Metropolis-Hastings algorithm for extracting periodic gravitational wave signals from laser interferometric detector data

Nelson Christensen,[1,*] Réjean J. Dupuis,[2,†] Graham Woan,[2,‡] and Renate Meyer[3,§]
[1]*Physics and Astronomy, Carleton College, Northfield, Minnesota 55057, USA*
[2]*Department of Physics and Astronomy, University of Glasgow, G12 8QQ, United Kingdom*
[3]*Department of Statistics, University of Auckland, Auckland, New Zealand*
(Received 8 January 2004; published 1 July 2004)

The Markov chain Monte Carlo methods offer practical procedures for detecting signals characterized by a large number of parameters and under conditions of low signal-to-noise ratio. We present a Metropolis-Hastings algorithm capable of inferring the spin and orientation parameters of a neutron star from its periodic gravitational wave signature seen by laser interferometric detectors.

DOI: 10.1103/PhysRevD.70.022001 PACS number(s): 04.80.Nn

I. INTRODUCTION

The worldwide network of laser interferometric gravitational wave detectors has begun to acquire scientifically significant data [1–4] and rapidly rotating neutron stars are an important potential source of signals (we will reserve the term ''pulsar'' to refer to the observed pulsating radio sources). Although a spinning spherically symmetric neutron star will not produce gravitational waves, a number of mechanisms have been proposed that are capable of producing quasi-periodic gravitational waves from biaxial or triaxial neutron stars [5,6]. Any gravitational waves from these neutron stars will likely be seen at Earth as weak continuous wave signals.

The data analysis task of identifying such a signal in the output of a laser interferometer is challenging and difficult, both because of the weakness of the signal and because its time evolution is characterized by a relatively large number of parameters. Radio observations can provide the sky location, rotation frequency and spindown rate of known pulsars, but the problem of looking for unknown (or poorly parameterized) neutron star sources is significantly more challenging. SN1987A is a good example of a poorly parameterized source for which the sky location in approximately known but also for which there is a large uncertainty in the frequency and spindown parameters of the putative neutron star [7].

Much work has already gone into all-sky hierarchical methods for searching for continuous gravitational waves [8,9]. Here we address the specific problem of a ''fuzzy'' parameter space search, in which a restricted volume of the space needs to be thoroughly investigated. We take a Bayesian approach to this problem and use Markov chain Monte Carlo (MCMC) techniques which have been shown to be especially suited to similar problems involving numerous parameters [10]. In particular, the Metropolis-Hastings (MH) algorithm [11,12] has been used for estimating cosmological parameters from cosmic microwave background data [13–15], and the applicability of the MH routine has been demonstrated in estimating astrophysical parameters for gravitational wave signals from coalescing compact binary systems [16,17]. MCMC methods have also provided Bayesian inference for noisy and chaotic data [18,19].

Here we demonstrate that a MH algorithm also offers great promise for estimating neutron star parameters from their continuous gravitational wave signals. This work builds on the development (by two of us) of an end-to-end robust Bayesian method of searching for periodic signals in gravitational wave interferometer data [20], summarized in Sec. II. Starting with this Bayesian approach we apply a similar MH routine to that used in [13,17]. The description of the Bayesian MH method is given in Sec. III. In Sec. IV we present the results of this study, using synthesized data, for four and five parameter problems. We believe that this method offers great hope for signal extraction as more parameters are included, and this point is discussed in Sec. V.

II. SIGNAL CHARACTERISTICS

We will initially consider searching for signals from known radio pulsars, and then expand the method to account for an uncertainty in the frequency of the gravitational wave signal. As gravitational waves from pulsars are certainly weak at Earth, long integration periods are required to extract the signal, and we must take account of the antenna patterns of the detectors and the Doppler shift due to the motion of the Earth.

As in the previous study [20,21] we consider the signal expected from a non-precessing triaxial neutron star. The gravitational wave signal from such an object is at twice its rotation frequency, $f_s = 2f_r$, and we characterize the amplitudes of each polarization with overall strain factor, h_0. The measured gravitational wave signal will also depend on the polarization antenna patterns of the detector $F_{\times,+}$ giving a signal

$$s(t) = \frac{1}{2}F_+(t;\psi)h_0(1+\cos^2\iota)\cos\Psi(t) + F_\times(t;\psi)h_0\cos\iota\sin\Psi(t), \quad (1)$$

*Electronic address: nchriste@carleton.edu
†Electronic address: rejean@astro.gla.ac.uk
‡Electronic address: graham@astro.gla.ac.uk
§Electronic address: meyer@stat.auckland.ac.nz

where ψ is the polarization angle of the radiation (which depends on the position angle of the spin axis in the plane of the sky) and ι is the inclination of the pulsar with respect to the line of sight.

Using a simple slowdown model, the phase evolution of the signal can be usefully parameterized as

$$\Psi(t)=\phi_0+2\pi\left[f_s(T-T_0)+\frac{1}{2}\dot{f}_s(T-T_0)^2+\frac{1}{6}\ddot{f}_s(T-T_0)^3\right], \tag{2}$$

where

$$T=t+\delta t=t+\frac{\mathbf{r}\cdot\mathbf{n}}{c}+\Delta_T. \tag{3}$$

Here, T is the time of arrival of the signal at the solar system barycenter, ϕ_0 is the phase of the signal at a fiducial time T_0, $\mathbf{r}$ is the position of the detector with regard to the solar system barycenter, $\mathbf{n}$ is a unit vector in the direction of the pulsar, c is the speed of light, and Δ_T contains the relativistic corrections to the arrival time [22].

