Gravitational wave detection using multiscale chirplets

E J Candés[1], P R Charlton[2] and H Helgason[3]

[1,3]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: [1]emmanuel@acm.caltech.edu, [2]pcharlton@csu.edu.au, [3]hannes@acm.caltech.edu

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 near-optimal statistic for detection of chirping signals in coloured noise. In this paper we present some results from using this method to detect binary black hole coalescences in simulated LIGO noise.

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 unmodeled or poorly modeled sources of gravitational waves. For detecting the inspiral of a binary system, the standard technique is *matched filtering* using a bank of templates parametrised by the component masses of the system. For low-mass binaries, the time evolution of the inspiral is well-modeled 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}–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 intermediate-mass black hole coalescences. Modeling the inspiral and ringdown

is relatively straightforward, whereas modeling 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 second) and are collected under the heading of *burst sources*. These include events such as supernovae, the final stages of black hole coalescence, and 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 time-frequency methods for detecting bursts have been proposed [12, 13, 14, 15], and some are currently being applied to interferometer data [16].

In this paper we examine a non-parametric detection scheme which we call the *best path (BP) method*. The terminology comes from the study of weighted graphs and refers to a path between two vertices of a graph which is of maximum total weight, subject to a constraint on it's length. The BP method is applied to the detection of quasi-periodic signals of the form

$$h(t) = A(t)\cos\phi(t) \tag{1}$$

where the amplitude $A(t)$ varies slowly with time and the phase $\phi(t)$ obeys some regularity conditions. Signals of this form have a well-defined *instantaneous frequency* $f(t) = \dot{\phi}(t)/2\pi$.

2. Chirplet path pursuit

Given detector output

$$u(t) = n(t) + \rho\, h(t) \tag{2}$$

where $n(t)$ is Gaussian coloured noise with 1-sided power spectral density $S(f)$, we seek a test statistic which will discriminate between the two hypotheses

$$\begin{aligned} H_0 &: \rho = 0 \\ H_1 &: \rho \neq 0\,. \end{aligned} \tag{3}$$

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

$$\langle u, v\rangle_S = 4\int_0^\infty \frac{\tilde{u}(f)\tilde{v}^*(f)}{S(f)}\,df\,. \tag{4}$$

Locally, chirps with smoothly-varying phase have a very simple structure. Over short times their frequency evolution can be approximated well by a line segment, leading to a good local approximation of the signal. For longer duration, local approximations can be joined together so that the instantaneous frequency of the original signal is approximated by a piecewise polynomial function. In the following we outline the methodology for obtaining a test statistic via best path pursuit – details may be found in [17].

2.1. Multiscale chirplets

Consider a signal on the interval $[0, 1)$. 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 unit *multiscale chirplets* of the form

$$c_{I,a,b}(t) \propto \mathrm{e}^{i2\pi(at+bt^2/2)}, \qquad t \in I \tag{5}$$

where a, b are frequency and chirp rate parameters, supported on dyadic intervals $I = [k2^{-s}, (k+1)2^{-s}]$. Here $s = 0, 1, 2, \ldots$ represents a scale index and defines the length of the dyadic interval. This dictionary, as a time-frequency portrait would reveal, has elements of various durations, locations, initial frequencies and chirp-rates. It is convenient to think of a chirplet as a line segment $a + bt$, $t \in I$ 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 total correlation to the signal. To achieve this we discretise the TF plane and consider coordinates (t_k, f_l) as vertices in a directed graph. Here the discrete times t_k correspond to the start time of dyadic intervals at the finest scale. The frequency intervals may be chosen as convenient – for example, to coincide with bins of a discrete Fourier transform. Fixing a TF discretisation also fixes the the discretisation of chirplet parameters since we think of chirplets as arcs connecting vertices of the graph, supported on dyadic time intervals. Using the FFT we can quickly calculate the local correlations $|\langle u, c\rangle_S|^2$ of $u(t)$ with elements of the chirplet dictionary, which gives the weights of the arcs connecting each vertex in the graph. The total correlation over a chirplet path P is then $\sum_{c\,\in\,P} |\langle u, c\rangle_S|^2$ where $c(t)$ are connected, non-overlapping chirplets and the path starts at the left-hand edge $(t = 0)$ of the TF plane and ends at the right-hand edge $(t = 1)$.

