

Coherent Bayesian inference on compact binary inspirals using a network of interferometric gravitational wave detectors

Christian Röver and Renate Meyer

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

Nelson Christensen

Physics and Astronomy, Carleton College, Northfield, Minnesota, USA

(Received 27 September 2006; published 30 March 2007)

Presented in this paper is the description of a Markov chain Monte Carlo (MCMC) routine for conducting coherent parameter estimation for interferometric gravitational wave observations of an inspiral of binary compact objects using multiple detectors. Data from several interferometers are processed, and all nine parameters (ignoring spin) associated with the binary system are inferred, including the distance to the source, the masses, and the location on the sky. The data is matched with time-domain inspiral templates that are 2.5 post-Newtonian (PN) in phase and 2.0 PN in amplitude. We designed and tuned an MCMC sampler so that it is able to efficiently find the posterior mode(s) in the parameter space and perform the stochastic integration necessary for inference within a Bayesian framework. Our routine could be implemented as part of an inspiral detection pipeline for a world-wide network of detectors. Examples are given for simulated signals and data as seen by the LIGO and Virgo detectors operating at their design sensitivity.

DOI: [10.1103/PhysRevD.75.062004](https://doi.org/10.1103/PhysRevD.75.062004)

PACS numbers: 04.80.Nn, 02.70.Uu, 05.45.Tp, 07.05.Kf

I. INTRODUCTION

The era of gravitational wave astronomy is now close upon us as numerous interferometric detectors are operating. The Laser Interferometer Gravitational Wave Observatory (LIGO) [1–3] has now reached its target sensitivity, and there is the hope that a detection could come at any time [4]. Around the globe a world-wide network of detectors is coming online; Virgo in Italy [5–7], GEO in Germany [8,9], and TAMA in Japan [10,11] are operating alongside LIGO in the quest for gravity wave detection. These ground based laser interferometers are sensitive to gravitational radiation in the frequency band from 40 Hz up to 8 kHz.

Coalescing binaries containing neutron stars or black holes promise to be one of the cleanest and most probable sources of detectable radiation [12]. Observation of inspiral events could provide important information on the structure of neutron stars [13,14]. Even cosmological information can be extracted from the observation of inspiral events [15–19]. 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 [18,20,21]. The LIGO Scientific Collaboration (LSC) has been actively searching for binary inspiral events [22,23], as well as conducting searches in coincidence with TAMA [24]. Using the data from LIGO's S2 run, it was possible to set an upper limit on the neutron star coalescence rate of less than 50 per year per Milky Way equivalent galaxy [23]. The LSC has also conducted searches for binary inspiral signals from primordial black holes ($0.2\text{--}1.0M_{\odot}$) in the halo of our galaxy [25], plus more

massive black hole systems where component masses are in the $3\text{--}20M_{\odot}$ range [26].

Bayesian inferential methods provide a means to use data from interferometric gravitational wave detectors in order to extract the parameters of a binary inspiral event. Markov chain Monte Carlo (MCMC) methods are a powerful computation technique for parameter estimation within this framework; they are especially useful in applications where the number of parameters is large [27]. Nice descriptions of the positive aspects of a Bayesian analysis of scientific and astrophysical data are provided in [28–30]. In previous work we have developed MCMC routines for extracting five parameters associated with a binary inspiral event from data generated by a single interferometric detector [31–33]. Our MCMC code took data from a single interferometer, Fourier transformed it into the frequency domain, and then compared the result with 2.0 post-Newtonian (PN) stationary phase templates [34]. One of the new methods that we implement in this current study, presented in this paper, is an MCMC routine that takes time-domain interferometric data, and compares it to time-domain templates that are 2.5 PN in phase, and 2.0 PN in amplitude; a trivial modification of the code (though not implemented in the study presented here) is to extend the accuracy of the signal waveforms to 3.5 PN in phase and 2.5 PN in amplitude [35–39]. A critical task for a world-wide gravity wave detection network will be to not only detect a binary inspiral signal, but to say where it came from. For this purpose the LSC has developed a *coherent* binary inspiral search pipeline [40–42]. Coherent binary search pipelines and methods are also being developed within the Virgo collaboration [43] and TAMA [44]. Along similar lines, we have developed a coherent

MCMC parameter estimation routine, and in this present paper we describe it and provide results from a test on simulated data. The simulations involve binary neutron star inspirals observed by three well-separated interferometers: the 4 km LIGO detectors at Hanford, WA and Livingston, LA, plus the 3 km Virgo detector in Cascina, Pisa, Italy. The synthetic data for the LIGO and Virgo detectors has Gaussian noise with power spectral densities (PSD) that match their target sensitivities [45,46]. The MCMC code takes data from several interferometers, and estimates the two individual masses, time and phase at coalescence, distance to source, gravity wave polarization angle, angle of inclination of the binary system’s orbital plane, and sky position in right ascension and declination. The additional parameters of polarization, inclination angle, and sky location can only be estimated accurately when data from more than one interferometer are considered; they also greatly inflate the parameter space and therefore complicate the analysis. MCMC methods have also been tested in a similar context to recover the nine parameters of a binary black hole coalescence in LISA data [47]. However, the problem setting is different (due to the different instrument, longer observation period, lower frequencies being investigated, and the 2.0 PN phase model), and MCMC techniques were employed rather for optimization than for integration; the MCMC was also only applied on a subset of the parameters, while others were solved for analytically.

The organization of this paper is as follows. After a brief introduction to the analysis problem we describe our approach alongside more detail on the applied model in Sec. II. Section III provides practical directions how we implemented MCMC methods in order to analyze data within the described framework. Section IV eventually illustrates results of applications of our code to simulated data. We conclude the paper with a discussion and outlook in Sec. V.

II. ANALYSIS STRATEGY AND MODELLING

A. Measuring gravitational waves

In an inspiralling binary system, the two companions orbit around their center of mass with decreasing orbital distance and period, until the system eventually collapses. The gravitational radiation emitted by the system exhibits a “chirp” form, that is, an oscillation of increasing frequency and amplitude. A laser interferometer is sensitive to space distortions along the directions of its two orthogonal arms, as it monitors the phase difference between the two laser beams that travel along the arms. A gravitational wave is a quadrupole wave that is characterized by its direction of travel, polarization angle, and its two polarization amplitudes. Its effect on a laser interferometer’s measurements then is a linear combination of the effects associated with the two polarizations, depending on the orientation of the interferometer with respect to the passing

wave. Actual measurements, of course, are also exposed to noise.

Measurements of a binary inspiral’s chirp signal by a single interferometer will not be sufficient to infer all of the parameters that determine the signal’s waveform and the interferometer’s response. Measurements from several separate interferometers will, in general, be required to derive (for example) the wave’s direction of travel by matching possible mutual orientations as well as different arrival times of the signals at the different sites. Combining measurements from several interferometers will also enhance sensitivity and signal-to-noise ratio [40].

