

Bayesian inference on compact binary inspiral gravitational radiation signals in interferometric data

Christian Röver¹, Renate Meyer¹ and Nelson Christensen²

¹ Department of Statistics, The University of Auckland, Auckland, New Zealand

² Department of Physics & Astronomy, Carleton College, Northfield, MN, USA

E-mail: christian@stat.auckland.ac.nz, meyer@stat.auckland.ac.nz,
nchrste@carleton.edu

Abstract. Presented is a description of a Markov chain Monte Carlo (MCMC) parameter estimation routine for use with interferometric gravitational radiational data in searches for binary neutron star inspiral signals. Five parameters associated with the inspiral can be estimated, and summary statistics are produced. Advanced MCMC methods were implemented, including importance resampling and prior distributions based on detection probability, in order to increase the efficiency of the code. An example is presented from an application using realistic, albeit fictitious, data.

PACS numbers: 04.80.Nn, 02.70.Uu.

Submitted to: *Classical and Quantum Gravity*

1. Introduction

Great effort is currently being expended in the search for gravitational radiation. Detectors around the world [1, 2, 3, 4] are reaching sensitivities where an event may be seen in the near future. The field of research has grown immensely since the first prediction of gravitational radiation [5], or even from the time of the observations that confirmed its existence [6]. The inspiral of compact binary objects may provide the cleanest system for comparing observations with general relativistic predictions [7]. Observation of inspiral events could provide important information on the structure of neutron stars [8, 9]. Even cosmological information can be extracted from the observation of inspiral events [10, 11, 12, 13]. The characteristics of radiation in the post-Newtonian regime will provide insight into highly non-linear general relativistic effects, such as the observation of the formation of a Kerr black hole as the binary system decays [12, 14, 15]. LIGO has recently conducted searches for signals from the inspiral of binary neutron star systems in science run data [16, 17].

Detection pipelines for binary inspiral signals [16, 18, 19] in interferometer data

will likely provide the first indication of detection. These pipelines provide parameter estimation through various parameter estimation techniques. Other techniques, like those that we will describe in this paper, provide a means to examine the data and produce not just point estimates of parameters, but their complete distributions including summary statistics. The unique analysis methods that we bring to the examination of LIGO data are computationally intensive MCMC algorithms. Our method of parameter estimation is an exercise in Bayesian inference. These Bayesian statistical techniques offer great promise in problems where the number of parameters is large. Since its initial application in digital signal analysis [20] MCMC methods have revolutionized many areas of applied statistics [21]. A distinct advantage of the MCMC approach is that computational time does not grow exponentially with parameter number, as it does for other methods [22]. After their introduction to the cosmological parameter estimation problem [23, 24], MCMC methods were quickly adopted and widely used with CMB data [25, 26, 27], including WMAP [28, 29]. We anticipate a similar scenario with gravitational radiation; after initial detection attention will be focused on parameter estimation and associated astrophysical information extracted from the data. The techniques that we describe in this paper are intended to be applied to interferometer data after an initial event has been registered by an inspiral detection pipeline [16, 18, 19]; our MCMC method would allow the extraction of information about the physical parameters. The technique we present here provides parameter estimates from the data output of a single interferometer. An MCMC technique using the coherent addition of signals from multiple detectors will be presented in a future publication.

In the present paper we describe the MCMC routine we have developed to examine interferometer data for binary neutron star inspiral signals. We have implemented advanced MCMC methods in order to effectively and efficiently search a five-parameter space, find the signal, and produce summary statistics. There have been significant advances in the execution of the binary neutron star inspiral MCMC routine since the publication of our previous method [30]; presented here is a description of our routine, implemented in C, that employs advanced MCMC techniques in order to increase efficiency and performance.

The outline of the paper is as follows. Section 2 describes our analysis strategy and contains a description of Bayesian parameter estimation methods, and the methods we employ in the binary neutron star inspiral MCMC application. In Section 3 we present an example application of our method using fictitious data; a binary neutron star inspiral event is embedded within noise (that matches the LIGO design sensitivity). We conclude the paper with a summary of our results, and a description of our intentions to expand this work to more complex situations.

2. Analysis strategy

2.1. Bayesian modeling and MCMC posterior simulation

We begin with a short introduction to Bayesian inference, which is a quite different way to approach problems like parameter estimation compared to more commonly used “frequentist” methods, like e.g. Maximum-Likelihood estimation. The model setup includes the specification of a *prior distribution* for all involved parameters as well as the *sampling distribution* for observed quantities. Inference is done via the derived *posterior distribution* given the observed data; this allows one to make probabilistic statements about parameters of interest by treating them as random variables. Due to its setup the concept is sometimes dubbed *full probability modeling*. While the theory behind Bayesian modeling is straightforward, its practical application for inference often requires the use of numerical methods.

