Gravitational wave detection using multiscale chirplets

E J Candès¹, P R Charlton² and H Helgason¹

¹ Division of Applied and Computational Mathematics, California Institute of Technology, Pasadena, CA 91125, USA

² School of Computing and Mathematics, Charles Sturt University, Wagga Wagga, NSW 2678, Australia

E-mail: emmanuel@acm.caltech.edu, pcharlton@csu.edu.au and hannes@acm.caltech.edu

Received 28 March 2008, in final form 18 June 2008

Published 2 September 2008

Online at stacks.iop.org/CQG/25/184020

Abstract

A generic ‘chirp’ of the form $h(t) = A(t) \cos \phi(t)$ can be closely approximated by a connected set of *multiscale chirplets* with quadratically-evolving phase. The problem of finding the best approximation to a given signal using chirplets can be reduced to that of finding the path of minimum cost in a weighted, directed graph, and can be solved in polynomial time via dynamic programming. For a signal embedded in noise we apply constraints on the path length to obtain a statistic for detection of chirping signals in coloured noise. In this paper we present some results from using this test to detect binary black hole coalescences in simulated LIGO noise.

PACS numbers: 04.80.Nn, 07.05.Kf, 95.30.Sf, 95.85.Sz

1. Introduction

Despite having achieved unprecedented sensitivities, experiments for laser interferometric detection of gravitational waves such as LIGO [1] face significant challenges, not least of which is the problem of detecting unmodelled or poorly modelled sources of gravitational waves. For detecting the inspiral of a binary system, the standard technique is *matched filtering* using a bank of templates parametrized by the component masses of the system. For low-mass binaries, the time evolution of the inspiral is well-modelled by post-Newtonian approximations, however for high-mass binaries the models are considerably less certain [2]. Furthermore, as the binary mass increases, the spin of the two bodies becomes a significant factor in the evolution of the signal [3]. A complete description of a binary system including the spin of both bodies requires 17 parameters, making the set of templates to be searched over infeasibly large. Even when some parameters are neglected, estimates of the number of templates needed to detect, for example, spinning extreme mass ratio inspirals using a space-based detector such as LISA range from 10^{15} to 10^{40} templates [4, 5]. Methods have been proposed to reduce the number of templates required, such as by using detection template

families which cover the expected range of gravitational wave signals [6, 7], but these still require $\sim 10^5$ templates [8].

Template methods for detecting binary coalescence events mostly focus on the inspiral component or the ringdown component [9] and do not attempt to match the merger component, believed to be a major contribution to the gravitational signature for black hole coalescences. Modelling the inspiral and ringdown is relatively straightforward, whereas modelling the merger requires robust techniques for solving the full Einstein equations numerically under extreme conditions. Much progress has been made in achieving this goal but the problem is far from solved [10, 11].

A number of potential gravitational wave signals are of short duration (less than 1 s) and are collected under the heading of *burst sources*. These include events such as supernovae, the final stages of binary black hole coalescence, and other potential sources of gravitational waves such as gamma-ray bursts. Generally, models for these sources are either non-existent or insufficient for constructing matched filters, and we must rely on *non-parametric* methods. Various methods for detecting bursts have been proposed [12–17], and some have been applied to interferometer data [18].

In this paper we apply a non-parametric detection scheme called the *best path (BP) test* introduced in [19] to the detection of binary black hole coalescences in simulated LIGO noise. The terminology comes from the study of weighted graphs and refers to the path between two vertices of a graph which is of maximum total weight, subject to a constraint on its length. The BP test is applicable to the detection of quasi-periodic signals of the form

$$h(t) = A(t) \cos \phi(t), \quad (1)$$

where the amplitude $A(t)$ varies slowly with time and the unknown phase $\phi(t)$ obeys some regularity conditions. Signals of this form have a well-defined *instantaneous frequency* $f(t) = \dot{\phi}(t)/2\pi$ (to avoid confusion, we note that there is an unrelated method called the fast chirp transform which is applicable to the detection of signals of the form (1) where the phase function is known [20]).