B. The Bayesian approach

We apply a *Bayesian* approach to this inferential problem, that is, the term “*probability*” is used in a broader sense than in the more common “*frequentist*” interpretation [30,48,49]. Probability calculus here is applied to process and infer *states of incomplete information* that are reflected by probability distributions, and that are conditional on prior knowledge and/or the data at hand. This allows one to treat unknown parameters as random variables that follow a *prior distribution* representing the researcher’s pre-experimental knowledge and uncertainty. The gain in information through observation of data then follows in a straight-forward fashion through the application of Bayes’ theorem, yielding the *posterior distribution* of the parameters. The posterior distribution, which is essentially the product of the parameters’ prior distribution and the likelihood of the data, then poses the basis for inference [50].

Inference through the posterior distribution usually involves the solution of integrals, since one is typically interested in figures such as the marginal (posterior) expectations of individual parameters, marginal (posterior) densities, or (posterior) probabilities of certain events, which are derived from the posterior distribution by integrating over the parameter space. In many cases when analytic integration is not possible, numerical methods are employed, usually (pseudo-) stochastic techniques like Markov chain Monte Carlo (MCMC) methods that simulate random draws from the posterior distribution, then allowing one to approximate the desired integrals by sample statistics [27,50].

C. Parameters

The waveform that is *measured* at a certain interferometer depends on the characteristics of the inspiral event as well as the orientation of the source relative to the interferometer. The nine “global” parameters determining the response of Earth-bound interferometers are:

- individual masses ($m_1, m_2 \in \mathbb{R}^+$; $m_1 \leq m_2$),
- luminosity distance ($d_L \in \mathbb{R}^+$),
- inclination angle ($\iota \in [0, \pi]$),
- coalescence phase ($\phi_0 \in [0, 2\pi]$),

coalescence time at geocenter ($t_c \in \mathbb{R}$),
 declination ($\delta \in [-\frac{\pi}{2}, \frac{\pi}{2}]$),
 right ascension ($\alpha \in [0, 2\pi]$) and
 polarization ($\psi \in [0, \pi]$),

the latter four of which affect the measurement at the I -th detector in terms of the “local” parameters

local coalescence time ($t_c^{(I)} \in \mathbb{R}$),
 altitude ($\vartheta^{(I)} \in [0, \pi]$),
 azimuth ($\varphi^{(I)} \in [0, 2\pi]$) and
 polarization ($\psi^{(I)} \in [0, \pi]$).

These “local” parameters are derived from the locations/orientations of the source and the individual interferometers with respect to each other. For more specific definitions and conventions see e.g. [51]. In the following we will refer to the two parameter sets as the *global* parameter vector

$$\theta^{\oplus} = (m_1, m_2, d_L, \iota, \phi_0, t_c, \delta, \alpha, \psi), \quad (1a)$$

and the *local* parameter vector

$$\theta^{(I)} = (m_1, m_2, d_L, \iota, \phi_0, t_c^{(I)}, \vartheta^{(I)}, \varphi^{(I)}, \psi^{(I)}) \quad (1b)$$

with respect to a specific interferometer I . Not all of the above parameters will usually be of primary interest; especially coalescence phase ϕ_0 , polarization ψ or inclination ι might be regarded as nuisance parameters.

D. Network likelihood

An interferometer’s data output z is assumed to be the sum of the actual signal $s(\theta)$, depending on the true parameters θ , and (interferometer-specific) colored noise. The (real-valued) data z and signal waveform $s(\theta)$ enter the likelihood function in terms of their (complex-valued) Fourier transforms \tilde{z} and $\tilde{s}(\theta)$, the noise is specified through its power spectral density (PSD) S_n . The likelihood function for a specific interferometer I then is (up to a normalizing constant) proportional to the following expression

$$\mathcal{L}^{(I)}(\theta^{(I)}) \propto \exp\left(-2 \int_0^\infty \frac{|\tilde{z}(f) - \tilde{s}(f, \theta^{(I)})|^2}{S_n(f)} df\right) \quad (2)$$

[52]. For actual data, discretized and measured over a finite interval of length δ_t , it is computed as the sum of squared and normalized differences between the Fourier transforms of the observed signal (\tilde{z}) and the signal template ($\tilde{s}(\theta)$) over the discrete set of Fourier frequencies $\{(i \times \Delta_f): i_L \leq i \leq i_U\}$

$$\mathcal{L}^{(I)}(\theta^{(I)}) \propto \exp\left(\frac{-2}{\delta_t} \sum_{i=i_L}^{i_U} \frac{\overbrace{|\tilde{z}(i \times \Delta_f) - \tilde{s}(i \times \Delta_f, \theta^{(I)})|^2}^{\text{data}}}{\underbrace{S_n(i \times \Delta_f)}_{\text{noise PSD}}}\right) \quad (3)$$

where $i_L \times \Delta_f$ and $i_U \times \Delta_f$ are the lower and upper bounds of the examined frequency range. Note that,

although not labeled as such here, data z , noise spectrum S_n , etc. are specific for the different interferometers I . Assuming that noise is independent across different interferometers, the network likelihood then is the product of the individual interferometer likelihoods

$$\mathcal{L}(\theta^{\oplus}) = \prod_I \mathcal{L}^{(I)}(\theta^{(I)}). \quad (4)$$

E. Prior specification

The prior information is specified in a straight-forward fashion for the “geometrical” parameters that determine the location and orientation of the inspiral event. *A priori*, the event is assumed to be equally likely across all possible directions; this leads to a uniform prior for the right ascension α , and a prior density

$$f(\delta) = \frac{1}{2} \cos(\delta) \quad (5)$$

that is proportional to the circumference of the corresponding “circle of latitude” for the declination δ . Analogously, the prior density of the inclination ι is

$$f(\iota) = \frac{1}{2} \sin(\iota), \quad (6)$$

and the remaining parameters, polarization ψ and coalescence phase ϕ_0 , have uniform priors. The prior specifications for these parameters may also be regarded as *Maximum Entropy* choices [49,53].

The coalescence time t_c is assumed to be known in advance up to a certain accuracy through preprocessing of the data [22,54,55]; for now we set its prior to be uniform across ± 5 ms around the true value, which of course is known for simulated data.

