Improving the Search for Gravitational Waves from the Coalescence of High Mass Black Hole Binaries

LIGO SURF Final Report

Johnathon Lowery
Advisors: Surabhi Sachdev, Tjonnie Li, Kent Blackburn and Alan Weinstein
Indiana University, Bloomington, IN 47405
(Dated: September 23, 2015)

Solutions to Einstein's field equations predict gravitational waves: disturbances in space-time that propagate at the speed of light. Detecting gravitational waves is challenging because the signals are very weak and so a very sensitive instrument is required. The Laser Interferometric Gravitational-Wave Observatory (LIGO) is a pair of detectors that search for these disturbances by looking for small length changes caused by passing gravitational waves. One promising class of sources of gravitational waves is binary black holes. Signals from such sources are searched for in the data from the detectors with an analysis pipeline called `gstlal` using an advanced form of matched filtering which helps to pull out small signals from a background of noise. This report details steps taken toward optimizing the `gstlal` pipeline for Advanced LIGO in the context of binary black hole detection. Specifically, I describe a novel method for the generation of stochastic template banks which is significantly faster than the traditional method and discuss attempts to improve the autocorrelation χ^2 used to veto glitches in the LIGO data.

I. BACKGROUND

A. Gravitational Waves and Their Sources

Gravitational waves (GWs) were first predicted by Einstein in a 1916 paper [1] where he solved the field equations of general relativity (GR) using the weak-field approximation and predicted that accelerating bodies would produce ripples in space-time that would propagate at the speed of light. Gravitational waves offer another way to test general relativity, but as of yet they have not been directly detected.

Strong evidence for the existence of gravitational waves was provided by Hulse and Taylor who noticed that the energy loss of a binary system containing a pulsar matched the predictions of general relativity [2]. A plot showing the curve from the quadrupole formula of GR against the pulsar observations is shown in Fig. 1. While the Hulse-Taylor system gives compelling indirect evidence for the existence of gravitational waves, physicists are eager to make a direct detection (i.e. an observation of the effects on spacetime from a passing gravitational wave) of gravitational waves to provide further support for GR and to use GWs to study the energetic astrophysical systems that emit them, such as neutron star and black hole binaries.

The production of gravitational waves requires a quadrupole source in contrast to electromagnetic waves which only require a non-zero dipole moment. This is because mass only comes in one variety as opposed to charge which can be positive and negative. Unlike the exchange of two charges, the exchange of two masses leaves the gravitational field the same. This means that binary systems are a natural place to begin searches for gravitational waves because they have a large mass quadrupole moment and we know that they exist. Binary systems containing neutron stars (NS) and black holes (BH) are important examples that have been studied extensively [4].

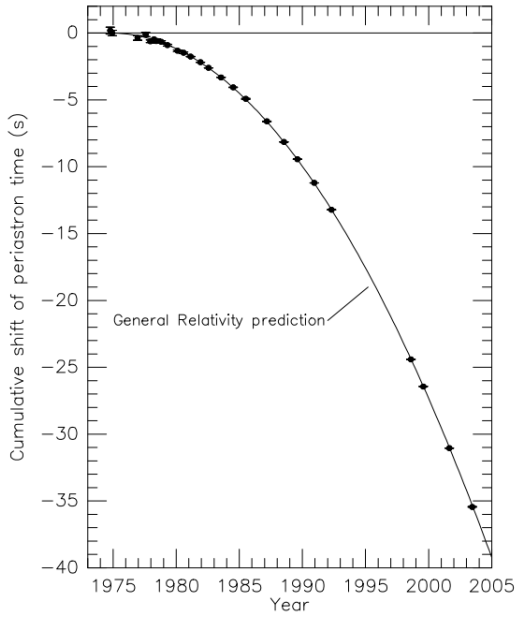

FIG. 1. Accumulated shift in the orbital phase relative to an assumed orbit with constant period caused by energy loss to gravitational waves. The straight line represents the prediction with no losses and the curve is the prediction from general relativity. Plot taken from [3].

Binary systems evolve through 3 phases: an inspiral phase where the orbital dynamics are driven by the re-

lease of energy through gravitational waves, a merger phase where the binary companions combine into a single highly-perturbed black hole, and a final phase where the resulting black hole rings down [5]. During this entire process, called compact binary coalescence (CBC), the binary system is emitting gravitational waves. While the three phases of a binary inspiral are treated equally by nature, the computation of the waveform is normally divided into these distinct regimes which are handled using separate techniques. As an example, consider Fig. 2 which shows a gravitational waveform for the inspiral of a pair of 50 $M_\odot$ black holes at a distance of 1 Mpc from the detectors.

If the binary companions are spinning, the gravitational waveforms become even more complicated. In the simpler case of aligned spins (spins pointing in the same direction as the orbital angular momentum), the phase evolution of the waveforms change slightly. In the more general case, the spins precess leading to modulation of the phase and amplitude of the gravitational waves.

The form of gravitational waves generated in compact binary coalescence depends at least 15 parameters of the system although many of them only enter into the overall amplitude of the signal (the sky position angles, binary plane orientation angles, and luminosity distance) [6]. Even more parameters are involved when considering the nuclear equation of state for neutron stars and extensions to general relativity. Table I lists these parameters and gives descriptions and the plots in Fig. 3 show how varying different parameters of the binary system impacts the waveform.

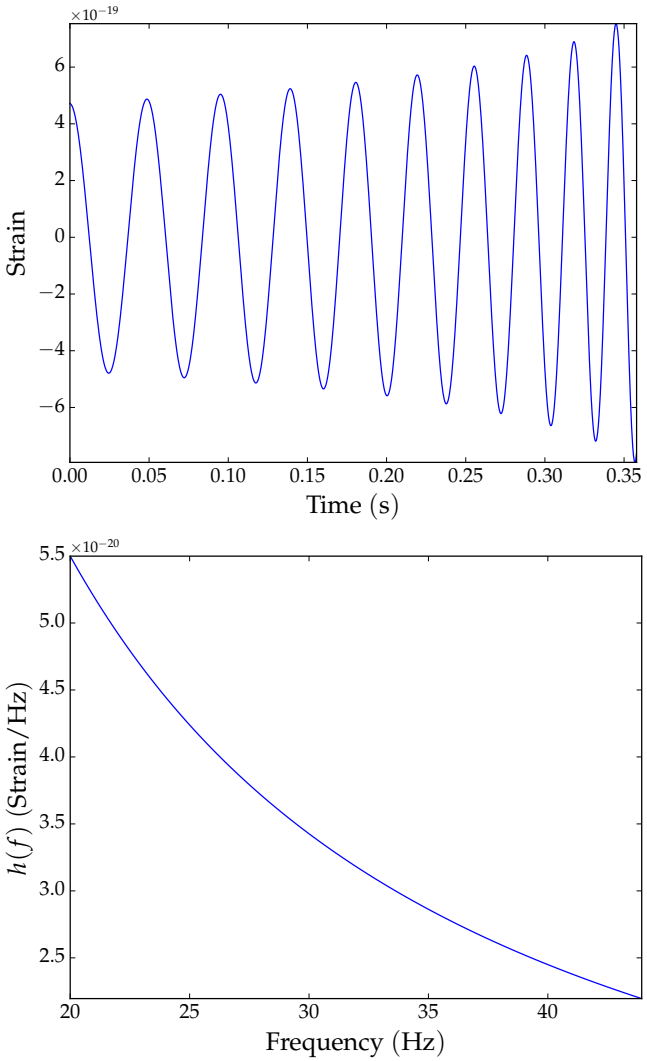

FIG. 2. A time domain waveform (top) and a frequency domain waveform (bottom) generated using the quadrupole and stationary phase approximations for the inspiral of a pair of non-spinning 50 $M_\odot$ black holes at a distance of 1 Mpc from the detectors. Full waveforms showing the inspiral, merger, and ringdown are shown in Fig. 15.

B. LIGO

The Laser Interferometric Gravitational-Wave Observatory (LIGO) is a part of a global effort to make the first direct detection of gravitational waves. The two LIGO detectors (shown in Fig. 4) are 4 km Michelson interferometers that search for small strains in their arms that result from passing gravitational waves. A GW detection would open the door to a new technique in astrophysics based on studying gravitational waves and would offer a test of general relativity in the most extreme and highly-dynamical regime of gravity that has ever been studied [4].

Advanced LIGO (aLIGO), an upgraded version of the initial LIGO detectors, is poised to begin a run with unprecedented sensitivity and bandwidth. It is possible that these upgrades will enable a detection in the near future. Compared to the initial LIGO detectors, Advanced LIGO will be 10 times more sensitive and will push the frequency band for gravitational wave searches down to 10 Hz compared to the previous 40 Hz [7]. These increases in sensitivity and bandwidth offer the possibility of hundreds of detections (see Table II for rates and detectable distances in Advanced LIGO).

The Advanced LIGO detectors (along with the Advanced Virgo detector which will come online in ∼2017) will be used to look for gravitational waves created by the coalescence of binary systems. For smaller mass systems like binary neutron stars, only the inspiral phase is in a frequency range detectable by LIGO. On the other hand, for large mass binary black hole systems the entire process of coalescence is detectable including the inspiral, merger, and ringdown phases [8]. Many example waveforms in the frequency domain for compact binary systems are shown against the early Advanced LIGO noise curve in Fig. 3 [**?**]. It is expected that Advanced LIGO will detect 40 neutron star mergers per year and between 30 and 100 black hole mergers [9].

Current limits on the rate of binary black hole mergers are set by combined searches from the LIGO and

TABLE I. 15 of the parameters specifying a compact binary system. The sky position angles, orientation relative to the line of sight (if the spins are aligned), and luminosity distance only enter into the overall amplitude of the signal and the coalescence phase and time can be efficiently determined external to the parameter search. The remaining parameters (the masses and spins) are intrinsic and must be determined using parameter estimation. Parameter list taken from [6].

