Method for estimation of gravitational-wave transient model parameters in frequency–time maps

2014 Class. Quantum Grav. 31 165012

(http://iopscience.iop.org/0264-9381/31/16/165012)

IOP Publishing Classical and Quantum Gravity
Class. Quantum Grav. **31** (2014) 165012 (14pp) doi:10.1088/0264-9381/31/16/165012

Method for estimation of gravitational-wave transient model parameters in frequency–time maps

M Coughlin[1,2], N Christensen[3], J Gair[2], S Kandhasamy[4,5] and E Thrane[6]

[1] Department of Physics, Harvard University, Cambridge, MA 02138, USA
[2] Institute of Astronomy, University of Cambridge, Cambridge, CB3 0HA, UK
[3] Physics and Astronomy, Carleton College, Northfield, MN 55057, USA
[4] School of Physics and Astronomy, University of Minnesota, Minneapolis, MN 55455, USA
[5] Physics and Astronomy, University of Mississippi, University, MS 38677-1848, USA
[6] LIGO Laboratory, California Institute of Technology, MS 100-36, Pasadena, CA 91125, USA

E-mail: coughlin@physics.harvard.edu

Received 17 April 2014, revised 25 June 2014
Accepted for publication 7 July 2014
Published 5 August 2014

Abstract

A common technique for detection of gravitational-wave (GW) signals is searching for excess power in frequency–time (*ft*)-maps of GW detector data. In the event of a detection, model selection and parameter estimation will be performed in order to explore the properties of the source. In this paper, we develop a Bayesian statistical method for extracting model-dependent parameters from observed GW signals in *ft*-maps. We demonstrate the method by recovering the parameters of model GW signals added to simulated advanced LIGO noise. We also characterize the performance of the method and discuss prospects for future work.

Keywords: gravitational waves, parameter estimation, time–frequency maps
PACS numbers: 95.75.-z, 04.30.-w

(Some figures may appear in colour only in the online journal)

1. Introduction

The Laser Interferometer Gravitational-wave Observatory (LIGO) [1], Virgo [2], and GEO600 [3] detectors are part of a network of gravitational-wave (GW) detectors seeking to

make direct observations of GWs. Previous analyses of the data have included searches targeting the coalescence of neutron stars or black holes [4, 5], short-duration bursts [6], isolated neutron stars [7], and a stochastic background of GWs [8, 9]. LIGO and Virgo are currently upgrading to Advanced LIGO (aLIGO) and Advanced Virgo (AdV), which will improve their strain sensitivities by one order of magnitude over the strain sensitivites achieved during previous science runs [10, 11]. These will be joined by GEO-HF [12] and KAGRA [13]. To date, none of the above searches have resulted in a GW detection, although with the current upgrades, the chances will increase significantly. In the event of a detection, one can perform model selection and parameter estimation in order to further explore the properties of the sources. Model selection and parameter estimation are topics of great interest in the GW community (see e.g. [14–20]).

The binary coalescence of compact objects are well-studied sources of GWs, and the most up-to-date models for the waveforms produced in these systems include most of the physical effects that influence the signals, including tidal and spin effects [21]. Searches and parameter estimation for these sources rely on matched filtering of the signal seen by detectors using models of the signals. Because models for these sources are thought to be reliable, a full Bayesian analysis utilizing matched filtering is possible for these sources, and the ability to precisely estimate injected waveform parameters for these sources has been demonstrated [18].

On the other hand, GW bursts cannot be modeled precisely (by assumption). GW emission by core-collapse supernovae is one such example. A number of competing models for the mechanism that drives the core-collapse exist, and each model produces qualitatively different waveforms. Logue *et al* demonstrated that a principal component analysis can be used to determine the correct model of injected GW waveforms by the computation of the Bayesian odds ratio [19]. Principle component analysis has also been used to reconstruct the stellar core-collapse GW signal after finding the amplitude of the individual principle components and arrival times [20].

Parameter estimation of signal models requires, at first, GW detection with high significance. For cases of well-modeled burst sources, such as cosmic strings, matched filtering has been used to perform searches [22]. A common technique for detection of unmodeled bursts is searching for excess power in frequency–time (ft)-maps of GW detector data [23–25]. Matched filtering is not used for these signal types because the precise waveforms are unknown. Excess power searches provide an effective alternative to matched filtering for such signals. Although these pipelines use generic search algorithms when searching for GWs, they utilize signal models that have the features of expected GW signals to tune their analyses. Some signals can be well-approximated by parameterized spectrogram curves which incorporate the salient features of the signals, and these curves can be used to focus the search with a 'phase-less template bank' [26, 27]. These unmodeled burst analyses mostly use ft-maps. However, there are modeled burst searches such as the cosmic string search that use other methods.

In this paper, we present a method for parameter estimation using GW tracks in ft-maps. We explore the possibility of performing parameter estimation and model selection, assuming that a search has been performed and a signal detected. We seek to address the question of how to fit the model parameters. As a concrete example, we show the recovery of parameters of an r-mode signal injected into simulated detector data. These GW sources are unstable oscillation modes which dampen the rotation of neutron stars by the emission of GWs [28]. We show how to estimate parameters such as the r-mode saturation amplitude, which is the amplitude above which the emitting neutron star will collapse into a black hole.

The remainder of this paper is organized as follows. We discuss the methods used to extract waveform parameters from tracks in *ft*-maps in section 2. To demonstrate the method and performance of parameter recovery, we perform sample injections into simulated aLIGO colored Gaussian noise and recover their parameters in section 3. We conclude with a discussion of topics for further study in section 4.

2. Formalism

In this section, we discuss the methods used to extract waveform parameters from tracks in *ft*-maps. To begin, we review the data products used in the detection of unmodeled GW transients.