The joint prior of the remaining parameters, masses m_1, m_2 and luminosity distance d_L , is set in order to reflect the distribution of parameters *given* the event has been detected in the first place. Initially, the prior for the two inspiral companions’ individual masses is assumed to be uniform between 0.6 and $3.0M_\odot$ (solar masses: $M_\odot \approx 2 \times 10^{30}$ kg), which effectively covers the range of values expected for neutron stars. 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). But inspiral signals obviously cannot originate from arbitrarily great distances, since at some point their signals become too faint to be detected. We incorporated this restriction by taking into consideration the *detection probability* D , which we assume to depend on the signal’s *amplitude*, which is roughly proportional to

$$\mathcal{A}(m_1, m_2, d_L) = \ln\left(\frac{\sqrt{m_1 m_2}}{d_L(m_1 + m_2)^{1/6}}\right), \quad (7)$$

neglecting for simplicity the effects of orientation parameters (\mathcal{A} is actually the logarithmic amplitude) [33]. We could have set a threshold amplitude below which inspirals

would be assumed undetectable, but favored a “smoother” transition that does not explicitly apply zero probability to parts of the parameter space. Instead we model the dependence between signal amplitude and detection probability using a (sigmoidal) logistic function of the form

$$D_{a,b}(x) = \frac{1}{1 + \exp(\frac{x-a}{b})} \quad (8)$$

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

$$a := \frac{x_L + x_U}{2} \quad \text{and} \quad b := \frac{x_U - x_L}{2 \ln(\frac{1-p}{p})}. \quad (9)$$

So the density of the resulting (proper) prior distribution of individual masses and distance is

$$f(m_1, m_2, d_L) = c \times d_L^2 \times D_{a,b}(\mathcal{A}(m_1, m_2, d_L)) \quad (10)$$

for some normalizing constant $c \in \mathbb{R}^+$ [33]. For the examples shown here, we set $x_U := \mathcal{A}(2.0M_\odot, 2.0M_\odot, 45 \text{ Mpc})$, $x_L := \mathcal{A}(2.0M_\odot, 2.0M_\odot, 50 \text{ Mpc})$ and $p := 0.1$, so an optimally oriented 2.0–2.0 M_\odot inspiral is assumed to be detectable out to 45 and 50 Mpc with probabilities of 90% and 10%, respectively.

As a “side effect” of this prior definition, larger masses have a greater prior probability, since inspirals involving large masses may originate from farther distances while low-mass inspirals need to be close in order to be observable at all. This feature is also known as the *Malmquist effect*; incorporating it into the prior will compensate for selection bias that would otherwise affect the results [56,57]. The definition of priors, especially for coalescence time t_c , individual masses m_1 and m_2 and their relation to the luminosity distance d_L , and possibly also for the sky location (δ, α) , may be refined at a later stage when e.g. some substantiated knowledge is available about the performance of the upstream detection pipeline, which might provide rough estimates of some of the parameters together with the detection [22,54,55]. For now we aim for simple and general formulations.

III. IMPLEMENTATION

A. General

In order to analyze data in terms of the above framework, we implemented an MCMC sampler in *C* that is supposed to both *find* the global mode(s) of the posterior distribution and then “*explore*” the distribution, i.e. simulate random draws from the posterior. Data is imported from the *Frame format* using the *Frame Library* [58]. Prior to the analysis, the data is filtered and downsampled by a

factor of 4 [59,60]. (Pseudo-) random number generation within the MCMC sampler was implemented using *Randlib* [61].

The MCMC sampler writes only every 25th of the drawn samples to a text file, in order to reduce the effects of subsequently correlated samples, and also to keep the data volume at a reasonable level. The MCMC output then is imported into *R*, a statistical software, for eventual analysis [62].

The marginal densities that are shown in this paper are *kernel density estimates* [63]. Two-dimensional densities are illustrated by greyscale plots (with darker areas corresponding to greater densities), and in addition the contour line enclosing the most probable region (accumulating 95% probability) is shown.

B. Likelihood implementation

In order to compute the coherent network-likelihood, first the individual interferometer-likelihoods need to be determined. The primary “ingredients” for the interferometer-likelihood are

the Fourier-transform of the data \tilde{z} ,

the noise spectrum S_n ,

the “local” parameter set $\theta^{(l)}$ and

the (Fourier-transformed) signal template $\tilde{s}(\theta^{(l)})$,

the first two of which only need to be determined once at the very beginning of the analysis, while the latter two (in general) need to be recomputed for each likelihood evaluation.

For all (discrete) Fourier-transformations we use the freely available FFTW library [64]. The noise spectrum is estimated from a section of data that is disjoint from the actually analyzed data set [65]. In order to minimize undesirable leakage effects, the data is “windowed” before Fourier-transformation; using a Hann window for spectrum estimation, and a Tukey window for data and template transformations [66].

Internally, along with the noise spectra, data Fourier transforms etc. corresponding to each of the interferometers I , a set of vectors defining the interferometer’s location and orientation is stored. This allows to derive the interferometer-specific parameters (local coalescence time $t_c^{(l)}$, altitude $\vartheta^{(l)}$, azimuth $\varphi^{(l)}$ and polarization $\psi^{(l)}$) with respect to the galactic and Earth coordinate systems from the global parameter vector θ^\oplus via vector operations like rotations, orthogonal projections, etc. [67–69].

C. Time-domain waveform generation

Template waveforms $\tilde{s}(\theta)$ are generated in the time-domain and then (numerically) Fourier transformed to the frequency-domain. Here we used waveform approximations that are 2.5 PN in phase and 2.0 PN in amplitude. The rather complex expressions for these are omitted here and can be found in [51]. We preferred working in the time-

domain, since frequency-domain templates might introduce discrepancies because they are exact analytic Fourier transforms of the “parametric” waveforms, while the actual data is of finite extent and affected by leakage introduced through the numerical discrete Fourier transformation. When using time-domain templates and Fourier transforming these, the resulting frequency-domain templates are the more accurate match to the Fourier-transformed data. Another advantage of generating templates in the time-domain is the availability of a wider range of signal waveforms; the extension to higher PN approximations (e.g. 3.5 PN phase and 2.5 PN amplitude [35–39]) or the consideration of additional parameters (e.g. spin effects [70]) would be straight forward to implement.

D. MCMC implementation

In order to enhance the MCMC sampler’s performance we applied reparametrizations to some of the parameters. The individual masses (m_1, m_2) are highly correlated in their posterior distribution [18], making sampling from the original parameters extremely difficult. Re-expressing the masses in terms of *chirp mass* m_c and the (symmetric) *mass ratio* η , where

$$m_c = \frac{(m_1 m_2)^{3/5}}{(m_1 + m_2)^{1/5}} \quad \text{and} \quad \eta = \frac{m_1 m_2}{(m_1 + m_2)^2}, \quad (12)$$

yields a posterior that is much easier to sample from. We then reparametrized the luminosity distance from d_L to $\ln(d_L)$, which implicitly yields an unbounded parameter space and proposal step widths that are proportional to the current distance d_L . Declination δ and inclination ι were transformed to $\sin(\delta)$ and $\cos(\iota)$, which leads to uniform prior distributions over the new parameters.