The *sampling distribution* is specified in terms of the corresponding probability density function $p(y|\theta)$, defining how the observations y come about given a (fixed) parameter setting θ . Viewed as a function of θ , it also defines the *likelihood function*. The use of the likelihood function, which may be thought of it as ‘linking observations to parameters’, is also common in frequentist analyses.

In the frequentist approach θ is viewed as fixed but unknown. With a Bayesian approach one treats θ as a random variable with a probability distribution that reflects the researcher’s uncertainty about the parameters. In addition to the likelihood a prior distribution $p(\theta)$ is used, reflecting the pre-experimental knowledge and uncertainty about the parameters. The availability of both prior $p(\theta)$ and sampling distribution $p(y|\theta)$ then allows to derive the *posterior distribution*

$$p(\theta|y) = \frac{p(\theta)p(y|\theta)}{p(y)} \propto p(\theta)p(y|\theta) \quad (1)$$

by applying (a generalisation of) Bayes’ theorem. In general, Bayes’ theorem deals with conditional probabilities, here it provides the probability density of θ *conditional on the observed data* y [31]. Note that the denominator $p(y) = \int p(y, \theta)d\theta = \int p(y|\theta)p(\theta)d\theta$ is a normalising constant that does not depend on θ , so the posterior $p(\theta|y)$ is essentially proportional to the product of prior and likelihood.

The posterior distribution expresses the knowledge about the parameter(s) given model (prior and sampling distribution) *and* data. Inference then aims at deriving probabilistic statements about parameters (or related quantities) from the posterior; one might for example be interested in posterior expectations, quantiles or marginal densities of quantities of interest, or probabilities of certain events.

Note that the impact of the prior specification is limited, since asymptotic theory shows that the importance of the prior distribution diminishes as the sample size increases, so the posterior is eventually dominated by the likelihood [31, 32]. When MCMC methods are used for posterior analysis as in this paper, the results (MCMC samples) yielded using a certain prior may easily be recycled to generate posterior

samples assuming different priors by using importance sampling [31].

Frequentist and *Bayesian* analyses are fundamentally different in their approach to given problems—Finn [33] puts the difference as “guessing nature’s state” in case of frequentist analysis as opposed to “learning from observation” for Bayesian analysis. For more detailed discussions see e.g. [21, 31, 33].

For posterior computations we use a simulation-based approach as the posterior distribution is high-dimensional and the calculation of marginal summaries like marginal posterior means would require high-dimensional integration. The simulation-based approach is to generate a sample from the posterior and then approximate the desired integrals by sample averages. For example, for some random variable X with probability density $p_X(x)$ its expectation $E[X] = \int x p_X(x) dx$ can be approximated through a sample average $\bar{x} = \frac{1}{N} \sum_{i=1}^N x_i$ where x_1, \dots, x_N are random draws from $p_X(x)$. Analogously, marginal densities, quantiles etc. can be estimated from samples from the distribution of interest.

Markov chain Monte Carlo (MCMC) methods in general may be used to generate random sequences of numbers that have a specified stationary distribution and in which each random step in the sequence only depends on the previous stage of the chain. The Metropolis-algorithm in particular can be applied if the target distribution is specified in terms of its density which only needs to be known up to a normalising constant. While a Metropolis sampler is guaranteed to function as the number of samples goes towards infinity, its proper behaviour and efficiency within finite time heavily depends on the sensible specifications of its *starting point* and *proposal distribution* [31].

2.2. Likelihood, signal templates & priors

The observed data are a (real-valued) time series of length N and sampling rate $\frac{1}{\Delta_t}$ indicating the phase shift between the two interferometer arms over time at evenly spaced time points. In order to analyse a data set, it is Fourier-transformed, from

$$\{z(t) \in \mathbb{R} : t = 0, \Delta_t, 2\Delta_t, \dots, (N-1)\Delta_t\} \quad (2)$$

into

$$\{\tilde{z}(f) \in \mathbb{C} : f = 0, \Delta_f, 2\Delta_f, \dots, (\frac{N}{2}-1)\Delta_f\}, \quad (3)$$

where $\Delta_f = \frac{1}{N\Delta_t}$ and

$$\tilde{z}(f) = \Delta_t \sum_{j=0}^{N-1} z(j\Delta_t) \exp(-2\pi i j f). \quad (4)$$

The likelihood depends on the data and a *signal template*, describing the theoretically derived detector response s_{ϑ} for a given parameter set ϑ .

The likelihood of some parameter set ϑ then is proportional to the sum of the squared and normalised differences between Fourier transforms of observed signal (\tilde{z}) and

signal template (\tilde{s}_ϑ) over the discrete set of Fourier frequencies $\{(j \times \Delta_f) : j_L \leq j \leq j_U\}$ [34, 35]:

$$p(z|\vartheta) = K \times \exp\left(-\frac{2}{N\Delta_t} \sum_{j=j_L}^{j_U} \frac{\overbrace{|\tilde{z}(j \times \Delta_f) - \tilde{s}_\vartheta(j \times \Delta_f)|^2}^{\text{data}}}{\underbrace{S_n(j \times \Delta_f)}_{\text{noise PSD}}}\right) \quad (5)$$

where $j_L \times \Delta_f$ and $j_U \times \Delta_f$ are the lower and upper bounds of the examined frequency range, $S_n(\cdot)$ is the (estimated) noise power spectral density (PSD), $N\Delta_t$ is the length of the analysed data segment, and K is a normalising constant.

