

A time-domain MCMC search and upper limit technique for gravitational waves of uncertain frequency from a targeted neutron star

John Veitch¹, Richard Umstätter³, Renate Meyer³, Nelson Christensen²
and Graham Woan¹

¹ Department of Physics and Astronomy, Kelvin Building, University of Glasgow,
Glasgow G12 8QQ, UK

² Physics and Astronomy, Carleton College, Northfield, MN 55057, USA

³ Department of Statistics, University of Auckland, Auckland, New Zealand

E-mail: jveitch@astro.gla.ac.uk

Received 9 April 2005, in final form 25 May 2005

Published 6 September 2005

Online at stacks.iop.org/CQG/22/S995

Abstract

It is computationally expensive to search the large parameter space associated with a gravitational wave signal of uncertain frequency, such as might be expected from the possible pulsar generated by SN1987A. To address this difficulty we have developed a Markov Chain Monte Carlo method that performs a time-domain Bayesian search for a signal over a 4 Hz frequency band and a spindown of magnitude of up to 1×10^{-9} Hz s⁻¹. We use Monte Carlo simulations to set upper limits on signal amplitude with this technique, which we intend to apply to a gravitational wave search.

PACS numbers: 04.80.Nn, 02.70.Uu, 06.20.Dk

(Some figures in this article are in colour only in the electronic version)

1. Introduction

Recent targeted searches for gravitational waves from radio pulsars have needed to consider only four signal parameters: h_0 , the signal amplitude, $\cos \iota$, the inclination angle of the pulsar towards the line of sight, ψ , the polarization angle of the gravitational wave and ϕ_0 , the initial phase of the signal. In the absence of a detection, it is possible to set an upper limit on h_0 , and therefore on the ellipticity of the pulsar in question. This is accomplished by defining a Bayesian posterior probability density function (PDF) of the four signal parameters and marginalizing it numerically on a grid. In the time domain, this search technique relies on precise heterodyning of data from the gravitational wave detector, which itself requires accurate information on the phase evolution of the signal. For known pulsars this is easily obtained

by monitoring the timing of their radio wave pulses, which directly provides the frequency and frequency derivatives of the gravitational signal (emitted at twice the object's rotation frequency) [1]. However, there are objects whose frequencies are not known accurately, such as the reported remnant of Supernova 1987A [3]. To perform a search over a range of frequencies would require a grid on frequency and frequency derivative with resolutions approximately $1/T_{\text{obs}}$ and $1/T_{\text{obs}}^2$ respectively, which makes the number of search templates very large ($\sim 4.5 \times 10^{11}$ for LIGO S3 data). We present an adaptation of this search based on a Markov Chain Monte Carlo (MCMC) method which does not require an exhaustive examination of the parameter space and therefore is able to search a range of frequencies in a reasonable time.

2. MCMC parameter estimation

After heterodyning at close to the expected signal frequency, the gravitational wave signal from a rotating triaxial neutron star has the form

$$y(t_k; \mathbf{a}) = \left(\frac{1}{4} F_+(t_k; \psi)(1 + \cos^2 \iota) - \frac{i}{2} F_\times(t_k; \psi) \cos \iota \right) h_0 e^{i\Phi(t_k)} \quad (1)$$

where \mathbf{a} represents the parameters responsible for generating the signal, F_+ and F_\times are the beam patterns for the + and \times polarizations, and $\Phi(t)$ is the phase of the signal given to the first two Taylor expansion terms by

$$\Phi(t) = \phi_0 + 2\pi \left[\delta f (T - T_0) + \frac{1}{2} \delta \dot{f} (T - T_0)^2 \right]. \quad (2)$$

Here δf is the deviation of the signal frequency from the heterodyne frequency, and $\delta \dot{f}$ is the deviation from first derivative of the heterodyne frequency [7]. By including these two parameters we can search around the heterodyne frequency for a signal. The heterodyned data are reduced to one complex sample per minute, B_k , with variance σ_k , allowing a frequency range of $[-1/120, 1/120]$ Hz around the heterodyne frequency to be searched. The search in $\delta \dot{f}$ is limited to $[-1 \times 10^{-9}, 1 \times 10^{-9}]$ Hz s^{-1} which is likely to include any possible pulsar spindown rates. The search proceeds by heterodyning 480 separate frequency bands at intervals of 1/120 Hz, and running parallel searches on each band. In this way we complete a search over a 4 Hz range—suitable for the putative SN1987A pulsar.