2.1. *ft*-maps

Many searches for GW bursts rely on searching for excess power in *ft*-maps of GW detector data [23–25]. The maps are computed by dividing detector strain time series into segments and computing a Fourier transform of each segment. Each column in the map corresponds to one of these segments. Searches for long-duration GW bursts in particular use the cross-correlation of two GW strain channels from spatially separated detectors to construct *ft*-maps of *cross-power* signal-to-noise ratio (SNR), $\rho(t; f)$ [25]

$$\rho(t; f) \equiv \hat{Y}(t; f)/\hat{\sigma}(t; f), \tag{1}$$

where t is the time of the segment, f is the frequency, $\hat{Y}(t; f)$ is an unbiased estimator for GW power and $\hat{\sigma}^2(t; f)$ is its variance. Arrays of $\rho(t; f)$ are visualized as *ft*-maps.

GWs appear as tracks or blobs on *ft*-maps. The morphology of the GW track depends on the source. If the signal is sufficiently loud, compact binaries appear as chirps of increasing frequency. Continuous-wave isolated neutron star sources appear as narrowband, nearly horizontal lines with a small change in frequency with time due to Doppler shifting. Figure 1 shows *ft*-maps of example sine-Gaussian injections with different durations (top row) and r-mode injections with different saturation amplitudes (bottom row).

Given an *ft*-map, GW searches employ pattern recognition algorithms to identify potentially significant clusters of pixels [25]. Next, the pattern-recognition algorithms are run repeatedly on noise-only maps to generate background statistics. These noise-only maps are created using GW detector strain data with a time-shift that removes any potential GW signal. Using time shifts to study noise and injections to study detection efficiency, false alarm and false dismissal rates can be estimated, and detections can potentially be made.

2.2. Waveform models

A metric perturbation, h_{ab}, can be written as a combination of two polarizations, h_+ and $h_\times$.

$$h_{ab} = \begin{pmatrix} h_+ & h_\times \\ h_\times & -h_+ \end{pmatrix}. \tag{2}$$

Far from an elliptically polarized source, we can write the metric pertubation as

$$h_+(t) = h_{\rm amp}(t)\left(1 + \cos^2\iota\right)\cos\psi(t), \tag{3}$$

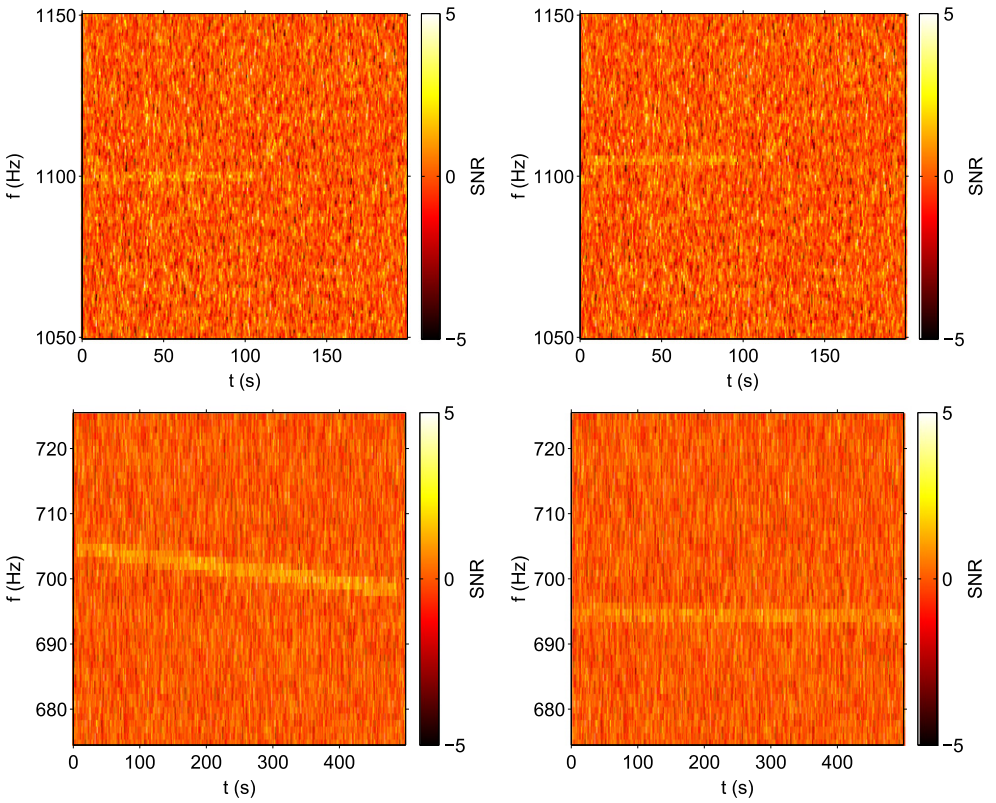

Figure 1. *ft*-maps of $\rho(t; f)$ with injected signals. The top row consists of sine-Gaussian injections [6]; on the left is an injection with $f_0 = 1100$ Hz and $\tau = 100$ s and on the right is $f_0 = 1105$ Hz and $\tau = 80$ s. The bottom row consists of r-mode injections [29]; on the left is an injection with f_0 of 705 Hz and a saturation amplitude $\alpha = 0.3$, and on the right is an injection with f_0 of 695 Hz and a saturation amplitude $\alpha = 0.1$. The injections are performed at a distance at which a GW signal can be observed above threshold with false alarm probability = 0.1% and false dismissal probability = 50% using a seed-based clustering algorithm [30]. This corresponds to a matched filter SNR of about 20 for the sine-Gaussian injections and about 30 for the r-modes [31].

and

$$h_\times(t) = 2h_{\rm amp}(t)\cos\iota\sin\psi(t), \tag{4}$$

where $h_{\rm amp}$ is the strain amplitude, ι is the inclination of the angular momentum axis of the object to the line of sight of the observer and ψ is the polarization angle. In the analysis below, we assume that we have a face-on source, so $\iota = 0$. We also assume that $\psi = 0$. This is for the sake of simplicity; in theory, one may estimate ι and ψ as well, but an analysis involving these parameters is beyond the scope of this paper.

