana (2012) have developed a method that can resolve multiple SMBHB sources in the sky using a maximum likelihood approach.

The aim of this paper it to determine the most sensitive *practical* detection technique for continuous GW sources in the PTA data making no assumptions about the S/N of the signal. This is done by comparing multiple detection techniques, using the minimum detectable amplitude as our figure of merit. The paper is ordered as follows. In Section 2 we introduce the formalism and notation that we will use in the paper. In Section 2.1 we define the GW signal from an SMBHB and derive the resulting GW induced pulsar timing residuals. Section 2.2 reviews methods of matched filtering, maximum likelihood detection techniques, and power spectral summing. In Section 3 we describe the simulated data sets that are used in this work. Section 4 describes the different detection techniques and discusses the main results of the paper. Finally, in Section 5 we summarize our work and mention prospects for future work.

2. METHOD

2.1. PTA Response to a Continuous GW

While PTAs are poised to detect a stochastic GW background due to SMBHBs in the next five years; single, resolvable sources may also be detected at expected five-year sensitivity limits (Sesana et al. 2009). A GW is defined as a metric perturbation to flat space time,

$$h_{ab}(t, \hat{\Omega}) = e^+_{ab}(\hat{\Omega})h_+(t, \hat{\Omega}) + e^\times_{ab}(\hat{\Omega})h_\times(t, \hat{\Omega}), \quad (1)$$

where $\hat{\Omega}$ is the unit vector pointing from the GW source to the Solar System Barycenter (SSB) and h_+, $h_\times$ and e^A_{ab} ($A = +, \times$) are the polarization amplitudes and polarization tensors, respectively (see the Appendix for more details). The GW will cause a fractional shift in frequency ν that can be defined by a redshift in the times of arrival (TOAs):

$$\frac{\delta\nu(t, \hat{\Omega})}{\nu} = -e^A_{ab}(\hat{\Omega})\frac{1}{2}\frac{\hat{p}^a\hat{p}^b}{1 + \hat{\Omega}\cdot\hat{p}}\Delta h_A(t, \hat{\Omega}), \quad (2)$$

where

$$\Delta h_A(t, \hat{\Omega}) = h_A(t_e) - h_A(t_p). \quad (3)$$

Note that we use the standard Einstein summing convention over A. Here t_e and t_p denote the time at which the GW passes the Earth and pulsar, respectively, and $\hat{p}$ is the unit vector pointing from the SSB to the pulsar. Henceforth, we will drop the subscript "e" denoting the Earth time unless otherwise noted. From geometry we can write[6]

$$t_p = t - D(1 - \cos\mu), \quad (4)$$

where $\cos\mu = -\hat{\Omega}\cdot\hat{p}$ and D is the distance to the pulsar. To simplify our notation we introduce the pulsar "antenna pattern functions":

$$F^A(\hat{\Omega}) = -\frac{1}{2}\frac{\hat{p}^a\hat{p}^b}{1 + \hat{\Omega}\cdot\hat{p}}e^A_{ab}(\hat{\Omega}), \quad (5)$$

where the redshift is now written as

$$\frac{\delta\nu(t, \hat{\Omega})}{\nu} = F^A(\hat{\Omega})\Delta h_A(t, \hat{\Omega}). \quad (6)$$

For this work we will only consider circular, non-precessing, monochromatic SMBHB systems. Astrophysical justification for some of the above approximations can be found in Sesana & Vecchio (2010).[7] For this work we will operate in the regime where the frequency evolution of the system is slow enough that we can make the approximation that $f(t_p) = f(t_e) = $ const and the pulsar term will just contribute sinusoids at a pulsar-dependent phase. It should be noted that we use the observed redshifted values. For example, the chirp mass and frequency in the rest frame are $\mathcal{M}_r = \mathcal{M}/(1 + z)$ and $f_r = f_0(1 + z)$, respectively, where z is the cosmological redshift. Assuming a monochromatic system with a circular orbit and an orbital angular frequency of $\omega_{\mathrm{orb}} = \pi f_0$, where f_0 is the GW frequency, we can now write the polarization amplitudes as

$$h_+(t) = h[(1 + \cos^2\iota)\cos 2\phi_n\cos(2\pi f_0 t - \phi_0)$$
$$- 2\cos\iota\sin 2\phi_n\sin(2\pi f_0 t - \phi_0)] \quad (7)$$

$$h_\times(t) = h[(1 + \cos^2\iota)\sin 2\phi_n\cos(2\pi f_0 t - \phi_0)$$
$$+ 2\cos\iota\cos 2\phi_n\sin(2\pi f_0 t - \phi_0)], \quad (8)$$

where ϕ_0 is the orbital phase of the binary at $t = 0$, ι is the inclination angle, and ϕ_n is the angle to the line of nodes. We define the amplitude h,

$$h = 2\frac{\mathcal{M}^{5/3}(\pi f_0)^{2/3}}{D_c}. \quad (9)$$

The magnitude of the polarization amplitudes is proportional to the SMBHB chirp mass $\mathcal{M} = (M_1 M_2)^{3/5}/(M_1 + M_2)^{1/5}$, the comoving distance to the source D_c, and the GW frequency f_0 (twice the orbital frequency), and can be written as

$$h \sim 8 \times 10^{-15}\left(\frac{\mathcal{M}}{10^9 M_\odot}\right)^{5/3}\left(\frac{D_c}{100\,\mathrm{Mpc}}\right)^{-1}$$

$$\times \left(\frac{f_0}{5 \times 10^{-8}\,\mathrm{Hz}}\right)^{2/3}. \quad (10)$$