The signal is *heterodyned* by multiplying the data by $\exp[-i\Psi(t)]$ so that the only time varying quantity remaining is the antenna pattern of the interferometer (which varies over the day). For convenience, the result is low-pass filtered and resampled. We are left with a simple model with four unknown parameters: the overall amplitude of the gravitational wave signal (h_0), its polarization angle (ψ), its phase at time T_0 (ϕ_0), and the angle between the spin axis of the pulsar and the line of sight (ι).

A detailed description of the heterodyning procedure is presented elsewhere [20,21]; here we just provide a summary of this standard technique. The raw signal, $s(t)$, is centered near twice the rotation frequency of the pulsar, but is Doppler modulated due to the motion of the Earth and the orbit of the pulsar if it is in a binary system. The modulation bandwidth is typically 10^4 times less than the detector bandwidth, so one can greatly reduce the effective data rate by extracting this band and shifting it to zero frequency. In its standard form the result is one binned data point, B_k, every minute, containing all the relevant information from the original time series but at only 2×10^{-6} the original data rate. If the phase evolution has been correctly accounted for at this heterodyning stage then the only time-varying component left in the signal will be the effect of the antenna pattern of the interferometer, as its geometry with respect to the neutron star varies with Earth rotation. Any small error, Δf, in the heterodyne frequency will cause the signal to oscillate at Δf, and for the second part of our study we have Δf as our fifth parameter. For both these studies we estimate the noise variance, σ_k^2, in the bin values, B_k, from the sample variance of the contributing data. It is assumed that the noise is stationary over the 60 s of data contributing to each bin.

III. THE METROPOLIS-HASTINGS ALGORITHM

This section presents a brief review of the Bayesian MH approach to parameter estimation. Comprehensive descriptions of MCMC methods and the MH algorithm can be found elsewhere [10,13,17].

We will denote the output from the above heterodyning procedure as $\{B_k\}$, with a joint probability distribution function (PDF) denoted by $p(\{B_k\}|\boldsymbol{a})$ conditional on unobserved parameters $\boldsymbol{a}=(a_1,\ldots,a_d)$. The PDF $p(\{B_k\}|\boldsymbol{a})$ is referred to as the *likelihood* and regarded as a function of the parameters $\boldsymbol{a}$. The parameters of interest for our four parameter study are $\boldsymbol{a}=(h_0,\psi,\phi_0,\iota)$, while for the five parameter study they are $\boldsymbol{a}=(h_0,\psi,\phi_0,\iota,\Delta f)$.

From Eq. (1), the (now complex) heterodyned signal is

$$y(t_k;\boldsymbol{a})=\frac{1}{4}F_+(t_k;\psi)h_0(1+\cos^2\iota)e^{i\phi_0} - \frac{i}{2}F_\times(t_k;\psi)h_0\cos\iota e^{i\phi_0}, \tag{4}$$

and the binning procedure should, by the central limit theorem, give the noise a near-Gaussian probability density characterized by a variance σ_k^2 for the kth bin. The likelihood that the data in this bin, taken at time t_k, is consistent with the above model is

$$p(B_k|\boldsymbol{a})\propto\exp\left(\frac{-|B_k-y(t_k;\boldsymbol{a})|^2}{2\sigma_k^2}\right), \tag{5}$$

and the joint likelihood that the data in all the bins (taken as independent) are consistent with a particular set of model parameters is

$$p(\{B_k\}|\boldsymbol{a})\propto\prod_k\exp\left(\frac{-|B_k-y(t_k;\boldsymbol{a})|^2}{2\sigma_k^2}\right). \tag{6}$$

Bayesian inference requires the specification of a prior PDF for $\boldsymbol{a}$, $p(\boldsymbol{a})$, that quantifies the researcher's pre-experimental knowledge about $\boldsymbol{a}$. The phase and polarization priors are flat in their space, and are set uniform for ϕ_0 over $[0,\pi]$, and for ψ over $[-\pi/4,\pi/4]$. The prior for ι is uniform in $\cos\iota$ over $[-1,1]$, corresponding to a uniform prior per unit solid angle of pulsar orientation. Finally, in the present study we take a prior for h_0 that is uniform for $0<h_0<1000$ (in our normalized units for which $\sigma_k=1$), and zero for all other values.

Using Bayes' theorem, the post-experimental knowledge of $\boldsymbol{a}$ is expressed by the *posterior* PDF of $\boldsymbol{a}$:

$$p(\boldsymbol{a}|\{B_k\})=\frac{p(\boldsymbol{a})p(\{B_k\}|\boldsymbol{a})}{p(\{B_k\})}\propto p(\boldsymbol{a})p(\{B_k\}|\boldsymbol{a}), \tag{7}$$

where $p(\{B_k\})=\int p(\{B_k\}|\boldsymbol{a})p(\boldsymbol{a})d\boldsymbol{a}$ is the marginal PDF of $\{B_k\}$ which can be regarded as a normalizing constant as it is independent of $\boldsymbol{a}$. The posterior PDF is thus proportional to the product of prior and likelihood.