2. Chirplet path pursuit

Given detector output

$$u(t) = n(t) + \rho h(t), \quad (2)$$

where $n(t)$ is Gaussian coloured noise with two-sided power spectral density $S(f)$, we seek a test statistic which will discriminate between the two hypotheses

$$H_0 : \rho = 0 \quad H_1 : \rho \neq 0. \quad (3)$$

The null hypothesis is that the data are pure noise, while the alternative is that the data contain a chirp-like signal of the form (1), normalized with respect to the inner product derived from $S(f)$,

$$\langle u, v \rangle = \int_{-\infty}^{\infty} \frac{\tilde{u}^*(f) \tilde{v}(f)}{S(f)} df. \quad (4)$$

The parameter ρ may be interpreted as the expectation of the SNR,

$$\text{SNR} = \frac{\langle u, h \rangle}{\text{rms}\langle n, h \rangle}. \quad (5)$$

Note that a one-sided PSD is more commonly used in the literature, equivalent to $2S(|f|)$. We use the two-sided PSD here to simplify the discretized form of (4).

Locally, chirps with smoothly-varying phase have a very simple structure. Over short times their frequency evolution is approximately linear. For longer duration, local approximations can be joined together so that the instantaneous frequency of the signal is approximated by a piecewise linear function. In the following we outline the methodology for obtaining a test statistic via chirplet path pursuit—details may be found in [19].

2.1. Multiscale chirplets

Consider a signal on the interval $I = [0, T)$. The preceding discussion suggests we should examine functions which will correlate well locally with signals of the form (1). Our detection method uses a dictionary of normalized *multiscale chirplets* of the form

$$c_{s,j,a,b}(t) \propto e^{i2\pi(at+bt^2/2)}, \quad t \in I_{s,j} \subseteq I \quad (6)$$

that is, a collection of chirplets supported on intervals $I_{s,j}$ and parametrized by length scale s , location j , initial frequency a and chirp rate b . The intervals are taken to be *dyadic* of the form $I_{s,j} = [j2^{-s}T, (j+1)2^{-s}T]$. Here $s = 0, 1, 2, \dots$, represents a scale index and defines the length of the dyadic interval. The dictionary has elements of various durations, locations, initial frequencies and chirp rates. It is convenient to think of a chirplet as a line segment $a + bt$ supported on $I_{s,j}$ in the time–frequency plane.

Our test statistic is constructed by looking for a connected ‘path’ of chirplets in the time–frequency plane that gives a good overall correlation with the signal. To achieve this we notionally discretize the time–frequency plane and consider points (t_i, f_k) as vertices in a directed graph. The frequency intervals may be chosen as convenient—for example, to coincide with bins of a discrete Fourier transform. Fixing a time–frequency discretization also fixes the discretization of the chirp parameter, since we think of chirplets as arcs connecting vertices of the graph supported on dyadic intervals. Using the FFT we can quickly calculate the local correlations $|\langle u, c_{s,j,a,b} \rangle|^2$ of $u(t)$ with elements of the chirplet dictionary, which we use as the weights of the arcs connecting each vertex in the graph. Given a connected, non-overlapping chirplet path $P = \{c_1, c_2, \dots, c_p\}$ supported on a partition $\mathcal{P} = \{I_1, I_2, \dots, I_p\}$ of I the total weight of the path is $\sum_p |\langle u, c_p \rangle|^2$. A description of our discretization scheme may be found in the appendix.

Simply maximizing $\sum_p |\langle u, c_p \rangle|^2$ over all chirplet paths will naively overfit the data. In the limit of small chirplets, such a statistic would simply fit $u(t)$ rather than a hidden signal. Instead we use a *multivariate* statistic obtained as the solution of the optimization problem

$$T_\ell^* = \max_P \sum_p |\langle u, c_p \rangle|^2 \quad \text{subject to} \quad |P| \leq \ell. \quad (7)$$

Here ℓ is a constraint on the path length ie. the number of chirplets in the path. To be adaptive, we calculate T_ℓ^* for several different path lengths, $\ell \in L = \{\ell_1, \ell_2, \dots\}$. While there are a vast number of possible paths, using a variant of Dijkstra’s algorithm, calculating T_ℓ^* reduces to a constrained dynamic programming problem which can be solved in $O(|L| \times \#\text{arcs})$ [21]. The number of arcs depends on such things as choice of discrete frequencies and chirp rates, but is typically not more than $N^2 \log_2 N$.