The MCMC algorithm was implemented as a Metropolis-sampler [27,50] that was extended to a parallel tempering algorithm. The idea of tempering is to consider a smoothed (“tempered”) version of the actual objective function (here the posterior distribution), or, analogously, do its exploration (optimization, MCMC sampling,...) following “relaxed” rules. In optimization contexts, the tempering is often faded out over time, the result being a “simulated annealing” algorithm, which starts off at a high temperature in order to find the global mode amongst other minor modes, but eventually ends up optimizing the original objective function. Parallel tempering is a special case of the “Metropolis-coupled MCMC” (MCMCMC) algorithm, in which several “tempered” chains are run in parallel, each having a different (constant) temperature. So the tempering does not vary *over time*, but instead is realized *across parallel chains*, with additional proposals allowing for swapping between chains. Instead of sampling from the regular posterior distribution with density function f (which is essentially the product of prior π and likelihood \mathcal{L} : $f(\theta) \propto \pi(\theta)\mathcal{L}(\theta)$), the tempered chains sam-

ple from a modified distribution

$$f_T(\theta) \propto \pi(\theta)\mathcal{L}(\theta)^{1/T}, \quad (13)$$

where $T \geq 1$ denotes the “temperature”, and for which in the extreme cases $T = 1$ and $T \rightarrow \infty$ the tempered distribution f_T equals posterior and prior, respectively. Chains running at higher temperatures can be considered as sampling from a “relaxed” or “widened” posterior which is then used as proposal distribution (through the swapping between chains) for “cooler” chains, thus improving both convergence and mixing. The draws from the “coolest” chain with $T = 1$ are the only ones that are eventually used for posterior inference.

The starting values for the sampler are determined using *importance resampling* [50]. In a simplified setting of the problem (considering only 5 parameters and one interferometer [33]), this was sufficient to yield reasonable posterior samples that were close enough to the main mode so the sampler then converged reliably and fast. Because of the much larger parameter space and the computationally more expensive likelihood, ensuring convergence through good starting values is not feasible any more. Instead, convergence is now supported through the use of parallel tempering, which enables the sampler to cross gaps between (local) modes and eventually find the posterior’s global mode(s).

As proposal distribution for the MCMC sampler we used a multivariate Student’s t -distribution with 3 degrees of freedom. In addition to these “regular” proposals, sometimes draws from the prior are proposed for some parameters in order to enhance convergence, or steps to “related” parts of the parameter space, like a step from phase ϕ_0 to $\phi_0 \pm \pi$, which corresponds to an (almost) equally likely parameter combination if the two masses are (almost) equal. Proposals like these are valid as long as a certain symmetry is maintained, i.e. every proposed step is as likely as the reverse step; otherwise one would need to switch to a Metropolis-Hastings sampler that is able to account for asymmetric steps [27,50].

E. Signal-to-noise ratios

The signal-to-noise ratios (SNR) stated in subsequent sections are defined as follows. The interferometer-specific SNR of a certain signal $s(\theta^\oplus)$ received at interferometer I and embedded in noise with spectral density S_n is defined as

$$\varrho^{(I)} = 2\sqrt{\int_0^\infty \frac{|\tilde{s}(f, \theta^{(I)})|^2}{S_n(f)} df} \quad (14)$$

[32]. We computed it—in analogy to Eq. (3)—over the same frequency range that is relevant for the likelihood. The network SNR then is defined as

$$\rho_{\text{network}} = \sqrt{\sum_I (\rho^{(I)})^2} \quad (15)$$

[18].

IV. EXAMPLE APPLICATION USING SIMULATED DATA

A. Overview

In the following two sections we will illustrate results generated by our MCMC implementation on simulated data. Firstly, in Sec. IV B the evidence on parameters in the measured data is illustrated for one simulated signal by looking at the results of an MCMC run in detail. In the following Sec. IV C the effect of varying signal properties on the evidence in the data is demonstrated and some peculiarities are pointed out. The posterior distributions of parameters are compared across simulated signals that are observed at different distances (and with that, SNRs), but which are otherwise identical. While the *individual* SNRs (at each interferometer) for each of these examples are similar, a more extreme example with an almost zero SNR at one of the interferometers is considered as well.

B. Recovering the inspiral's parameters

We simulated an inspiral event that is measured at three interferometer sites, namely, Hanford (LHO), Livingston (LLO) and Pisa (V). Because of their different noise characteristics, the frequency ranges of the likelihoods were set to 30–1600 Hz for the Virgo observatory, and 40–1600 Hz for the other two LIGO interferometers (cp. Equation (3)).

The amount of data to consider was set to be the 40 seconds before coalescence for Virgo, and 20 s for the other two. This matches the time an inspiral of this kind spends emitting radiation within the above frequency ranges, and would in a realistic search need to be set either with respect to worst-case considerations, or based on prior information supplied by the detection pipeline. The original sampling rates of the data were 20000 Hz (V) and 16384 Hz (LHO, LLO), and the signals were superimposed with Gaussian noise matching the corresponding interferometer's design sensitivities [45,46]. The example inspiral had parameter values of $d_L = 10$ Mpc for the distance and masses of $m_1 = 1.5M_\odot$ and $m_2 = 2.0M_\odot$ (chirp mass $m_c = 1.505M_\odot$, mass ratio $\eta = 0.245$). The resulting SNRs at the three interferometer sites are: LHO: 16.4, LLO: 21.2 and V: 12.6 (network SNR: 29.6).

Six parallel MCMC chains were run within the parallel tempering scheme. With this amount of data the MCMC code generated some 80 samples per minute on a 3.2 GHz Pentium desktop PC, so considering the parallel chains (six) and the thinned output (every 25th), an actual posterior sample is generated every 113 seconds. The first chain of the parallel tempering MCMC sampler converged after some 75000 iterations, after “thinning out” of the samples and discarding the burn-in phase, the resulting posterior sample size was 12500 samples.