We now use the redshift given in Equation (6) to compute the GW induced pulsar timing residuals:

$$r(t, \hat{\Omega}) = \int_0^t \frac{\delta\nu(t, \hat{\Omega})}{\nu}dt = \frac{1}{2\pi f_0}F^A(\hat{\Omega})\Delta h_A(t, \hat{\Omega})$$

$$= \frac{1}{2\pi f_0}F^A(\hat{\Omega})[h_A(t_e) - h_A(t_e - D(1 + \hat{\Omega}\cdot\hat{p}))].$$

$$(11)$$

---

[6] Note that we use units in which $c = G = 1$.

[7] Recent work by Roedig & Sesana (2011) shows that many individual sources may have large eccentricities ($e > 0.1$) in the PTA band. However, for this work, we will only consider sources with circular orbits.

Here we have written out the explicit dependence on the pulsar distance and sky location. The first term in square brackets is the so-called Earth term because it refers to the metric perturbation at Earth and is correlated in all sets of pulsar timing residuals. The second term in the square brackets is the so-called pulsar term because it refers to the metric perturbation at the pulsar and is uncorrelated in all sets of pulsar timing residuals. Note that the pulsar term carries a dependence on the GW sky location $\hat{\Omega}$, and *all* of the dependence on the pulsar distance, D. As with the strain amplitude h, we can express the approximate amplitude of the GW induced timing residuals as a function of the SMBHB source parameters (Sesana et al. 2009),

$$r \sim 25.7\,\mathrm{ns} \left(\frac{\mathcal{M}}{10^9\,M_\odot}\right)^{5/3} \left(\frac{D_c}{100\,\mathrm{Mpc}}\right)^{-1}$$
$$\times \left(\frac{f_0}{5 \times 10^{-8}\,\mathrm{Hz}}\right)^{-1/3}. \tag{12}$$

The residuals can be written as a function of $7 + M$ parameters:

$$\boldsymbol{\lambda} = \{\theta, \phi, \zeta, f_0, \iota, \phi_n, \phi_0, \boldsymbol{D}\}, \tag{13}$$

where $\zeta = \mathcal{M}^{5/3} D_c^{-1}$ and $\boldsymbol{D}$ is a vector of the M pulsar distances. Here we have written ζ as a single parameter since it is not possible to disentangle the chirp mass $\mathcal{M}$ from the comoving distance D_c without a measurement of $\dot{f}$ (see Corbin & Cornish 2010 for more details). This measurement is only possible if the Earth and pulsar terms are at measurably different frequencies, which we do not consider in this work. Detecting and characterizing a signal that is a function of many parameters can be quite difficult and will be discussed in future papers. In this work we aim to give a baseline to the problem, and we will assume that all parameters are known and simply assess how well a particular detection method can confidently detect the signal.

2.2. Matched Filtering and Power Spectral Summing

Here we will outline matched filtering in terms of the log-likelihood, and a power spectral summing technique. First, we will review matched filtering basics in the context of a PTA. The problem of detecting a signal in noisy data is well studied (Wainstein & Zubakov 1962). We assume that the noise in each pulsar is additive, stationary, and Gaussian. For this case, the data for each set of pulsar timing residuals $x_\alpha(t_i) = x_{i\alpha}$ can be written as

$$x_{i\alpha} = r_{i\alpha} + n_{i\alpha}, \tag{14}$$

where $r_{i\alpha} = r_\alpha(t_i)$ and $n_{i\alpha} = n_\alpha(t_i)$ are the signal and the noise in each data set. Here i refers to the time index and α refers to the pulsar number. As is the method in matched filtering, we want to compare our data to a signal template of known form. Here we define $r_{i\alpha} = r_\alpha(t_i, \boldsymbol{\lambda})$ as our template of known form where $\boldsymbol{\lambda}$ is the vector of search parameters given in Equation (13).

We define the inner product of two functions of time $x(t_i)$ and $y(t_i)$ as

$$(x|y) = \sum_{i,j} x_i (C)_{ij}^{-1} y_j, \tag{15}$$

where C is the covariance matrix of the noise. With this framework in place we can specialize to the case of our PTA data and templates with white noise. We find the inner product

of our full set of residual data with the corresponding set of templates for that data set as

$$(x|r(\boldsymbol{\lambda})) = \sum_\alpha \sum_i \frac{x_{i\alpha} r_{i\alpha}}{\sigma_\alpha^2}, \tag{16}$$

where σ_α^2 is the rms of the residuals from the αth pulsar. A Wiener optimal statistic can be defined as

$$\rho(\boldsymbol{\lambda}) = \frac{(x|r(\boldsymbol{\lambda}))}{\sqrt{(r(\boldsymbol{\lambda})|r(\boldsymbol{\lambda}))}}. \tag{17}$$