Since T_ℓ^* is a multivariate statistic we use a *multiple comparison* rule for rejecting the null hypothesis [22]. Given data $u(t)$ we test H_0 at false alarm probability α using the following procedure:

- (1) For each $\ell \in L$, calculate T_ℓ^* and find the corresponding p -value under H_0 , p_ℓ .
- (2) Compare the minimum p -value $p^* = \min_\ell p_\ell$ with the distribution of minimum p -values under H_0 .

Figure 1. The best path found for a binary inspiral signal with total mass $16 M_{\odot}$ in white noise, indicated by the dashed curve. Vertical lines delimit the support of individual chirplets in the path. Note that the BP test uses long chirplets when the frequency is changing slowly, and short chirplets when it is changing rapidly.

(This figure is in colour only in the electronic version)

- (3) If p^* is small enough to lie in the α -quantile of the distribution, reject H_0 —we conclude a signal is present.

In this procedure we are choosing the ordinate of the multivariate test statistic that gives the greatest evidence against the null hypothesis. We then compare this p -value with what one would expect under the null hypothesis. Although there do not exist analytic expressions for the distributions of T_{ℓ}^* and the minimum p -value, we can estimate them using Monte Carlo simulations. We call T_{ℓ}^* the *best path (BP) statistic*. As an example, figure 1 shows the path obtained for an inspiral signal in white noise.

3. Simulations

3.1. Noise model

To estimate the statistical power of the BP test we have studied the detection of certain gravitational wave signals in simulated LIGO noise. Discretely sampled Gaussian noise is produced via the following method. We generate two sequences of white noise a_k and b_k , then construct a discrete Fourier representation of an instance of coloured noise $\tilde{n}_k$ using the PSD as follows:

$$\begin{aligned} \tilde{n}_0 &= \left[\frac{NS_k}{\Delta t} \right]^{\frac{1}{2}} a_0 \\ \tilde{n}_k &= \left[\frac{NS_k}{\Delta t} \right]^{\frac{1}{2}} \frac{a_k + ib_k}{2} \quad k = 1, \dots, N/2 - 1 \end{aligned}$$

$$\tilde{n}_{N/2} = \left[\frac{N S_{N/2}}{\Delta t} \right]^{\frac{1}{2}} a_{N/2}$$

$$\tilde{n}_k = \tilde{n}_{N-k}^* \quad k = N/2 + 1, \dots, N - 1. \quad (8)$$

By construction, the inverse DFT n_k is real Gaussian noise with PSD S_k . The PSD used is the polynomial fit given in ([23], table 5). This fit is only valid for frequencies above the LIGO-I seismic wall frequency $f_s = 40$ Hz. Seismic noise renders region below f_s inaccessible to gravitational wave searches. For the purposes of simulation, we mimic high-pass filtered data by rolling off $S(f)$ below 20 Hz. When calculating the BP statistics we only search over paths with instantaneous frequencies above f_s . At this time we have not included non-Gaussian features such as instrumental bursts in our noise model.

3.2. Signal model

Since the object of the exercise is to detect ‘real’ gravitational waves, we will use as our test signals a collection of physically realistic waveforms for binary black hole coalescence. We use a modification of the method in [12] to model a complete coalescence waveform. The signal consists of an inspiral component, a merger component and a ringdown component. While the inspiral and ringdown models are reasonable, the simulated merger should not be taken to be physically realistic. Instead, it is meant to approximate the overall time and frequency characteristics of a real merger.

The test signals are parametrized by the total mass $M = m_1 + m_2$ of the two bodies and the symmetric mass ratio $\eta = m_1 m_2 / M^2$. The full waveform is obtained by combining the components in such a way that the instantaneous frequency and amplitude are continuous up to first derivatives:

$$h(t) = \begin{cases} A^{\text{insp}}(t) \cos \phi^{\text{insp}}(t) & t \leq 0 \\ A^{\text{merge}}(t) \cos \phi^{\text{merge}}(t) & 0 < t \leq t_m \\ A^{\text{ring}}(t) \cos \phi^{\text{ring}}(t) & t_m < t. \end{cases} \quad (9)$$

Here we have arranged for the inspiral component to end at $t = 0$ and the merger component to end at $t = t_m$. Following [24] we take the merger duration to be $t_m = 50M/M_\odot \times T_\odot$.