Component masses	m_1, m_2
Component spin vectors	$\vec{S_1}$, $\vec{S_2}$
Sky position angles	right ascension α, declination δ
Orientation relative to line of sight	inclination ι, polarization angle ψ
Luminosity distance	D
Coalescence phase	φ_{coal}
Coalescence time	t_{coal}

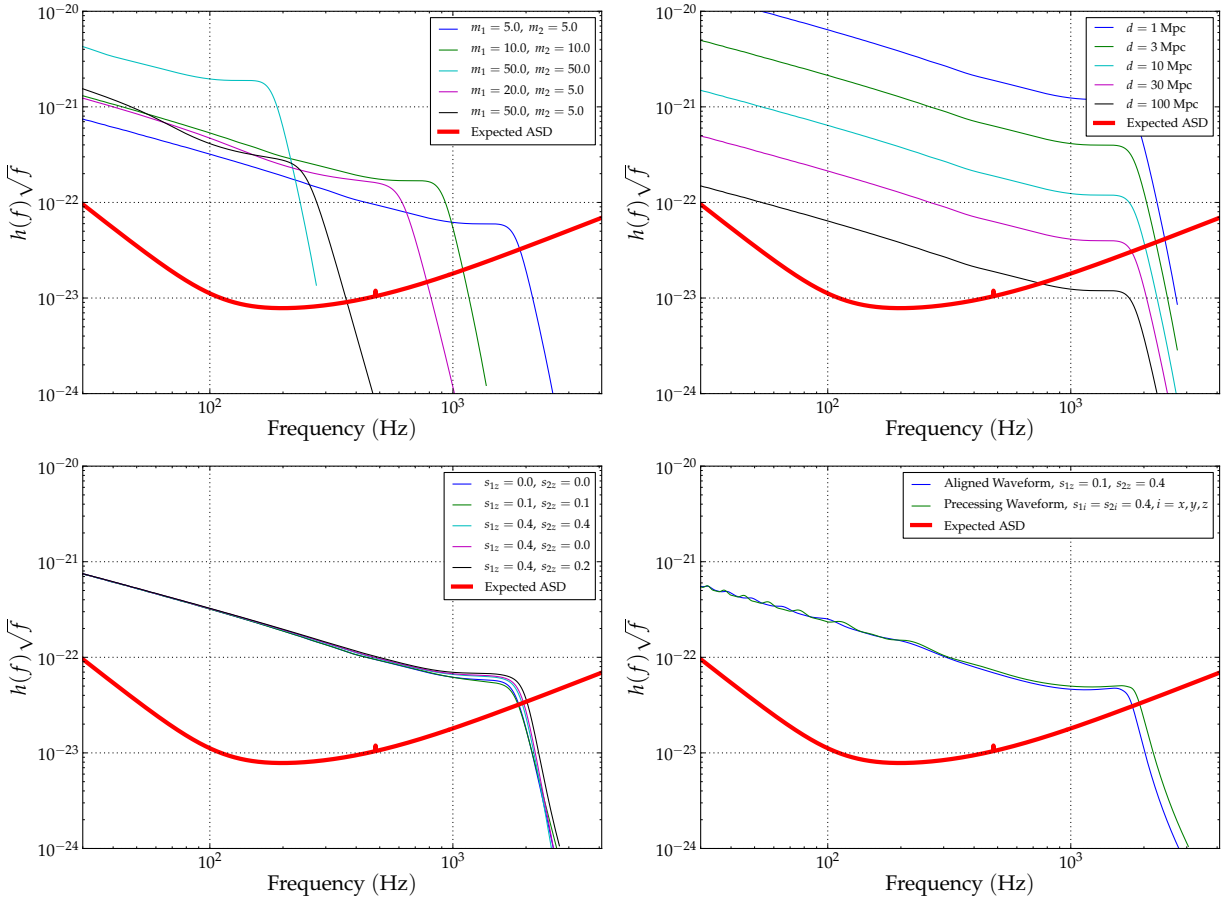

FIG. 3. Frequency domain waveforms for compact binary systems generated by varying the binary companions' masses (top left), distances from the detectors (top right), and spins in the aligned (z) direction (bottom left). The bottom right plot shows the effect of adding non-zero spins in the x- and y-directions to create a precessing waveform. Both binary companions have masses of 5 solar masses in the top right and bottom plots. The expected ASD is the amplitude spectral density of the detector noise for the early Advanced LIGO detector.

VIRGO scientific collaborations. By using large sets of data, these collaborations have looked for the occurrence of binary black hole coalescences for systems with total mass between 2 and 25 $M_\odot$ [10], systems with total mass between 25 and 100 $M_\odot$ [11], and systems with total mass between 100 and 450 $M_\odot$ [12]. No detections have been made yet, but the upgraded Advanced LIGO detectors may find the first definitive gravitational wave signal from these sources.

II. THE `gstlal` PIPELINE

When the LIGO detectors are in operation, there is a near-constant stream of data that must be analyzed to look for gravitational wave signals. The general flow

FIG. 4. Images showing the two LIGO sites in Livingston, Lousiana (top) and Hanford, Washington (bottom). The interferometer arms in each detector are 4 km long and contain vacuum pipes along the whole optical path.

TABLE II. Maximum detection distances, $\mathcal{D}$, and rates, $\mathcal{R}$, for various binary systems for Advanced LIGO. The masses are taken to be $\sim$1.4 $M_\odot$ for NS and $\sim$10 $M_\odot$ for BH. Distances and rates taken from [9].

	NS/NS	NS/BH	BH/BH
$\mathcal{D}$	445 Mpc	927 Mpc	2187 Mpc
$\mathcal{R}$, yr^{-1}	40	10	20

of the search pipelines used by LIGO are shown in Fig. 5. In an ideal case, the search pipelines used by LIGO would be able to identify candidate signals in the data in real time. This would allow for prompt follow-up by telescopes around the world that could detect an electromagnetic counterpart to the gravitational wave signal. In reality, some latency is incurred although the low-latency pipeline `gstlal` currently lags behind the incoming data on the order of 30 seconds making it close to real time [13]. This fast separation of signal and noise is accomplished through an advanced form of matched filtering [14]. In what follows, matched filtering will first be discussed followed by a description of some of the many additional techniques LIGO must use.

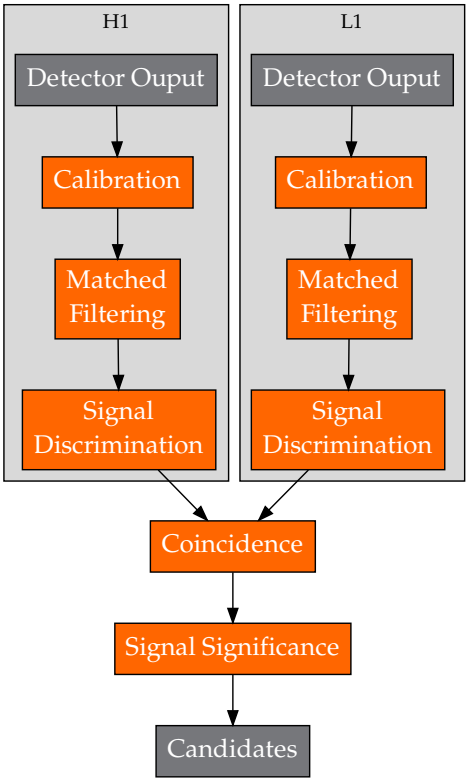

FIG. 5. A diagram showing the flow of data through the search pipelines. The output from each detector is initially filtered independently before being combined to look for coincidence.

A. Matched Filtering

Basic matched filtering correlates a template signal (which is known *a priori*) with detector output to search for a potential signal. This means that if the gravitational wave signal is known ahead of time, it is possible to use matched filtering to look for it in the data. This technique forms the basis of the LIGO CBC searches and so it is an important starting point. In this section, the matched filter will be described and I will show that it is optimal in the presence of stationary, Gaussian noise.

Let's start by assuming that the data coming from the detectors $x(t)$ can be written as a sum of some signal $s(t)$ and some noise $n(t)$. The noise is assumed to be stationary and Gaussian meaning the amplitude follows

a Gaussian probability distribution that is independent of time. Written out, the detector output takes the form

$$x(t) = s(t) + n(t). \tag{1}$$

Given a filter template $h(t)$, the correlation with the signal is defined to be

$$c(\tau) \equiv \int_{-\infty}^{\infty} x(t)h(t+\tau)dt \tag{2}$$

where τ is the time that the filter lags the detector output. This equation can be rewritten in the frequency domain as

$$c(\tau) = \int_{-\infty}^{\infty} \tilde{x}(f)\tilde{h}^*(f)e^{-2\pi i f\tau}df \tag{3}$$

where the tilde denotes the Fourier transform given by

$$\tilde{x}(f) = \int_{-\infty}^{\infty} x(t)e^{-2\pi i f t}dt. \tag{4}$$

Since the noise is assumed to be zero-mean, then the mean of c, called S, is given by

$$\langle c\rangle \equiv S = \int_{-\infty}^{\infty} \tilde{s}(f)\tilde{h}^*(f)e^{-2\pi i f\tau}df \tag{5}$$

where $\langle\ \rangle$ denotes an average over an ensemble of noise realizations. S characterizes the signal response generated by a filter $h(t)$ to an input signal which contains $s(t)$. Another quantity of interest is the variance of c which is given by

$$\left\langle (c - \langle c\rangle)^2\right\rangle \equiv N^2 = \int_{-\infty}^{\infty} S_n(f)|\tilde{h}(f)|^2df \tag{6}$$

where $S_n(f) = 2\langle|\tilde{n}(f)|^2\rangle$ is the one-sided noise power spectral density (PSD). This is the square of the noise amplitude spectral density (ASD) which is idealized as the red curve in the plots in Fig. 3. A plot showing the expected noise ASD for Advanced LIGO and all the contributions to it is shown in Fig. 6. The quantity we are interested in is the signal to noise ratio (SNR) which is given by

$$\rho = S/N. \tag{7}$$

This can be written in a more illuminating way if we first define the scalar product of two functions $a(t)$ and $b(t)$ to be

$$(a|b) = 2\int_0^{\infty} \frac{df}{S_n(f)}\left[\tilde{a}(f)\tilde{b}(f)^* + \tilde{a}(f)^*\tilde{b}(f)\right]. \tag{8}$$

Now, taking advantage of the fact that the Fourier transform of a real function $y(t)$ obeys $\tilde{y}(-f) = \tilde{y}^*(f)$, we can rewrite S and N in terms of inner products to give

$$\rho(\tau) = \frac{(xe^{2\pi i f\tau}|S_n h)}{\sqrt{(S_n h|S_n h)}}. \tag{9}$$

In practice, ρ is a function of τ and so the filtering process involves maximizing ρ over τ. From Eq. (8) it follows that ρ is maximized when the filter h takes the form

$$\tilde{h}(f) = \gamma\frac{\tilde{s}(f)e^{2\pi i f\tau}}{S_n(f)} \tag{10}$$

where γ is an arbitary constant that I will simply set to 1. The form of the filter we have derived is called the matched filter and is the unique linear filter which maximizes SNR.