The marginal posterior distribution for parameter a_i is the integral of the joint posterior PDF over all other components of $\boldsymbol{a}$ other than a_i, i.e.,

$$p(a_i|\{B_k\}) = \int \ldots \int p(\boldsymbol{a}|\{B_k\}) \times da_1 \ldots da_{i-1} da_{i+1} \ldots da_d, \quad (8)$$

and contains all the analysis has to say about the value of a_i alone. However it is often useful to summarize this in a single ''point estimate'' of a_i using, for example, the posterior mean:

$$\langle a_i \rangle = \int a_i p(a_i|\{B_k\}) da_i. \quad (9)$$

Calculating the normalization constant $p(\{B_k\})$ and calculating each marginal posterior PDF requires difficult d- and $(d-1)$-dimensional integrations, respectively, that can be evaluated using a sampling approach and MCMC methods [10,13,17]. Rather than sampling directly from $p(\boldsymbol{a}|\{B_k\})$, a sample from a Markov chain is generated which has $p(\boldsymbol{a}|\{B_k\})$ as its equilibrium distribution. Thus, after running the Markov chain for a certain ''burn-in'' period, these (correlated) samples can be regarded as samples from the limiting distribution, provided that the Markov chain has reached convergence. Despite their correlations, the ergodic theorem guarantees that the sample average is still a consistent estimate of the posterior mean Eq. (9) [23].

The specific MCMC technique used for this study was the MH algorithm [11,12]. The MH algorithm generates a sample from the target PDF $p(\boldsymbol{a}|\{B_k\})$ using a technique that is similar to the well-known simulation technique of *rejection sampling*. A candidate is generated from an auxiliary PDF and then accepted or rejected with some probability. Starting with an arbitrary initial state $\boldsymbol{a}_0$, at time n a new candidate $\boldsymbol{a}'$ is generated from the candidate generating PDF, $q(\boldsymbol{a}|\boldsymbol{a}_n)$, which can depend on the current state $\boldsymbol{a}_n$ of the Markov chain. This new candidate $\boldsymbol{a}'$ is accepted with a certain *acceptance probability* $\alpha(\boldsymbol{a}'|\boldsymbol{a}_n)$, also depending on the current state $\boldsymbol{a}_n$, given by

$$\alpha(\boldsymbol{a}'|\boldsymbol{a}_n) = \min\left\{\frac{p(\boldsymbol{a}')p(\{B_k\}|\boldsymbol{a}')q(\boldsymbol{a}_n|\boldsymbol{a}')}{p(\boldsymbol{a}_n)p(\{B_k\}|\boldsymbol{a}_n)q(\boldsymbol{a}'|\boldsymbol{a}_n)}, 1\right\}. \quad (10)$$

For good efficiency a multivariate normal distribution centered at the current state $\boldsymbol{a}_n$ is used for $q(\boldsymbol{a}'|\boldsymbol{a}_n)$. This then implies that if the posterior probability at $\boldsymbol{a}'$ is larger than at the current state $\boldsymbol{a}_n$, the proposed step to $\boldsymbol{a}'$ is always accepted. However, if the step is in a direction of lower posterior probability, then this step is accepted only with a certain probability given by the ratio of the posterior PDFs (since our multivariate normal generating function is symmetric in $\boldsymbol{a}'$ and $\boldsymbol{a}_n$ and therefore cancels out). If the candidate is accepted, the next state of the Markov chain is $\boldsymbol{a}_{n+1} = \boldsymbol{a}'$, otherwise the chain does not move, i.e. $\boldsymbol{a}_{n+1} = \boldsymbol{a}_n$.

The steps of the MH algorithm are therefore:

Step 0: Start with an arbitrary value $\boldsymbol{a}_0$;

Step $n+1$: Generate $\boldsymbol{a}'$ from $q(\boldsymbol{a}|\boldsymbol{a}_n)$ and u from $U(0,1)$;

If $u \leq \alpha(\boldsymbol{a}'|\boldsymbol{a}_n)$ set $\boldsymbol{a}_{n+1} = \boldsymbol{a}'$ (acceptance),

If $u > \alpha(\boldsymbol{a}'|\boldsymbol{a}_n)$ set $\boldsymbol{a}_{n+1} = \boldsymbol{a}_n$ (rejection).

$U(0,1)$ is the uniform distribution between 0 and 1. The efficiency of the MH algorithm depends heavily on the choice of the proposal density. The closer the proposal is to the target distribution, the faster convergence will be accomplished. This link between the closeness of the proposal to stationary distribution and speed of convergence has also been substantiated by Holden [24]. In the study presented here we dynamically altered the proposal distribution based on information from the chain's history. The approach, called *pilot adaptation*, is to perform a separate pilot run to gain insight about the target density and then tune the proposal accordingly for the successive runs. Such adaptation can be iterated but allowing it infinitely often will destroy the Markovian property of the chain and thereby often compromise the stationarity of the chain and the consistency of sample path averages ([25]; see [26] for an example).

Based on the central limit theorem, the posterior PDF should be well approximated by a multivariate normal distribution with mean equal to the posterior mode and covariance matrix equal to minus the Hessian evaluated at the posterior mode. Thus, we use a multivariate normal distribution for the proposal density $q(\boldsymbol{a}|\boldsymbol{a}_n)$. As the mode is unknown, we try to make use of pilot samples to estimate its covariance matrix. When we initially run the MH algorithm, we sample candidate parameters from a normal distribution with covariance matrix equal to the identity matrix and centered around the current state. After the completion of this pilot run we use the empirical covariance matrix of the sample as covariance matrix of the multivariate normal proposal density, again with mean equal to the current state.