GW detectors measure strain, $h_0(t)$,

$$h_0(t) = h_+(t)F_+(t) + h_\times(t)F_\times(t), \tag{5}$$

where $F_+(t)$ and $F_\times(t)$ are the detector antenna response functions to the two polarizations [32]. GW amplitudes are sometimes characterized by the root-sum-square amplitude, $h_{\rm rss}$, defined as

$$h_{\rm rss}^2 = \int \left[h_+^2(t) + h_\times^2(t) \right] {\rm d}t. \tag{6}$$

Each GW burst creates a specific pattern in *ft*-maps which depends on astrophysical parameters. In this study, three different models of GW signals are used. The models are designed to have the salient features of the types of burst signals expected to be detected, while being general enough to capture a variety of potential phenomena. The first is a sine-Gaussian, which is commonly used in searches for GW bursts [6]. This model depends on four parameters: the waveform duration, τ, the start frequency, f_0, the signal distance, D, and the time of the maximum of burst, t_0. It has the following form

$$h_0(t) = k \frac{\exp\left(-\left(\frac{(t-t_0)^2}{4\tau^2} + 2\pi {\rm i}(t-t_0) f_0 \right) \right)}{D\left(2\pi\tau^2\right)^{1/4}}, \tag{7}$$

$$h_+(t) = {\rm Re}\left[h_0(t)\right], \quad h_\times(t) = {\rm Im}\left[h_0(t)\right], \tag{8}$$

where k is a constant and i is the imaginary unit. The second waveform represents a simple r-mode model, based on a model by Owen *et al* [29]. This model depends on three parameters: the saturation amplitude, α, the start frequency, f_0, and the signal distance, D. It has the following form

$$f(t) = \left(\frac{1}{f_0^{-6} - 6kt} \right)^{1/6}, \tag{9}$$

$$h_+(t) = h_0 \cos\left(2\pi f(t) t\right), \quad h_\times(t) = h_0 \sin\left(2\pi f(t) t\right), \tag{10}$$

where

$$h_0 = 3.6 \times 10^{-23} \alpha \left(\frac{f(t)}{1000} \right)^3 \Big/ D, \tag{11}$$

$$k = -1.8 \times 10^{-21} \alpha^2 \ {\rm Hz}^{-5}. \tag{12}$$

The third waveform is a slowly varying sinusoid waveform with a time-varying frequency, $f = f_0 + \dot{f} t$. This model is chosen here as its morphology is similar to the r-modes. This model depends on three parameters: the time derivative of signal frequency, $\dot{f}$, the start frequency, f_0, and the signal distance, $D \propto 1/h_{\rm rss}$.

$$h_0(t) = c \frac{\exp\left(2\pi {\rm i} \left(f_0 + \dot{f} t \right) t \right)}{D}, \tag{13}$$

where c is a constant. h_+ and $h_\times$ are calculated in the same way as the sine-Gaussian.

2.3. Likelihood

We use the above models to illustrate our method. Figure 1 shows two pairs of *ft*-maps of cross-power with sine-Gaussian and r-mode injections. Our goal is to determine, based on the map structure, the parameters which best fit the models. In order to estimate the parameters, we employ a likelihood formalism.

The first step is to compute the probability distribution of $\rho_B(t; f)$ due to background, $f_B(\rho)$. A distribution valid for Gaussian and stationary noise is derived in appendix A. In cases for which an analytic distribution is impractical to construct, it can be estimated from

time-shifted data. In the analysis below, we assume Gaussian noise for simplicity; the use of time-shifted distributions will be explored in a future study. It is likely that this distribution would be necessary in the case of glitchy or non-stationary data. We also assume that $\rho_B(t; f)$ of each pixel is drawn from the same distribution. The second step is to determine the contribution to ρ from a signal. We denote the expected ρ value due to a signal with parameters θ by $\rho_S(\theta)$. We calculate the expected contribution using an approximation described in appendix B. The assumption leads to an approximate formulation, which can be made more accurate by performing injection studies and computing the distributions with signals present.

Armed with both the distribution of $\rho_B(t; f)$ due to background as well as that of the waveform models, $\rho_S(t; f)$, we are able to construct our likelihood. The idea is to subtract $\rho_S(\theta)$ from ρ, which would just leave detector noise if $\rho_S(\theta)$ was the correct waveform model. Minimizing the residuals maximizes the likelihood function. The probability density function describing the residuals is $p(\rho - \rho_S(\theta))$, which is calculated by finding the probability that $\rho - \rho_S(\theta)$ is due to noise, as given by $f_B(\rho)$. The likelihood is

$$L\Big(\{\rho_i\}\,|\,\theta\Big) = \Pi_{i=1}^{N} p\Big(\rho_i - \rho_{s_i}(\theta)\,|\,\theta\Big), \tag{14}$$

where i is the pixel index and N is the number of pixels in the *ft*-map. The goal is to maximize the likelihood in order to determine confidence intervals for θ.

Ideally, one would produce an *ft*-map of the GW signal and evaluate equation (14) for every set of possible parameters. In this way, we could generate the posterior density functions (PDFs) for the relevant model parameters. Because this is computationally intractable, we use algorithms that efficiently sample the posterior while minimizing the computational burden. In the examples below, we use flat, non-informative priors on the parameters of interest. This could be modified, for example, in the event of an r-mode detection, where models predict a small value of α.