The signal templates we use are 2.0-post-Newtonian stationary phase approximations of the Fourier-transformed inspiral waveform [36]. The model has five parameters: the individual masses of the two involved companions m_1 and m_2 , the coalescence time t_c , the coalescence phase ϕ_0 and the luminosity distance d_L . The luminosity distance is not the actual distance to the source, but reflects the effect on the amplitude of the gravity wave from other parameters. The other parameters affecting the gravity wave amplitude are the binary inspiral system's sky position and orbital plane inclination angle, the interferometer orientation, and the polarization of the wave. We also ignore any effects due to the spins of the compact objects. The template is then defined in terms of *total mass* $m_t = m_1 + m_2$ and *mass ratio* $\eta = \frac{m_1 m_2}{m_t^2}$ as:

$$\tilde{s}_\vartheta(f) = \frac{\sqrt{\eta} m_t^{\frac{5}{6}}}{d_L} \frac{\sqrt{5} G^{\frac{5}{6}}}{2\sqrt{6} \pi^{\frac{2}{3}} c^{\frac{3}{2}}} f^{-\frac{7}{6}} \exp\left(-i \underbrace{(\psi(f) + \phi_0 + 2\pi f t_c)}_{\text{phase evolution}}\right) \quad (6)$$

where

$$\psi(f) := \sum_{i=1}^4 a_i \zeta_i(f), \quad (7a)$$

$$a_1 = \frac{3}{128\eta} q^{-\frac{5}{3}}, \quad (7b)$$

$$a_2 = \frac{1}{384\eta} \left(\frac{3715}{84} + 55\eta \right) q^{-1}, \quad (7c)$$

$$a_3 = -\frac{1}{128\eta} 48\pi q^{-\frac{2}{3}}, \quad (7d)$$

$$a_4 = \frac{3}{128\eta} \left(\frac{15\,293\,365}{508\,032} + \frac{27\,145}{504} \eta + \frac{3085}{72} \eta^2 \right) q^{-\frac{1}{3}}, \quad (7e)$$

$\zeta_1(f) = f^{-\frac{5}{3}}$, $\zeta_2(f) = f^{-1}$, $\zeta_3(f) = f^{-\frac{2}{3}}$, $\zeta_4(f) = f^{-\frac{1}{3}}$, and $q = \pi G m_t c^{-3}$ [35, 36].

Priors are specified with respect to some preliminary considerations. The coalescence time t_c is assumed to be known in advance with a certain accuracy through preprocessing of the data [16, 18, 19]; we set its prior to be uniform across ± 5 ms around the true value (which of course is known for our simulated data). The prior for the coalescence phase ϕ_0 is uniform across its domain $[0, 2\pi]$.

The prior for masses m_1 , m_2 and luminosity distance d_L is set in order to reflect

the distribution of parameters of a neutron star inspiral *given* that it has been detected in the first place. Initially, the prior for two companions' individual masses (m_1 and m_2) is uniform between 0.6 and 3.0 M_\odot (solar masses: $M_\odot = 1.988919 \times 10^{30}$ kg), which effectively covers the range of values expected for binary neutron star systems. The prior for d_L is derived from the assumption that inspirals happen uniformly across space, so that $P(d_L \leq x) \propto x^3$. So far, this leads to an improper distribution (that has an infinite integral). For actual data, in order to be investigated at all, an inspiral event needs to emit a signal that is strong enough to produce a trigger in an inspiral search routine [16, 18, 19], and thus cannot originate from arbitrarily great distances. We incorporated this constraint by downweighing the low-mass/great-distance inspirals according to their lower *detection probability*. The probability of detection is assumed to depend simply on the signal's amplitude, which is affected by the inspiral's masses and distance and is proportional to:

$$\begin{aligned} \alpha(m_1, m_2, d_L) &= \frac{1}{2}(\log(m_1) + \log(m_2)) - \frac{1}{6} \log(m_1 + m_2) - \log(d_L) \quad (8) \\ &= \log\left(\frac{\sqrt{m_1 m_2}}{(m_1 + m_2)^{\frac{1}{6}} d_L}\right) = \log\left(\frac{\sqrt{\eta} m_t^{\frac{5}{6}}}{d_L}\right) \end{aligned}$$