If the noise is Gaussian and the signal is present, then the S/N is given as

$$\mathrm{S/N} = \langle \rho(\boldsymbol{\lambda}') \rangle = \sqrt{(r(\boldsymbol{\lambda}')|r(\boldsymbol{\lambda}'))}, \tag{18}$$

where $\boldsymbol{\lambda}'$ is the best estimate of the source parameters. For our data-analysis purposes we use the log-likelihood as our matched filtering statistic. Under the assumption of Gaussian noise, we define the likelihood function as the probability of the data $x_\alpha(t_i)$ given some set of model parameters $\boldsymbol{\lambda}$:

$$p(x|\boldsymbol{\lambda}) = C_{\mathrm{norm}} \exp\left[-\frac{1}{2}((x - r(\boldsymbol{\lambda}))|(x - r(\boldsymbol{\lambda})))\right], \tag{19}$$

where C_{norm} is a normalization constant. We now define the relative likelihood as $\Lambda(\boldsymbol{\lambda}) = p(x|\boldsymbol{\lambda})/p(x|0)$, where $p(x|0)$ is the probability of the data given the null hypothesis. From this we define the log-likelihood function:

$$\ln \Lambda(\boldsymbol{\lambda}) = \left((x|r(\boldsymbol{\lambda})) - \frac{1}{2}(r(\boldsymbol{\lambda})|r(\boldsymbol{\lambda}))\right). \tag{20}$$

By defining the log-likelihood in terms of the relative likelihood, we incorporate hypothesis testing (whether the GW signal is present or not) and parameter estimation into one statistic. We can define the S/N in terms of the log-likelihood as follows:

$$\langle \ln \Lambda(\boldsymbol{\lambda}') \rangle = \frac{1}{2} \langle \rho(\boldsymbol{\lambda}') \rangle^2 = \frac{1}{2} \mathrm{S/N}^2. \tag{21}$$

To determine if a signal is present and to determine the source parameters $\boldsymbol{\lambda}$, one would need to search parameter space to find the maximum value of the likelihood. This can be done through grid-based methods, Nested Sampling, or MCMC. For this work we are only concerned with detection of a source. In this case, to claim a detection, our statistic (log-likelihood) must be greater than a threshold value determined by a specified false alarm probability.

While the log-likelihood has the ability to simultaneously carry out detection, parameter estimation, and hypothesis testing, we now describe a method aimed at detection. In this method, we simply calculate the power spectrum of each set of pulsar timing residuals and then sum the power weighted by the variance of each data set and look for the maximum value over all frequency bins. We define our detection statistic as

$$\mathcal{P} = \max_f \sum_{\alpha=1}^M \frac{S_\alpha(f)}{\sigma_\alpha^2}, \tag{22}$$

where $S_\alpha(f)$ and σ_α^2 are the one-sided power spectrum and the variance of the αth pulsar data set, respectively. A detection is claimed when the value of $\mathcal{P}$ is greater than some threshold value $\mathcal{P}_0$ corresponding to a false alarm probability.

3. SIMULATED PTA DATA SETS

The simulated PTA data sets used for this work represent a best case scenario when it comes to *data quality*. While the quality of the data is optimal, the properties of the PTA (i.e., distances, sky location, rms, etc.) are meant to represent a realizable case. The array consists of up to 100 pulsars uniformly distributed in both the azimuthal angle ϕ and in the cosine of the polar angle $\cos\theta$. The pulsar distances are also drawn from a uniform random distribution ranging from 0.5 to 3 kpc. It should be noted that although there are known millisecond pulsars with distances less than 0.5 kpc, the simulated pulsar distances here are meant to represent an average for a typical PTA. The pulsar timing residuals are evenly spaced over 10 years with 250 TOAs for each pulsar simulating roughly bimonthly observations, and rms values drawn randomly from a uniform sample ranging from 100 to 300 ns. The noise is simulated to be white, Gaussian, additive, and stationary.

In real pulsar timing data, the residuals will be unevenly sampled and the noise may have various red components. Fortunately, recent work suggests that NANOGrav data appear to be mostly white with little to no red noise contributions (D. Perrodin et al. 2012, in preparation; J. Ellis et al. 2012, in preparation). In addition, the pulsar timing residuals will not be stationary, as a quadratic must be fit out of the data to account for the spin-down of the pulsar. Specifically, our definition of the inner product in Equation (16) no longer holds as we will need to include the covariance matrix of the data and incorporate a linear operator that takes into account this fitting. These are issues that will need to be addressed in order to make a fully functional data-analysis pipeline for continuous GW searches and will be addressed in future papers. However, here we will deal with the simple case to illustrate the efficacy of the studied search techniques on a data set of optimal quality.

4. ANALYSIS

In this section we compare four different detection techniques and determine their efficacy in terms of a minimum detectable amplitude as a function of the number of pulsars in the array. The four detection methods are the full matched filter, the Earth-term-only-matched filter, the pairwise matched filter, and a simple power spectral summing technique. For all of these methods we will assume that we know the parameters of the source exactly and will only be interested in the lowest amplitude that each method can detect. Though unrealistic in most astrophysical scenarios, this gives us a simple baseline for the comparison of the four methods.