Figure 1 shows marginal posterior densities of the nine parameters for our example problem, and Table 1 lists some numerical posterior estimates. Although correlations between parameters are already greatly reduced through the reparametrization, some correlation still remains. Figure 2

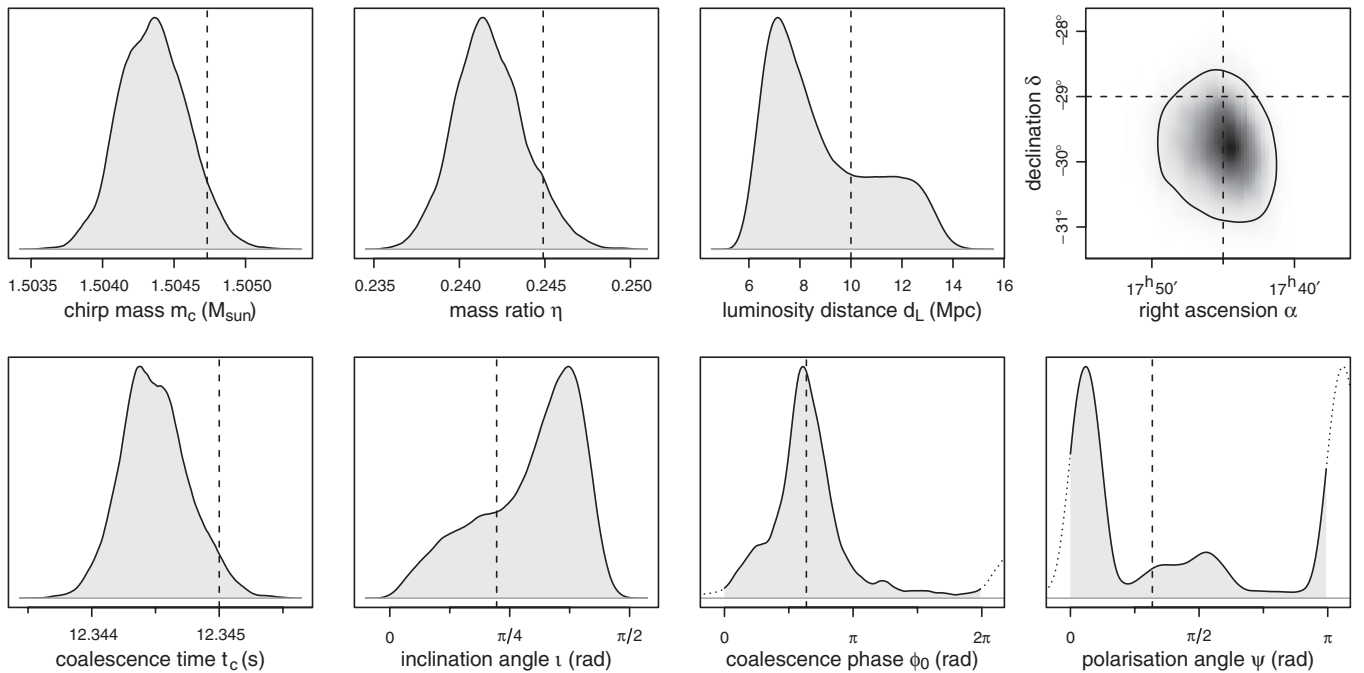


FIG. 1. Marginal posterior densities of the inspiral's parameters for our example problem. Dashed lines indicate the true values.

TABLE I. Some key figures summarizing the marginal posterior distributions of individual parameters, where meaningful. Mean and median characterize the distributions' centers. Given the observed (simulated) data, the parameters fall within the central posterior intervals with 95% probability. The true parameter values used to generate the data are shown as well.

| | Mean | Median | 95% c.p.i. | True | Unit |
|----------|----------------------|----------------------|--|----------------------|-----------|
| m_c | 1.5044 | 1.5044 | [1.5039, 1.5048] | 1.5047 | M_\odot |
| η | 0.2418 | 0.2417 | [0.2380, 0.2460] | 0.2449 | |
| t_c | 12.3445 | 12.3445 | [12.3440, 12.3450] | 12.3450 | s |
| d_L | 8.89 | 8.29 | [6.25, 13.1] | 10.00 | Mpc |
| δ | -29.76° | -29.77° | $[-30.74^\circ, -28.84^\circ]$ | -29.00° | |
| α | $17^{\text{h}}45.0'$ | $17^{\text{h}}44.9'$ | $[17^{\text{h}}42.1', 17^{\text{h}}48.9']$ | $17^{\text{h}}45.0'$ | |
| ι | 0.911 | 0.995 | [0.194, 1.354] | 0.700 | rad |
| m_1 | 1.446 | 1.442 | [1.389, 1.526] | 1.5 | M_\odot |
| m_2 | 2.080 | 2.084 | [1.965, 2.170] | 2.0 | M_\odot |
| m_t | 3.526 | 3.527 | [3.490, 3.559] | 3.5 | M_\odot |




FIG. 2. Marginal joint posterior distributions of two pairs of parameters. Dashed lines indicate the true values.

illustrates correlations between two such pairs of parameters, one can see that the “new” mass parameters m_c and η are still dependent (though orders of magnitude less than m_1 and m_2 were), and also that the uncertainty in the luminosity distance d_L is tied to the uncertainty in the inclination angle ι , since these two parameters can compensate or mimic each other’s effect to a certain degree.

Distributions of other variables derived from the parameters could be investigated, their distributions then depending on the *joint* distribution of the involved parameters [33]. Examples would be the individual masses (m_1 , m_2), or the total mass $m_t = m_1 + m_2$; the distribution of m_t is narrower than those of both m_1 and m_2 , due to their negative correlation (cp. Table I).

C. Results for varying signal characteristics

We performed additional runs with varying settings of the “true” parameters of the simulated signal. As one would expect, the precision of parameter estimation is proportional to the signal’s strength; Table II shows the standard deviations of some of the parameters’ posterior

distributions. The posterior is narrowest for a close-by inspiral of high masses, and gets wider for both lower mass or greater distance.

These results are in agreement with earlier estimates of the accuracy to be expected from such parameter estimates [18]. The great difference in relative accuracies of parameters related to phase evolution (like chirp mass m_c and reduced mass $\mu = \frac{m_1 m_2}{m_1 + m_2} = m_t \eta$) versus those affecting the signal’s amplitude (like distance d_L) is confirmed, and the correlation between m_c and μ is verified as well.

At decreasing SNRs, certain parameters cannot be determined unambiguously any more. One example is the inclination angle ι , which still has a “well-behaving” posterior distribution at 10 Mpc distance (see Fig. 2). For a weaker signal originating from 30 Mpc distance, the distribution then turns bimodal (Fig. 3). The “orientation” of the inclination angle is not clear any more, the result being two roughly equally likely “mirror image” solutions with $P(\iota < \frac{\pi}{2}) \approx \frac{1}{2} \approx P(\iota > \frac{\pi}{2})$. Note that the two solutions ι and $\pi - \iota$ correspond to opposite orbital directions (clockwise/counterclockwise), as seen from Earth, which might be of minor interest anyway.