IV. RESULTS

In the first part of our study we reproduced the results presented in [20] where the four unknown parameters were h_0, ι, ψ, and ϕ_0. The signal $s(t)$ was synthesized assuming a source at right ascension $= 4^{\mathrm{h}}41^{\mathrm{m}}54^{\mathrm{s}}$ and declination $= 18°23'32''$, as would be seen by the LIGO-Livingston interferometer. This was then added to white Gaussian noise, $n(t)$, which is a good approximation to the detector noise in our band. Our normalized data had a noise variance of $\sigma_k^2 = 1$ for each sample, and the amplitude of the signal used in our test runs was varied in the range $h_0 = 0.0$ to 10.0. We were able to detect signals for $h_0 > 0.1$. The length of the data set corresponded to 14 400 samples or 10 days of data at a rate of one sample per minute (which was the rate used for the LIGO/GEO S1 analysis described in [21]). Although we will work with strains normalized to $\sigma_k = 1$, the results can be cast into a more conventional form by multiplying σ_i and h_0 by $(S_h/60)^{1/2}$, where $(S_h)^{1/2}$ is the strain noise spectral density of the detector at the frequency of interest, in Hz^{-1}.

An example of the MH routine output is shown in Fig. 1. Displayed are the trace plots and the kernel densities (posterior PDFs). For this example the program ran for 10^6 iterations. The first 10^5 iterations were discarded as the burn-in. Short-term correlations in the chain were eliminated by

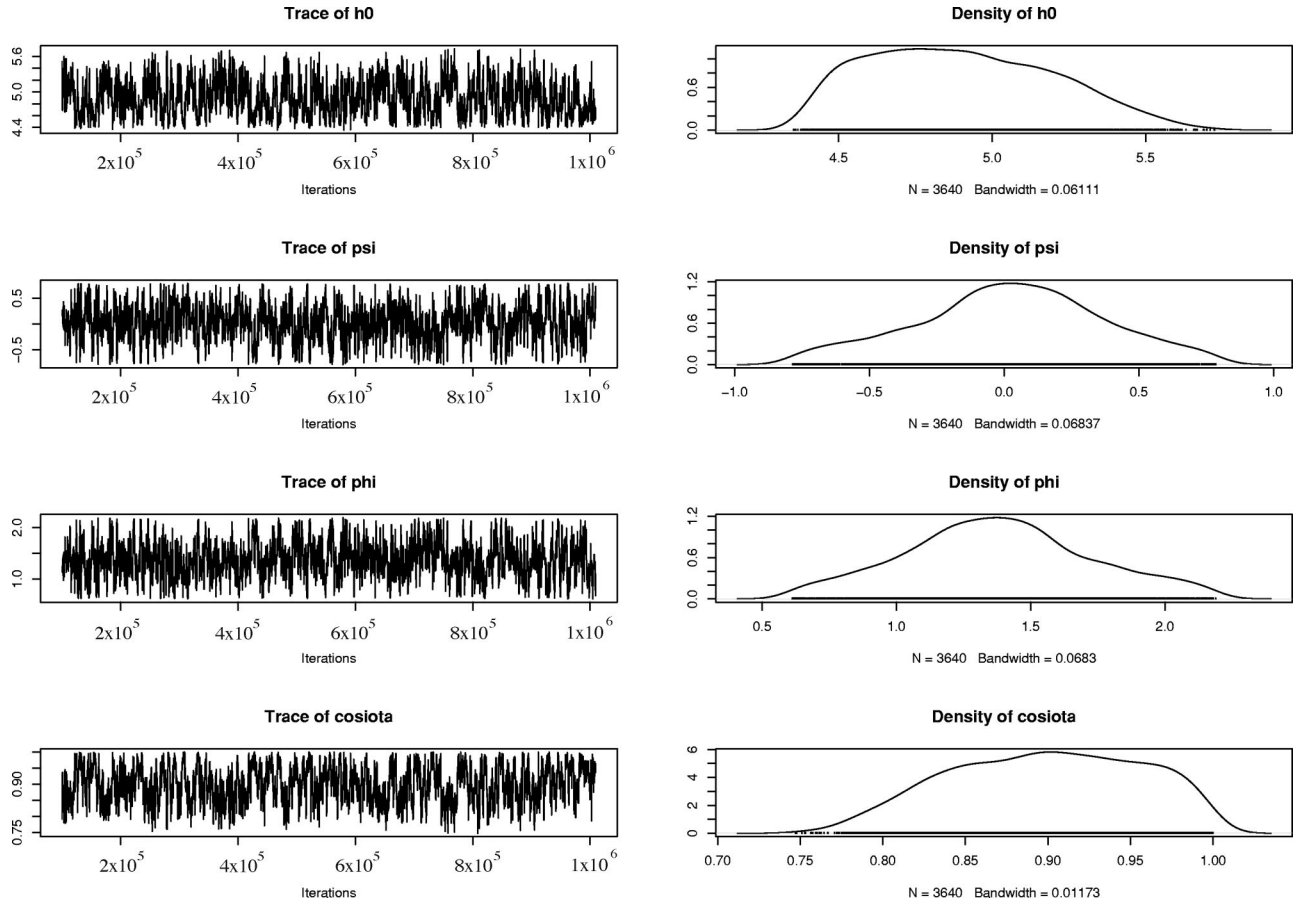

FIG. 1. Trace plot (left) and MCMC-estimated posterior PDFs (right) for the pulsar parameters h_0, ψ, ϕ_0 and $\cos\iota$. In this example the true parameters were $h_0=5.0$, $\phi_0=1.0$, $\psi=0.4$, and $\iota=0.5$ implying $\cos\iota=0.88$.