There are three main algorithms presently used to rapidly evaluate the posterior in GW parameter estimation and model selection: Markov Chain Monte Carlo [14, 18], Nested Sampling [15, 33], and MultiNest [16, 17, 34]. Nested Sampling and MultiNest calculate the Bayesian evidence for a given set of parameters, which can be used to assign relative probabilities to different models. We use a MATLAB implementation of Nested Sampling and MultiNest [35], which implements the MultiNest algorithm, as described in [34], and Nested Sampling, as described in [33].

3. Demonstration

In this section, we present two examples of parameter estimation of toy model waveforms. We inject GW signals into simulated aLIGO colored Gaussian noise and create *ft*-maps based on the resulting timeseries. We use the design sensitivity aLIGO noise curve [10]. We perform injections at the waveform models' detection distance, which we define as the distance at which a GW signal can be observed above threshold with false alarm probability = 0.1% and false dismissal probability = 50% using a seed-based clustering algorithm [30]. We have many *ft*-maps that contain both GW signals and noise. In order to construct parameter posterior distributions, we produce *ft*-maps containing only GW signals for various sets of waveform parameters. Equation (14) is evaluated repeatedly for each set of parameters by subtracting the GW signal maps from those that contain signals and noise. From the equation, the likelihood is maximized for those parameters that best minimize the residuals.

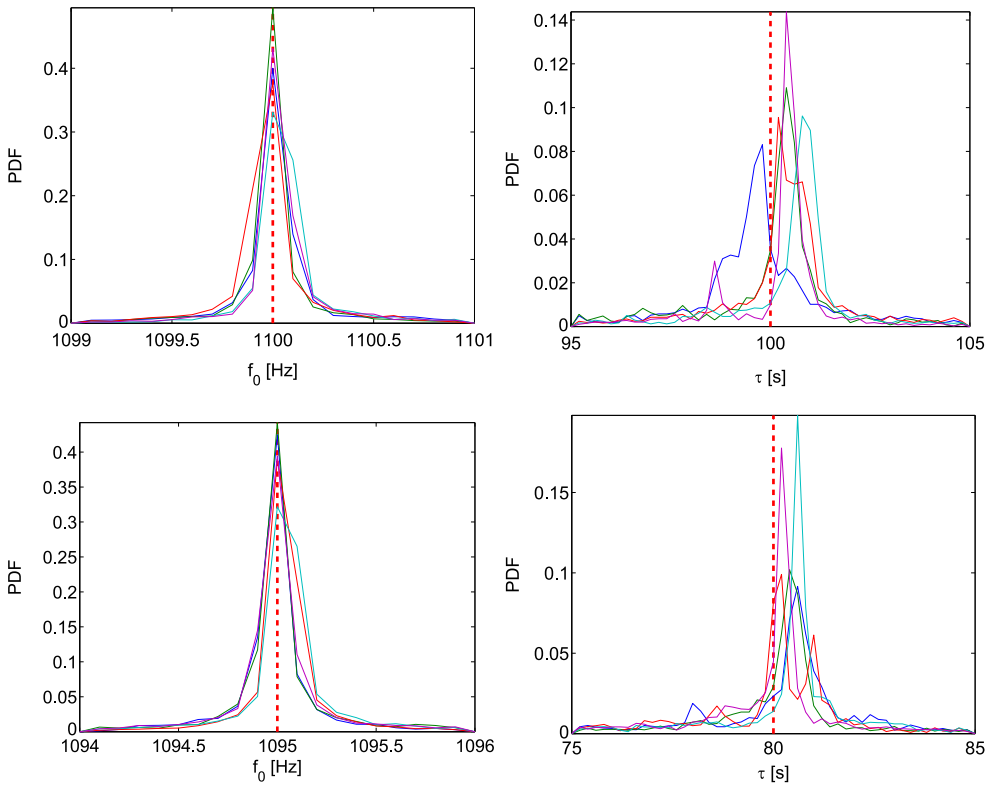

Figure 2. PDFs of distributions for injected sine-Gaussian signals. Each injection is performed into ten simulated aLIGO colored Gaussian noise realizations. The injections are performed at a distance at which a GW signal can be observed above threshold with FAP = 0.1% and FDP = 50%. This corresponds to a matched filter SNR of about 30 for these injections. The PDF for each injection is plotted in a different color. The red dotted line shows the true injected value. The plot on the left is the PDF of f_0. The plot on the right is the PDF of τ. The top row corresponds to an injection of $f_0 = 1100$ Hz and $\tau = 100$ s, while the bottom row corresponds to an injection of $f_0 = 1095$ Hz and $\tau = 80$ s. f_0 is recovered to within a fraction of a percent, while the τ sampling is within a few percent.

Parameter posterior distributions are constructed for parameter sets of equal likelihood that maximize this likelihood.

3.1. Sine-Gaussian burst

Figure 2 shows the distribution of posterior samples for both f_0 and τ for one of the injections. The posteriors of the parameters are consistent with the injected values. In general, the recoveries for f_0 are within a fraction of a percent of the true value for all injections. The recoveries for τ are within a few percent.

3.2. r-mode

We perform injections of the r-mode waveforms with $f_0 = 705$ Hz and $\alpha = 0.3$, as well as an injection of $f_0 = 695$ Hz and $\alpha = 0.1$. Figure 3 shows the performance of the parameter recoveries. In general, the recoveries for α and f_0 are within a few percent for all injections.