TABLE II. Individual and total SNRs for different signals, and some characteristics of the resulting posterior distributions. The accuracy of some of the parameters is illustrated by the posterior standard deviations for (δ, α) , t_c , d_L , m_c and μ (percentages refer to the true value). The correlation coefficient for m_c and μ shows the (posterior) interdependence between the two parameters. Our results are consistent with those presented in [18].

| Masses | Distance | Signal-to-noise ratios | | | | Posterior standard deviations | | | | | |
|-------------------|----------|------------------------|------|------|---------|-------------------------------|---------|-------|--------|-------|-------------------|
| | | LHO | LLO | V | Network | $(\delta, \alpha)^a$ | t_c | d_L | m_c | μ | Cor(m_c, μ) |
| 1.5–2.0 M_\odot | 10 Mpc | 16.4 | 21.2 | 12.6 | 29.6 | 0.011 rad | 0.26 ms | 20% | 0.016% | 0.35% | 0.95 |
| 1.5–2.0 M_\odot | 20 Mpc | 8.2 | 10.6 | 6.3 | 14.8 | 0.030 rad | 0.49 ms | 25% | 0.031% | 0.69% | 0.94 |
| 1.5–2.0 M_\odot | 30 Mpc | 5.5 | 7.1 | 4.2 | 9.9 | 0.207 rad | 1.04 ms | 25% | 0.074% | 1.33% | 0.91 |
| 2.0–2.0 M_\odot | 10 Mpc | 18.4 | 23.9 | 14.1 | 33.3 | 0.008 rad | 0.14 ms | 14% | 0.009% | 0.14% | 0.80 |
| 2.0–2.0 M_\odot | 20 Mpc | 9.2 | 11.9 | 7.1 | 16.7 | 0.017 rad | 0.28 ms | 18% | 0.014% | 0.23% | 0.73 |
| 2.0–2.0 M_\odot | 30 Mpc | 6.1 | 8.0 | 4.7 | 11.1 | 0.026 rad | 0.42 ms | 21% | 0.021% | 0.37% | 0.78 |

^aspherical standard deviation [71]

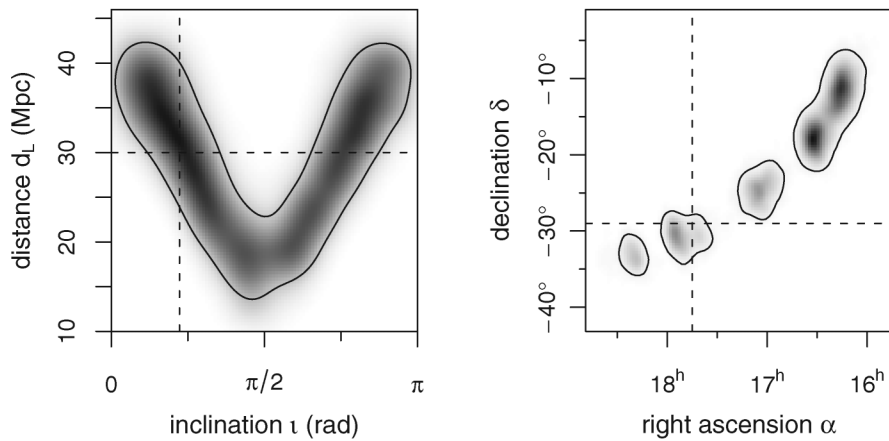


FIG. 3. At greater distance the “orientation” of the inclination angle ι cannot be resolved any more, both directions are roughly equally likely ($P(\iota < \frac{\pi}{2}) \approx \frac{1}{2} \approx P(\iota > \frac{\pi}{2})$). At the same time, with the lower SNR the sky location’s posterior turns multimodal. (Dashed lines indicate the true values.)

The sky location’s posterior also exhibits multiple modes for this weaker signal (Fig. 3). This illustrates some potential pitfalls of Maximum-Likelihood (ML) or Maximum-a-Posteriori (MAP) methods; these would advise picking the highest of the several modes, which might just be the narrowest one, but not necessarily the most likely. If one then proceeded by extrapolating the curvature at that mode and deriving error bounds from the Fisher Information matrix, the resulting estimates might not only be far off, but also associated with overestimated accuracies.

We also tried MCMC runs with a modified prior setting; we extended the prior for the coalescence time t_c from its original range of ± 5 ms around the true value to ± 27 ms, allowing for an additional margin of 22 ms, which is the time it takes a gravitational wave to travel from Earth’s surface to its center. This setting reflects the case where the inspiral detection pipeline received triggers from less than three interferometer sites, so the signal’s arrival time at the

geocenter could not be estimated to greater accuracy in advance. The MCMC algorithm is still able to find the mode in the enlarged time parameter range, but takes more iterations to converge.

One scenario in which such an approach would be necessary is when the SNR for one of the interferometers is almost zero. In such a case the data from the interferometer under consideration also would not (directly) contribute to the estimation of phase- and frequency-related parameters, but would still carry information about amplitude-related parameters—by implicitly “ruling out” those parameter combinations that *would* have resulted in a response at that interferometer. Figure 4 shows the sky location posteriors for such a signal, a 1.5–2.0 M_\odot inspiral at 10 Mpc distance, where the SNRs at the three interferometer sites are: LHO: 9.6, LLO: 13.9, V: 0.18 (network: 16.9). Including the data from the third interferometer (with almost zero SNR) into the analysis still yields a much more accurate estimate of the sky location.

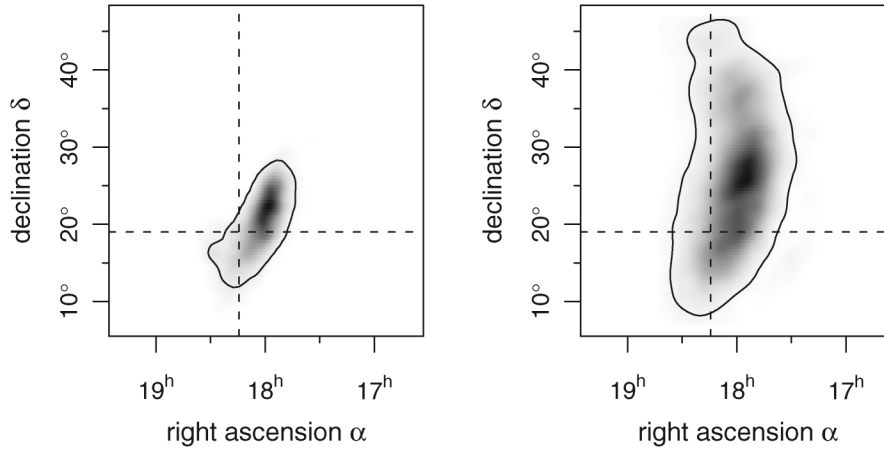


FIG. 4. Even if the SNR at one of the interferometers is almost zero, it still contributes to estimates' accuracies—the posterior is much narrower if its data is included (left plot) than if it is omitted (right plot). (Dashed lines indicate the true values.)