The probability of a particular combination of the six parameters \mathbf{a} representing a signal in the data $\{B_k\}$ is

$$p(\mathbf{a}|\{B_k\}) \propto p(\mathbf{a}) \exp \left[- \sum_k \frac{|B_k - y(t_k; \mathbf{a})|^2}{2\sigma_k^2} \right], \quad (3)$$

which is the prior probability of \mathbf{a} multiplied by the likelihood, $p(\{B_k\}|\mathbf{a})$. This proportional definition is adequate for our application, as we will only be evaluating ratios of probabilities.

We now have a joint posterior probability distribution over the six parameters in vector \mathbf{a} , but in the presence of a signal the probability will be strongly concentrated around the signal parameters. The Markov Chain Monte Carlo method takes advantage of this fact by preferentially sampling the areas of greatest probability density, in proportion to that probability [4, 2]. In this way we can approximate the posterior distribution $p(\mathbf{a}|\{B_k\})$ by allowing the Markov Chain to accumulate many samples, where the density of samples in an area is then proportional to the probability density in that area. In order to accomplish this,

a Markov Chain in state \mathbf{a} chooses a candidate sample \mathbf{a}' from proposal distribution $q_1(\mathbf{a}'|\mathbf{a})$, and then accepts the candidate as the next state of the chain with probability

$$\alpha_1(\mathbf{a}'|\mathbf{a}) = \min \left[1, \frac{p(\mathbf{a}')p(\{B_k\}|\mathbf{a}')q_1(\mathbf{a}|\mathbf{a}')}{p(\mathbf{a})p(\{B_k\}|\mathbf{a})q_1(\mathbf{a}'|\mathbf{a})} \right]. \quad (4)$$

If the proposal is not accepted, the current state \mathbf{a} is recorded as a sample again and the process repeats. The proposal distribution $q_1(\mathbf{a}'|\mathbf{a})$ is a multivariate normal distribution, with covariances determined by the correlation between parameters in trial runs. In practice, the highly correlated parameters h_0 and $\cos \iota$ are replaced with a re-parametrization, such that $a_1 = \frac{1}{4}h_0(1 + \cos^2 \iota)$ and $a_2 = \frac{1}{2}h_0 \cos \iota$. The frequency parameters are also changed to $f_{\text{start}} = \delta f + \frac{1}{2}\delta \dot{f} \times t_{\text{start}}$ and $f_{\text{end}} = \delta f + \frac{1}{2}\delta \dot{f} \times t_{\text{end}}$, being the signal frequency at the start and end times of the data.

It is important that the chain explore the parameter space adequately, in order to find the area of high probability. To this end, we have included a delayed rejection stage of the algorithm, so that if a proposal \mathbf{a}' is rejected, a new candidate is generated from distribution $q_2(\mathbf{a}''|\mathbf{a}', \mathbf{a})$, which has a narrower width and makes more conservative steps. These are accepted with probability

$$\alpha_2(\mathbf{a}''|\mathbf{a}) = \min \left[1, \frac{p(\mathbf{a}'')p(\{B_k\}|\mathbf{a}'')q_1(\mathbf{a}'|\mathbf{a}'')q_2(\mathbf{a}|\mathbf{a}', \mathbf{a}'')(1 - \alpha_1(\mathbf{a}'|\mathbf{a}''))}{p(\mathbf{a})p(\{B_k\}|\mathbf{a})q_1(\mathbf{a}'|\mathbf{a})q_2(\mathbf{a}''|\mathbf{a}, \mathbf{a}')(1 - \alpha_1(\mathbf{a}'|\mathbf{a}))} \right], \quad (5)$$

where the proposal distributions are included to satisfy the principle of detailed balance. This allows the chain to make small steps where a large step would be rejected.

In addition to these measures, there is also a burn-in period at the start of each MCMC run where the exponent in equation (3) is multiplied by an inverse temperature factor β in the proposed step. Initially $\beta = 0.01$, and gradually increases during the burn-in period to $\beta = 1$, where equation (3) is restored. This decreases the likelihood of the chain getting stuck in a local maximum of probability without exploring the space adequately. No samples from the burn-in period are used in calculating the final PDF, as they do not represent the target distribution $p(\mathbf{a}|\{B_k\})$. In our implementation the burn-in lasts for 1×10^6 iterations, followed by 100 000 iterations with every 50th used as a sample, so as to reduce correlation between samples. The program is implemented in C, using the LIGO Algorithm Library to calculate the beam patterns and time delays necessary for analysing a real signal [5].