For the inspiral component of the signal we use the non-spinning 2PN approximation for the phase in the form given by ([25], equation (15.24)). For amplitude we use the leading order (i.e. Newtonian) expression given in ([25], equations (15.27) and (15.28)). For simplicity we average over orientation (ι, β) and sky position to obtain

$$A^{\text{insp}}(t) = \frac{8}{5} \frac{T_\odot c}{D} \frac{\eta M}{M_\odot} \left[\frac{\pi T_\odot M f^{\text{insp}}(t)}{M_\odot} \right]^{2/3}, \quad (10)$$

where D is the distance to the source.

We model the inspiral component from the time the instantaneous frequency enters the sensitive band of the detector above f_s up to the commencement of the merger component. Deciding where the boundary between inspiral and merger lies is somewhat arbitrary. We follow [24] in making the transition at the point where post-Newtonian approximations begin to break down. It is convenient to fix this transition at $t = 0$. A conservative estimate [24] is that errors in the 2PN approximation become significant when the instantaneous frequency reaches

$$f_0 = \frac{M_\odot}{M} \times 4100 \text{ Hz}, \quad (11)$$

so we set the coalescence time t_c of the inspiral in ([25], equation (15.24)) by solving $f^{\text{insp}}(0) = f_0$.

The ringdown component is assumed to be an exponentially damped sinusoid with constant frequency f^{ring} as given in ([25], equation (18.3)). Our amplitude model, adapted from [26], is

$$A^{\text{ring}}(t) = \frac{\mathcal{A}}{\sqrt{20\pi}} \frac{T_{\odot} c}{D} \frac{M}{M_{\odot}} e^{-\pi f^{\text{ring}}(t-t_m)/Q}, \quad (12)$$

where a is the dimensionless spin parameter, $Q = 2(1-a)^{-0.45}$ is the quality factor,

$$\mathcal{A} = 4 \left[\frac{\pi \epsilon}{Q[1 - 0.63(1-a)^{0.3}]} \right]^{1/2} \quad (13)$$

and ϵ is the fraction of M radiated as gravitational waves during the ringdown. The factor of $1/\sqrt{20\pi}$ in (12) comes from averaging over orientations and sky positions. This is essentially the same amplitude model as given in ([25], equation (18.5)).

Our inspiral component has been arranged to terminate at $t = 0$, with ringdown commencing at $t = t_m$. Since no analytic models exist for the merger component, we fit the amplitude and phase functions to bridge the gap between inspiral and ringdown. Assuming that the merger waveform is of the form (1), a simple way to connect the inspiral and ringdown waveforms is to require that the amplitude be continuous to first derivatives, and the phase to be continuous up to second derivatives (thus ensuring that the instantaneous frequency is continuous up to first derivatives). This gives four conditions that must be satisfied by $f^{\text{merge}}(t)$ and $A^{\text{merge}}(t)$ at $t = 0$ and $t = t_m$, so we model $f^{\text{merge}}(t)$ and $A^{\text{merge}}(t)$ by cubic polynomials. Since we also require the phase to be continuous at $t = 0$, we obtain $\phi^{\text{merge}}(t)$ from the anti-derivative of $f^{\text{merge}}(t)$ with an appropriate constant of integration. We note that phenomenological templates for coalescing binaries have recently become available [27], however our waveforms are qualitatively very similar, and for testing purposes it is convenient to know the exact form of the instantaneous frequency and be able to set the precise times of transition from inspiral to merger to ringdown. Phenomenological templates will be examined in future work.