Using the expression for h given in (10) one can compute the optimal SNR and find that it is

$$\rho_{opt} = 2\left(\int_0^{\infty} \frac{|\tilde{s}|^2}{S_n}df\right)^{1/2} = (s|s)^{1/2}. \tag{11}$$

Again, in practice one must also maximize over the time and phase at coalescence along with the intrinsic source parameters to arrive at the optimal SNR.

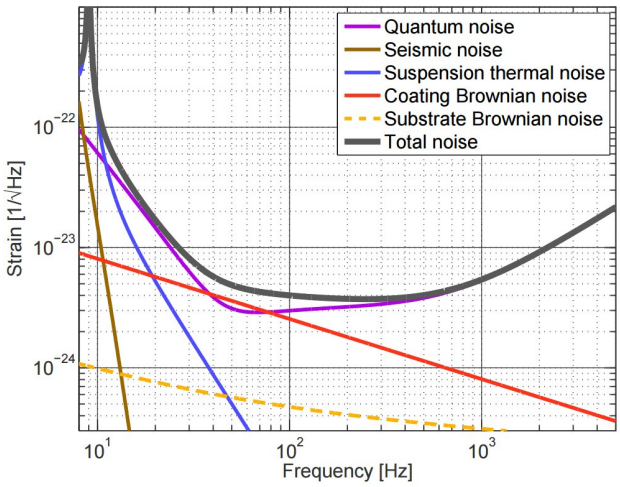

FIG. 6. The Advanced LIGO noise ASD with the different noise contributions shown. Plot taken from [15].

The SNR is an important quantity to distinguish signal from noise for LIGO and is used to identify candidate events. The process of matched filtering produces an SNR time series and peaks in the time series indicate possible GW signals. Those peaks that exceed a certain threshold (taken to be 4 in current `gstlal` analyses) are analyzed further and are called triggers.

While matched filtering forms the basis of LIGO data analysis, more advanced techniques must be employed in the LIGO searches because (i) the exact form of the signal is not known because the waveform shape depends on parameters such as the masses and (ii) the data contains non-stationary, non-Gaussian noise which cannot be accounted for in the standard matched filtering method. Below I will discuss some of the techniques used by `gstlal` to overcome these problems.

B. Template Banks

The discussion of matched filtering above assumed that the signal was known ahead of time, but in practice this is not the case. In order to account for the fact that the exact form of the gravitational wave signal is unknown, it is necessary to construct large template banks which span the compact binary parameter space. The intrinsic parameters (the masses and spins of the binaries) are the important parameters for determining the shape of the waveform and so in practice template banks are created by discretely sampling some subset of the mass and spin parameter spaces. Fig. 7 shows an example of a template bank spanning the parameter space of masses and the effective spin parameter $\chi_{\rm eff}$ which is given by

$$\chi_{\rm eff} = \frac{m_1 s_{1z} + m_2 s_{2z}}{m_1 + m_2}. \tag{12}$$

Since the parameter values are continuous, it is never possible to perfectly match a signal to a template, but template banks can be constructed to give arbitrarily small losses in detection rates and SNR. Typically, a minimal match between any signal and the nearest template of somewhere between 95% and 97% is chosen to give banks of manageable size which still represent the parameter space well.

The procedure of constructing an optimal template bank has been extensively studied and several methods have emerged. In the case of searches using inspiral-only templates over just the mass parameters (not as relevant now that spin is considered important in these searches), it has been shown that metric-based hexagonal template placement is optimal [16]. For higher-dimensional parameter spaces (e.g. those including spin), a stochastic placement algorithm is used which attempts to place random templates and discards those that have a sufficient minimal match with the templates around them [17]. This method has the advantage that it is scalable and does not require knowledge of the metric for the parameter space. In section IV, I discuss a novel method of generating stochastic template banks which is significantly faster than the traditional algorithm.

Substantial work has also gone into finding ways to use banks which have fewer templates. One method uses a modification of the Gram-Schmidt process to construct an optimal basis of templates [18]. Below I will discuss how singular value decomposition is used by the Advanced LIGO pipelines to compress the template banks to make the search procedure faster.

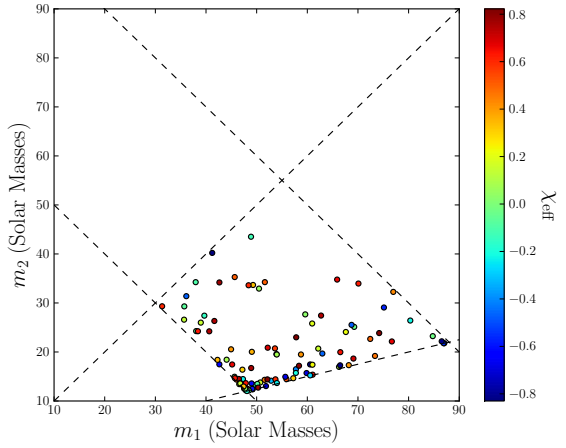

FIG. 7. An example of the parameter values used in a template bank where the color of the dots indicates the value of the effective spin parameter $\chi_{\rm eff}$. The dashed lines indicate the boundaries of the parameter space set by limits on the total mass and mass ratio.

1. *Singular Value Decomposition*

In order to make the `gstlal` pipeline run faster, it would be ideal if a smaller set of templates could be used to cover the same parameter space. This can be accomplished using a truncated singular value decomposition (SVD) which identifies a set of basis templates which can be used to reconstruct the entire template space with almost no losses [19]. In what follows, some of the details of this technique along with the computational benefits are discussed.

To begin, we construct the $N \times M$ template matrix $\mathbf{H}$. The rows of $\mathbf{H}$ are the templates (technically every two rows are a template where one is the real part and one is the imaginary part) and the columns are slices in time. Fig. 8 shows how this is done for an example set of templates.

In general, $\mathbf{H}$ will not be a square matrix, but it is possible to factor it into a product of three useful matrices. This is called the singular value decomposition and is written:

$$\mathbf{H} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^{\mathbf{T}} \tag{13}$$

where $\mathbf{U}$ is an $m \times m$ unitary matrix whose rows are orthonormal basis vectors, $\mathbf{V}$ is an $n \times n$ unitary reconstruction matrix, and $\mathbf{\Sigma}$ is an $m \times n$ diagonal matrix whose diagonal values σ_i are called the singular values of $\mathbf{H}$. This decomposition can be rewritten as

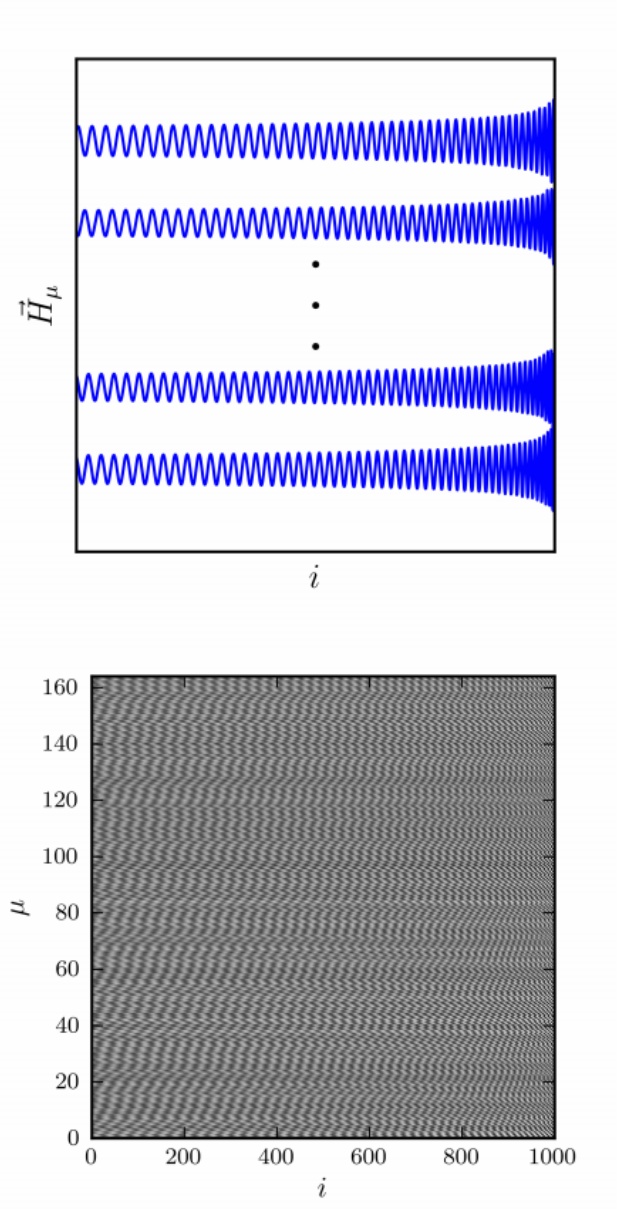

FIG. 8. Top: a series of templates in the time domain. Bottom: a template matrix **H** whose rows are templates and the columns are time samples. Fig. taken from [19].

$$H_{\mu j} = \sum_{\nu=1}^{N} v_{\mu\nu}\sigma_\nu u_{\nu j}. \tag{14}$$

In order to reduce the computational cost of filtering, only the basis vectors corresponding to the largest singular values are kept. These are the basis vectors that are most important to the reconstruction of **H**. As a common convention, the σ_i are listed in descending order so that choosing the largest N' of the N values of σ_i is equivalent to choosing the first N'. This truncation produces an approximate reconstruction of **H** given by

$$H_{\mu j} \approx H'_{\mu j} = \sum_{\nu=1}^{N'} v_{\mu j}\sigma_\nu u_{\nu j} \tag{15}$$

where $N' < N$. This reduces the number of basis templates from N to N' and allows for a reconstruction of the original bank from a smaller basis. It is possible to show that performing this truncation produces a fractional loss of SNR given by

$$\left\langle \frac{\delta\rho}{\rho} \right\rangle = \frac{1}{2N} \sum_{\nu=N'+1}^{N} \sigma_\nu^2. \tag{16}$$

In practice, the value of N' is chosen so that the fractional loss in SNR is less than .001. In the case of a small template bank considered in [19], the number of templates was reduced from $N = 912$ to $N' = 118$ under this condition. This represents almost an order of magnitude reduction in the number of templates and corresponds to huge reduction in computing costs. SVD is typically applied to small sub-banks which are close in parameter space so that most of the waveforms are similar and thus a smaller number of basis templates can be used.