(so α denotes the logarithmic amplitude; see also equation (6)). We could have set a threshold amplitude below which neutron star inspirals would be assumed to be undetectable, but favoured a 'smoother' transition that does not explicitly assign zero probability to parts of the parameter space. We model the dependence between a given (logarithmic) amplitude x and detection probability using a (sigmoidal) logistic function of the form:

$$d_{a,b}(x) = \frac{1}{1 + \exp\left(\frac{x-a}{b}\right)} \quad (9)$$

where a and b are set so that $d_{a,b}(x_U) = 1 - p$ and $d_{a,b}(x_L) = p$ for some $0 < p < 0.5$ (e.g.: $p := 0.1$) and upper and lower reference points x_U and x_L . So x_U denotes the amplitude at which the detection probability reaches $(1 - p)$, and x_L is the amplitude where the probability falls below p . In order to fit d through these points, its parameters are set to:

$$a := \frac{x_L + x_U}{2} \quad \text{and} \quad b := \frac{x_U - x_L}{2 \log(p)}. \quad (10)$$

So, given the above prior *occurrence* and *detection* probabilities, the resulting (proper) joint prior for individual masses m_1 and m_2 and distance d_L is:

$$p(m_1, m_2, d_L) \propto I_{[0.6,3.0]}(m_1) \times I_{[0.6,3.0]}(m_2) \times d_L^2 \times d_{a,b}(\alpha(m_1, m_2, d_L)) \quad (11)$$

where $I_{\mathcal{A}}(x) := \begin{cases} 1 & \text{if } x \in \mathcal{A} \\ 0 & \text{otherwise} \end{cases}$ is the indicator function for any set \mathcal{A} . For the examples given in this paper we have chosen values for the sigmoidal logistic function so that a 2-2 M_\odot inspiral would be detectable out to 90 and 95 Mpc with probabilities of 90% and 10% respectively. These values effectively produce a prior probability for the effective distance that smoothly covers the values of interest, and then falls to zero as the distance

gets too large.

Fig. 1 shows estimates of the marginal prior probability densities of masses and luminosity distance, in terms of individual masses (m_1, m_2) as well as in the alternative parametrisation of chirp mass (m_c) and mass ratio (η). Greater individual masses are

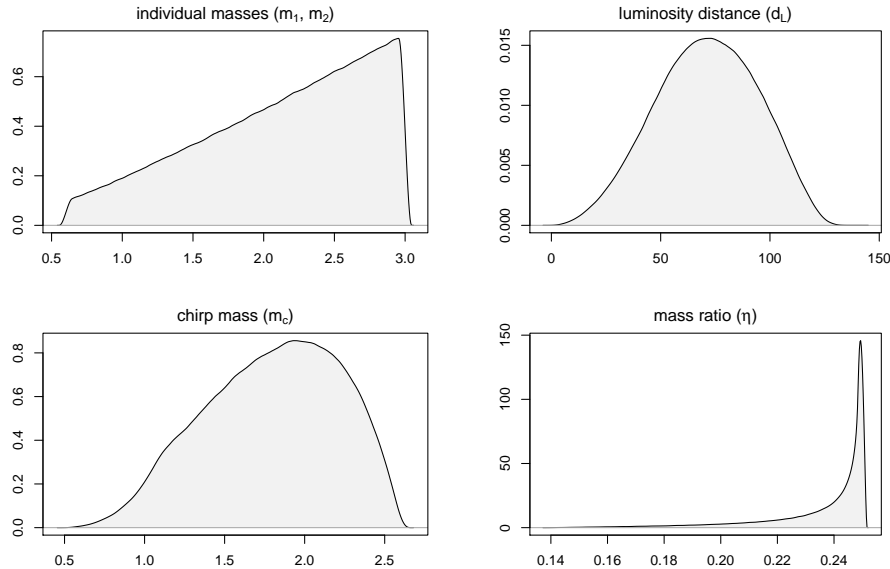


Figure 1. Marginal *prior* probability density estimates for masses and luminosity distance. The upper left plot shows the density of individual masses (m_1, m_2), the lower plots show densities in the alternative mass parametrisation of chirp mass (m_c) and mass ratio (η).

more likely to be observed since low-mass inspirals need to be close to be detectable, while high-mass inspirals may as well originate from greater distances: a feature that is also known as the *Malmquist effect* [37]. Note that the marginal prior for d_L extends well beyond 95 Mpc, because the thresholds x_L and x_U were specified for 2-2 M_\odot binary neutron star inspirals; inspirals involving greater masses still generate signals of observable amplitudes when happening further out.

2.3. Technical details

The MCMC sampler is implemented in C. Data is imported from the *Frame File format* using the *Frame Library* [38]. Before transformation, the data are low-pass-filtered and downsampled (from 16 384 to 4096 Hz) [39]; required Filter coefficients are determined by the ‘Parks-McClellan’-algorithm (see e.g. [40]), again using existing software [41]. The Discrete Fourier Transform of the data (for both likelihood computation and spectral density estimation) is carried out using the free *FFTW* library [42]. The noise spectrum is estimated from a data section that is disjoint from the actually investigated data [43]. (Pseudo-) random numbers are generated using the *Randlib* library [44].