3.3. Choice of signal parameters

To test detection efficiency we used signals of length $N = 512$, $N = 1024$ and $N = 2048$ sampled at 2048 Hz. Signals of roughly this duration are produced by BBH systems with total mass in the range 20–50 $M_{\odot}$. As most models for the ringdown waveforms assume equal mass binaries, we will only consider this case. The masses used were $m_1 = m_2 = 22.5, 15$ and 10 . Motivated by recent numerical experiments [10, 11], we take $a = 0.7$ and $\epsilon = 0.01$. While the procedure for producing a merger waveform is crude, it does produce a signal with frequency and amplitude characteristics similar to those seen in numerical relativity simulations. Figure 2 shows the strain and instantaneous frequency for these binary coalescences at a distance of 1 Mpc for the $M = 45$ and $30 M_{\odot}$ cases.

4. Results

To use the BP test we first need the distribution of T_{ℓ}^* under H_0 . There is no analytic expression for the distribution of the BP statistic so we have used a Monte Carlo simulation to estimate them. As our test signals have different lengths we generated three null distributions, one for each N . In each case we generated 10^5 instances of simulated LIGO noise and calculated T_{ℓ}^* for each of them with chirplet path lengths ℓ drawn from the set $L = \{1, 2, 4, 8, 16\}$. These random trials give an approximation to the distributions of T_{ℓ}^* under H_0 for each ℓ . Using our empirical distributions we can estimate the p -value for an observed T_{ℓ}^* .

Figure 2. $h(t)$ and instantaneous frequency for a binary coalescence with masses (a) $m_1 = m_2 = 22.5 M_\odot$ and (b) $m_1 = m_2 = 15 M_\odot$ at a distance of 1 Mpc.

To test detection efficiency, we first constructed normalized test signals using the model described in section 3.2. For each $\rho = 7, 8, 9, 10, 11$ and 12 we generated 10^5 instances of noise and injected the signal at that level. The BP statistic T_ℓ^* was calculated, as was p^* , and we determined the detection probability for a given α by counting the number of $p^* \leq \alpha$. Figure 3 gives the detection probabilities as a function of α (the receiver operating characteristic curve) for $\rho = 8, 10$ and 12 . For comparison, we also give an ROC curve obtained using matched filtering to detect the signal. For these curves ρ has been chosen to give a good match to the ROC curve obtained via the BP test. From this it can be seen that the BP statistic is about half as sensitive as matched filtering. Since the distance D to the source is inversely proportional to the overall signal amplitude, we can consider the BP test to have a seeing distance about half that of matched filtering.

In figure 3 we also give the detection probability as a function of ρ (or equivalently, inverse distance to the source) for $\alpha = 0.05, 0.01$ and 0.001 . The corresponding distances at $\rho = 10$ are $D = 100$ Mpc for $M = 45 M_\odot$, $D = 80$ Mpc for $M = 30 M_\odot$ and $D = 65$ Mpc for $M = 20 M_\odot$. This shows that, for example, at a false alarm probability of $\alpha = 0.001$ we can see an event out to ~ 100 Mpc with a false dismissal probability of about 10%. Note that since we have averaged the signal amplitude over sky positions and orientations, an optimally aligned and positioned source could be detected much farther away.

In the above comparison we are injecting a known signal into noise and using the same (normalized) signal as our template for matched filtering. Real signals in interferometer data will have unknown parameters, and a bank of templates using discrete values of the parameters (mass, spin etc) is needed to cover the range of physically plausible coalescences. Since a real signal has parameters drawn from a continuum there will usually be some degree of mismatch between the signal and templates in the bank. As such, the comparison above is very conservative in comparing the BP test with the most favourable matched filter detection scenario, one which is unlikely to be attained in practice. A more realistic benchmark is obtained by examining the performance of the BP test when the signal parameters are chosen at random from a range of values. Here we present a comparison of the BP test with detection via a bank of templates, and with another method employed in searches for unmodelled signals, the *excess power* statistic [13].