Simply maximising $\sum_{c\,\in\,P} |\langle u, c\rangle_S|^2$ 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 given by

$$T_\ell^* = \max_{P,\ |P|\le\ell} \sum_{c\,\in\,P} |\langle u, c\rangle_S|^2 \tag{6}$$

where 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, $\ell \in L = \{\ell_1, \ell_2, \ldots\}$. 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 $\mathrm{O}(|L| \times \#\text{arcs})$. 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 H_0. To test H_0 at significance α we use 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.
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 we do not have 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.

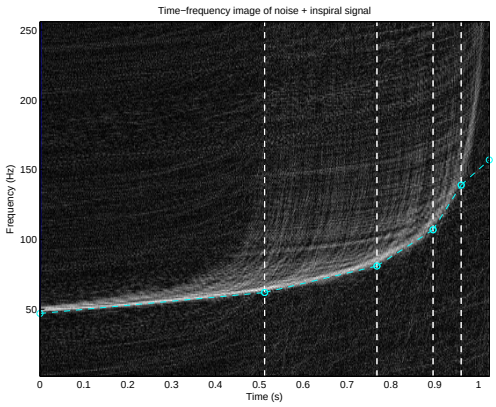

Figure 1. Best path found for a binary inspiral signal with total mass 16 $M_\odot$ in white noise. Note that the BP method by design uses long chirplets when the frequency is changing slowly, and short chirplets when it is changing rapidly.

3. Simulations

3.1. Noise model

To test the effectiveness of the BP method we have simulated the detection of a selection of gravitational wave signals in LIGO noise. Discretely sampled coloured 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:

$$\begin{aligned}
\tilde{n}_0 &= \left[\frac{NS_k}{\Delta t}\right]^{\frac{1}{2}} a_0 \\
\tilde{n}_k &= \left[\frac{NS_k}{\Delta t}\right]^{\frac{1}{2}} \frac{a_k + ib_k}{2} \qquad k = 1, \ldots, N/2 - 1 \\
\tilde{n}_{N/2} &= \left[\frac{NS_{N/2}}{\Delta t}\right]^{\frac{1}{2}} a_{N/2} \\
\tilde{n}_k &= \tilde{n}^*_{N-k} \qquad k = N/2 + 1, \ldots, N - 1. \qquad (7)
\end{aligned}$$

By construction, the inverse DFT n_k is real Gaussian noise with PSD S_k. The PSD used is the polynomial fit given in [18, Table 5]. This fit is only valid for frequencies above the LIGO-I seismic wall frequency $f_s = 40$ Hz. Below f_s, the seismic contribution causes the noise floor to rapidly becomes several orders of magnitude larger than S_0 and so this region is 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 for paths in the region where frequency exceeds f_s.

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 parametrised 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:

$$h(t) = \begin{cases} A^{\mathrm{insp}}(t) \cos \phi^{\mathrm{insp}}(t) & t \leq 0 \\ A^{\mathrm{merge}}(t) \cos \phi^{\mathrm{merge}}(t) & 0 < t \leq t_m \\ A^{\mathrm{ring}}(t) \cos \phi^{\mathrm{ring}}(t) & t_m < t\,. \end{cases} \tag{8}$$

Here we have arranged for the inspiral component to end at $t = 0$ and the merger component to end at $t = t_m$. Following [19] 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 post^2-Newtonian approximation for the phase in the form given by [20, eqn. 15.24]. For amplitude we use the leading order (ie. Newtonian) expression given in [20, eqn. 15.27–28]. For simplicity we average over orientation (ι, β) and sky position to obtain

$$A^{\mathrm{insp}}(t) = \frac{8}{5} \frac{T_\odot c}{D} \frac{\eta M}{M_\odot} \left[\frac{\pi T_\odot M f^{\mathrm{insp}}(t)}{M_\odot} \right]^{2/3} \tag{9}$$