''thinning'' the remaining terms; we kept every 250th item in the chain. The true parameter values for this run were $h_0=5.0$, $\psi=0.4$, $\phi_0=1.0$ and $\iota=0.5$ ($\cos\iota=0.88$). In the example displayed in Fig. 1 the MCMC yielded mean values and 95% posterior probability intervals of $h_0=4.9$ (4.43 to 5.50), $\psi=0.02$ (-0.68 to 0.69), $\phi_0=1.34$ (0.71 to 2.08), and $\cos\iota=0.90$ (0.79 to 0.99). The 95% posterior probability interval is specified by the 2.5% and 97.5% percentiles of $p(a_i|\{B_k\})$. In Fig. 2 we display the estimated posterior PDF of h_0 on an expanded scale, along with the real and estimated value for h_0.

It is crucial that our algorithm is sensitive to the true value of the gravitational wave amplitude, h_0, even under conditions of relatively low signal-to-noise ratio, and Fig. 3 shows injected h_0 values versus their values inferred by the MH routine. The error bars correspond to the 95% posterior probability interval, i.e. the lower and upper bound are specified by the 2.5% and 97.5% percentiles of $p(a_i|\{B_k\})$. The algorithm clearly is successful in finding and estimating h_0. While the error bars increase as the signal gets larger, the relative error $\Delta h_0/h_0$ does diminish as h_0 increases. The fact that the 95% posterior probability interval grows with h_0 for constant noise level would seem to be counterintuitive. In addition, the widths of the posterior probability distributions for h_0 are larger than would be naively expected from a search for a simple periodic signal. The reason is that these error bars represent the uncertainty in the parameter rather than just the level of the noise, and this is affected both by the noise level and the posterior covariance between all of the parameters. The MCMC technique also allows one to calculate cross-correlation coefficients from the Markov chains of the parameters, and the value between h_0 and $\cos\iota$

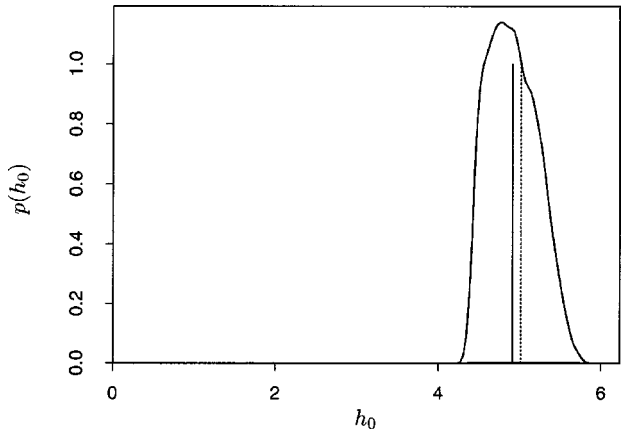

FIG. 2. An expanded view of the estimated posterior PDF based on the MCMC sample for parameter h_0. The vertical solid line shows the posterior mean of $h_0=4.9$, while the vertical dotted line marks the true parameter value of $h_0=5.0$.

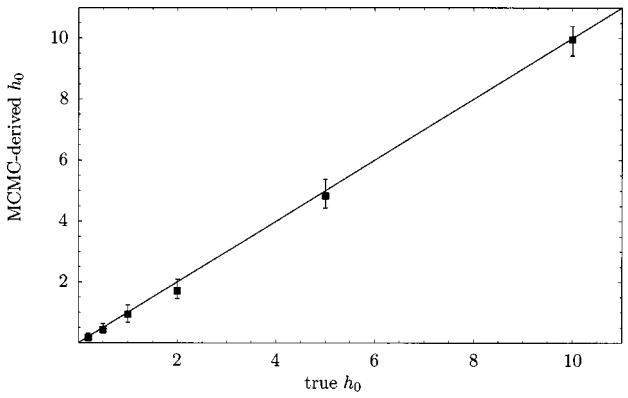

FIG. 3. The posterior mean based on the MCMC sample for the gravitational wave amplitude parameter h_0 versus the actual value of h_0 used in synthesizing the data. The error bars correspond to lower and upper bounds at the 2.5% and 97.5% percentiles of the posterior PDF. The solid line has a slope of 1. The calculations were performed over 14 000 data points, each with noise variance of $\sigma_k^2=1$.

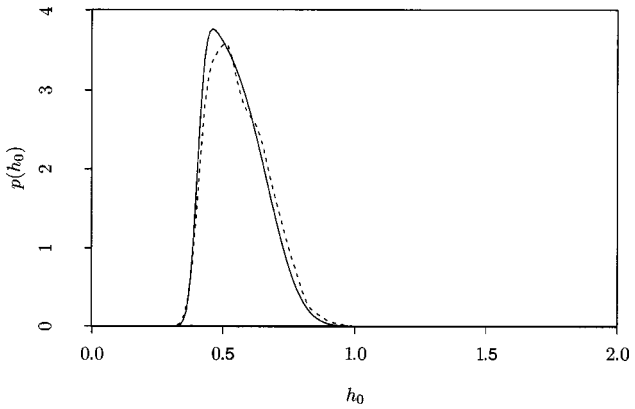

FIG. 4. The posterior mean based on the MCMC sample for the gravitational wave amplitude parameter h_0 (dotted line), along with that produced via the method presented in [20] (solid line). In this example the *true* value was $h_0=0.5$, while the other *true* parameter values were $\psi=0.4$, $\phi=1.0$, and $\iota=0.5$.

in all of our runs was ~ -0.95. As a result the data are consistent with a relatively broad range of combinations of the two parameters, making their *individual* values rather uncertain here—an effect evident from Eq. (1).