4.1. Detection Methods

The full matched filter includes both the Earth and pulsar terms from Equation (11) and is thus a coherent search technique. This is the optimal detection statistic for a continuous wave buried in Gaussian noise. However, this method has the major drawback that it is computationally expensive to carry out in practice as the pulsar distances must be added as search parameters. For example, if we have an array of 20 pulsars and want to search over just 100 trial distances for each pulsar, then we will need to use at least 10^{40} templates. If one does not use grid-based methods, this number will drop significantly. However, even using more advanced methods like Nested Sampling or MCMC, including the pulsar distance as a search parameter is still computationally expensive. In the low S/N regime, where the likelihood surface is relatively flat and noisy, it may
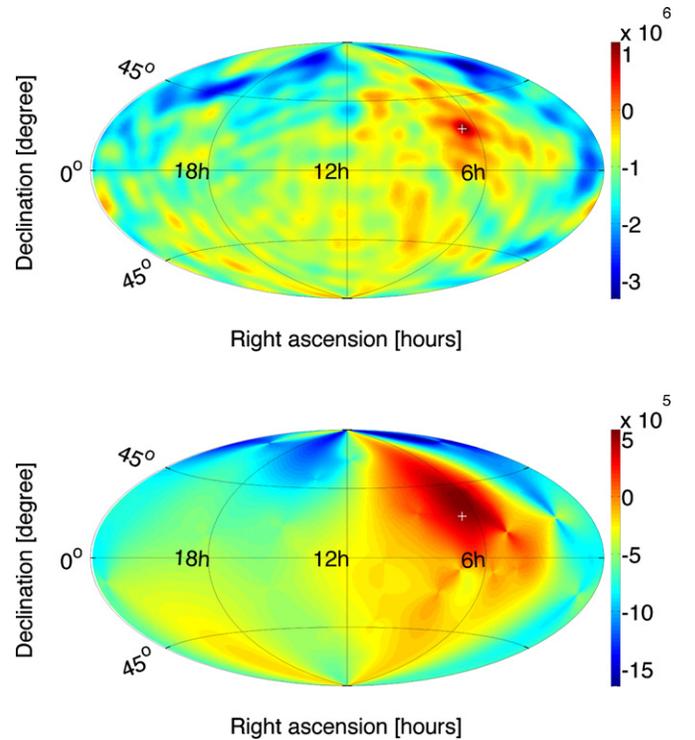

Figure 1. Skymaps created using the full matched filter (top) and the Earth-term-matched filter (bottom). The injected signal was very large (S/N = 1000) for illustration purposes. The "+" symbol indicates the injected sky location. We can see that the sky localization is biased and has a large uncertainty for the Earth-term-matched filter as compared to the full matched filter. In this particular case we observe a 46% loss in recovered S/N.

(A color version of this figure is available in the online journal.)

be impossible to get an accurate pulsar distance estimate with finite computational resources. The referee points out that it may be possible to analytically maximize over the unknown phase induced by the pulsar term. In effect this would allow us to conduct a search that includes the pulsar term without the extremely large overhead of searching over the pulsar distances. However, this would also increase our false alarm threshold, so the benefits of this method are unknown at this time. An exploration of this problem will be discussed in future work.

The Earth-term-matched filter uses templates that only depend on the coherent Earth term and treats the pulsar term as a noise source. The major advantage of using only the Earth term from Equation (11) for this detection method is that one does not need to include the pulsar distance in the search, making this method much less computationally expensive. Since we are not including the pulsar term in the analysis, we will always measure an S/N that is lower than the injected value. Also, there is a strong correlation between the pulsar distance and the sky location of the GW source. This will cause the recovered sky location to be biased (in the case where the Earth and pulsar term frequencies fall within the same bin) and have larger uncertainties than the full matched filter case. This is illustrated in Figure 1 in which a large GW (S/N = 1000) is injected into the data from 40 pulsars. Both the full matched filter and the Earth-term-matched filter search over sky location, and a skymap is created where the color scale indicates the log-likelihood. We can see that the full matched filter does a good job of localization whereas the Earth-term-matched filter is biased with a much larger error box in the sky. For this realization, using the Earth-term-matched filter also results in a 46% decrease in

S/N compared to the full matched filter. It is important to note that when we use an Earth-term-only-matched filter, we are essentially labeling the pulsar term as noise. The inner product defined in Equation (16) assumes that the noise is uncorrelated among pulsars. When the pulsar term is considered noise, this assumption no longer holds. It is unclear how this will affect the likelihood since the pulsar term is a deterministic signal and not realization dependent; however, since we are calculating *minimum* detectable amplitudes, the pulsar term will be very small compared to the intrinsic white noise. Although this is an important distinction, the effect will be small for our purposes and can be safely ignored.

The pairwise matched filter is a method that will allow us to take advantage of the pulsar term without the hindrance of an unworkable number of templates. This is done by constructing the full matched filter for each pair of pulsars, including the distance as a search term, and then adding the likelihoods as defined in Equation (20) in a pairwise fashion:

$$\ln \Lambda^{\mathrm{PW}} = \sum_{\alpha < \beta} (x|r(\boldsymbol{\lambda}))_{\alpha\beta} - \frac{1}{2}(r(\boldsymbol{\lambda})|r(\boldsymbol{\lambda}))_{\alpha\beta}, \qquad (23)$$

where the $\alpha\beta$ subscript denotes an inner product using two pulsars and the sum indicates a sum over all $M(M-1)/2$ unique pulsar pairs. While we still need the same number of templates for the intrinsic parameters of the source, we need far fewer distance templates. For example, in the case of M pulsars and 100 trial distances for each pulsar, the full matched filter requires at least 10^{2M} distance templates. However, for the pairwise matched filter we only perform a coherent search each pulsar pair, therefore for each pair we only require 10^4 distance templates making the total number of distance templates $M(M-1)/2 \times 10^4$, which is *significantly* less than the full matched filter. This method is still negatively affected by the strong degeneracy between pulsar distance and sky location but by using many pulsar pairs, this degeneracy will be greatly reduced.