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 [19] 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 [19] is that errors in the 2PN approximation become significant when the instantaneous frequency reaches

$$f_0 = \frac{M_\odot}{M} \times 4100 \text{ Hz} \tag{10}$$

so we set the coalescence time t_c of the inspiral in [20, eqn. 15.24] by solving $f^{\mathrm{insp}}(0) = f_0$.

The ringdown component is assumed to be an exponentially damped sinusoid with constant frequency f^{ring} as given in [20, eqn. 18.3]. Our amplitude model, adapted from [21], is

$$A^{\mathrm{ring}}(t) = \frac{\mathcal{A}}{\sqrt{20\pi}} \frac{T_\odot c}{D} \frac{M}{M_\odot} \mathrm{e}^{-\pi f^{\mathrm{ring}}(t - t_m)/Q} \tag{11}$$

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

$$\mathcal{A} = 4 \left[\frac{\pi \epsilon}{Q\,[1 - 0.63(1-a)^{0.3}]} \right]^{1/2} \tag{12}$$

and ϵ is the fraction of M radiated as gravitational waves during the ringdown. The factor of $1/\sqrt{20\pi}$ in (11) comes from averaging over orientations and sky positions. This is essentially the same amplitude model as given in [20, eqn. 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 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 $f^{\mathrm{merge}}(t)$ and $A^{\mathrm{merge}}(t)$ at $t = 0$ and $t = t_m$, so we model $f^{\mathrm{merge}}(t)$ and $A^{\mathrm{merge}}(t)$ by cubic polynomials. Since we also require the phase to be continuous at $t = 0$, we obtain $\phi^{\mathrm{merge}}(t)$ from the anti-derivative of f^{merge} with an appropriate constant of integration.

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 BH-BH binary 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, respectively. Motivated by recent numerical experiments [10, 11], we take $a = 0.7$ and $\epsilon = 0.01$. While our naive procedure for producing a merger waveform seems very crude, remarkably 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.

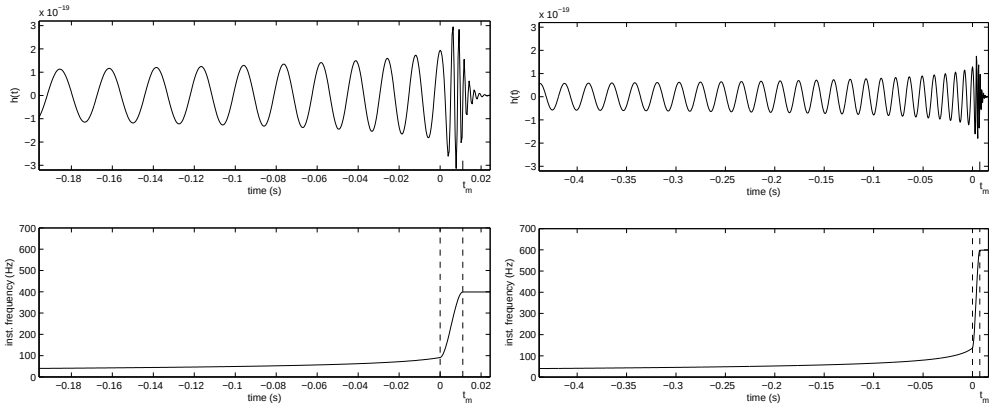

Figure 2. $h(t)$ and instantaneous frequency for a binary coalescence with masses (a) $m_1 = m_2 = 22.5\ M_\odot$ and (b) $m_1 = m_2 = 15\ M_\odot$ at a distance of 1 Mpc.

4. Results

To use the BP method we first need the distributions of T_ℓ^* and p^* 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 coloured 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^*.