The frequency range over which the likelihood is computed (see equation (5)) was

set to 40–1800 Hz. In the original mass parametrisation of individual masses m_1 and m_2 the two masses showed high posterior correlation, making sampling from the posterior difficult. In order to improve sampling we expressed the likelihood in terms of *chirp mass* (m_c) and *mass ratio* (η) instead, where $m_c = \frac{(m_1 m_2)^{0.6}}{(m_1 + m_2)^{0.2}}$ and $\eta = \frac{m_1 m_2}{(m_1 + m_2)^2}$. In addition we reparameterised the luminosity distance from d_L to $\log(d_L)$, thus implicitly yielding an unbounded parameter space and proposal step widths that are proportional to the actual distance (d_L).

Approximate posterior samples as starting values for the MCMC chains are generated using *importance resampling*, by first generating a large sample from a distribution covering the whole prior, and then drawing the actual sample out of these with correspondingly assigned weights depending on the posterior density [31, 45]. The Metropolis sampler’s proposal distribution was chosen to be a Multivariate Student-t distribution. The t-distribution has a similar shape to the Normal distribution, but possesses an additional degrees-of-freedom parameter ν , and for $\nu \rightarrow \infty$ approaches a Normal distribution. It has ‘heavier tails’ than a Normal distribution, which means that extreme values are more likely under a t-distribution than they were under a Normal, making it a more robust choice as a proposal distribution [31]. The degrees of freedom were set to $\nu := 3$, the lowest possible integer value for which the distribution’s variance is finite. Starting off from an initial setting, the covariance parameter of the proposal distribution is recursively adapted to the sample covariance of generated samples during an initial burn-in phase [46, 47]. The scale of the proposal covariance is set to $\frac{1}{10}$ of the sample covariance estimate, yielding a reasonable acceptance rate for the sampler.

Convergence of the Markov chains is monitored by using several chains that are run simultaneously from different starting points, so one may e.g. verify whether these eventually end up in the same mode [48]. In order to minimise the effect of correlated draws from the MCMC output, only every 50th MCMC draw is written to a file which is then imported into R, a statistical software, for eventual analysis [49]. The C code generates some 2500 MCMC samples per minute on a 3.2 GHz Pentium 4 desktop PC.

The univariate marginal densities that are shown in this paper are *kernel density estimates* (for more details about these see e.g. [50]). The bivariate densities (Fig. 3) are 2-dimensional histograms, the greyscale plots showing relative densities normalised to the mode; the actual density values are not relevant here.

3. Example application with simulated data

We illustrate the results of a run of the MCMC sampler on simulated data for which the true parameter values are known. Six parallel chains were run; the starting points of the chains were generated by importance resampling of 100 000 draws, a number that proved to yield enough eventual draws that were sufficiently close to the main posterior mode to ensure reliable and fast convergence of the Metropolis algorithm. The first 30 000 iterations of each chain were considered the burn-in-phase, during which the iterations 15 000–30 000 were used to tune the proposal covariance. The code was then run for

2 million iterations in total, which after thinning out of the samples (every 50th was kept) and discarding the burn-in yielded a sample of size 236 400 from the posterior. The ‘multivariate potential scale reduction factor’ \hat{R}^p was close to 1 ($\hat{R}^p = 1.0034$), indicating convergence of all chains [51].

Figures 2–5 illustrate estimates of *marginal* posterior densities, that is, densities with respect to the concerned parameters, with all other parameters integrated out. Firstly, Fig. 2 shows posterior densities for the five individual parameters. Most of them exhibit a mode near the true parameter value (indicated by dashed lines). One