We first created a bank of templates using discrete values for the parameters. Although our complete signal model contains a large number of free parameters, for simplicity we chose

Figure 3. Detection probability as a function of ρ and false alarm probability for BBH coalescences with total mass (a) $M = 45 M_\odot$, (b) $M = 30 M_\odot$ and (c) $M = 20 M_\odot$. In each case we give an estimate for the ρ which gives a similar curve using matched filtering.

to only vary m_1, m_2 and a , as these have the greatest effect on the waveform. For the same reason we have used equal spacing in all parameters, rather than attempting to construct a template bank spaced to give equal overlap between adjacent templates. While methods exist

Figure 4. Detection probability as a function of false alarm probability for random signals in the mass ranges (a) $20.5\text{--}30.4 M_{\odot}$, (b) $13.5\text{--}20.4 M_{\odot}$ and (c) $9.5\text{--}13.4 M_{\odot}$. In each case we give an estimate for the ρ which gives a similar curve using matched filtering with a bank of templates and the excess power statistic.

to construct optimally-spaced template banks, our templates have the additional complication of including merger and ringdown components.

For each of the signal lengths $N = 512, 1024$ and 2048 we generated a bank of normalized templates using the criteria that

- (1) The range of masses m_1 and m_2 is chosen so that the length of the signals ranges from $N/2$ to N samples.
- (2) The spin ranges from $a = 0.18$ to 0.98 .
- (3) The spacing between masses and spins is chosen so that the minimal match of a signal with parameters drawn from the range of parameters is at least 0.97 .

For each N we then generated 1000 test signals with mass and spin parameters drawn at random from the appropriate range, and injected them into simulated LIGO noise with $\rho = 10$. The resulting data were used to calculate a BP statistic for each segment. Comparing the BP statistics with the empirical null distribution as above, we obtained the ROC curves shown in figure 4. Searching for the same signals via matched filtering, we found that the ROC curves

matched well when the signals were injected with ρ around 6.5—in other words, the BP test sees about 2/3 as far as the template bank.

We performed a similar analysis using the excess power statistic, which is optimal when the only known features of the signal are the duration and bandwidth [13]. The excess power statistic is simply the power $\langle u, u \rangle$ calculated using (4) where the integration is performed over the bandwidth of the expected signals, taken to be 40–1024 Hz in this instance. Under H_0 this has a χ^2 distribution with degrees of freedom twice the number of frequency bins. Using the excess power statistic, we found that the ROC curves matched those of the BP test in figure 4 well when ρ was around 14.

5. Conclusion

Chirplet path pursuit has previously been shown to be effective at detecting a broad class of chirp-like but otherwise unmodelled signals in coloured noise [19]. In this paper we have demonstrated that the method can be successfully applied to the problem of detecting test signals with similar characteristics to those expected from binary black hole coalescence. The method is able to detect a range of signals of modest strength hidden in simulated LIGO noise, and exhibits somewhat better statistical power than the excess power test.

As with other methods for detecting bursts, in real LIGO noise there is the difficulty of distinguishing genuine gravitational wave signals from instrumental and environmental events. For matched filter searches the χ^2 discriminator can be used to reject signals that do not have the correct distribution of power across frequency bands, however this requires that the gravitational waveform be known [28]. This discriminator is not applicable to chirplet path pursuit since the signal is not known and we do not impose any assumptions on the distribution of power. Instead we would rely on the methods being employed in current searches: requiring events to be coincident across multiple detectors, vetoing events based on environmental channels and testing if waveforms measured in different detectors are consistent [18, 29].

As expected for a non-parametric method, chirplet path pursuit is not as sensitive as matched filtering using a template bank, nevertheless our comparison shows that the method has similar effectiveness to matched filtering for a signal that is roughly 1.5 times as strong. Significantly, since the method is sensitive to a wide range of chirp-like signals, an exact model of the signals to be detected is not necessary. This makes the method particularly of interest in situations where the signal is unmodelled or poorly modelled, as is the case for the late inspiral and merger components of intermediate mass black hole coalescences.

Acknowledgments

EC was partially supported by National Science Foundation grants DMS 01-40698 (FRG) and ITR ACI-0204932. We thank Warren Anderson for supplying his Maple code for simulating BBH coalescence. This paper has LIGO document number LIGO-P080017-00-Z.

Appendix

In this appendix we present the scheme used for calculating chirplet coefficients of discretized data. The data $u[n] = u(n\Delta t)$ are discretely sampled at $N = 2^S$ intervals of duration Δt . Notionally, this discretizes the time–frequency plane into points (t_i, f_k) , where $t_i = i\Delta t$ and $f_k = k\Delta f = k/(N\Delta t)$. Points in the time–frequency plane are considered to be vertices

in a directed graph where the weight of the arc connecting two vertices is given by the local correlation of $u(t)$ with the corresponding chirplet.