The most simple and computationally inexpensive method that one could use to detect a continuous GW is power spectral summing as described in Equation (22). A method very similar to this was used in Yardley et al. (2010) on real data to produce sensitivity curves for the PPTA pulsars. This method is relatively robust in that it does not depend on any signal model templates. The disadvantages of this method are that it is incoherent because it does not keep track of phase information and it gives no indication of the true parameters of the source. With real data that is irregularly sampled and may contain red noise processes, going into the Fourier domain may pose problems that will be addressed in future work. However, for this analysis we are using this method as a baseline to compare the matched filtering statistics.

4.2. Efficacy of Detection Statistics

Here we will outline our Monte Carlo simulations and present our results. A good figure of merit for a detection statistic is the minimum detectable amplitude, that is, we wish to find the amplitude that can be detected above some threshold in 95% of the simulated realizations. To find the minimum detectable amplitude, we must first define a false alarm rate, that is, that rate at which we expect to make a detection when no signal is present. Assuming Gaussian statistics, we want a 4σ detection significance corresponding to a false alarm rate of 1/15,787. All four of our detection methods will have a different threshold

value corresponding to this false alarm rate. To calculate these thresholds we perform the following simulation. We choose a template at random for each realization and calculate the detection statistic for 15,787 realizations of noise and record the maximum value. Statistically, this means that we would expect to get a value for our statistic that is above this maximum value $\sim 0.016\%$ of the time if our data were pure Gaussian noise.

A Monte Carlo simulation was run to determine the minimum detectable amplitude as a function of the number of pulsars in the array for four detection statistics. The steps of the simulation are as follows. (1) Choose the number of pulsars in the array and simulate residual data with white noise. (2) Fix[8] the chirp mass $\mathcal{M}$, frequency f_0, and distance D_c to construct the amplitude h given in Equation (9). Then create a GW source in the sky with a given set of parameters $\boldsymbol{\lambda} = (\theta, \phi, \iota, \phi_n)$ drawn from random distributions. (3) Add the GW induced residuals into the simulated PTA data. (4) Run the detection statistic code assuming all parameters are known exactly[9] and output the log-likelihood. If the log-likelihood is above a given threshold value, count as a detection, otherwise, count as a non-detection. (5) Draw a new parameter vector $\boldsymbol{\lambda}$ and repeat steps (3) and (4). Repeat this for 10,000 GW source realizations (different realizations of $\boldsymbol{\lambda}$) and record the percent of sources detected. (6) Keep $\mathcal{M}$ and f_0 fixed and change D_c to obtain a new amplitude h, repeat step v until 95% of the realizations are detected. (7) Change the number of pulsars in the array and repeat the entire procedure. In practice, a bisection root finding method is used to determine the detection probability instead of linearly increasing h until the 95% level is reached.

4.2.1. Single Template Case

We ran the simulation described above for each of our detection methods for PTAs with 15–100 pulsars. In this case we assume that there are no search templates except for the one exactly matching the data. This means that we will obtain the lowest possible false alarm probability. The results are shown in Figure 2. It is obvious that that the unrealizable full matched filter (thin dashed line) can detect the lowest amplitudes of the four detection methods tested. However, if one uses the pairwise matched filter (thin dotted line), very little sensitivity is lost for PTAs with up to 30 pulsars. Also, by using only the Earth term in a matched filter search (thin solid line), the resulting minimum detectable amplitude is only a factor of two higher than the optimal method of filtering for the entire signal. It is also important to note that these two matched filtering methods (full and Earth term) scale roughly the same with the number of pulsars in the array, therefore, this factor of two is independent of the number of pulsars. The incoherent power spectral summing method obtains roughly the same sensitivity as the Earth-term-matched filter up to ~ 55 pulsars. It is known that the S/N scales as $\sqrt{M}$ for coherent methods and scales as $M^{1/4}$ for incoherent methods, where M is the number of "detectors." However, when using the Earth-term-matched filter we are ignoring the pulsar term which contributes roughly a factor of two to the S/N. Therefore, the Earth-term-matched