C. Multibanding

Another way that the `gstlal` pipeline improves the filtering speed is by taking advantage of the form of compact binary coalescence signals. These signals, known as chirps, monotonically increase in frequency throughout the inspiral phase. In addition, a binary system spends most of the time during inspiral at low frequencies and the merger phase is relatively short compared to the inspiral. It is possible to take advantage of this predictable pattern in the signal by using different sampling rates. Dividing the signal into different pieces that are sampled at different rates is called multibanding and is a technique used to efficiently analyze LIGO data [20].

The Nyquist-Shannon Theorem states that if a function is band-limited so that the frequencies that comprise it satisfy $|f| < B$, then it can be completely determined by taking samples at a rate of $2B$ Hz. This rate is called the Nyquist rate and it represents the minimum sampling rate that can be used to determine a signal without aliasing.

This theorem can be applied to LIGO waveforms by dividing each waveform into a series of time slices which are all sampled at different rates. Since the binary spends most of the inspiral at low frequencies, a small sampling frequency can be used. It is only during the more dynamic high-frequency portion of the waveform that a high sampling frequency is necessary. This greatly reduces the amount of data that must be processed in the the pipeline.

Since the discreteness of the signals is important in this discussion, I will represent templates as discrete functions and I will reserve the letter k to indicate a time index. If the original signal was sampled at a rate f^0 and is divided into S non-overlapping (and thus orthogonal) time slices, it can be written as a sum of those slices:

$$h[k] = \sum_{s=0}^{S-1} \begin{cases} h^s[k] & \text{if } t^s \leq k/f^0 < t^{s+1} \\ 0 & \text{otherwise} \end{cases} \tag{17}$$

for the S integers $\{f^0 t^s\}$ such that $0 = f^0 t^0 < ... < f^0 t^S$. For `gstlal` , the time slice boundaries are chosen such that each interval $[t^s, t^{s+1})$ is sub-critically sampled by a power-of-two sample rate f^s. Once the time slice boundaries are selected, the templates can be downsampled without aliasing. The resulting downsampled signal is made up of time slices given by:

$$h^s[k] = \begin{cases} h[k\frac{f^0}{f^s}] & \text{if } t^s \leq k/f^s < t^{s+1} \\ 0 & \text{otherwise} \end{cases} . \tag{18}$$

An image showing how this downsampling is typically done for a chirp signal is shown in Fig. 9. As with SVD, time slicing is typically performed over sub-banks which contain many templates that are close in parameter space. Since the waveforms in the sub-banks are all quite similar, the time-slicing can be performed identically for all of them.

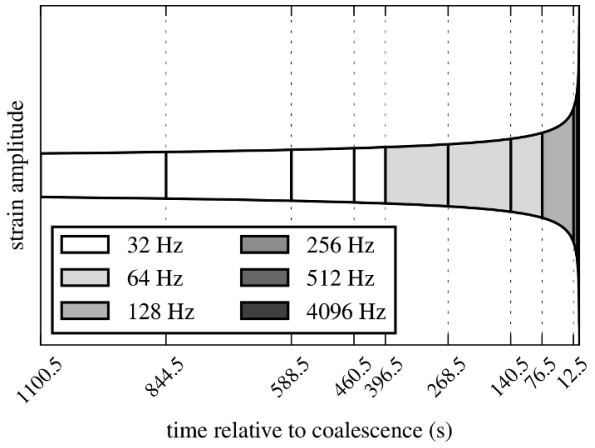

FIG. 9. A time domain chirp waveform with time slices chosen. The sampling frequency of each time slice is indicated by the shading. Fig. taken from [20].

D. χ^2 Test

Once the peaks in the SNR are found and counted as triggers, there is still a possibility that they correspond to very loud noise in the detector. These short-duration, high-amplitude noise events are called glitches and unfortunately they can create responses in many templates at once. Glitches represent non-Gaussian noise in the detector that is very hard to deal with. One method that has been developed to handle glitches is the use of a χ^2 test. Many different types of χ^2 tests have been developed [21, 22]) and below I will describe two of them and how they help to veto glitches.

1. *Traditional* χ^2

In the traditional χ^2 test, the interval $[0, \infty)$ is divided into p disjoint sub-intervals $\Delta f_1, ..., \Delta f_p$. The intervals are chosen so that the expected signal contributions from a chirp in each interval are equal. This condition can be written more succinctly by first defining a set of p Hermitian inner products

$$(a(f)|b(f))_p = \int_{-\Delta f_p \cup \Delta f_p} \frac{a^*(f)b(f)}{S_n(f)} df. \tag{19}$$

Then the frequency bands are chosen so that for a normalized template, h,

$$(\tilde{h}|\tilde{h})_p = \frac{1}{p} \tag{20}$$

Now, suppose the signal takes the form

$$s(t) = h(t - t_0) \tag{21}$$

where t_0 is the coalescence time. This is the basic form of a chirp signal where the phase is known. Then, as shown above, the signal to noise ratio in the optimal case is

$$\rho = (s|s)^{1/2} = (h|h)^{1/2} = 1 \tag{22}$$

Now define the contribution to the SNR from a specific frequency interval as ρ_j. One can show that

$$\langle \rho_j \rangle = \frac{1}{p}, \quad \langle \rho_j^2 \rangle = \frac{1}{p} + \frac{1}{p^2} \tag{23}$$

where here $\langle\ \rangle$ corresponds to an average over many noise ensembles. In the absence of signal ($d \to \infty$) we have $\langle \rho_j \rangle = 0$, and $\langle \rho_j^2 \rangle = \frac{1}{p}$. Now define

$$\Delta\rho_j \equiv \rho_j - \frac{\rho}{p}. \tag{24}$$

Then the χ^2 is defined to be [23]:

$$\chi^2 \equiv p \sum_{j=0}^{p} (\Delta\rho_j)^2 \tag{25}$$

In the case of stationary, Gaussian noise, this quantity is a classical χ^2 distribution with $p-1$ degrees of freedom, hence the name χ^2. Low values of χ^2 indicate potential signal while high values indicate probable glitches. A plot demonstrating how the χ^2 can be used to separate signal from noise is shown in Fig. 10. The long tail of background events into the high SNR region is non-Gaussian and makes rejection of noise much harder without the χ^2.

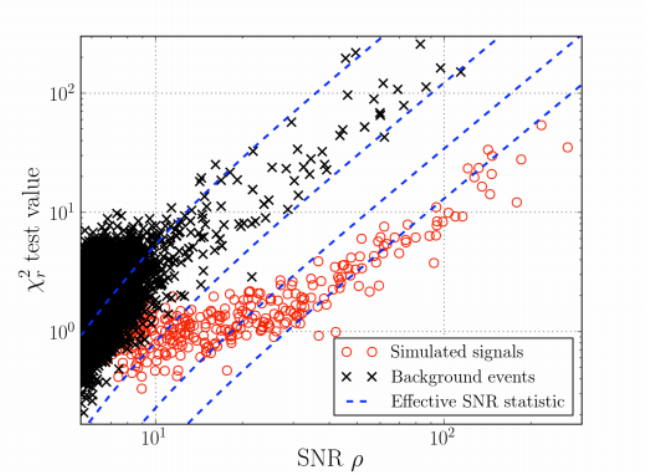

FIG. 10. A plot showing signal (simulated) and background events with different SNR and χ^2 values. It is clear that a simple SNR threshold is not enough to separate signal from background and so a combined cut on ρ and χ^2 such as an "effective SNR statistic", shown as dashed lines, is necessary. Fig. taken from [6].

One can intuitively understand the way this test works by looking at what it does. For each of the frequency intervals the χ^2 looks at the contribution to ρ and compares the contribution to the expectation. For a glitch, the contributions to the SNR will be dominant in a small number of frequency intervals leading to a large χ^2 value. For a normal chirp signal, the contributions to the SNR will be evenly spread out among the frequency intervals and so the χ^2 will be low. This allows for the rejection of unwanted glitches.

2. *Autocorrelation* χ^2

Another form of χ^2 that one can use to reject glitches is called the autocorrelation χ^2 and is obtained by comparing the SNR time series to the auto-correlation of the template. This statistic has the advantage that once the SNR time series is computed, all of the necessary pieces are already in memory and so it is extremely computationally efficient [6].