To test detection efficiency, we first constructed unit test signals using the model described in Section 3.2. For each SNR $\rho = 7, 8, 9, 10, 11$ and 12 we generated 10^5 instances of noise and injected the signal at that SNR. The BP statistic T_ℓ^* was calculated, as was p^*, and we determined the detection probability for a given false alarm probability α by counting the number of $p^* \leq \alpha$. Figure 3 gives the detection probabilities as a function of α (the so-called Receiver Operator Characteristic curve) for SNRs $\rho = 8$, 10 and 12. For comparison, we also give the curve obtained when using matched filtering to detect the signal. For these curves the SNR of the injected signal has been chosen to give a reasonable match to the BP ROC curve. 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 SNR, another way to think of it is that the BP statistic has a seeing distance half that of matched filtering.

In Figure 3 we also give the detection probability as a function of SNR (or equivalently, inverse distance to the source) for false alarm levels $\alpha = 0.05$, 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 f about 10%. Note that since we have averaged the signal amplitude over sky positions and orientations, D is an average effective distance. An optimally aligned and positioned source could be detected much farther away.

It must be emphasised that in this comparison we are injecting a known signal into noise and using the exact 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, we are being very conservative in comparing the BP statistic with the most favourable matched filter detection scenario, one which is unlikely to be attained in practise.

A more realistic benchmark is to compare the performance of the BP method with matched filtering when the signal parameters are unknown. To test this we 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 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 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 normalised templates using the criteria that