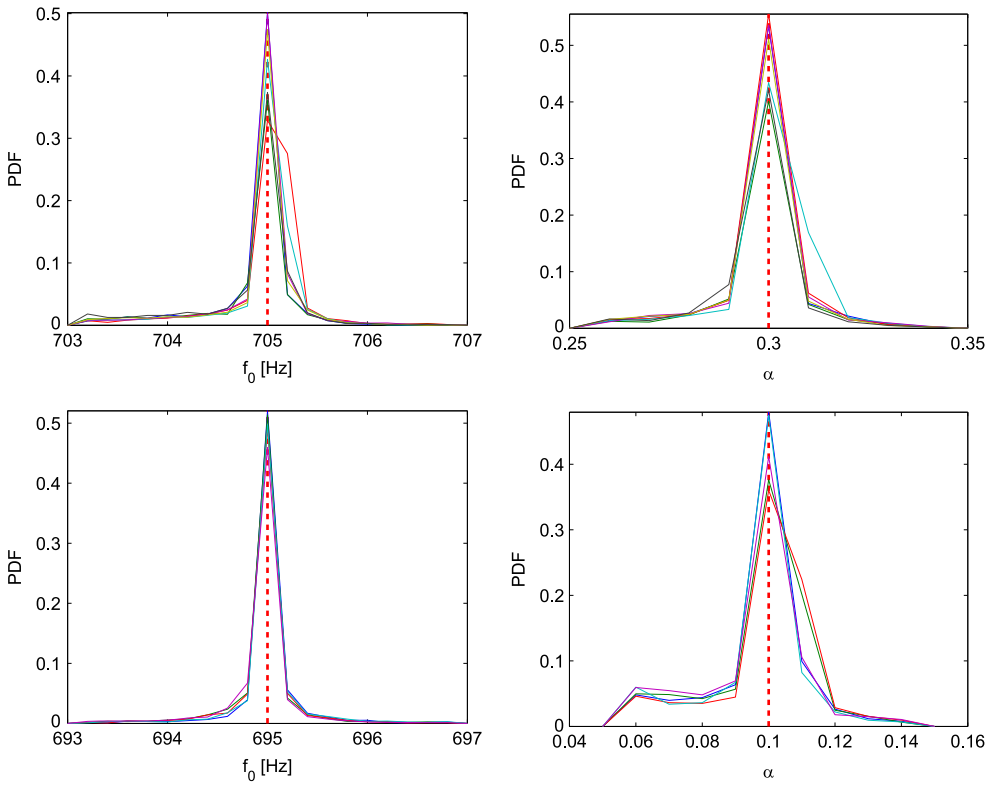

Figure 3. PDFs of distributions for injected r-modes signals. Each injection is performed into ten simulated aLIGO colored Gaussian noise realizations. The injections are performed at a distance at which a GW signal can be observed above threshold with FAP = 0.1% and FDP = 50%. This corresponds to a matched filter SNR of about 20 for these injections. The PDF for each injection is plotted in a different color. The red dotted line shows the true injected value. The plot on the left is the PDF of f_0. The plot on the right is the PDF of α. The top row corresponds to an injection of $f_0 = 705$ Hz and $\alpha = 0.3$, while the bottom row is an injection of $f_0 = 695$ Hz and $\alpha = 0.1$.

We can ask whether the r-mode or the varying sinusoidal model is a better description of the *ft*-map for the r-mode signal. This can be done by evaluating the Bayes factor, which is the ratio of the evidences for the two models. The evidence computed by the search algorithm for the r-mode model was 278, while the evidence for the CW model was 12, meaning the Bayes factor is 23. The evidences are similar across noise realizations as well as injection parameters. This implies that the r-mode model is strongly favored over the varying sinusoidal model. This is despite the fact that the varying sinusoid is a good fit for the linear portion of the r-mode parameter space.

It is a well-documented fact that in the parameter estimation of compact binary coalescences non-Gaussian noise can significantly affect the posterior recoveries [36]. Therefore, it is worthwhile to test the algorithm when the noise background is non-Gaussian and non-stationary and thus violates the approximations that go into deriving the noise model used in this analysis. For this reason, we repeat the test with initial LIGO noise which has been recolored to match the design sensitivity of aLIGO [1, 10]. We introduce an artificial time-