3. Setting upper limits

In the standard time-domain search, if a signal is not detected it is useful to set an upper limit on gravitational radiation emitted by the source under examination. This translates as an upper limit on the ellipticity of the pulsar which can be used to constrain physical models of a neutron star. This is achieved by marginalizing the posterior PDF over $\cos \iota$, ψ and ϕ_0 , leaving a distribution on h_0 which can then be integrated upward from $h_0 = 0$ until 95% of the probability is included in the interval. The upper limit of the interval is then the 95% upper limit on h_0 for the target.

In the MCMC routine however, if the signal is below a certain threshold then the chain may not converge on the correct mode of the signal in the PDF [6]. This is partly due to the narrowness of the mode in δf and $\delta \dot{f}$, where the bulk of the probability lies within a single frequency bin of width $1/T_{\text{obs}}$ with a very small attraction area.

A weak signal will raise probability only slightly in this bin, so that even if the chain were to find the mode, it has a probability of jumping out and not returning. These factors therefore preclude the evaluation of an upper limit simply by marginalizing the posterior PDF when

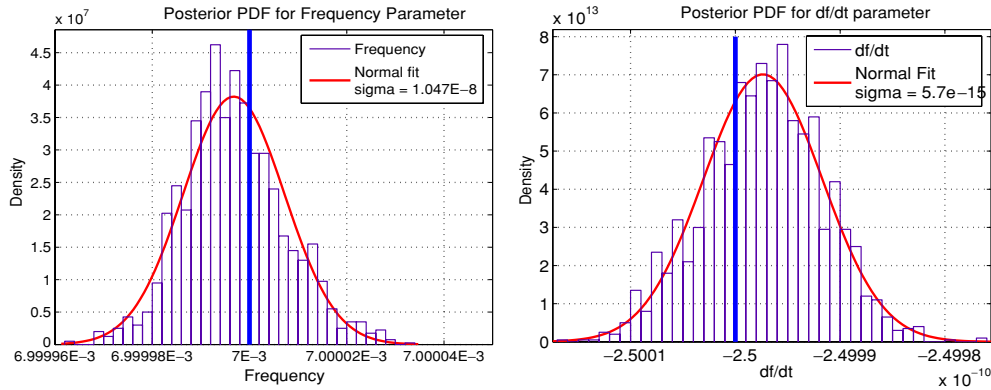


Figure 1. The posterior PDFs (bars) for a 2000 sample chain which has converged on the injected values of $\delta f = 7 \times 10^{-3}$ Hz and $\delta \dot{f} = -2.5 \times 10^{-10}$ Hz s $^{-1}$ (vertical lines), and the best normal fits, having standard deviations 1.047×10^{-8} Hz and 5.7031×10^{-15} Hz s $^{-1}$ respectively.

there is no detection. Instead, we performed a Monte Carlo simulation by injecting signals of known parameters into noise, and running the MCMC code to try and recover them. Signals are injected with known parameters by calculating the value of the signal at each timestamp t_k using LAL and adding it to each sample B_k in the input data file. Real data would have to be heterodyned prior to this step, but the artificial noise is generated at 1/60 Hz so heterodyning is unnecessary in the simulations. It can be shown that the probability of detecting a signal using our method depends strongly on the injected values of h_0 and $\cos \iota$, and much more weakly on the value of ψ . To determine our upper limit for a particular set of data, we inject signals of varying h_0 and $\cos \iota$ into the noise σ_k , then analyse the results of the MCMC routine to determine if it has detected the signal. This is accomplished by observing the output PDFs in δf and $\delta \dot{f}$, and determining if the chain has converged or not. If the injected value lies within three standard deviations of the mean, and the standard deviations themselves are less than one frequency bin in size, the chain is judged to have converged. If the chain has not converged, the samples are distributed widely over the entire δf range, and the standard deviation is typically five orders of magnitude higher. Figure 1 shows the posterior PDFs in the δf and $\delta \dot{f}$ parameters for a chain converged on a signal with injected parameters $h_0 = 1.3$, $\cos \iota = 0$, $\psi = 0.281$, $\phi_0 = 4.234$, $\delta f = 7 \times 10^{-3}$ Hz and $\delta \dot{f} = -2.5 \times 10^{-10}$ Hz s $^{-1}$, where the noise level was $\sigma = 1.0$. Figure 2 shows the same type of plot for a signal with $h_0 = 0.7$ and otherwise identical parameters; this chain failed to converge, so there is no concentrated mode in the PDF. The attempted fit to a normal distribution to test convergence is therefore very wide, and poorly fitted.