1. The range of masses m_1 and m_2 is chosen so that the signal lengths range from

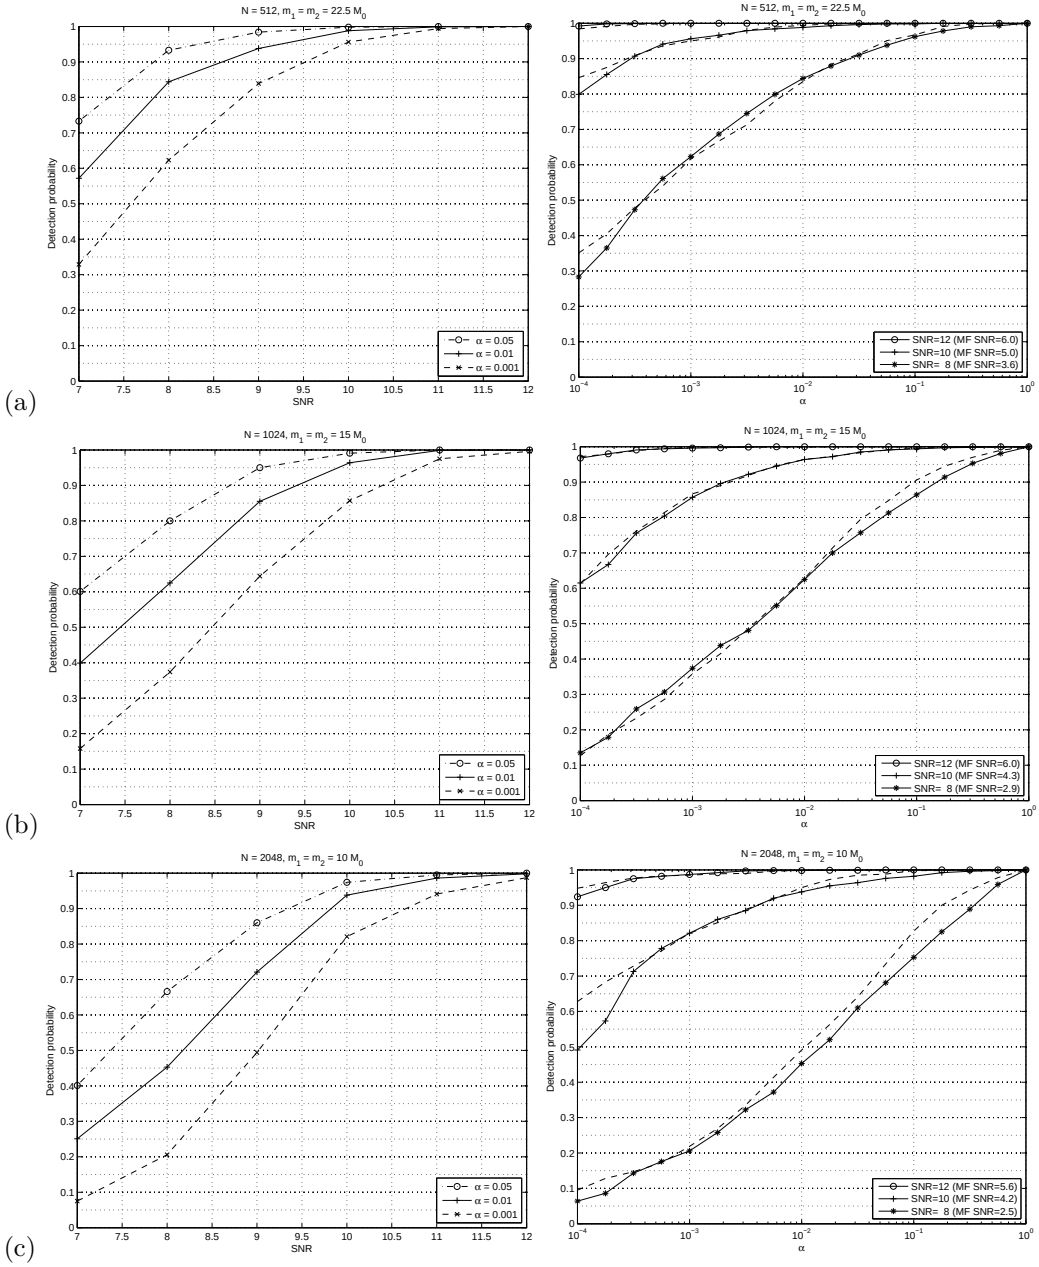

Figure 3. Detection probability as a function of SNR and false alarm probability for BBH coalescences with total mass (a) $M = 45\ M_\odot$ (b) $M = 30\ M_\odot$ and (c) $M = 20\ M_\odot$. In each case we give an estimate for the SNR which gives a similar ROC using matched filtering, represented by the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

$N/2$ 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 greater than 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 at SNR 10. The resulting data was passed through our code to obtain a BP statistic for each segment. Comparing the BP statistics with the empirical null distribution as before, 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 had an SNR of around 6.5 – in other words, the BP method sees about 2/3 as far as the template bank.