The effect of the other unknown parameters (particularly ι) on the posterior PDF for h_0 can be clearly shown by repeating the analysis for Fig. 3 but with h_0 as the only unknown, namely, all of the other parameters set to their *actual* values in the MCMC routine. Under these circumstances the widths of all 95% posterior probability intervals are 0.116, independent of the value of h_0. Comprehensive analyses have investigated detection statistics for a periodic

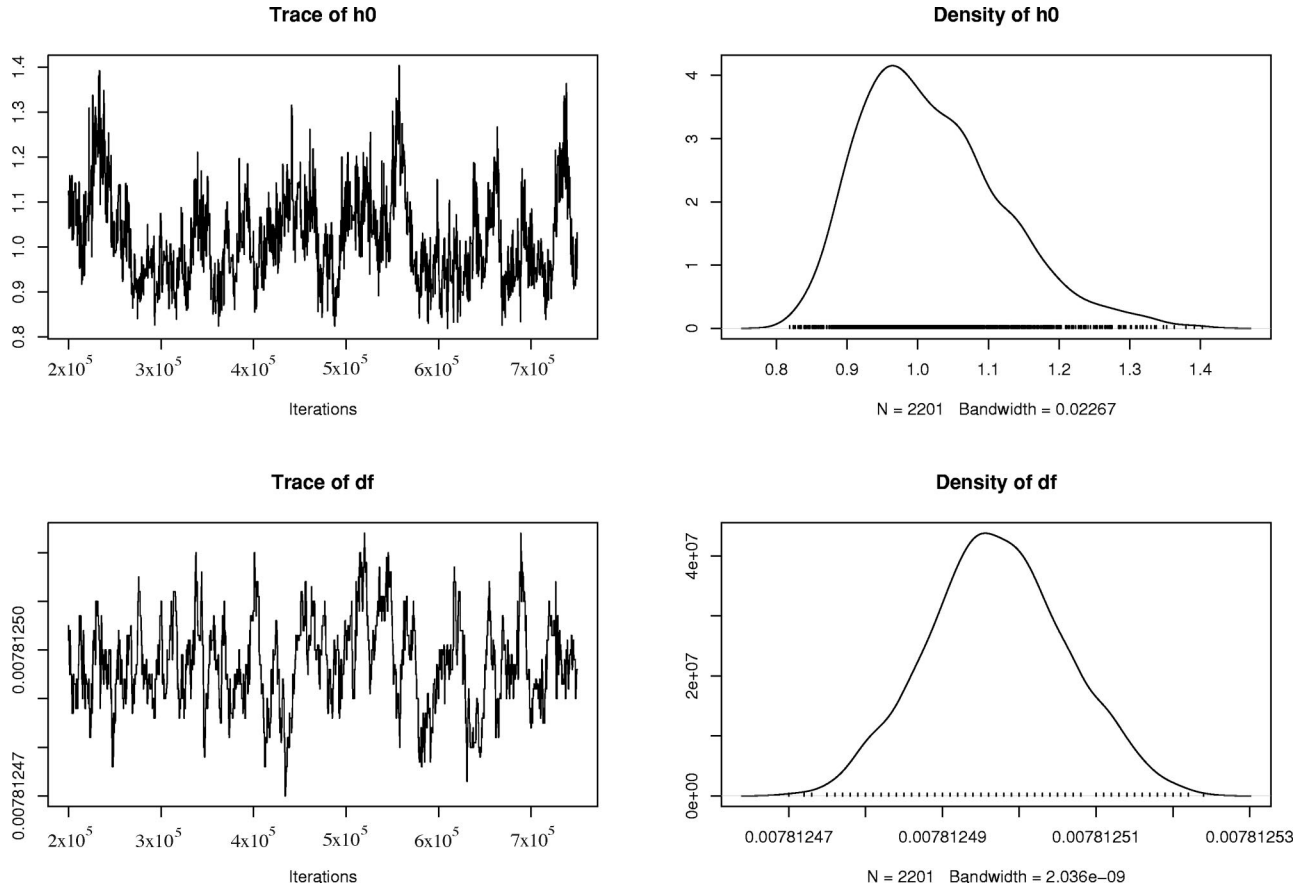

FIG. 5. Trace plot (left) and posterior PDFs (right) for the pulsar parameters h_0 and Δf. In this example from the five parameter problem the true values for these critical parameters were $h_0=1.0$ and $\Delta f=0.0078125$ Hz.