---

[8] For this work we fix the frequency to the lowest detectable frequency of $f_0 = 1/T_{\mathrm{obs}}$ and the chirp mass to a reasonable value of $\mathcal{M} = 5 \times 10^8 \, M_\odot$.
[9] This step is somewhat different for the Earth-term-matched filter because the largest S/N does not correspond to the case where the filter signal parameters are the same as the input signal parameters since we are not including the pulsar term. For this case we carry out a search over source sky location for each iteration in order to obtain the maximum possible log-likelihood.
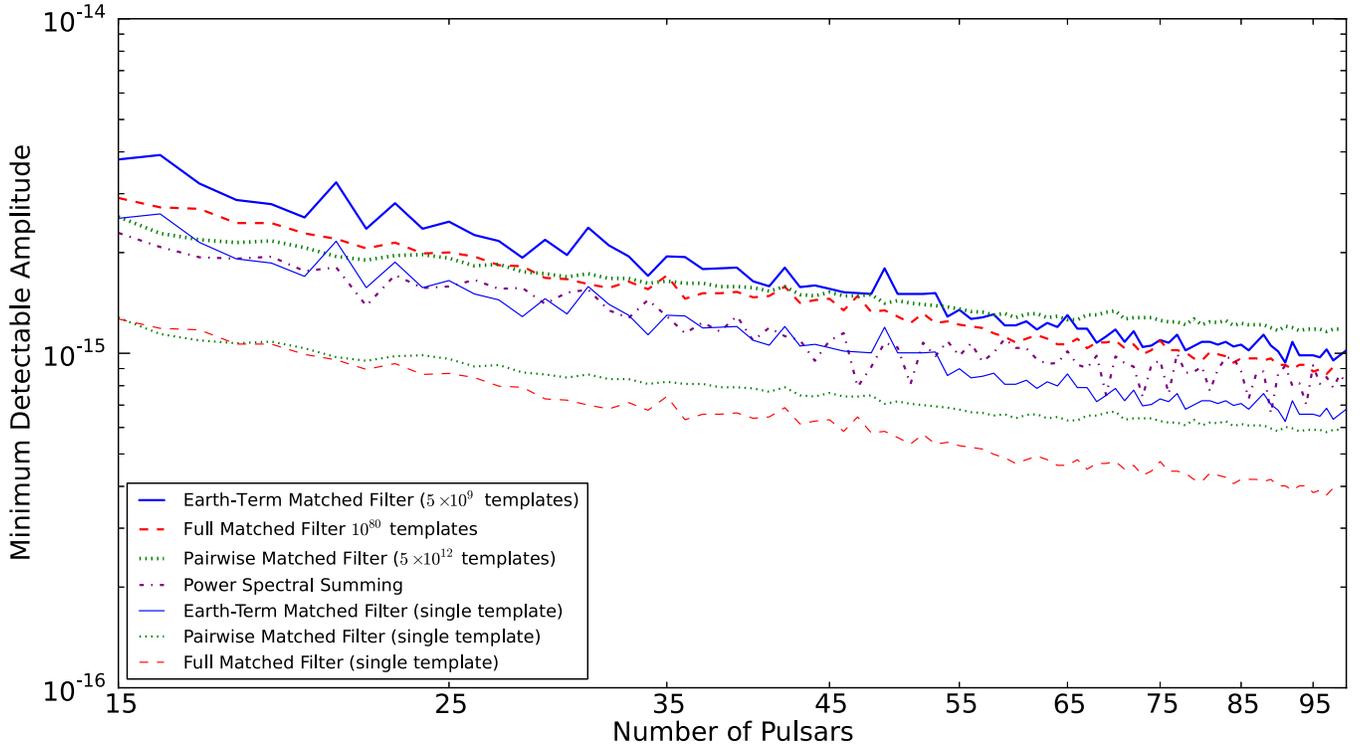
Figure 2. Plot of minimum detectable amplitude vs. number of pulsars for the full matched filter (thin dashed line), Earth-term-matched filter (thin solid line), power spectral summing (thick dash-dotted line), and the pairwise matched filter (thin dotted line). Also plotted are the full matched filter (thick dashed line), the Earth-term-matched filter (thick solid line), and the pairwise matched filter (thick dotted line) with a realistic number of search templates. To make these plots, a Monte Carlo simulation was run to find the amplitude at which 95% of GW source realizations were detected for a given number of pulsars.

(A color version of this figure is available in the online journal.)

filter scales as $\sqrt{M}/2$ whereas, the power spectral summing technique contains both the Earth and pulsar terms and scales as $M^{1/4}$. This results in roughly equal sensitivity as the ratio of the Earth-term-matched filter to the power spectral summing method scaling like $M^{1/4}/2$. Note that this is only the case when the Earth and pulsar terms are in the same frequency bin. It is likely that these results will change significantly when we are not in this regime. An exploration of this problem will be the subject of future work. The pairwise matched filter is also incoherent so the S/N has a flatter slope versus M than when using the other two matched filtering methods. For a real PTA, the S/N will not scale exactly as mentioned above because each set of pulsar timing residuals do not have the exact same characteristics (different noise, sky location, distance, etc.). As a sanity check, we ran the simulation on PTAs where every pulsar had the same sky location, distance, and rms residual. In this case the curves in Figure 2 scale exactly as mentioned above.

4.2.2. Multiple Template Case

The above section deals with the case of one template. In reality, the act of searching over many templates will serve to increase the false alarm probability. In the case of Gaussian noise the false alarm rate can be calculated analytically as (Maggiore 2007)

$$p_{\mathrm{FA}} = 2 \int_{\rho_0}^{\infty} d\rho \, e^{-\rho^2/2\sigma^2} = 2 \operatorname{erfc}(\rho_0/\sqrt{2}\sigma), \qquad (24)$$

where ρ_0 is a threshold value of the S/N, ρ is the output of the Wiener filter in Equation (17), σ is the standard deviation of the probability distribution function, and $\operatorname{erfc}(z)$ is the

complementary error function. Formally, this false alarm rate is derived based on a Wiener filter defined in Equation (17); however, since the expectation value of the log-likelihood is proportional the expectation value of the Wiener filter (see Equation (21)), this equation is still valid for the log-likelihood method that we use here. Because we deal with only white Gaussian noise in this paper, our simulated thresholds are in agreement with Equation (24). However, when calculating a false alarm probability for a search with N templates, the total false alarm probability is $P_{\mathrm{FA}} = Np_{\mathrm{FA}}$, when p_{FA} is much less than unity. This implies that if one wants to keep the same detection significance (less than 1/15,787 chance of occurring in noisy data), then the threshold value must be increased.