The injections are repeated, using a different starting point in parameter space for the MCMC chain each time, and an empirical probability of signal recovery $P(\text{detection}|h_0, \cos \iota, \sigma_k)$ for each point on the $(h_0, \cos \iota)$ plane is obtained from the set of results. Figure 3 shows the result on the $(h_0, \cos \iota)$ plane of performing this procedure on 64 000 samples of white noise with a variance of $\sigma = 1.0$. The dark region shows where no signal is detectable, above which is a small transition zone where there is a finite probability of detection. With increasing iterations of the chain, the width of this zone reduces—we are using a chain of 1100 000 iterations, of which the first 1000 000 are burn-in time and the last 100 000 used for sampling the PDF. The white area above this transition zone represents signals that are detected with very high probability. These are strong enough to cause the chain to converge on them in every instance of the Monte Carlo simulation. It is clear from

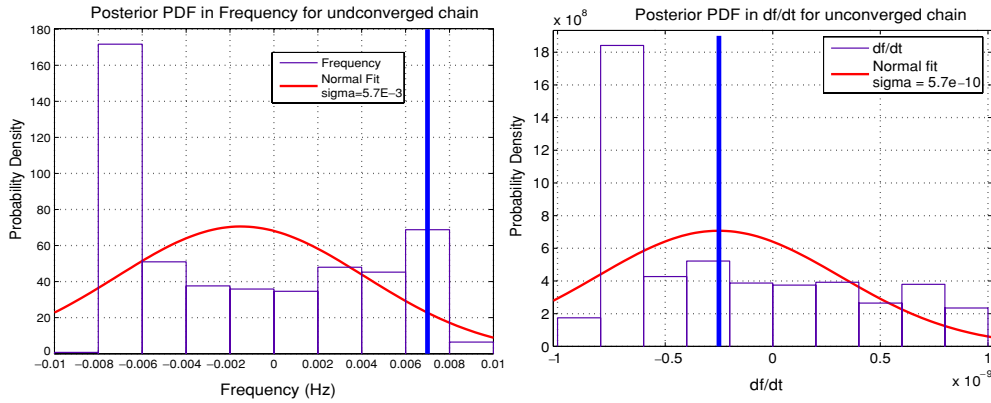


Figure 2. The δf and $\delta \dot{f}$ posterior PDFs (bars) for a 2000 sample chain which failed to converge on the injected parameters (vertical lines). This distribution is clearly distinguishable from the converged case using the criteria defined above.

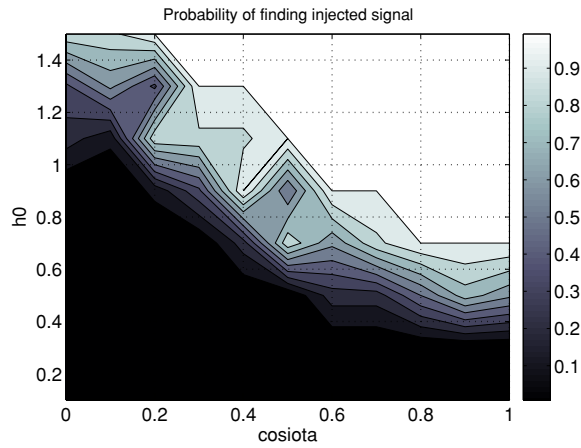


Figure 3. Results of Monte Carlo evaluation of probability of detection of a signal injected with varying values of h_0 and $\cos \iota$. This shows the strong dependence on both these parameters. The transition from detectable to undetectable occurs over a narrow region for a particular value of $\cos \iota$, but decreases in h_0 as $\cos \iota$ tends to 1.

the figure that this is strongly dependent on $\cos \iota$, as this acts a weighting factor between the power in the real and imaginary parts of the heterodyned signal. The distribution is symmetric about $\cos \iota = 0$, so only the positive half was calculated.

This result can then be marginalized by summing over $\cos \iota$ and normalizing to give a distribution on h_0 , as required for an upper limit. The probability of there being an undetected gravitational wave is then $1 - P(\text{detection}|h_0, \sigma_k)$, and the upper limit $h_{95\%}$ satisfies $1 - P(\text{detection}|h_{95\%}, \sigma_k) = 0.95$. Note that $P(\cdot|\cdot)$ is a probability rather than a probability density. The upper limit determined from this stage is applicable to each of the 480 bands in the 4 Hz frequency interval because the values of σ_k are determined by estimating the noise over the entire 4 Hz band at the heterodyne stage. The heterodyne then performs filtering of this band to calculate the signal B_k in each 1/60 Hz interval, from which the 480 bands are derived.