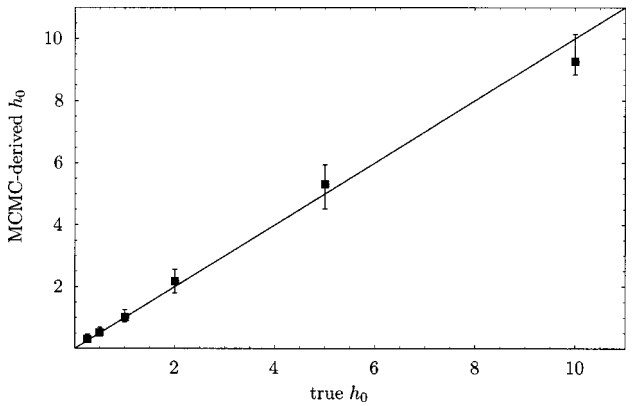

FIG. 6. The posterior mean based on the MCMC sample for the gravitational wave amplitude parameter h_0 versus the actual value of h_0 used in synthesizing the data. This example is from the five parameter problem. The error bars correspond to the lower and upper bounds being specified by the 2.5% and 97.5% percentiles of the posterior PDF. The solid line has a slope of 1.

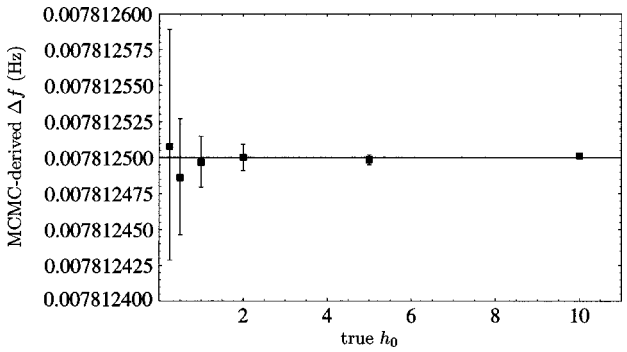

FIG. 7. The posterior mean based on the MCMC sample for the uncertainty in the frequency, Δf, versus the actual value of h_0 used in synthesizing the data. This example is from the five parameter problem. The error bars correspond to the lower and upper bounds being specified by the 2.5% and 97.5% percentiles of the posterior PDF. The horizontal line corresponds to the real value of $\Delta f = 0.007\,812\,500$ Hz.

signal in a gravity wave detector [27]. However, these statistics are concerned only with the amplitude of the periodic signal, and not with parameter estimation (as described above). If we write Eq. (1) as $s(t)=A\cos(\Psi+\Phi)$ (with A being the periodic signal amplitude and Φ a phase term) then the detection statistic of [27] would apply to finding a signal amplitude A in the presence of the detector noise. In terms of Eq. (1), the amplitude of the periodic signal would be

$$A=\{[F_+(t;\psi)h_0(1+\cos^2\iota)/2]^2+[F_\times(t;\psi)h_0\cos\iota]^2\}^{1/2}. \tag{11}$$

It is clear that A has a complicated dependence on h_0 and $\cos\iota$. We will never know, *a priori*, the value of all the pulsar parameters. Our study here is about parameter estimation, and not knowing the values of all the pulsar parameters ultimately increases the width in the posterior PDF for the gravity wave magnitude h_0.

As the magnitudes of the signals are diminished there comes a point when one is no longer able to *confidently* claim a detection. This threshold is somewhat arbitrary, and dependent on the statistics and interpretation. In the study presented here we *claim* that a signal is detected when the $h_0=0$ point is more that two standard deviations from the mean value of the MCMC generated posterior PDF for h_0. For the synthesized signals we investigated this corresponded to a threshold for detection of $h_0=0.1$; in this case the measured mean of the posterior PDF for h_0 was 2.1 standard deviations away from zero. For an initial detection of gravitational radiation it is likely that the scientific community will demand a significantly larger signal-to-noise ratio. However, the performance of the MCMC routine is still very good for these relatively low signal levels.

Although 10^6 Monte Carlo iterations were used in this study (taking 1 day on a 1 GHz processor) adequate distributions can be generated from 10^5 iterations after the burn-in, so good results can be achieved after just a few hours. In fact the marginalizations discussed above can be tackled more quickly using simple summing methods as performed by [20], and the result of a comparison of the two is shown in Fig. 4. The great advantage of the MCMC method for us is its demonstrated ability to deal with problems that have a large number of parameters [10], where other numerical integration techniques (such as employed by [20]) are not feasible. The ultimate goal of our research is to expand this pulsar parameter estimation work to include more parameters. The next step in increasing the complexity of the pulsar signal is to consider potential sources of known location, but with unknown rotation frequency. In order to start this investigation we added a new parameter, the uncertainty in the frequency of the source, Δf. In this example the exact value of the pulsar's gravitational wave signal is uncertain to within 1/60 Hz. In the study we present here there is a difference, Δf, between the gravitational wave signal frequency and the heterodyne frequency. The addition of this new parameter did not significantly increase the rate at which the code ran, but did (by about 20%) increase the length of the burn-in time. If one wanted to increase this frequency range to 5 Hz then this could be done by running the MCMC code on 300 processors, with each run differing in center frequency by 1/60 Hz. The Markov chain using the *correct* frequency would converge, while the other 299 chains would not. This will be a future research project for us.