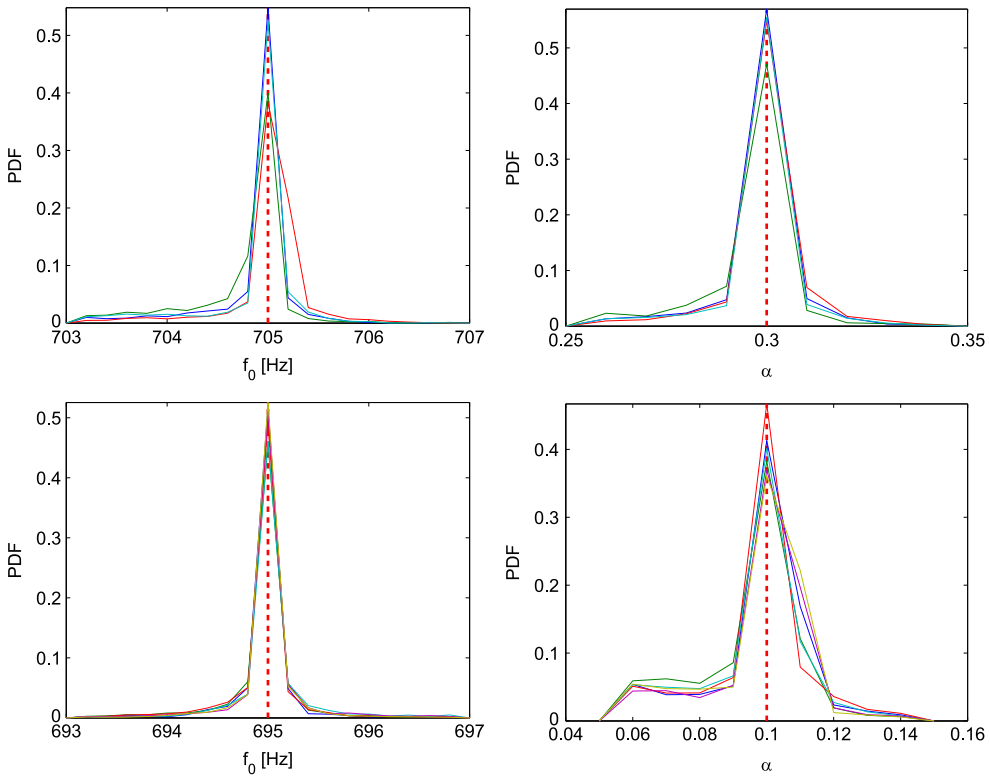

Figure 4. Same as figure 3 for initial LIGO noise recolored to advanced LIGO noise. The parameter recoveries are similar to the Gaussian noise case.

shift in the initial LIGO data to remove any potential GW signals present[7]. This has the benefit of preserving non-stationary noise artifacts while using the expected spectrum. Figure 4 shows the performance of the recoveries for injections into the recolored noise, using the noise model that assumes Gaussian noise. We find that for these injections, the performance is similar to that of the Gaussian noise case. To fully understand the effects of glitches, it will be necessary to perform more comprehensive injection studies. Because the data segments analyzed and waveforms are long, it is likely that our pipeline is less susceptible to noise transients than the compact binary case, which has most of its SNR in the last second before coalescence.

4. Conclusion

In this paper, we have demonstrated the ability to perform basic parameter estimation on GW signals from their signature in *ft*-maps. We described the likelihood method used and showed that these methods correctly recovered the parameters of waveform models and were able to differentiate between two similar models.

In the future, we will move beyond the generic models presented here to more complicated models. This will be necessary to identify the physics underlying a particular GW

[7] The data are taken in between GPS times 822917487 and 847549782.

source by distinguishing between different variations of similar models. Also, further studies of the assumptions made in the paper will be conducted. We have assumed that the noise is Gaussian and stationary and ignore the correlation between pixels in the maps. We have also assumed that the cross terms when multiplying the noise and waveform signals are zero. There are three complicating assumptions involved in the use of the likelihood here. The first is that in reality the cross-terms of the noise and signal are non-zero (see appendix B). The second is that the cross-power statistic uses adjacent PSD's for the purpose of estimating σ (from quation (1)), meaning there is a correlation between adjacent pixels (see appendix A). As such, the multiplication of the pixel probabilities in the *ft*-map, which requires that the probabilities are all independent if one wants a true cumulative probability, is not valid. The third is that real detectors have noise transients and non-stationary noise, which violate some of the approximations used here. One way to rectify this is to perform many injections and measure $f_S(\theta)$ empirically (this distribution would change for each signal model). These issues will be explored in the future.

The use of *ft*-maps to perform parameter estimation has the natural advantage over matched filtering in terms of the speed at which it can be done. Because we fit the amplitude of the waveform to the track in the *ft*-map (removing the phase information), it also means that our signal models do not need to be quite as exact as for parameter estimation relying on matched filtering. It is easier to match the amplitude of the signal than the phase, which is required by matched filtering. We can estimate the potential performance of matched filtering parameter estimation using the Fisher Information Matrix (FIM), which is a tool that has been used to estimate the potential accuracy of parameter estimates for GW signals [37, 38]. In the limit of high SNR, the inverse of the FIM is the variance–covariance matrix of the estimated signal parameters. It provides a first-order estimate of the errors when measuring parameters. Applying this technique to the r-mode model discussed above, this technique finds that the errors would be of order 0.1% for f_0 and 1% for α, which is about an order of magnitude better than for the *ft*-map based technique. This seems reasonable as the matched filtering technique includes phase information and includes none of the approximations. We also expect this to be a reasonable estimate for the signals used above because of the high matched filtering SNRs. This is important as it has been previously shown that the FIM is biased for near-threshold SNR signals [39]. Further study may include a detailed comparison between matched filtering and *ft*-map perform parameter estimation.

Acknowledgments

The authors would like to thank Matthew Pitkin for help running the likelihood sampler. MC was supported by the National Science Foundation Graduate Research Fellowship Program, under NSF grant number DGE 1144152. ET is a member of the LIGO Laboratory, supported by funding from United States National Science Foundation. NC's work was supported by NSF grant PHY-1204371. JG's work is supported by the Royal Society. LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the National Science Foundation and operates under cooperative agreement PHY-0757058. The authors thank the LIGO and Virgo Collaborations for providing the data plotted in figure 4. This paper has been assigned LIGO document number LIGO-P1400043.

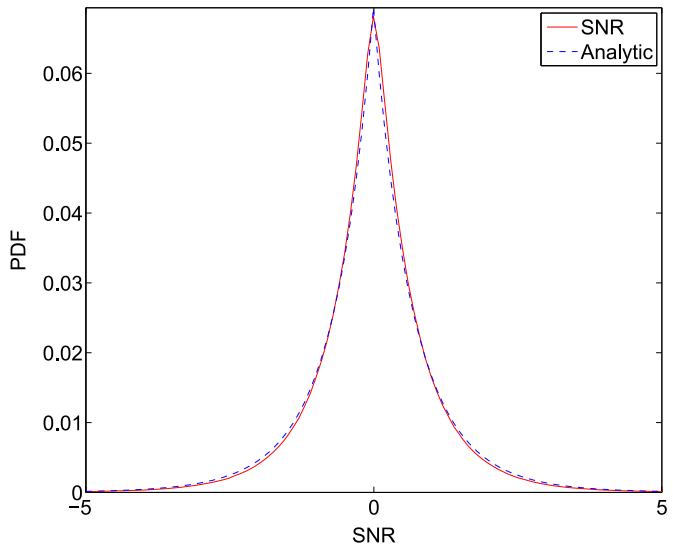

Figure A1. Probability density function of the pixels in a *ft*-map of cross-power SNR with the theoretical distribution given by equation (A.7) overlaid.