Suppose that the detector output is of the form $x(t) = n(t) + Ah(t)$ where A is some amplitude and h is unit-normalized so that $(h|h) = 1$. Then, the SNR time series is:

$$\rho(\tau) = (n|he^{2\pi i f\tau}) + A(h|he^{2\pi i f\tau}) \quad (26)$$
$$= (n|he^{2\pi i f\tau}) + A\alpha(\tau) \quad (27)$$

where $\alpha(\tau)$ is the autocorrelation of the template. The time $\tau = 0$ is chosen to be the point when the SNR is at a maximum. Maximizing in time and taking an ensemble average so that the noise term disappears gives $\langle\rho_{max}\rangle \approx A$. The quantities $\rho_{\max}$, $\rho(\tau)$, and $\alpha(\tau)$ are easily computable from the templates and data while the quantity $(n|he^{2\pi i f\tau})$ will be Gaussian distributed when the noise is Gaussian. Thus, it is possible to compute a χ^2 of the form

$$\chi^2 = \int_0^{T_{\max}} |\rho(\tau) - \rho_{\text{peak}}\alpha(\tau)|^2 d\tau. \quad (28)$$

$T_{\max}$ is a tunable parameter called the autocorrelation length and determines the number of degrees of freedom in the χ^2 distribution (e.g. if the SNR time series is computed with a time interval Δt, the number of degrees of freedom is $N = T_{\max}/\Delta t$).

E. Non-stationary Noise

One final hurdle for the LIGO detectors is non-stationary noise (i.e. noise that evolves over time). The non-stationary nature of the LIGO noise can best be seen by looking at the evolution of the noise ASD shown in Fig. 11.

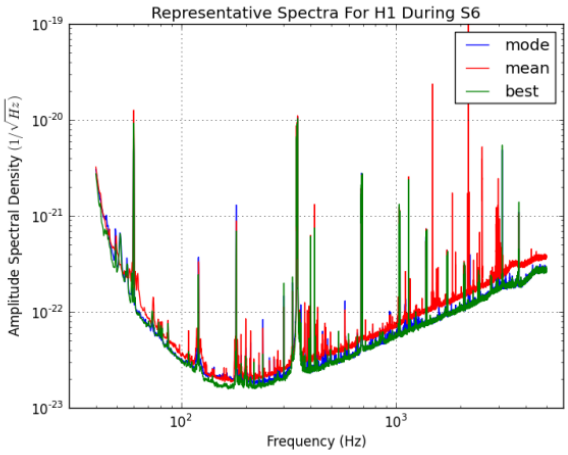

FIG. 11. Noise ASD curves for the H1 detector during S6 (May 9, 2010 - November 4, 2010). The ASD for the mode (blue), mean (red), and best (green) operation are shown. Plot taken from [24].

For forms of non-stationary noise that evolve slowly over time (i.e. those due to changes to the detector hardware), it is possible to simply filter the data over time intervals that are small compared to the time scale of the

non-stationary noise. Non-stationary noise that evolves on a shorter time-scale is more difficult to deal with (such as anthropogenic noise) and impacts the detector sensitivity.

An obvious way to reduce the backgrounds from non-stationary glitches is the requirement of coincidence between detectors. Candidate gravitational wave events with high enough signal-to-noise ratios for a specific template waveform in one detector are counted as triggers. Coincidence occurs when identical triggers (those corresponding to the same template waveforms and close in time) are found in multiple detectors [6]. It is unlikely that glitches will occur at the same time in multiple detectors and so the coincidence requirement makes the LIGO searches more robust against them.

III. A NEW STOCHASTIC TEMPLATE BANK ALGORITHM

A. Stochastic Template Banks

In early searches for gravitational waves from compact binary coalescence (CBC), the only parameters under consideration were the binary masses. Using this restriction, it was possible to treat the parameter space as a manifold with a metric that gave the separation between points in terms of the match [25]. Methods using this metric treatment of the parameter space were developed which placed templates on a lattice that covered the parameter space [26]. In the case of the hexagonal lattice, the template placement is actually optimal meaning that the space is covered and a minimum number of templates is used [16].

While optimal methods have been developed and implemented for low-dimensional spaces, in the higher-dimensional cases it is not known how to create optimal template banks for curved spaces of dimension $D > 2$ because the metric is not known [27]. In fact, even in the case of only two-dimensional template banks, the metric is only known for inspiral waveforms and so the inclusion of merger and ringdown is not possible.

As a result, we must instead use stochastic methods which randomly place templates, check if they are covered by templates already in the bank, and keep them if they are not covered [17]. Stochastic techniques have the advantage that they do not require knowledge of the metric and so they can be implemented for parameter spaces of arbitrary dimension and in those without a metric. These methods have found use in searches for compact binary coalescence where the necessity of including spin demands at least a 3-dimensional template bank [28, 29].

One problem with stochastic methods is that they are not especially computationally efficient. Randomly generating the coordinates for templates uniformly in the parameter space and checking the match with neighboring templates causes many extra expensive match calculations to be performed. These extra calculations are incurred because: (1) the parameter space is not flat and so uniformly chosen coordinates are not uniformly separated in match and (2) when the bank begins to become full, a lot of time is spent placing the last few templates to fill in the remaining holes in the bank. Metric methods do not suffer from the same pitfalls because they exploit knowledge of the metric to place templates optimally.

An obvious way to combat the issue of the curved parameter space is by choosing coordinates where the metric is approximately flat. In the case of searches for compact binary coalescence, the chirp time coordinates τ_0, τ_3 provide a nearly Cartesian metric that alleviates the difficulties due to the parameter space curvature [30]. In terms of the masses, τ_0, τ_3 can be written

$$\tau_0 = \frac{5}{256\pi f_L \eta}(\pi M f_L)^{-5/3} \tag{29}$$

$$\tau_3 = \frac{1}{8 f_L \eta}(\pi M f_L)^{-2/3} \tag{30}$$

where $\eta = m_1 m_2/M^2$ is the symmetric mass ratio, $M = m_1 + m_2$ is the total mass, and f_L is the lower bound on the frequency of the waveform. Physically τ_0, τ_3 correspond to the Newtonian and 3PN contributions to the time it takes for the gravitational wave frequency to go from f_L to infinity. In Appendix A, the boundaries of the τ_0, τ_3 space are computed and an efficient method to sample from this space is described.

The problem of placing the last few templates is an inherent difficulty in current stochastic placement algorithms that stems from their lack of knowledge about where previous templates were placed. In order to address this issue, we propose a new stochastic template bank creation algorithm which maintains a memory of where templates have been placed to guide the placement of proposal templates.

B. New Algorithm

Following the notation in [17], let $\mathcal{M}$ be a signal manifold of dimension D with a positive-definite distance $d(x, y)$ where $x, y \in \mathcal{M}$. A template bank T is a set of n points $\{x_1, \ldots, x_n \mid x_i \in \mathcal{M}\}$. T covers the signal manifold $\mathcal{M}$ with radius Δ if $\forall\ y \in \mathcal{M},\ \exists\ x_i \in T$ such that $d(x_i, y) < \Delta$.

In this paper, the following algorithm is proposed as an improvement to the stochastic template placement algorithm presented in [17] including the improvements from [28]:

1. Cover $\mathcal{M}$ with k^D equally sized bins B_i, $i = 1, \ldots, k^D$ where k is the number of divisions per dimension. To construct the bins:

   (a) Suppose the binary source parameters are given by $\theta_i, i = 1, \ldots, D$ and that each parameter takes values in the interval $[\theta_{i,min}, \theta_{i,max}]$

on $\mathcal{M}$. Divide this interval into k equal length sub-intervals $\theta_{i,\alpha}$ where $\alpha = 1, \ldots, k$.

(b) Each bin is then given by the Cartesian product $\theta_{1,\alpha} \times \theta_{2,\beta} \times \ldots \times \theta_{D,\omega}$ where all the greek indices go from 1 to k. One could imagine allowing each dimension to be binned differently (i.e. with more or less than k bins), but for simplicity the case of k bins in each dimension is considered.

Note that using this construction, not all the bins will necessarily intersect $\mathcal{M}$. An example gridding of two dimensions of a three-dimensional space is shown in Fig. 12.

2. For each bin B_i, assign two values: the number of accepted templates A_i and the number of rejected templates R_i. Initially these values are set to $A_i = 1, R_i = 0$ although other choices are possible.

3. Choose a proposal template y uniformly from the signal manifold and determine which bin, B_i it falls into. Then compute the rejection probability P_i for that bin given by:

$$P_i = \frac{R_i}{R_i + A_i} \tag{31}$$

4. Choose a random number r uniformly from $[0, 1]$. If $r < P_i$, discard y and return to step 3.

5. If $d(y, x_i) > \Delta \; \forall \; x_i \in T$, add y to T and increment A_i. Otherwise, discard y, increment R_i, and return to step 3.

6. Continue 3-5 until $P_i > P^* \; \forall \; i$ where P^* is a cutoff probability that can be varied to change the performance of the algorithm.

Currently the dominant contribution to the time required for template bank generation is the computation of the match between the proposal template and templates already in the bank. This algorithm increases the speed of template bank construction by keeping track of where templates have already been placed and reducing the number of match calculations. As regions of the signal manifold are covered and P_i increases for those regions, it is less likely that proposal templates will be placed there. In effect, this algorithm guides template bank placement to regions of the bank that are not full and eliminates match calculations in regions that are already sufficiently covered.

In practice some computational cost is introduced by having to check if all the $P_i > P^*$ for each new proposal. This cost can be reduced by vectorizing the computation so that the additional burden is negligible.