Consider chirplets supported on the interval $[0, 2^{-s}T)$. At scale s , $0 \leq s < S$ this interval has length $N_s = 2^{-s}N$ samples. While there are many ways to discretize chirplets on this interval, it is convenient to choose the spacing of the frequency parameter to correspond with the bins of a discrete Fourier transform, and choose the spacing of the chirp parameter so that at the end of the interval the instantaneous frequency has changed by a whole number of bins. Thus our dictionary of chirplets is indexed by scale index s , frequency index k and chirp index l , and the (unnormalized) discrete chirplet is given by

$$c_{s,k,l}[n] = e^{i2\pi\phi_{s,k,l}[n]} \quad 0 \leq n < N_s, \quad 0 \leq k \leq N/2, \quad (A.1)$$

where the phase is

$$\phi_{s,k,l}[n] = k \frac{n}{N} + l \frac{n^2}{2NN_s}. \quad (A.2)$$

The discretized instantaneous frequency is

$$\dot{\phi}_{s,k,l}[n] = k + l \frac{n}{N_s}. \quad (A.3)$$

Such a chirplet has initial frequency $k\Delta f$ Hz and rises to frequency $(k+l)\Delta f$ Hz at a rate of $l\Delta f/(N_s\Delta t)$ Hz s^{-1} . Since we only deal with real signals, the range of the chirp index l is chosen to restrict the chirplets to non-negative frequencies up to Nyquist, thus $-k \leq l \leq N/2 - k$.

In general, the inner product (4) for a noise process with covariance matrix Σ is $u^*\Sigma^{-1}v$, where u^* is the conjugate transpose of u . For our noise model the Fourier matrices $F_{mn} = e^{-i2\pi mn/N}$ diagonalize Σ , and so $\langle u, v \rangle = u^*F^*D^{-1}Fv = \tilde{u}^*D^{-1}\tilde{v}$ where $D = \text{diag}(\sigma_0^2, \sigma_1^2, \dots, \sigma_{N-1}^2)$ and $\sigma_k^2 = \langle |\tilde{n}_k|^2 \rangle$ are the eigenvalues of Σ . Calculating $\langle u, c \rangle$ for a chirplet supported on a dyadic interval $I_{s,j} = [j2^{-s}T, (j+1)2^{-s}T)$ is equivalent to calculating the inner product of $u(t + j2^{-s}T)$ with a chirplet supported on $[0, 2^{-s}T)$. As the time index of the first sample in $I_{s,j}$ is jN_s , let $u_{s,j} = (u[jN_s], u[jN_s+1], \dots, u[(j+1)N_s-1])$ be the samples of $u(t)$ restricted to $I_{s,j}$. Then to find $\langle u, c \rangle$ we pad $u_{s,j}$ and c to length N with zeros and FFT. In discrete form, the inner product then reduces to

$$\langle u, c \rangle = \frac{\Delta t}{N} \sum_{n=0}^{N-1} \frac{\tilde{u}_{s,j}^*[n]\tilde{c}[n]}{S[n]} \quad (A.4)$$

where $S[n] = S(n\Delta f)$. If c has indices k, l then after normalizing, $|\langle u, c \rangle|^2/|\langle c, c \rangle|^2$ is the weight of the arc connecting (t_{jN_s}, f_k) to $(t_{(j+1)N_s}, f_{k+l})$.

To calculate the BP statistic we must find the total weight of connected, non-overlapping chirplet paths in the time–frequency plane starting at $t = 0$ and ending at $t = T$. To keep the number of arcs manageable we further restrict our chirplet paths to those supported on a *recursive dyadic partition* (RDP) of I constructed using the following definition [30]:

- (1) The trivial partition $\mathcal{P} = \{I\}$ is an RDP.
- (2) If $\mathcal{P} = \{I_1, I_2, \dots, I_p\}$ is an RDP, then so is the partition obtained by splitting any interval I_j into two adjacent dyadic intervals.