Appendix A. Distribution of cross-power SNR

Cross-power SNR ρ is defined as

$$\rho \approx \frac{\mathrm{Re}\left(s_1(f)s_2(f)\right)}{\sqrt{P1_{\mathrm{adj}}(f)P2_{\mathrm{adj}}(f)}}, \tag{A.1}$$

where $s_1(f)$ and $s_2(f)$ are Fourier transforms of times series from two detectors (hence complex numbers) and $P1_{\mathrm{adj}}(f)$ and $P2_{\mathrm{adj}}(f)$ are the (averaged) PSDs calculated from adjacent segments [25]. Up to a scaling factor, the numerator is known as Y, the signal estimate, and the denominator σ_Y, its error. $s_1(f)$ and $s_2(f)$ are each Gaussian variables with mean 0 and variance σ^2. The distribution of a new variable z defined as $z = s_1(f)s_2(f)$, known as a normal product, is given by the expression [40]

$$f(z) = \frac{1}{\sigma^2}K_0\left(\frac{|z|}{\sigma^2}\right), \tag{A.2}$$

where $K_0(x)$ is the modified Bessel function of the second kind. As $s_1(f)$ and $s_2(f)$ are complex vectors, this is actually the sum of two normal products, which is known as a double exponential or Laplace distribution. This has a distribution of the form

$$f_Y(y) = \frac{1}{2\sigma^2}\exp\left(\frac{-|y|}{\sigma^2}\right). \tag{A.3}$$

This is the distribution of Y. The second step is to calculate the distribution of σ_Y. $P1_{\mathrm{adj}}$ and $P2_{\mathrm{adj}}$ are the average PSDs calculated from segments on either side of the segment used to calculate Y

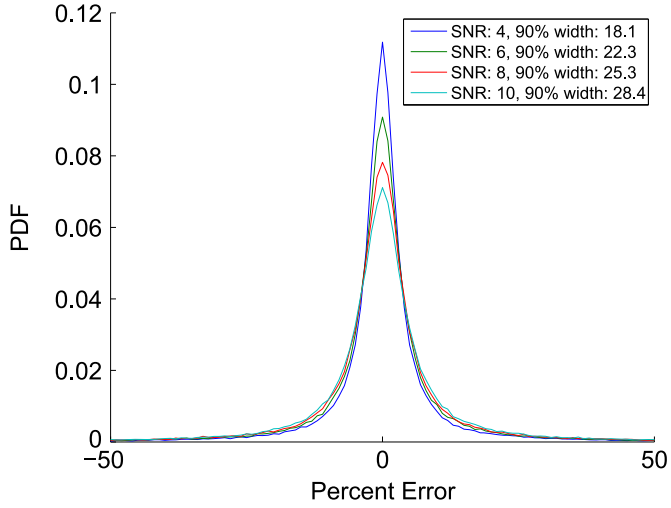

Figure B1. Probability density function of the percentage error of using $s_1 s_2 + n_1 n_2$ as an approximation for $h_1 h_2$. The width is the percentage error within which 90% of the distribution is contained.

$$P1_{\rm adj} = \frac{1}{N}\sum_{j=1}^{N} P1_j, \quad P2_{\rm adj} = \frac{1}{N}\sum_{j=1}^{N} P2_j. \tag{A.4}$$

In the frequency domain, each $P1_j$ and $P2_j$ are chi-squared distributed variables, the sum of which have distributions of the form

$$f_P(z) = \frac{N^N}{2^N \sigma^{2N} \Gamma(N)} \exp\left(\frac{Nz}{2\sigma^2}\right) z^{N-1} \quad (z \geqslant 0). \tag{A.5}$$

The distribution for σ_Y is then

$$f_{\sigma_Y}(z) = \frac{N^{2N}}{2^{2N-2}\sigma^{4N}(\Gamma(N))^2} K_0\left(\frac{Nz}{\sigma^2}\right) z^{2N-1}. \tag{A.6}$$

The final step is to combine the distributions for Y and σ_Y

$$f_{\rm SNR}(z) = \frac{N^{2N}}{2^{2N-1}\sigma^{4N+2}(\Gamma(N))^2} \int_0^\infty |x| e^{-\left|\frac{xz}{\sigma^2}\right|} K_0\left(\frac{Nx}{\sigma^2}\right) x^{2N-1} {\rm d}x. \tag{A.7}$$

Figure A1 shows the probability density function of the pixels of cross-power SNR used in this analysis overlaid with a distribution of $\rho(t; f)$ calculated from actual data.

Appendix B. Error approximation

If the timeseries of two detectors, h_1 and h_2, are composed of the sum of a signal and a noise part, (i.e., $h_1 = s_1 + n_1$ and $h_2 = s_2 + n_2$), when the two data streams are multiplied, the result will be in the form of

$$h_1 h_2 = s_1 s_2 + s_1 n_2 + s_2 n_1 + n_1 n_2. \tag{B.1}$$

The quantity $h_1 h_2$ is proportional to $\rho(t; f)$. The expectation values of the cross-terms, $s_1 n_2$ and $s_2 n_1$, are 0 because signal and noise are uncorrelated. To test the approximation that on

average the cross-terms, $s_1 n_2$ and $s_2 n_1$, will sum to zero, sets of 100 pixels with different total SNRs associated with them are generated. The SNR for each set is computed by performing a sum of the individual pixel SNRs. This process seeks to imitate the total error accumulated due to the assumption above. In this case, the total error is the sum of $s_1 n_2$ and $s_2 n_1$ for all of the pixels.

