Gravitational Wave Detection using Multiscale Chirplets

Emmanuel Candès[1], Philip Charlton[2] and Hannes Helgason[1]

[1]Division of Applied and Computational Mathematics, Caltech
[2]School of Computing and Mathematics, Charles Sturt University

1. Background

A generic 'chirp' 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 a 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[1]. In this poster we present some results from using this method to detect binary black hole coalescences in simulated LIGO noise.

[1]Candès, Charlton and Helgason, "Detecting highly oscillatory signals by chirplet path pursuit", *Appl. Comput. Harmon. Anal.* **24** (2008)

2. Detection problem

We want to test for the presence of a *chirp-like* but otherwise *unknown* signal of the form

$$h(t) = A(t)\cos\varphi(t)$$

under some mild conditions:

- $A(t)$ slowly varying
- $|\dot\varphi(t)| < \pi$ and $|\dot\varphi(t)|^2 >> |\ddot\varphi(t)|$

Such a signal has a well-defined instantaneous frequency and is well-localised along the curve

$$f(t) = \dot\varphi(t)/2\pi$$

in the time-frequency plane. Given detector output

$$u(t) = n(t) + \rho h(t)$$

where $n(t)$ is coloured noise, we seek a statistic which will discriminate between the hypotheses

$$\mathrm{H}_0 : \rho = 0$$

$$\mathrm{H}_1 : \rho \neq 0$$

3. Chirplets

This suggests we should examine functions which will correlate well *locally* with $h(t)$. We define a family of *multiscale chirplets* of the form

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

defined on dyadic subintervals $I = [k2^{-s}, (k+1)2^{-s}]$ where s = 0, 1, 2, ... represents a scale index. Each chirplet is normalised according to the inner product

$$\langle u, v\rangle_\Sigma = u^*\Sigma^{-1}v$$

where Σ is the covariance of $n(t)$. Chirplets have linearly-evolving instantaneous frequency $a + bt$ and form line segments in the TF plane.

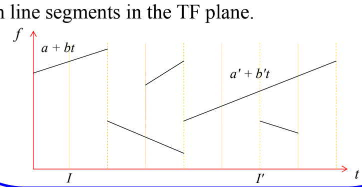

4. Test statistic

Our test statistic is calculated by looking for a connected 'path' of chirplets in the TF plane that gives a good match to the signal. However, simply maximising the sum $\Sigma_{p\in P} |\langle u, c_p\rangle|^2$ over all chirplet paths P will naively overfit the data. In the limit of small chirplets, such a statistic would simply fit u itself rather than a hidden signal. Instead, we use a *multivariate* statistic given by

$$T_\ell^* = \max_{|P|\le\ell} \sum_{p\in P} \left|\langle u, c_p\rangle_\Sigma\right|^2$$

where ℓ is a constraint on the path length $|P|$. In other words, for each possible path length ℓ we find the path of that length which gives the largest total correlation. Although there are a vast number of paths, if we discretise the TF plane and consider points (t_k, f_l) as nodes in a graph with arcs between them having weight $|\langle u, c\rangle|^2$, calculating T_ℓ^* reduces to a *constrained* dynamic programming problem which can be solved in polynomial time, approximately $O(\#\ell \times \#\text{arcs})$.

5. Example of a chirplet path

The figure below shows the time-frequency image of a binary black hole system with $m_1 = m_2 = 8$ solar masses. The best chirplet path constrained to $\ell = 5$ is overlaid, with some representative nodes and arcs of the chirplet graph.

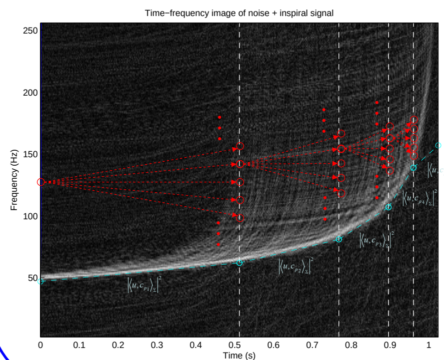

6. Multiple comparison

Since T_ℓ^* is multivariate we have a complex decision rule for rejecting H_0. One approach is to use the Bonferroni approximation: to achieve an overall type I error α we test each T_ℓ^* at significance α/k where k is the number of path lengths used. However, this is known to be conservative. We use the following more powerful multiple comparison:

1. Calculate the p-value for each T_ℓ^* and find the minimum p-value p^*.
2. Compare p^* 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 step 1, we choose the coordinate of T_ℓ^* that gives the greatest evidence against H_0. In step 2, we compare p^* to what we would expect under H_0.

7. Simulated BBH coalescence

We test our method using simulations of binary black hole coalescences with total mass in the range 20–45 solar masses. These signals are good candidates for chirplet analysis because they are

- chirp-like but otherwise poorly modeled
- short: 0.5–2 s

The test signals have three components:

$$h(t) = \begin{cases} h^{\text{insp}}(t) & t \le 0 & \text{(inspiral)} \\ h^{\text{merge}}(t) & 0 < t \le t_{\text{merger}} & \text{(merger)} \\ h^{\text{ring}}(t) & t > t_{\text{merger}} & \text{(ringdown)} \end{cases}$$

Inspiral and ringdown components use standard models from the literature. The "merger" component is simply a chirp signal where amplitude $A(t)$ and instantaneous frequency $\varphi'(t)/2\pi$ have been smoothly connected across the gap using cubic polynomials[2].

[2]We thank Warren Anderson for providing us his Maple code to generate BBH coalescences

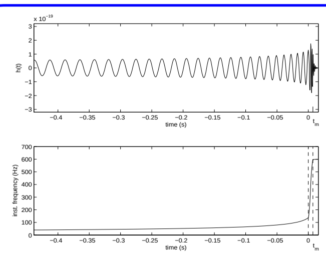

The figure above shows an example of simulated $h(t)$ for the coalescence of a $m_1 = m_2 = 15$ BBH system at 1 MPc. The lower plot is the instantaneous frequency.

8. Results

For the signal above, the figures opposite show the detection rate in simulated LIGO noise for (1) fixed false alarm rates α as a function of SNR (top left) and (2) fixed SNRs as a function of α (top right). A signal at SNR 10 (~80 MPc) has about an 85% chance of being detected.

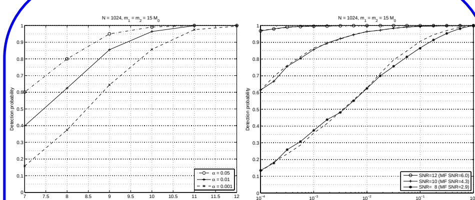

We also show the curve corresponding to the SNR that gives a similar detection rate via matched filtering as if the the signal were known exactly – typically this is about half the SNR required by the chirplet path method. In other words:

The chirplet path method can see a signal about half the distance that a matched filter would see if the signal was known.

This is the cost of a non-parametric detection method that is not targeted at a specific signal – however, the method can detect a *much* larger set of signals than a bank of matched filters.