This means that, for example, $\{[0, 1/4), [1/4, 1/2), [1/2, 1)\}$ is a recursive dyadic partition of $[0, 1)$, but $\{[0, 1/4), [1/4, 1)\}$ is not. The total weight of a chirplet path $P = \{c_1, c_2, \dots, c_p\}$ supported on $\mathcal{P} = \{I_1, I_2, \dots, I_p\}$ is then

$$T_P = \sum_p \frac{|\langle u, c_p \rangle|^2}{|\langle c_p, c_p \rangle|^2}. \quad (A.5)$$

References

- [1] <http://www.ligo.caltech.edu>
- [2] Damour T, Iyer B R and Sathyaprakash B S 2001 *Phys. Rev. D* **60** 044023
- [3] Apostolatos T A, Cutler C, Sussman G J and Thorne K S 1994 *Phys. Rev. D* **49** 6274
- [4] Sathyaprakash B S and Schutz B F 2003 *Class. Quantum Grav.* **20** S209–18
- [5] Gair J R, Barack L, Creighton T, Cutler C, Larson S L, Phinney E S and Vallisneri M 2004 *Class. Quantum Grav.* **21** S1595–606
- [6] Buonanno A, Chen Y and Vallisneri M 2003 *Phys. Rev. D* **67** 024016
- [7] Pan Y, Buonanno A, Chen Y and Vallisneri M 2004 *Phys. Rev. D* **69** 104017
- [8] Buonanno A, Chen Y, Pan Y, Tagoshi H and Vallisneri M 2005 *Phys. Rev. D* **72** 084027
- [9] Goggin L 2006 *Class. Quantum Grav.* **23** S709–13
- [10] Pretorius F 2005 *Phys. Rev. Lett.* **95** 121101
- [11] Baker J G, Centrella J, Choi D I, Koppitz M and van Meter J 2006 *Phys. Rev. D* **73** 104002
- [12] Anderson W G and Balasubramanian R 1999 *Phys. Rev. D* **60** 102001
- [13] Anderson W G, Brady P R, Creighton J D E and Flanagan É É 2001 *Phys. Rev. D* **63** 142003
- [14] Sylvestre J 2002 *Phys. Rev. D* **66** 102004
- [15] Klimentenko S and Mitselmakher G 2004 *Class. Quantum Grav.* **21** S1819–30
- [16] Chatterji S 2005 The search for gravitational wave bursts in data from the second LIGO science run *PhD Thesis* Massachusetts Institute of Technology
- [17] Chassande-Mottin É and Pai A 2006 *Phys. Rev. D* **73** 042003
- [18] Abbott A *et al* 2007 *Class. Quantum Grav.* **24** 5343–69
- [19] Candès E J, Charlton P R and Helgason H 2008 *Appl. Comput. Harmon. Anal.* **24** 14–40
- [20] Jenet F A and Prince T A 2000 *Phys. Rev. D* **62** 122001
- [21] Joksche H C 1966 *J. Math. Anal. Appl.* **14** 191–7
- [22] Benjamini Y and Hochberg Y 1995 *J. R. Statist. Soc. B* **57** 289–300
- [23] Grishchuk L P, Lipunov V M, Postnov K A, Prokhorov M E and Sathyaprakash B S 2001 *Phys.-Uspekhi* **44** 1–51
- [24] Flanagan É É and Hughes S 1998 *Phys. Rev. D* **57** 4535
- [25] The LIGO Scientific Collaboration, LAL Software Documentation <http://www.lsc-group.phys.uwm.edu/lal/slug/nightly/doc/lsc-nightly.pdf>
- [26] Thorne K S 1987 *300 Years of Gravitation* ed S W Hawking and W Israel (Cambridge: Cambridge University Press)
- [27] Ajith P *et al* 2007 *Class. Quantum Grav.* **24** S689–99
- [28] Allen B 2005 *Phys. Rev. D* **71** 062001
- [29] Abbott A *et al* 2005 *Phys. Rev. D* **72** 122004
- [30] Candès E J 2002 Multiscale chirplets and near-optimal recovery of chirps *Tech. Rep.* Stanford University