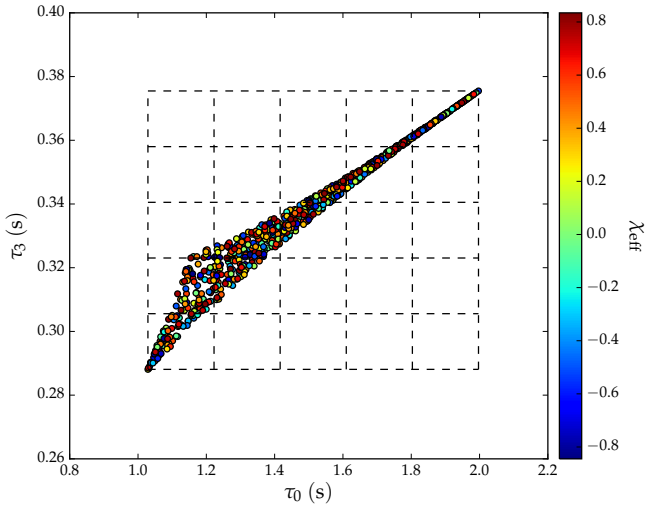

FIG. 12. A 3-dimensional template bank with a grid overlaid in the τ_0 and τ_3 directions. The values of τ_0 and τ_3 shown correspond to systems with total mass between 20 and 30 solar masses and component masses between 10 and 15 solar masses. The third dimension is $\chi_{\rm eff}$ which is not gridded in this case. Notice that some of the bins may not intersect the signal manifold at all.

C. Results

Since the match calculations required to compare a proposal template to the existing templates in the bank is the most expensive part of stochastic template bank generation, one way to benchmark the performance of the new algorithm is to compare the number of match calculations it requires to the traditional algorithm. A plot of this comparison is shown for a specific template bank in Fig. 13. It is clear that the new algorithm requires far fewer match calculations to make a bank of the same size and this is reflected in the amount of time is required to generate the banks.

Perhaps more important than speed is the coverage of the banks. It is crucial that a template bank have no holes (i.e. that it covers all possible signals). To ensure that the new algorithm is generating banks that are robust, bank simulations were performed where 5000 injections were tested against the bank to see if they were covered. The results of these bank simulations are shown in Fig. 14.

As is clear from these figures, the new algorithm creates banks which cover the parameter space just as well as the old algorithm. In contrast however, the new algorithm requires less than half the number of match calculations and consequently requires less than half of the time. This is a substantial reduction that will allow for new template banks to be generated more regularly.

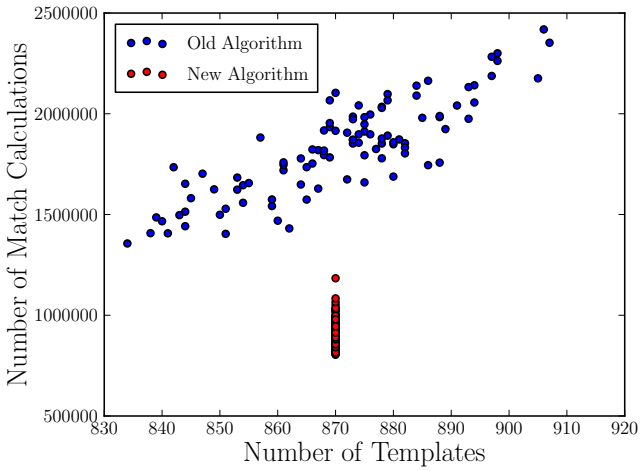

FIG. 13. A plot showing the number of match calculations required to make a template bank of a certain size for both the new and old stochastic template bank algorithms. The number of match calculations in the new algorithm was fixed to get a fair comparison between the algorithms.

IV. IMPROVING THE AUTOCORRELATION χ^2

In Section IID, two different types of χ^2 discriminator tests were described. Below, I describe efforts to improve the second of these tests: the autocorrelation χ^2.

Recall from Eq. (28) that the autocorrelation χ^2 test compares the autocorrelation of the best-matching template to the SNR time series. Fig. 15 shows an example time domain waveform and the accompanying autocorrelation.

When the detector output contains only Gaussian noise, the SNR time series will be very similar to the autocorrelation and so the χ^2 will be small. In the presence of a glitch, however, the SNR time series is the same as the template (or a time-reversed version of it depending on the order of the functions in the correlation) which will not match the autocorrelation well. Fig. 16 shows a glitch on top of a waveform and the correlation between the waveform and the glitch demonstrating that the correlation is nothing like the autocorrelation of the template.

The current autocorrelation χ^2 test used in the `gstlal` pipeline is computed on an interval centered on the autocorrelation peak and only uses this region for the computation of the autocorrelation χ^2. The length of this region is called the autocorrelation length and in current analyses it is set to be the same for all templates. Below I propose and discuss the results of using two different methods to modify the calculation of the autocorrelation χ^2 so that it is adaptive to the templates being used to filter the detector output.

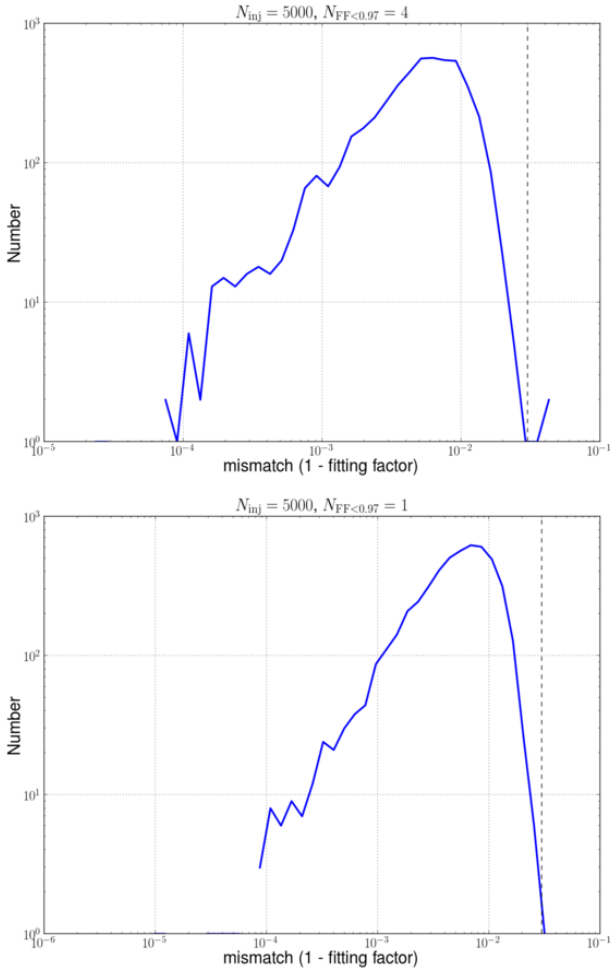

FIG. 14. The number of injections in the bank simulation with a given value of mismatch for the old (top) and new (bottom) algorithms. The dashed line represents the maximum mismatch desired for a template bank.

A. Windowing in Autocorrelation Power

One way to modify the autocorrelation χ^2 is to make the window where it is computed depend on the region where most of the autocorrelation power is. For each template, the autocorrelation is computed and squared to get the power and then the small region centered on 0 containing some chosen fraction of the autocorrelation power is determined. This is the region used to compute the autocorrelation χ^2 for triggers corresponding to the given template. Choosing to compute the autocorrelation χ^2 only over the region where the autocorrelation is highest should lower all χ^2 values since the region of integration is shrinking. If it lowers the χ^2 values for signal more than it does for background the detection efficiency will improve.

This method of defining the autocorrelation length takes into account differences in the templates. Very long templates corresponding to low mass systems have very tight autocorrelations and so only small autocorrelation

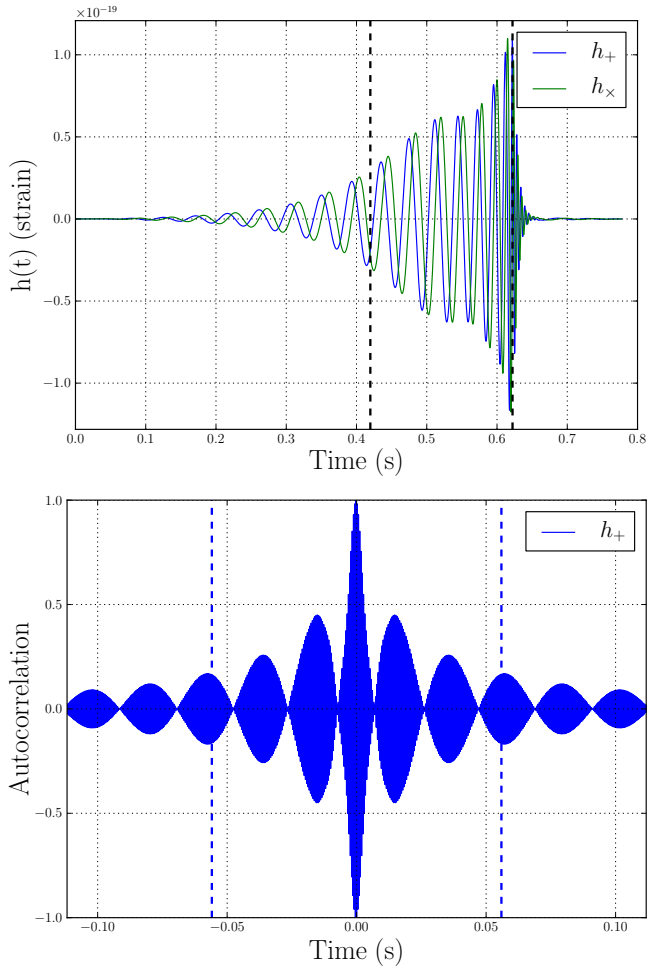

FIG. 15. Time domain waveforms for a pair of 50 $M_\odot$ black holes (top) and the autocorrelation of the '+' polarization (bottom). In both plots the region between the dashed lines corresponds to the part of the waveforms (top) and autocorrelation (bottom) where 90% of the power falls.