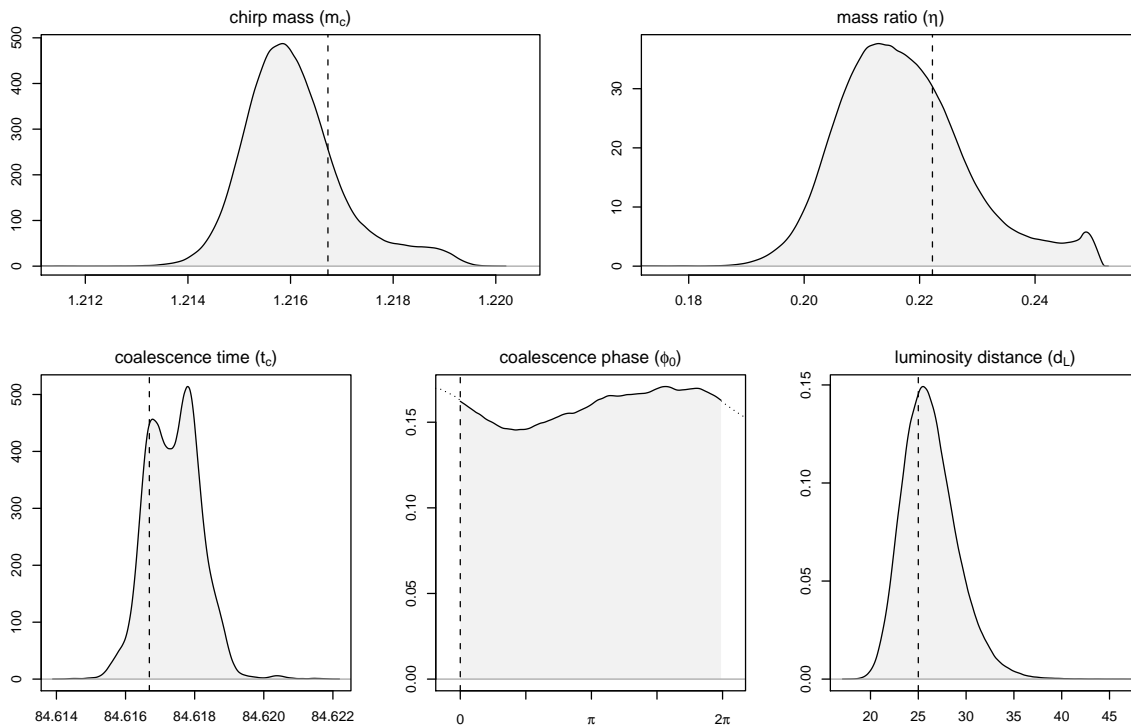


Figure 2. Marginal posterior densities of the five parameters. Dashed lines indicate true parameters.

can see that the relative precision of parameter estimation varies significantly between different parameters. For example, the posterior of the chirp mass covers a range of about $0.01 M_{\odot}$, while the prior range initially was some $2.1 M_{\odot}$ (cp. Fig. 1). The coalescence phase’s posterior, on the other hand, still covers the complete prior domain.

Fig. 3 allows for some insight into joint posterior distributions of some of the parameters. The joint density of chirp mass and mass ratio (3a) shows a positive correlation between the two parameters. Fig. 3b shows interaction between two parameters (ϕ_0 and η), and in particular demonstrates that although the marginal density of ϕ_0 alone is almost uniform (see Fig. 2), this does not imply that its effect on the posterior was negligible.

The MCMC sampler internally works with chirp mass (m_c) and mass ratio (η) instead of individual masses (m_1, m_2). A posterior sample of the individual masses still




Figure 3. Bivariate marginal posterior densities for two pairs of parameters. Dashed lines indicate the true values (the true coalescence phase is $\phi_0 = 0$).

can easily be obtained by back-transforming each pair of (m_c, η) samples. Fig. 4 shows these two marginal densities combined into one plot.




Figure 4. Combined plot of marginal posterior densities of the two companions' individual masses (w.l.o.g. assuming $m_1 \leq m_2$). Dashed lines indicate true parameters.

Analogously, other functions of the parameters can be derived and distributional features investigated; if e.g. one was interested in whether the masses differ ‘significantly’ or are ‘almost equal’, we can estimate: $P(m_2 > 3m_1) = 0.11\%$ or $P(m_2 < 1.5m_1) = 4.84\%$. Fig. 5 shows the posterior distribution of the logarithmic amplitude $\alpha(m_1, m_2, d_L)$ (8). Comparing it to the prior density you can see that, since it is significantly above the reference points x_U and x_L , the particular specification of the lower bound of the parameter space does not affect our conclusions.




Figure 5. Posterior density of the signal’s logarithmic amplitude $\alpha(m_1, m_2, d_L)$; the dashed line indicates the true value. The prior density (dotted line) and x_U and x_L are shown as well.

Table 1 shows numerical estimates of the posterior distributions of the inspiral’s parameters.

Table 1. Posterior estimates: Means, medians and 95% central posterior intervals for several parameters.

parameter	mean	median	95% c.p.i.	true	unit
chirp mass (m_c)	1.2161	1.2159	[1.2145, 1.2186]	1.2167	M_\odot
mass ratio (η)	0.2174	0.2162	[0.1987, 0.2457]	0.2222	
coalescence time (t_c)	84.6174	84.6174	[84.6160, 84.6189]	84.6167	s
coalescence phase (ϕ_0)	— <i>not meaningful</i> —			0.0	radian
luminosity distance (d_L)	26.28	25.99	[21.55, 32.68]	25.00	Mpc
mass 1 (m_1)	0.980	0.964	[0.876, 1.229]	1.0	M_\odot
mass 2 (m_2)	2.062	2.085	[1.600, 2.327]	2.0	M_\odot

4. Discussion

We have developed an efficient method for examining gravitational wave interferometer data for binary neutron star inspiral signals. Our MCMC code searches for signals described by five parameters (two masses, effective distance, time, and phase at coalescence). This code can be applied at the end of a binary neutron star detection pipeline [16, 18, 19], thereby producing parameter estimates and summary statistics.