TABLE III. Relative accuracies of different parameters (in analogy to Table II) when considering/not considering the Virgo data (where the example situation is such that the SNR is nearly zero).

| Virgo data ... | (δ, α) | t_c | d_L | m_c | μ |
|----------------|--------------------|---------|-------|--------|-------|
| ...Included | 0.071 rad | 0.81 ms | 21% | 0.023% | 0.33% |
| ...Excluded | 0.150 rad | 2.38 ms | 23% | 0.022% | 0.31% |

Table III compares the resulting parameter accuracies of these two settings. The posteriors for sky location (δ, α) and coalescence time t_c , which are closely related, are much narrower when the Virgo data is considered in the analysis, while estimates of the rather phase- and frequency-related parameters m_c and μ do not gain from the additional information.

On the one hand, not only a high (total) SNR is desirable but also one that is rather “evenly spread” over different interferometers. On the other hand even a near-zero SNR at one of the interferometers does not make its measurement useless. Inference on different parameters will be affected to different degrees by such an unbalanced SNR arrangement.

V. DISCUSSION

We have developed a new MCMC program for estimating the nine parameters associated with an inspiral of compact binary objects from the data coming from a network of gravitational wave interferometers. The determination of the sky location of the source is an important consequence of the procedure. Numerous new features have been implemented in this binary inspiral MCMC. The MCMC uses waveform approximations that are 2.5 PN in phase and 2.0 PN in amplitude [51] (and by the time of final submission of this paper a version using

waveforms that are 3.5 PN in phase and 2.5 PN in amplitude [35–39] was running as well). The data from multiple interferometers (two or more) are coherently analyzed in order to produce posterior probability distributions for all nine parameters.

Advanced MCMC techniques were implemented in our program in order to maximize the efficiency of converging to the correct parameter values in the large, 9-dimensional, parameter space. The initial parameter values for the sampler were determined using importance resampling [50]. We recently extended (though not with the results presented in this paper) the parallel tempering algorithm to an Evolutionary MCMC algorithm [72]. This MCMC variety implements proposals that are motivated by genetic algorithms [73], and so recombinations of parameter samples from different MCMC chains are used as proposals in order to improve convergence and mixing.

Another current related research effort is the application of a version of this MCMC code to burst waveforms. This problem is by orders of magnitude computationally less expensive, due to the much shorter duration of the signals. But it appears that on the other hand convergence, i.e. finding the main posterior mode in the parameter space, still poses a major problem. The theoretical background of the various potential burst sources is rather vague, so realistic waveforms and sensible specifications of parameterizations and priors also need to be identified. The results of this study on the MCMC parameter estimation of burst signals using the coherent analysis of multi-interferometer data will be presented in a forthcoming publication.

We are also extending the MCMC techniques used in this study to the application of data analysis for LISA detection of binary inspiral signals. While our present program coherently analyzes data from multiple ground based interferometers, we have found it is a straightforward extension of the code so that we can coherently analyze the time delay interferometry data from LISA. These results are also forthcoming.

Presently LIGO is at its target sensitivity. Virgo is fast approaching its design sensitivity. Using the LIGO-Virgo network it will be possible to observe neutron star binary inspirals out to a distance of 35 Mpc [45,46]. A detection of such an inspiral could occur at any time [4]. As displayed in this paper, our MCMC routine is capable of coherently analyzing the data from the multiple interferometers, and then using it to estimate the nine parameters associated with such a signal.

ACKNOWLEDGMENTS

This work was supported by The Royal Society of New Zealand Marsden Fund grant No. UOA-204, National Science Foundation grant No. PHY-0553422, and the Fulbright Program.

-
- [1] B.C. Barish, in *Gravitational Wave Detection: Proceedings of the TAMA International Workshop on Gravitational Wave Detection*, edited by K. Tsubono *et al.* (Universal Academy Press, Tokyo, 1997), p. 155.
- [2] B.C. Barish and R. Weiss, *Phys. Today* **52**, 44 (1999).
- [3] B. Abbott *et al.*, *Nucl. Instrum. Methods Phys. Res., Sect. A* **517**, 154 (2004).
- [4] V. Kalogera, K. Belczynski, R. O’Shaughnessy, and B. Willms, *astro-ph/0612144*.
- [5] A. Brilliet, in *Gravitational Wave Detection: Proceedings of the TAMA International Workshop on Gravitational Wave Detection*, edited by K. Tsubono *et al.* (Universal Academy Press, Tokyo, 1997), p. 163.
- [6] B. Caron *et al.*, *Nucl. Phys. B, Proc. Suppl.* **48**, 107 (1996).
- [7] F. Acernese *et al.*, *Class. Quant. Grav.* **23**, S63 (2006).
- [8] J. Hough *et al.*, in *Gravitational Wave Detection: Proceedings of the Tama International Workshop on Gravitational Wave Detection*, edited by K. Tsubono *et al.* (Universal Academy Press, Tokyo, 1997), p. 175.
- [9] H. Lück *et al.*, *Class. Quant. Grav.* **23**, S71 (2006).
- [10] K. Tsubono *et al.*, in *Gravitational Wave Detection: Proceedings of the Tama International Workshop on Gravitational Wave Detection*, edited by K. Tsubono *et al.* (Universal Academy Press, Tokyo, 1997), p. 163.
- [11] R. Takahashi *et al.*, *Class. Quant. Grav.* **20**, S593 (2003).
- [12] K. S. Thorne, in *300 Years of Gravitation*, edited by S. W. Hawking and W. Israel (Cambridge University Press, Cambridge, 1987), chap. 9, pp. 330–358.
- [13] C. Cutler *et al.*, *Phys. Rev. Lett.* **70**, 2984 (1993).
- [14] S. A. Hughes, *Phys. Rev. D* **66**, 102001 (2002).
- [15] B. F. Schutz, *Nature (London)* **323**, 310 (1986).
- [16] D. F. Chernoff and L. S. Finn, *Astrophys. J.* **411**, L5 (1993).
- [17] D. Marković, *Phys. Rev. D* **48**, 4738 (1993).
- [18] C. Cutler and E. E. Flanagan, *Phys. Rev. D* **49**, 2658 (1994).
- [19] L. S. Finn, *Phys. Rev. D* **53**, 2878 (1996).
- [20] E. E. Flanagan and S. A. Hughes, *Phys. Rev. D* **57**, 4535 (1998).
- [21] E. E. Flanagan and S. A. Hughes, *Phys. Rev. D* **57**, 4566 (1998).
- [22] B. Abbott *et al.*, *Phys. Rev. D* **69**, 122001 (2004).
- [23] B. Abbott *et al.*, *Phys. Rev. D* **72**, 082001 (2005).
- [24] B. Abbott *et al.*, *Phys. Rev. D* **73**, 102002 (2006).
- [25] B. Abbott *et al.*, *Phys. Rev. D* **72**, 082002 (2005).
- [26] B. Abbott *et al.*, *Phys. Rev. D* **73**, 062001 (2006).
- [27] W. R. Gilks, S. Richardson, and D. J. Spiegelhalter, *Markov chain Monte Carlo in practice* (Chapman & Hall/CRC, Boca Raton, 1996).
- [28] P. C. Gregory, in *Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, edited by A. Mohammad-Djafari, *AIP Conf. Proc.* No. 568 (American Institute of Physics, New York, 2001), pp. 557–568.
- [29] T. J. Loredo, in *Statistical Challenges in Modern Astronomy*, edited by E. D. Feigelson and G. J. Babu (Springer-Verlag, New York, 1992), chap. 12, pp. 275–297.
- [30] L. S. Finn, *gr-qc/9709077*.
- [31] N. Christensen and R. Meyer, *Phys. Rev. D* **64**, 022001 (2001).
- [32] N. Christensen, R. Meyer, and A. Libson, *Class. Quant. Grav.* **21**, 317 (2004).
- [33] C. Röver, R. Meyer, and N. Christensen, *Class. Quant. Grav.* **23**, 4895 (2006).
- [34] T. Tanaka and H. Tagoshi, *Phys. Rev. D* **62**, 082001 (2000).
- [35] L. Blanchet, G. Faye, B. R. Iyer, and B. Joguet, *Phys. Rev. D* **65**, 061501 (2002), note the erratum [36].
- [36] L. Blanchet, G. Faye, B. R. Iyer, and B. Joguet, *Phys. Rev. D* **71**, 129902 (2005) (See also [35]).
- [37] K. G. Arun, L. Blanchet, B. R. Iyer, and M. S. S. Qusailah, *Class. Quant. Grav.* **21**, 3771 (2004), note the erratum [38].
- [38] K. G. Arun, L. Blanchet, B. R. Iyer, and M. S. S. Qusailah, *Class. Quant. Grav.* **22**, 3115 (2005) (See also [37]).
- [39] L. Blanchet, T. Damour, G. Esposito-Farèse, and B. R. Iyer, *Phys. Rev. Lett.* **93**, 091101 (2004).
- [40] A. Pai, S. Dhurandhar, and S. Bose, *Phys. Rev. D* **64**, 042004 (2001).
- [41] A. Pai, S. Bose, and S. Dhurandhar, *Class. Quant. Grav.* **19**, 1477 (2002).
- [42] S. Bose, *Class. Quant. Grav.* **19**, 1437 (2002).
- [43] S. Birindelli *et al.* (unpublished).
- [44] H. Mukhopadhyay, N. Sago, H. Tagoshi, S. Dhurandhar, H. Takahashi, and N. Kanda, *Phys. Rev. D* **74**, 083005 (2006).
- [45] D. Sigg, *Class. Quant. Grav.* **21**, S409 (2004).
- [46] F. Acernese *et al.*, *Class. Quant. Grav.* **22**, S869 (2005).