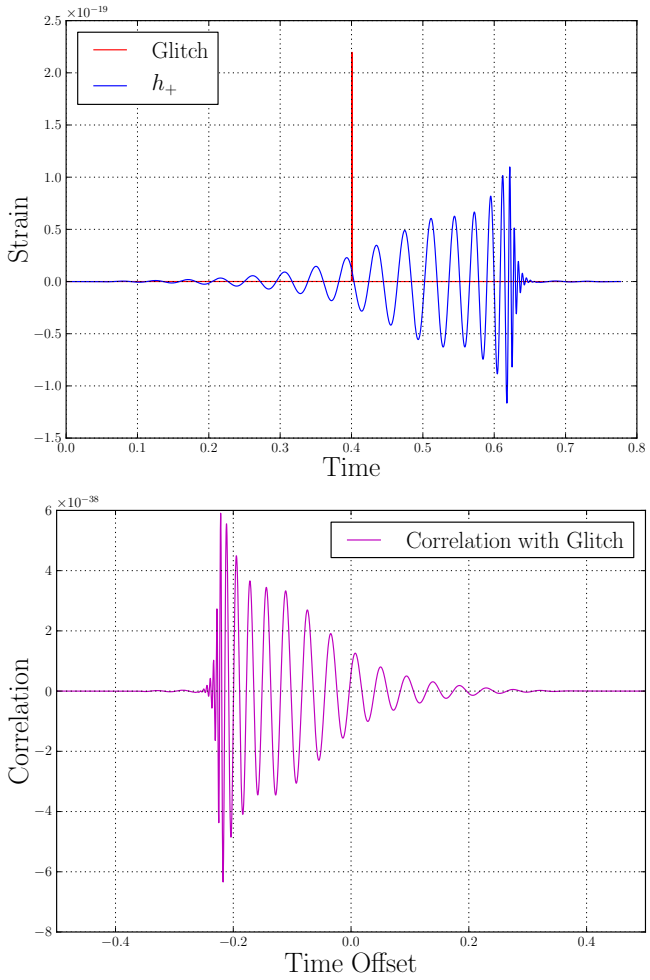

FIG. 16. A glitch (approximated as a very narrow Gaussian) shown against the '+' polarization of the waveform from Fig. 15 (top) and the correlation between the glitch and the waveform (bottom).

lengths are needed to capture most of the autocorrelation power. In contrast, high mass systems produce short waveforms that have wider autocorrelations and so larger autocorrelation lengths are needed.

This method was implemented in the `gstlal` code and tested using a 90% power interval. This interval is indicated by the dashed blue lines in Fig. 15. The results of filtering using an autocorrelation χ^2 defined this way are compared to the standard method in Fig. 17.

As described above, all the χ^2 values were significantly reduced by using a smaller autocorrelation length, but the background χ^2 values were reduced so much that the high SNR signal population became contaminated with background. Of the ∼180,000 total background triggers, about 160,000 were in the signal region using the old χ^2 calculation and an additional 200 were added using the new calculation. While this is only a small contamination, this change has still made the separation of the signal and background worse and so it is not viable.

One could imagine trying this technique using a larger portion of the autocorrelation (i.e. by using 99% of the power to define the autocorrelation length), but improvement does not seem likely except in the high mass region where the relevant length scales of the autocorrelation and template are similar. The failure of this method seems to be due to the fact that for low mass systems, the template is much longer than the autocorrelation meaning that the relevant part of the correlation with the glitch may not fall in the window specified by the autocorrelation power.

B. Windowing in Template Power

An alternative approach to determining an appropriate autocorrelation length for each template is to choose it based on template itself rather than the autocorrelation. In this case, the region of 90% power for the waveform (indicated by the dashed black lines in Fig. 15) is used to

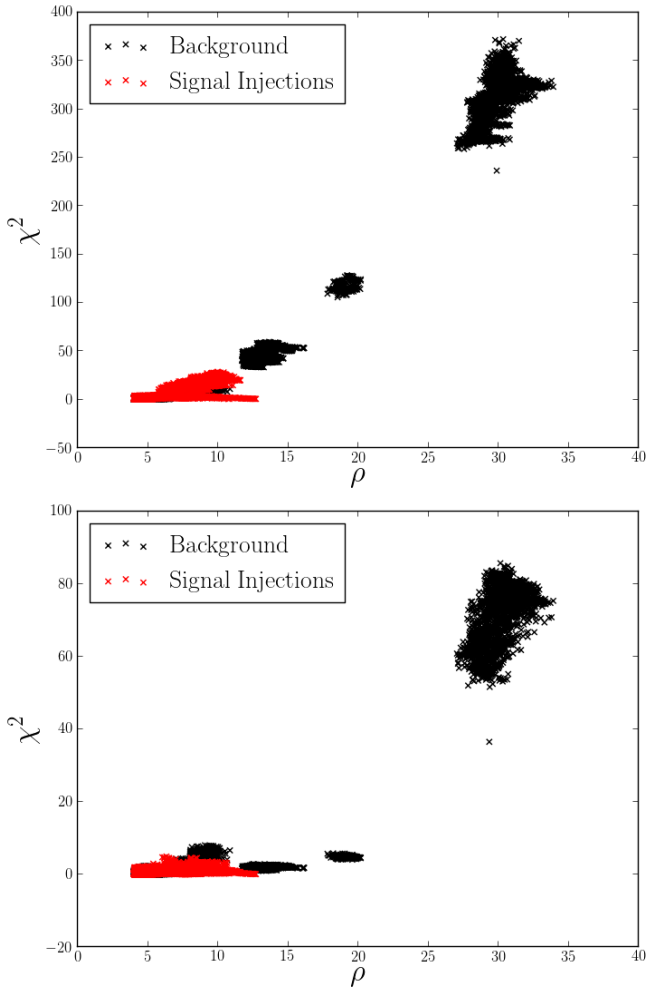

FIG. 17. Plots showing the SNR vs. the χ^2 for the standard method of computing the autocorrelation χ^2 (top) and the new method using a window defined by the autocorrelation power (bottom). Signal injections are shown in red and background is shown in black. The difficult background events to deal with are those that overlap with the signal region (i.e. those with low ρ values). The signal injections were for systems with total mass between 4 $M_\odot$ and 10 $M_\odot$.

define the autocorrelation length. This method has the advantage that in the case of a glitch, most of the power of the glitch will be picked up by the autocorrelation χ^2 and so glitches will be assigned high χ^2 values.

A comparison of this method to the one described above is shown in Fig. 18. One can clearly see the tradeoffs between the two methods of choosing the autocorrelation length. When choosing the autocorrelation length based on the autocorrelation power, a signal will have a very low autocorrelation χ^2, but background events will as well. On the other hand, choosing the autocorrelation length based on the template power will ensure that glitches have very high χ^2 values, but will also increase the χ^2 of signal events.

Results from choosing the autocorrelation length for each template based on the template power are shown in

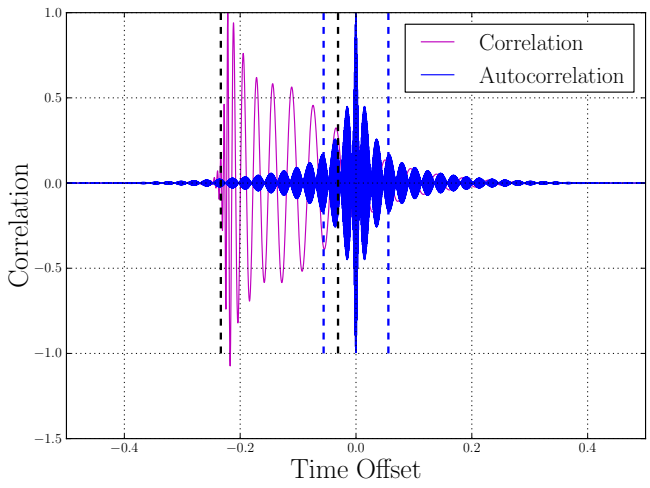

FIG. 18. The autocorrelation from Fig. 15 shown with the correlation of the signal with the glitch from Fig. 16. The blue dashed lines indicate the region that would be used when windowing on the autocorrelation power and the black dashed lines indicate the region that would be used when windowing on the template power.

Fig. 19 for the same background and signal populations used to evaluate the previous technique. In this case, of the ∼160,000 background triggers that overlapped the signal region, about 200 were elevated out of this region. It is unclear how large of an impact this elevation of background events will have on the sensitivity of LIGO searches. Further investigation with larger data sets will be necessary.

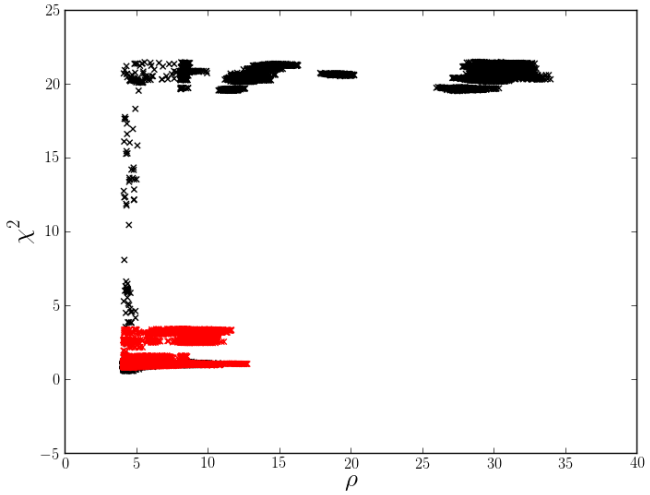

FIG. 19. A plot showing the SNR vs. the χ^2 for autocorrelation χ^2 values calculated by windowing on the template power. This should be compared to the top plot on Fig. 17. As above signal injections (total mass between 4 $M_\odot$ and 10 $M_\odot$) are in red and background triggers are in black.

One could in principle use other percentages of the template power to tune the autocorrelation χ^2 in order to determine if more than 200 events could be elevated out

of the signal region. One must also be concerned with the additional computational burden of using a larger autocorrelation length. For long templates, a 90% power window corresponds to many time samples leading to very long autocorrelation lengths (sometimes larger by 2 orders of magnitude). This will increase the size of SVD banks and possibly filtering time due to increased cost in the χ^2 calculation.

V. CONCLUSION

The `gstlal` pipeline is an elaborate search pipeline with many optimizations for speed and sensitivity in CBC searches. I have presented some methods which seek to improve the performance and efficiency of CBC searches by allowing for significant improvements in the speed of template bank generation. In addition, a framework for modifying and tuning the autocorrelation χ^2 test has been developed which has the potential create better separation of signal and background.