When this is done, we can see from Figure 2 that the minimum detectable amplitudes for the Earth-term-matched filter, pairwise matched filter, and full matched filter are nearly the same because of the significantly smaller number of templates for the Earth term and pairwise matched filters. Since the power spectral addition method is not a template-based search, its false alarm threshold remains the same. Therefore, for this simple analysis it appears that the power spectral summing technique is the most sensitive when considering a monochromatic wave in a large dimensional parameter space. However, recent work by Babak & Sesana (2012) shows that one can maximize the likelihood over many of the parameters while retaining all of the S/N, resulting in just a three-dimensional search space. This would result in a dramatic decrease in the number of templates required in the Earth-term-matched filter search, thereby increasing its sensitivity. This prospect is actively being explored and will be the subject of future work. It is also important to note the values of h on the y-axis. Since we are dealing with

the best case scenario in terms of data quality (white Gaussian noise, evenly spaced data, no timing model subtraction), this plot shows us that, at current levels of timing precision (rms $\sim$100 ns), we could never confidently detect any signal with an amplitude below $h = 9 \times 10^{-16}$ even with a PTA of 100 pulsars. At 20 pulsars, we could only possibly detect a source with $h > 2 \times 10^{-15}$.

5. SUMMARY

In this work we have tested the efficacy of four detection techniques on simulated data sets when searching for continuous GW signals from SMBHBs. We have shown that a matched filter using only the correlated Earth term in the search templates results in a minimum detectable amplitude that is two times higher than the optimal matched filter using both the correlated and uncorrelated terms. We have also shown that when using a pairwise matched filter it is possible, in principle, to obtain nearly the same sensitivity as the full matched filter. However, when performing a real search, the strong correlations between sky location of the GW source and pulsar distance may cause problems with S/N recovery, a problem that does not affect the Earth-term-matched filter. We have also shown that an incoherent power spectra summing method results in a minimum detectable amplitude that is quite comparable to that of the Earth-term-matched filter. When the approximate number of search templates is taken into account, we find that using an Earth-term-matched filter is nearly as sensitive as the full matched filter, but the power spectral summing technique results in the lowest minimum detectable amplitude because it is independent of the number of search templates used in a matched filter search. Moreover, this work gives an idea of the prospects of detecting a continuous GW with PTAs, by placing lower limits on the detectable amplitude for data of optimal quality. The advantages and disadvantages of the various detection methods have been discussed and it has been shown that using a full matched filter with the pulsar distances explicitly included as search parameters is very computationally expensive (maybe even impossible for some cases). Because of this and the relatively low computational cost, an Earth-term-matched filter or power spectral summing technique appear to be practical choices for a detection method in a data-analysis pipeline for use on real pulsar timing data.

This work gives some insight into what detection techniques should be used in a fully functional pipeline. More sophisticated data-analysis methods will have to include the effects of irregularly sampled data, red noise (both correlated noise in the form of the stochastic GW background and uncorrelated noise in the form of intrinsic timing noise and interstellar medium effects), and timing model parameter fits. The methods described here give basic detection algorithms that can be modified for use with real data. After a detection is made, the next step is parameter estimation. This will require fast, efficient algorithms to find the correct parameters in a large parameter space. Both of these challenges are currently being studied and will be the subject of future papers.

APPENDIX

POLARIZATION TENSOR IN THE SSB REFERENCE FRAME

Here we will show how one can convert the polarization tensors into the SSB coordinates. Once again, a GW is defined as a metric perturbation to flat space time,

$$h_{ab}(t, \hat{\Omega}) = e_{ab}^{+}(\hat{\Omega})h_{+}(t, \hat{\Omega}) + e_{ab}^{\times}(\hat{\Omega})h_{\times}(t, \hat{\Omega}), \quad (A1)$$

where $\hat{\Omega}$ is the unit vector pointing from the GW source to the SSB and h_{+}, $h_{\times}$ and e_{ab}^{A} ($A = +, \times$) are the polarization amplitudes and polarization tensors, respectively. The polarization tensors can be converted to the SSB by the following transformation. Following Wahlquist (1987) we write the polarization tensors in terms of the wave principal axes described by unit vectors $\hat{m}$ and $\hat{n}$:

$$e_{ab}^{+}(\hat{\Omega}) = \hat{m}_a\hat{m}_b - \hat{n}_a\hat{n}_b, \quad (A2)$$

$$e_{ab}^{\times}(\hat{\Omega}) = \hat{m}_a\hat{n}_b + \hat{n}_a\hat{m}_b. \quad (A3)$$