Figure B1 shows the percent difference between $h_1 h_2$ and $s_1 s_2 + n_1 n_2$. For pixel sets with moderate total SNR, 90% of the time, this approximation is within 25% of its true value. Extremely high SNR events, which are an order of magnitude larger, yield errors on the order of 100%. Examining the cross-terms, their contribution becomes more significant as the magnitude of the signal increases and the approximation breaks down in the high SNR regime. Conversely, the bias when using signals of moderate SNR is shown to be small in section 3.

References

[1] Abbott B *et al* (LIGO Scientific Collaboration) 2009 *Rep. Prog. Phys.* **72** 076901
[2] Accadia T *et al* 2012 *J. Instrum.* **7** P03012
[3] Grote H (LIGO Scientific Collaboration) 2010 *Class. Quantum Grav.* **27** 084003
[4] Aasi J *et al* (LIGO Scientific Collaboration and Virgo Collaboration) 2013 *Phys. Rev.* D **87** 022002
[5] Abadie J *et al* LIGO Scientific Collaboration and Virgo Collaboration 2012 *Phys. Rev.* D **85** 082002
[6] Abadie J *et al* (The LIGO Scientific Collaboration and The Virgo Collaboration) 2012 *Phys. Rev.* D **85** 122007
[7] Abbott B *et al* (The LIGO Scientific Collaboration and The Virgo Collaboration) 2010 *Astrophys. J.* **713** 671
[8] Abbott B *et al* (LIGO Scientific Collaboration and Virgo Collaboration) 2009 *Nature* **460** 990–4
[9] Abadie J *et al* LIGO Scientific Collaboration and Virgo Collaboration 2012 *Phys. Rev.* D **85** 122001
[10] Harry G (LIGO Scientific Collaboration) 2010 *Class. Quantum Grav.* **27** 084006
[11] Acernese F *et al* (VIRGO Scientific) 2009 *Virgo Internal Report* VIR027A09 (https://tds.ego-gw.it/itf/tds/file.php?callFile=VIR-0027A-09.pdf)
[12] Willke B *et al* 2006 *Class. Quantum Grav.* **23** S207
[13] Kuroda K LCGT Collaboration 2010 *Class. Quantum Grav.* **27** 084004
[14] Christensen N and Meyer R 1998 *Phys. Rev.* D **58** 082001
[15] Veitch J and Vecchio A 2010 *Phys. Rev.* D **81** 062003
[16] Feroz F, Gair J, Graff P, Hobson M and Lasenby A 2010 *Class. Quantum Grav.* **27** 075010
[17] Feroz F, Gair J, Hobson M and Porter E 2009 *Class. Quantum Grav.* **26** 215003
[18] Aasi J *et al* LIGO-Virgo Scientific Collaboration 2013 *Phys. Rev.* D **88** 062001
[19] Logue J *et al* 2012 *Phys. Rev.* D **86** 044023
[20] Rover C *et al* 2009 *Phys. Rev.* D **80** 102004
[21] Abadie J *et al* (LIGO Scientific Collaboration and Virgo Collaboration) 2010 *Class. Quantum Grav.* **27** 173001
[22] Aasi J *et al* (LIGO Scientific Collaboration and Virgo Collaboration) 2014 *Phys. Rev. Lett.* **112** 131101
[23] Sutton P *et al* 2010 *New J. Phys.* **12** 053034
[24] Klimenko S, Yakushin I, Mercer A and Mitselmakher G 2008 *Class. Quantum Grav.* **25** 114029
[25] Thrane E *et al* 2011 *Phys. Rev.* D **83** 083004
[26] Thrane E and Coughlin M 2013 *Phys. Rev.* D **88** 083010
[27] Thrane E and Coughlin M 2014 *Phys. Rev.* D **89** 063012
[28] Alford M *et al* 2012 *Phys. Rev.* D **85** 044051
[29] Owen B *et al* 1998 *Phys. Rev.* D **58** 084020
[30] Prestegard T and Thrane E 2012 *LIGO DCC* L1200204 (https://dcc.ligo.org/cgi-bin/DocDB/ShowDocument?docid=93146)
[31] De Rosa R, Forte L, Garufi F and Milano L 2012 *Phys. Rev.* D **85** 042001

[32] Sathyaprakash B S and Schutz B 2009 *Living Rev. Relativ.* **12** 2
[33] Skilling J 2004 *AIP Conf. Proc. of the 24th Int. Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering* **735** 395–405
[34] Feroz F, Hobson M P and Bridges M 2009 *Mon. Not. R. Astron. Soc.* **398** 1601–14
[35] Pitkin M and Romano J 2013 http://ccpforge.cse.rl.ac.uk/gf/project/multinest/frs/
[36] Littenberg T *et al* 2013 *Phys. Rev.* D **88** 084044
[37] Finn L S 1992 *Phys. Rev.* D **46** 5236–49
[38] Cutler C and Flanagan E E 1994 *Phys. Rev.* D **49** 2658–97
[39] Rodriguez C L, Farr B, Farr W M and Mandel I 2013 *Phys. Rev.* D **88** 084013
[40] Durrett R 2010 *Probability: Theory and Examples (Cambridge Series on Statistical and Probabilistic Mathematics)* (Cambridge: Cambridge University Press)