We are extending our MCMC research on binary neutron star inspiral signals. Active research is in progress for more complex signals from massive black hole systems, and binaries with significant amounts of spin for each mass. Another on-going research

project is in applying MCMC methods to the multiple detector binary neutron star inspiral problem, where sky position parameters can also be estimated. We expect that MCMC methods will prove useful with more complex binary inspiral scenarios.

Acknowledgments

This work was supported by the The Royal Society of New Zealand Marsden Fund grant UOA-204 and National Science Foundation grant PHY-0244357.

References

- [1] B. Abbott et al. Detector description and performance for the first coincidence observations between LIGO and GEO. *Nuclear Instruments and Methods in Physics Research A*, 517:154–179, January 2004.
- [2] F. Acernese et al. Status of VIRGO. *Classical and Quantum Gravity*, 22(18):S869–S880, September 2005.
- [3] M. Ando et al. Current status of the TAMA300 gravitational-wave detector. *Classical and Quantum Gravity*, 22(18):S881–S889, September 2005.
- [4] M. Hewitson. Preparing GEO 600 for gravitational wave astronomy—a status report. *Classical and Quantum Gravity*, 22(18):S891–S900, September 2005.
- [5] A. Einstein. Näherungsweise Integration der Feldgleichungen der Gravitation. *Sitzungsberichte der Preußischen Akademie der Wissenschaften*, pages 688–696, June 1916.
- [6] J. H. Taylor and J. M. Weisberg. Further experimental tests of relativistic gravity using the binary pulsar PSR1913+16. *The Astrophysical Journal*, 345:434–450, October 1989.
- [7] K. S. Thorne. Gravitational radiation. In S. W. Hawking and W. Israel, editors, *300 years of gravitation*, chapter 9, pages 330–358. Cambridge University Press, Cambridge, 1987.
- [8] C. Cutler et al. The last three minutes: Issues in gravitational-wave measurements of coalescing compact binaries. *Physical Review Letters*, 70(20):2984–2987, May 1993.
- [9] S. A. Hughes. Tuning gravitational-wave detector networks to measure compact binary mergers. *Physical Review D*, 66(10), November 2002.
- [10] B. F. Schutz. Determining the Hubble constant from gravitational wave observations. *Nature*, 323:310–311, September 1986.
- [11] D. Marković. Possibility of determining cosmological parameters from measurements of gravitational waves emitted by coalescing compact binaries. *Physical Review D*, 48(10):4738–4756, November 1993.
- [12] C. Cutler and É. É. Flanagan. Gravitational waves from merging compact binaries: How accurately can one extract the binary’s parameters from the inspiral waveform? *Physical Review D*, 49(6):2658–2697, March 1994.
- [13] L. S. Finn. Binary inspiral, gravitational radiation, and cosmology. *Physical Review D*, 53(6):2878–2894, March 1996.
- [14] É. É. Flanagan and S. A. Hughes. Measuring gravitational waves from binary black hole coalescences. I. Signal to noise for inspiral, merger, and ringdown. *Physical Review D*, 57(8):4535–4565, April 1998.
- [15] É. É. Flanagan and S. A. Hughes. Measuring gravitational waves from binary black hole coalescences. II. The waves’ information and its extraction, with and without templates. *Physical Review D*, 57(8):4566–4587, April 1998.
- [16] B. Abbott et al. Analysis of LIGO data for gravitational waves from binary neutron stars. *Physical Review D*, 69(12), June 2004.