- [47] N. J. Cornish and E. K. Porter, *Phys. Rev. D* **75**, 021301 (2007).
- [48] E. T. Jaynes, in *Maximum-Entropy and Bayesian Methods in Applied Statistics*, edited by J. H. Justice (Cambridge University Press, Cambridge, 1986).
- [49] E. T. Jaynes, *Probability Theory: The Logic of Science* (Cambridge University Press, Cambridge, 2003).
- [50] A. Gelman, J. B. Carlin, H. Stern, and D. B. Rubin, *Bayesian Data Analysis* (Chapman & Hall/CRC, Boca Raton, 1997), 1st ed.
- [51] L. Blanchet, in *Gravitational Waves: Proceedings of the Como School on Gravitational Waves in Astrophysics*, edited by I. Ciufolini, V. Gorini, U. Moschella, and P. Fré (Institute of Physics Publishing, Bristol, 2001); see also gr-qc/0104084.
- [52] L. S. Finn and D. F. Chernoff, *Phys. Rev. D* **47**, 2198 (1993).
- [53] E. T. Jaynes, *IEEE Trans. Syst. Sci. Cybern.* **SEC-4**, 227 (1968).
- [54] F. Marion *et al.*, *Proceedings of the Rencontres de Moriond 2003* (Editions Frontières, Gif-sur-Yvette, France, 2004).
- [55] P. Amico *et al.*, *Comput. Phys. Commun.* **153**, 179 (2003).
- [56] P. Teerikorpi, *Annu. Rev. Astron. Astrophys.* **35**, 101 (1997).
- [57] A. Sandage, in *Encyclopedia of Astronomy and Astrophysics*, edited by P. Murdin (Institute of Physics Publishing, Bristol, 2001), Vol. 2, pp. 1645–1650.
- [58] Frame library (Fr), version v6r19 (2005), <http://lappweb.in2p3.fr/virgo/FrameL>.
- [59] R. E. Crochiere, in *Programs for Digital Signal Processing*, edited by A. C. Schell *et al.* (IEEE Press, New York, 1979), chap. 8.2.
- [60] J. Janovetz, Parks-McClellan algorithm for FIR filter design (C version) (1998), <http://www.janovetz.com/jake>.
- [61] B. W. Brown, J. Lovato, K. Russell, and J. Venier, Randlib - Library of C routines for random number generation (1997), <http://biostatistics.mdanderson.org/SoftwareDownload>.
- [62] R Development Core Team, *R: A Language and Environment for Statistical Computing* (R Foundation for Statistical Computing, Vienna, Austria) (2004).
- [63] D. W. Scott, *Multivariate Density Estimation: Theory, Practice and Visualization* (Wiley & Sons, New York, 1992).
- [64] M. Frigo and S. G. Johnson, FFTW 3.0.1, a C subroutine library for computing the discrete Fourier Transform (DFT) (2003), <http://www.fftw.org>.
- [65] P. D. Welch, *IEEE Trans. Audio Electroacoust.* **AU-15**, 70 (1967).
- [66] F. J. Harris, *Proc. IEEE* **66**, 51 (1978).
- [67] B. Allen, gr-qc/9607075.
- [68] K. R. Lang, *Astrophysical Formulae* (Springer-Verlag, Berlin, 1999), 3rd ed., Vol. II.
- [69] WGS 84 implementation manual, *European Organization for the Safety of Air Navigation (EUROCONTROL)* (Institute of Geodesy and Navigation (IfEN), Brussels/Munich, 1998), especially Appendices B and E.
- [70] A. Buonanno, Y. Chen, and M. Vallisneri, *Phys. Rev. D* **67**, 104025 (2003).
- [71] K. V. Mardia and P. E. Jupp, *Directional statistics* (Wiley & Sons, Chichester, 2000), Chap. 9.2.
- [72] F. Liang and H. W. Wong, *J. Am. Stat. Assoc.* **96**, 653 (2001).
- [73] D. E. Goldberg, *Genetic Algorithms in Search, Optimization, and Machine Learning* (Addison-Wesley, Reading, Mass., 1989).