In the SSB coordinate center we define

$$\hat{\Omega} = -(\sin\theta\cos\phi)\hat{x} - (\sin\theta\sin\phi)\hat{y} - (\cos\theta)\hat{z}, \quad (A4)$$

$$\hat{m} = -(\sin\phi)\hat{x} + (\cos\phi)\hat{y}, \quad (A5)$$

$$\hat{n} = -(\cos\theta\cos\phi)\hat{x} - (\cos\theta\sin\phi)\hat{y} + (\sin\theta)\hat{z}. \quad (A6)$$

In this coordinate system, $\theta = \pi/2 - \delta$ and $\phi = \alpha$ are the polar and azimuthal angles of the source, respectively, where δ and α are the declination and right ascension in usual celestial coordinates. In this coordinate system, the polarization tensors can be written as

$$e^{+}(\theta, \phi) = \begin{pmatrix} \sin^2\phi - \cos^2\phi\cos^2\theta & -\cos^2\phi(1+\cos^2\theta)\sin\phi & \cos\phi\cos\theta\sin\theta \\ -\cos\phi(1+\cos^2\theta)\sin\phi & \cos^2\phi - \cos^2\theta\sin^2\phi & \cos\theta\sin\theta\sin\phi \\ \cos\phi\cos\theta\sin\theta & \cos\theta\sin\phi\sin\theta & -\sin^2\theta \end{pmatrix} \quad (A7)$$

$$e^{\times}(\theta, \phi) = \begin{pmatrix} \cos\theta\sin 2\phi & -\cos 2\phi\cos\theta & -\sin\phi\sin\theta \\ -\cos 2\phi\cos\theta & -2\cos\phi\cos\theta\sin\phi & \cos\phi\sin\theta \\ -\sin\phi\sin\theta & \cos\phi\sin\theta & 0 \end{pmatrix}. \quad (A8)$$

REFERENCES

Anholm, M., Ballmer, S., Creighton, J. D. E., Price, L. R., & Siemens, X. 2009, Phys. Rev. D, 79, 084030
Babak, S., & Sesana, A. 2012, Phys. Rev. D, 85, 044034
Corbin, V., & Cornish, N. J. 2010, arXiv:1008.1782
Damour, T., & Vilenkin, A. 2001, Phys. Rev. D, 64, 064008
Demorest, P. B., Ferdman, R. D., Gonzalez, M. E., et al. 2012, arXiv:1201.6641
Finn, L. S., & Lommen, A. N. 2010, ApJ, 718, 1400
Foster, R. S., & Backer, D. C. 1990, ApJ, 361, 300
Hobbs, G., Archibald, A., Arzoumanian, Z., et al. 2010, Class. Quantum Grav., 27, 084013
Janssen, G. H., Stappers, B. W., Kramer, M., et al. 2008, in AIP Conf. Ser. 983, 40 Years of Pulsars: Millisecond Pulsars, Magnetars and More, ed. C. Bassa, Z. Wang, A. Cumming, & V. M. Kaspi (Melville, NY: AIP), 633
Jenet, F., Finn, L. S., Lazio, J., et al. 2009, arXiv:0909.1058

Jenet, F. A., Hobbs, G. B., Lee, K. J., & Manchester, R. N. 2005, ApJ, 625, L123
Jenet, F. A., Lommen, A., Larson, S. L., & Wen, L. 2004, ApJ, 606, 799
Lazio, J. 2009, in Panoramic Radio Astronomy: Wide-field 1-2 GHz Research on Galaxy Evolution, ed. G. Heald & P. Serra, http://pos.sissa.it/cgi-bin/reader/conf.cgi?confid=89
Leblond, L., Shlaer, B., & Siemens, X. 2009, Phys. Rev. D, 79, 123519
Lee, K. J., Wex, N., Kramer, M., et al. 2011, MNRAS, 414, 3251
Maggiore, M. 2007, Gravitational Waves: Theory and experiments, Gravitational Waves (Oxford: Oxford Univ. Press)
Manchester, R. N. 2008, in AIP Conf. Ser. 983, 40 Years of Pulsars: Millisecond Pulsars, Magnetars and More, ed. C. Bassa, Z. Wang, A. Cumming, & V. M. Kaspi (Melville, NY: AIP), 584
Roedig, C., & Sesana, A. 2011, arXiv:1111.3742

Romani, R. W. 1989, in Timing Neutron Stars, ed. H. Ögelman & E. P. J. van den Heuvel (New York: Kluwer), 113
Sesana, A., & Vecchio, A. 2010, Phys. Rev. D, 81, 104008
Sesana, A., Vecchio, A., & Volonteri, M. 2009, MNRAS, 394, 2255
Siemens, X., Mandic, V., & Creighton, J. 2007, Phys. Rev. Lett., 98, 111101
Vallisneri, M. 2008, Phys. Rev. D, 77, 042001
van Haasteren, R., Levin, Y., McDonald, P., & Lu, T. 2009, MNRAS, 395, 1005
Wahlquist, H. 1987, Gen. Rel. Grav., 19, 1101
Wainstein, L. A., & Zubakov, V. D. 1962, in Extraction of Signals from Noise (Englewood Cliffs, NJ: Prentice-Hall), 86
Wyithe, J. S. B., & Loeb, A. 2003, ApJ, 590, 691
Yardley, D. R. B., Coles, W. A., Hobbs, G. B., et al. 2011, MNRAS, 414, 1777
Yardley, D. R. B., Hobbs, G. B., Jenet, F. A., et al. 2010, MNRAS, 407, 669