In our MH code we used a uniform prior for the uncertainty in the frequency, Δf, over $\pm 0.016\,67$ Hz. The injection parameters used were $\psi=0.4$, $\phi=1.0$, $\Delta f = 0.007\,812\,5$ Hz, and $\iota=0.5$ ($\cos\iota=0.88$). h_0 was again injected with a number of values between 0.25 and 10.0. In Fig. 5 we show sample trace plots and posterior PDFs for Δf and h_0 when the injected value of h_0 was 1.0. For this example the MCMC algorithm yielded mean values and 95% posterior probability intervals of $h_0=1.02$ (0.86 to 1.26) and $\Delta f=0.007\,812\,497$ Hz (0.007 812 480 Hz to 0.007 812 515 Hz). The frequency PDF is quite narrow, which was responsible for the increase in the burn-in time as the Markov chain must find this narrow region of parameter space. In Fig. 6 we display the estimate for the gravitational wave amplitude

(h_0) predicted by the five parameter MH routine versus the actual h_0. In Fig. 7 we display the estimate for the difference in frequency Δf predicted by the five parameter MH routine versus the injected h_0.

V. DISCUSSION

Recent applications of MCMC techniques have provided a Bayesian approach to estimating parameters in a number of physical situations. These include cosmological parameter estimation from cosmic microwave background data [13–15], estimating astrophysical parameters for gravitational wave signals from coalescing compact binary systems [16,17], and parameter estimation of a chaotic system in the presence of noise [18,19]. An all sky survey for periodic gravitational waves from neutron stars must explore a very large parameter space, and this has partially been addressed in [8]. Generically, the signal from a neutron star in a binary system will be characterized by at least 13 parameters. It is our hope that the MCMC techniques will prove fruitful in dealing with these complex signals.

In this paper we have demonstrated that the success of the MH routine for the five parameter problem: h_0, ψ, ϕ_0, ι and Δf. Our longer term plans are to account for other parameters, such as spindown rate, pulsar wobble, and possibly location of the signal in the sky. This research is currently in progress.

ACKNOWLEDGMENTS

This work was supported by National Science Foundation grants PHY-0071327 and PHY-0244357, the Royal Society of New Zealand Marsden fund award UOA204, the Natural Sciences and Engineering Research Council of Canada, Universities U.K., and the University of Glasgow.

[1] B. Barish and R. Weiss, Phys. Today **52** (10), 44 (1999).
[2] B. Willke *et al.*, Class. Quantum Grav. **19**, 1377 (2002).
[3] B. Caron *et al.*, Nucl. Phys. B (Proc. Suppl.) **54**, 167 (1996).
[4] K. Tsubono, in *1st Edoardo Amaldi Conference on Gravitational Wave Experiments*, edited by E. Coccia, G. Pizella, and F. Ronga (World Scientific, Singapore, 1995), p. 112.
[5] C. Cutler, Phys. Rev. D **66**, 084025 (2002).
[6] L. Bildsten, Astrophys. J. Lett. **501**, L89 (1998).
[7] J. Middleditch, J.A. Kristan, W.E. Kunkel, K.M. Hill, R.D. Watson, R. Lucinio, J.N. Imamura, T.Y. Steiman-Cameron, A.S. Shearer, R. Butler, M. Redfern, and A.C. Danks, New Astron. **5**, 243 (2000).
[8] P. Jaranowski, A. Królak, and B.F. Schutz, Phys. Rev. D **58**, 063001 (1998).
[9] P. Brady and T. Creighton, Phys. Rev. D **61**, 082001 (2000).
[10] W.R. Gilks, S. Richardson, and D.J. Spiegelhalter, *Markov Chain Monte Carlo in Practice* (Chapman and Hall, London, 1996).
[11] N. Metropolis, A.W. Rosenbluth, M.N. Rosenbluth, A.H. Teller, and E. Teller, J. Chem. Phys. **21**, 1087 (1953).
[12] W.K. Hastings, Biometrika **57**, 97 (1970).
[13] N. Christensen, R. Meyer, L. Knox, and B. Luey, Class. Quantum Grav. **18**, 2677 (2001).
[14] L. Knox, N. Christensen, and C. Skordis, Astrophys. J. Lett. **563**, L95 (2001).
[15] L. Verde, H.V. Peiris, D.N. Spergel, M. Nolta, C.L. Bennett, M. Halpern, G. Hinshaw, N. Jarosik, A. Kogut, M. Limon, S.S. Meyer, L. Page, G.S. Tucker, E. Wollack, and E.L. Wright, Astrophys. J., Suppl. Ser. **148**, 195 (2003).
[16] N. Christensen and R. Meyer, Phys. Rev. D **64**, 022001 (2001).
[17] N. Christensen, R. Meyer, and A. Libson, Class. Quantum Grav. **21**, 317 (2004).
[18] R. Meyer and N. Christensen, Phys. Rev. E **62**, 3535 (2000).
[19] R. Meyer and N. Christensen, Phys. Rev. E **65**, 016216 (2002).
[20] Réjean J. Dupuis and Graham Woan, ''A Bayesian Method to Search for Periodic Gravitational Waves.''
[21] B. Abbott *et al.*, Phys. Rev. D **69**, 082004 (2004).
[22] J.H. Taylor, Phys. Rev. D **66**, 084025 (2002).
[23] L. Tierney, Ann. Stat. **22**, 1701 (1994).
[24] L. Holden, ''Adaptive chains,'' http://www.statslab.cam.ac.uk/mcmc
[25] W.R. Gilks, G.O. Roberts, and S.K. Sahu, J. Am. Stat. Assoc. **93**, 1045 (1998).
[26] A.E. Gelfand and S.K. Sahu, J. Comput. Graph. Stat. **3**, 261 (1994).
[27] B. Allen, M.A. Papa, and B.F. Schutz, Phys. Rev. D **66**, 102003 (2002).