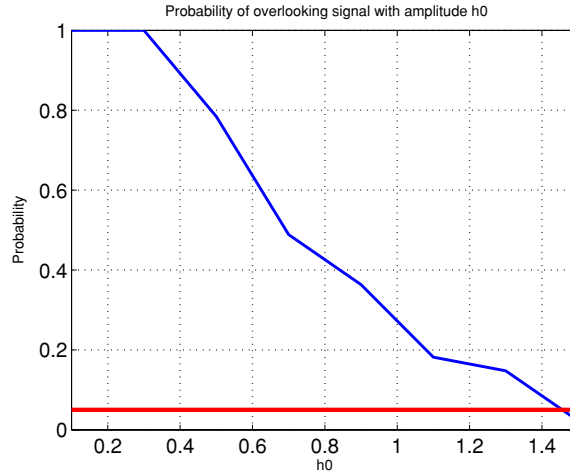


Figure 4. The marginalization of figure 3 onto the h_0 axis. This shows that the 95% confidence (represented by the horizontal line) may be drawn at $h_{95\%} = 1.44$.

Figure 4 shows this final distribution on h_0 , where the horizontal line represents 95% confidence in detection, required for a 95% upper limit, which is achieved at $h_0 = 1.44$. It is useful to express this in terms of noise power spectral density, S_h , yielding the relation

$$h_{95\%} = 515.2 \sqrt{\frac{S_h}{T_{\text{obs}}}}, \quad (6)$$

where the 512.2 is an indicator of the sensitivity of the search, empirically derived from the simulation.

For a realistic example, consider a noise power spectral density of $S(f) = 6.6 \times 10^{-44} \text{ Hz}^{-1}$, or $h(f) = 2.57 \times 10^{-22}$, and a data run of 44 days, the upper limit would then be approximately

$$h_{95\%} = 515.2 \sqrt{\frac{6.6 \times 10^{-44}}{3801600}} = 6.8 \times 10^{-23}. \quad (7)$$

4. Summary

This MCMC search method is capable of reliably detecting a signal above a certain threshold and estimating its parameters including frequency and spindown. The computing time it takes to perform the search is dependent only on the amount of data and the number of MCMC iterations desired, which provides significant improvement over the previous time-domain targeted search [8]. With the current version of the pipeline, to analyse the entire 4 Hz window would take about 17 000 CPU hours on a 1.8 GHz Athlon processor, which is approximately two days on the 366-CPU Merlin cluster at AEI Golm. In the absence of a signal an upper limit can be set using Monte Carlo injections, which require a similar amount of processing on top of the search itself. Additional work is under way to tune the algorithm further, which may lead to improvements in sensitivity and speed of execution.

We hope that this technique could be applied to a search for gravitational radiation from the possible pulsar in the remnant of SN1987A at $935 \pm 2 \text{ Hz}$, and in the absence of detection

place upper limits on the gravitational radiation emitted by this object if it were triaxial. The method could also be easily applied to other similar objects.

Acknowledgments

This work was supported by National Science Foundation grant PHY-0244357, the Royal Society of New Zealand Marsden fund award UOA204, Universities UK, PPARC and the University of Glasgow. We would like to also thank Peter Saulson for reading and commenting on the manuscript.

References

- [1] Abbot B *et al* (The LIGO Scientific Collaboration) 2004 *Phys. Rev. D* **69** 082004
- [2] Umstätter R, Meyer R, Dupuis R J, Veitch J, Woan G and Christensen N 2004 *Class. Quantum Grav.* **21** S1655–65
- [3] Middleditch J *et al* 2000 Rapid photometry of supernova 1987A: a 2.14 ms pulsar? *New Astron.* **5** 243–83
- [4] Gamerman D 1997 *Markov Chain Monte Carlo: Stochastic Simulation for Bayesian Inference* (London: Chapman and Hall)
- [5] <http://www.lsc-group.phys.uwm.edu/lal/>
- [6] Gilks W R, Richardson S and Spiegelhalter D J 1996 *Markov Chain Monte Carlo in Practice* (London: Chapman and Hall)
- [7] Jaranowski P, Królak A and Schutz B 1998 *Phys. Rev. D* **58** 063001
- [8] Abbott B *et al* (The LIGO Scientific Collaboration), Kramer M and Lyne A 2005 *Phys. Rev. Lett.* **94** 181103 (Preprint gr-qc/0410007)