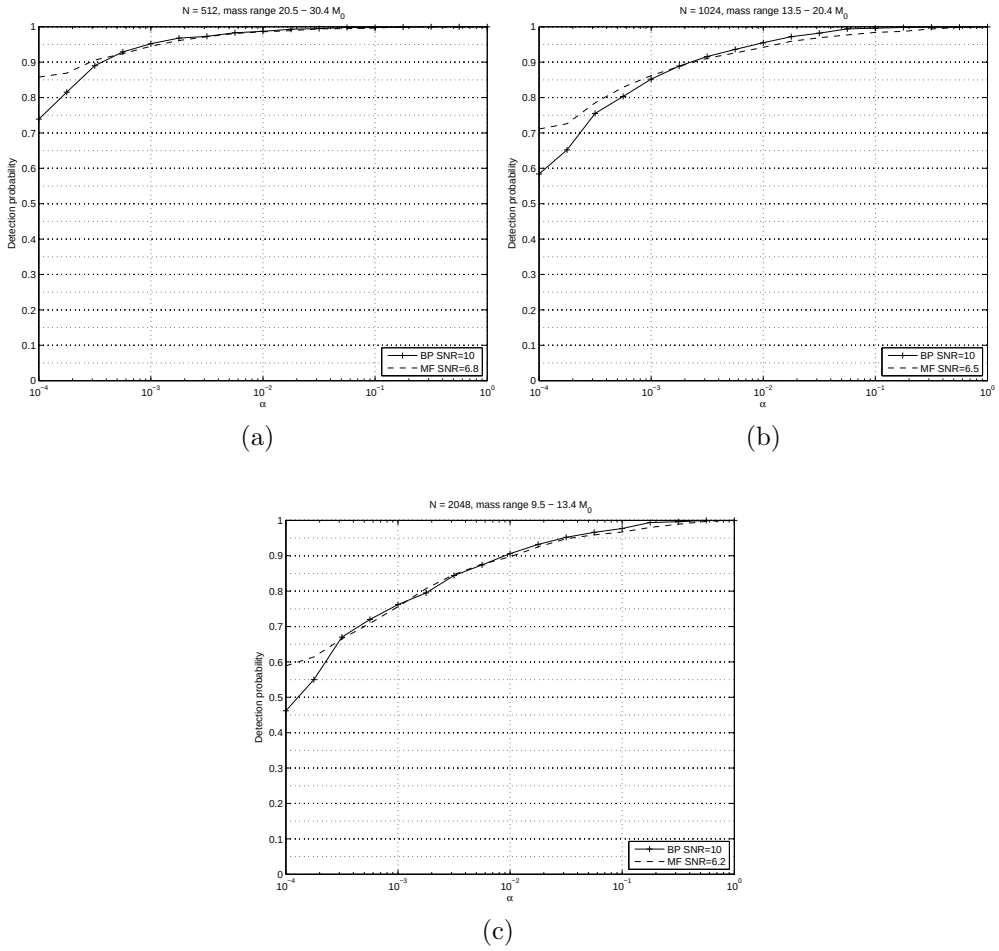

Figure 4. Detection probability as a function of significance for random signals in the mass ranges (a) 20.5–30.4 $M_\odot$ (b) 13.5–20.4 $M_\odot$ and (c) 9.5–13.4 $M_\odot$.

5. Conclusion

Chirplet path pursuit has previously been shown to be effective at detecting a broad class of chirp-like but otherwise unmodeled signals in coloured noise [17]. In this

paper we have applied the method to the problem of detecting signals with similar to characteristics to those expected of gravitational waves. We have demonstrated that the method is effective in detecting a range of test signals of modest strength embedded in simulated LIGO noise.

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 our method has similar effectiveness to matched filtering for a signal that is roughly 1.5 times as loud. 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, as is required in matched filtering. This makes the method particularly of interest in situations where the signal is unmodeled or poorly modeled, 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. PC was partially supported by NSF grant PHY-0107417. We would like to thank Warren Anderson for supplying us with his Maple code for simulating BBH coalescence. This article has LIGO Document Number LIGO-P080017-00-Z.

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–S218
[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–S1606
[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–S713
[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] Sylvestre J 2002 *Phys. Rev. D* **66** 102004
[14] Chatterji S 2005 *The search for gravitational wave bursts in data from the second LIGO science run* Ph.D. thesis Massachusetts Institute of Technology
[15] Chassande-Mottin É and Pai A 2006 *Phys. Rev. D* **73** 042003
[16] A Abbott *et al* 2007 *Class. Quantum Grav.* **24** S1595–S1606
[17] Candés E J, Charlton P R and Helgason H 2008 *Appl. Comput. Harmon. Anal.* **24** 14–40
[18] Grishchuk L P, Lipunov V M, Postnov K A, Prokhorov M E and Sathyaprakash B S 2000 *Physics-Uspekhi* **73**(3)
[19] Flanagan É É and Hughes S 1998 *Phys. Rev. D* **57** 4535
[20] The LIGO Scientific Collaboration LAL Software Documentation `http://www.lsc-group.phys.uwm.edu/lal/slug/nightly/doc/lsd-nightly.pdf`
[21] Thorne K S 1987 *300 Years of Gravitation* ed Hawking S W and Israel W (Cambridge University Press)