- [17] B. Abbott et al. Search for gravitational waves from galactic and extra-galactic binary neutron stars. *Physical Review D*, 72(8), October 2005.
- [18] F. Marion et al. Gravitational waves and experimental gravity. In *Proceedings of the Rencontres de Moriond 2003*, 2004.
- [19] P. Amico et al. A parallel Beowulf-based system for the detection of gravitational waves in interferometric detectors. *Computer Physics Communications*, 153(2):179–189, June 2003.
- [20] S. Geman and D. Geman. Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6:721–741, 1984.
- [21] T. J. Loredo. The promise of Bayesian inference for astrophysics. In E. D. Feigelson and G. J. Babu, editors, *Statistical challenges in modern astronomy*, chapter 12, pages 275–297. Springer-Verlag, New York, 1992.
- [22] W. R. Gilks, S. Richardson, and D. J. Spiegelhalter. *Markov chain Monte Carlo in practice*. Chapman & Hall / CRC, Boca Raton, 1996.
- [23] N. Christensen and R. Meyer. Bayesian methods for cosmological parameter estimation from cosmic microwave background measurements. *Astrophysics*, 06(401), 2000. url: <http://arxiv.org/abs/astro-ph/0006401> (preprint).
- [24] N. Christensen, R. Meyer, L. Knox, and B. Luey. Bayesian methods for cosmological parameter estimation from cosmic microwave background measurements. *Classical and Quantum Gravity*, 18(14):2677–2688, July 2001.
- [25] L. Knox, N. Christensen, and C. Skordis. The age of the universe and the cosmological constant determined from cosmic microwave background anisotropy measurements. *The Astrophysical Journal Letters*, 563(2):L95–L98, December 2001.
- [26] J. A. Rubiño-Martín et al. First results from the Very Small Array—IV. cosmological parameter estimation. *Monthly Notices of the Royal Astronomical Society*, 341(4):1084–1092, June 2003.
- [27] J. Dunkley et al. Fast and reliable Markov chain Monte Carlo technique for cosmological parameter estimation. *Monthly Notices of the Royal Astronomical Society*, 356(3):925–936, January 2005.
- [28] D. N. Spergel et al. First-year Wilkinson Microwave Anisotropy Probe (WMAP) observations: Parameter estimation methodology. *The Astrophysical Journal Supplement Series*, 148(1):175–194, September 2003.
- [29] L. Verde et al. First-year Wilkinson Microwave Anisotropy Probe (WMAP) observations: Determination of cosmological parameters. *The Astrophysical Journal Supplement Series*, 148(1):195–211, September 2003.
- [30] N. Christensen, R. Meyer, and A. Libson. A Metropolis-Hastings routine for estimating parameters from compact binary inspiral events with laser interferometric gravitational radiation data. *Classical and Quantum Gravity*, 21(1):317–330, January 2004.
- [31] A. Gelman, J. B. Carlin, H. Stern, and D. B. Rubin. *Bayesian data analysis*. Chapman & Hall / CRC, Boca Raton, 1997.
- [32] L. Le Cam. On some asymptotic properties of maximum likelihood estimates and related Bayes' estimates. *University of California Publications in Statistics*, 1(11):277–329, January 1953.
- [33] L. S. Finn. Issues in gravitational wave data analysis. *General Relativity and Quantum Cosmology*, 09(077), September 1997. url: <http://arxiv.org/abs/gr-qc/9709077> (preprint).
- [34] L. S. Finn and D. F. Chernoff. Observing binary inspiral in gravitational radiation: One interferometer. *Physical Review D*, 47(6), March 1993.
- [35] N. Christensen and R. Meyer. Using Markov Chain Monte Carlo methods for estimating parameters with gravitational radiation data. *Physical Review D*, 64(2), July 2001.
- [36] T. Tanaka and H. Tagoshi. Use of new coordinates for the template space in a hierarchical search for gravitational waves from inspiraling binaries. *Physical Review D*, 62(8), September 2000.
- [37] A. Sandage. Malmquist bias and completeness limits. In P. Murdin, editor, *Encyclopedia of Astronomy and Astrophysics*. Institute of Physics Publishing, Bristol, 2001. url: <http://eaa.iop.org>.
- [38] Frame library (Fr), version v6r19, May 2005. url: <http://lappweb.in2p3.fr/virgo/FrameL>.

- [39] R. E. Crochiere. A general program to perform sampling rate conversion of data by rational ratios. In A. C. Schell et al., editors, *Programs for digital signal processing*, chapter 8.2. IEEE Press, New York, 1979.
- [40] E. C. Ifeachor and B. W. Jervis. *Digital signal processing: A practical approach*, chapter 6, pages 285–309. Addison-Wesley, Wokingham, England, 1993.
- [41] J. Janovetz. Parks-McClellan algorithm for FIR filter design (C version), 1998. url: <http://www.janovetz.com/jake>.
- [42] M. Frigo and S. G. Johnson. FFTW 3.0.1, a C subroutine library for computing the discrete Fourier Transform (DFT), 2003. url: <http://www.fftw.org>.
- [43] P. D. Welch. The use of Fast Fourier Transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms. *IEEE Transactions on Audio and Electroacoustics*, AU-15(2):70–73, June 1967.
- [44] B. W. Brown, J. Lovato, K. Russell, and J. Venier. Randlib - Library of C routines for random number generation, 1997. url: <http://biostatistics.mdanderson.org/SoftwareDownload>.
- [45] A. F. M. Smith and A. E. Gelfand. Bayesian statistics without tears: A sampling-resampling perspective. *The American Statistician*, 46(2):84–88, May 1992.
- [46] B. P. Welford. Note on a method for calculating corrected sums of squares and products. *Technometrics*, 4(3):419–420, August 1962.
- [47] T. F. Chan, G. H. Golub, and R. J. Le Veque. Algorithms for computing the sample variance: Analysis and recommendations. *The American Statistician*, 37(3):242–247, August 1983.
- [48] A. Gelman and D. B. Rubin. Inference from interative simulation using multiple sequences. *Statistical Science*, 7(4):457–472, March 1992.
- [49] R Development Core Team. *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria, 2004. url: <http://www.R-project.org>.
- [50] D. W. Scott. *Multivariate density estimation: theory, practice and visualization*. Wiley & Sons, New York, 1992.
- [51] S. P. Brooks and A. Gelman. General methods for monitoring convergence of iterative simulations. *Journal of Computational and Graphical Statistics*, 7(4):434–455, December 1998.