In the future, it is possible to expand on this work to create even larger improvements to the performance of `gstlal` . In its current incarnation, `gstlal` must convert the τ_0, τ_3 coordinates to component masses to check that they satisfy the given parameter space constraints which is costly. An alternative would be to analytically compute the boundaries of the τ_0, τ_3 space given constraints on the component masses so that the sampling could be directly done in this coordinate system. Some work on this issue is presented in Appendix A although some subtleties still need to be worked out.

In addition, tuning still needs to be done to determine the optimal power windows for separating signal and background in the `gstlal` pipeline using the autocorrelation χ^2 test. Once these windows are found, it would also be prudent to determine the increases in computing costs, if any, that are introduced by adding this complexity to the search process.

VI. ACKNOWLEDGEMENTS

I would like to acknowledge the constant help of my mentors. They helped me immensely in learning about the analysis of LIGO data as it relates to my projects. I would also like to thank my fellow SURF students who helped me talk through my work on many occasions. In addition, I am grateful for computational resources provided by the Leonard E Parker Center for Gravitation, Cosmology and Astrophysics at University of Wisconsin-Milwaukee. Finally, I would like to acknowledge the Caltech SURF Program and the NSF which funded my summer research.

Appendix A: Boundaries of the τ_0, τ_3 Space

Below I show in detail the calculation of the boundaries of the τ_0, τ_3 space given some constraints on the masses of the binary companions. In addition, I will show how this knowledge can be used to efficiently sample the τ_0, τ_3 space without many discarded trials. This derivation follows Appendix B of [16] with some corrections and extensions.

Suppose that a parameter space of interest has masses constrained by $m_1 \in [m_{1,\min}, m_{1,\max}]$, $M \in [M_{\min}, M_{\max}]$, and $q \in [q_{\min}, q_{\max}]$ where $q = m_1/m_2$ is the mass ratio. Notice that the conditions on m_1 and M are equivalent to putting constraints on m_1 and m_2.

It is useful to rewrite the expressions for τ_0 and τ_3 as

$$\tau_0 = \frac{A_0}{\eta} M^{-5/3} \tag{A1}$$

$$\tau_3 = \frac{A_3}{\eta} M^{-2/3} \tag{A2}$$

where f_L is the lower cutoff frequency and A_0, A_3 are constants given by

$$A_0 = \frac{5}{256(\pi f_L)^{8/3}}, \quad A_3 = \frac{\pi}{8(\pi f_L)^{5/3}}. \tag{A3}$$

Eliminating M from (A1) and (A2) yields

$$\tau_3 = \frac{A_3}{\eta}\left(\frac{\eta\tau_0}{A_0}\right)^{2/5} = \frac{A_3}{A_0^{2/5}}\frac{\tau_0^{2/5}}{\eta^{3/5}}. \tag{A4}$$

Using the bounds on q it is possible to constrain η by realizing that

$$\eta = \frac{q m_2^2}{m_2^2(q+1)^2} = \frac{q}{(q+1)^2} \tag{A5}$$

meaning η takes values between $\eta_{\min}$ and $\eta_{\max}$ which are given by

$$\eta_{\max} = \frac{q_{\min}}{(q_{\min}+1)^2}, \quad \eta_{\min} = \frac{q_{\max}}{(q_{\max}+1)^2}. \tag{A6}$$

Substituting these values into (A4) gives two boundary curves for the space:

$$\tau_3 < \frac{A_3}{A_0^{2/5}}\frac{\tau_0^{2/5}}{\eta_{\min}^{3/5}} \tag{A7}$$

$$\tau_3 > \frac{A_3}{A_0^{2/5}}\frac{\tau_0^{2/5}}{\eta_{\max}^{3/5}}. \tag{A8}$$

These conditions on τ_0, τ_3 ensure that the original conditions on the mass ratio are satisfied. Eliminating η from (A1) and (A2) instead gives

$$\tau_3 = \frac{A_3}{A_0}\tau_0 M. \tag{A9}$$

Applying the constraints on M gives two more boundary curves:

$$\tau_3 > \frac{A_3}{A_0}\tau_0 M_{\rm min} \tag{A10}$$

$$\tau_3 < \frac{A_3}{A_0}\tau_0 M_{\rm max}. \tag{A11}$$

Finally, we need to apply the constraints on m_1. Suppose that $m_1 = m_e$ where m_e is one of the extreme values of m_1. In this case,

$$\eta = \frac{m_1 m_2}{M^2} = \frac{m_e(M-m_e)}{M^2} \Rightarrow \tau_0 = \frac{A_0 M^{1/3}}{m_e(M-m_e)}. \tag{A12}$$

Rearranging this equation gives

$$M - \frac{A_0}{\tau_0 m_e} M^{1/3} - m_e = 0. \tag{A13}$$

Letting $M = x^3$, $-A_0/(\tau_0 m_e) = p$, and $-m_e = q$ turns this into an equation of the form

$$x^3 + px + q = 0. \tag{A14}$$

The solution to this cubic equation can be obtained from the more general cubic equation. The solutions for x are:

$$x = \left(\frac{-q}{2} + \sqrt{\frac{q^2}{4} + \frac{p^3}{27}}\right)^{1/3} + \left(\frac{-q}{2} - \sqrt{\frac{q^2}{4} + \frac{p^3}{27}}\right)^{1/3}. \tag{A15}$$

Substituting for p, q in the above for both values of m_e gives two different values of x. Each value of x can then be used to give $M = x^3$ and the two values of M can be substituted into (A9) to give the final two boundaries:

$$\tau_3 > \left.\frac{A_3}{A_0}\tau_0 x^3\right|_{m_e = m_{1,\rm min}} \tag{A16}$$

$$\tau_3 < \left.\frac{A_3}{A_0}\tau_0 x^3\right|_{m_e = m_{1,\rm max}}. \tag{A17}$$

Fig. 20 shows an example τ_0, τ_3 parameter space with the boundaries drawn in. Notice that not all of the boundaries constrain the space at one time.

Now that we know the boundaries of the τ_0, τ_3 space, it is possible to develop an efficient method to sample from it. To start, notice from equations (1) and (2) that τ_0, τ_3 are maximized when η, M take their minimum values and are minimized when η, M take their maximum values. This gives a simple way to determine the extreme values of τ_0 and τ_3.

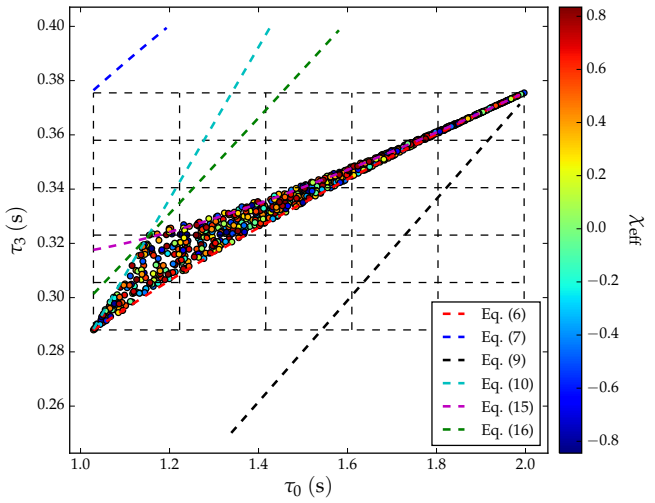

FIG. 20. Points in a template bank with the boundary curves in τ_0, τ_3 drawn in. Each of the curves corresponds to one of the boundary equations above.

To sample in τ_0, τ_3 space, start by pulling a value of τ_0 uniformly from $[\tau_{0,\rm min}, \tau_{0,\rm max}]$. Then, compute all of the constraints on τ_3 using the equations above and choose the most restrictive of them (i.e. the lowest maximum and the highest minimum). Finally, pull a value of τ_3 uniformly from the values allowed by the constraints.

Currently sampling in τ_0, τ_3 is done by choosing τ_0, τ_3 randomly in the calculated bounds and then converting to m_1, m_2 to check if the conditions on the component masses are satisfied. This method has the disadvantage that it requires many trials because most of the points do not satisfy the conditions on m_1, m_2. The method proposed above eliminates these trials because it guarantees that each point it generates satisfies the conditions on m_1, m_2.

[1] A. Einstein, Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften (Berlin), Seite 688-696. , 688 (1916).

[2] J. H. T. Jr., Binary pulsars and relativistic gravity, Nobel Prize lecture by Joseph Taylor., 2014.

[3] D. R. Lorimer, Living Reviews in Relativity **8** (2005).

[4] B. Sathyaprakash and B. F. Schutz, Living Reviews in Relativity **12** (2009).

[5] C. Cutler and K. S. Thorne, page 72 (2013).

[6] S. Privitera, *The Importance of Spin for Observing Gravitational Waves from Coalescing Compact Binaries with LIGO and Virgo*, PhD thesis, California Institute of Technology, 2014.

[7] G. M. Harry and the LIGO Scientific Collaboration, Classical and Quantum Gravity **27**, 084006 (2010).

[8] J. Abadie et al., Phys.Rev. **D83**, 122005 (2011).

[9] LIGO Scientific Collaboration et al., ArXiv e-prints (2013).

[10] J. Abadie et al., Classical and Quantum Gravity **27**, 173001 (2010).

[11] J. Abadie et al., Phys.Rev. **D85**, 082002 (2012).

[12] J. Aasi et al., Phys.Rev. **D87**, 022002 (2013).

[13] J. Abadie et al., Phys.Rev. **D85**, 102004 (2012).

[14] K. C. Chad Hanna, gstlal_inspiral: Code performance and scaling, Technical report, LIGO Scientific Collaboration, 2015.

[15] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. Creighton, Phys.Rev. **D85**, 122006 (2012).

[16] S. J. Waldman, The Advanced LIGO Gravitational Wave Detector, in *22nd Rencontres de Blois on Particle Physics and Cosmology Blois, Loire Valley, France, July 15-20, 2010*, 2011.

[17] T. Cokelaer, Phys. Rev. D **76**, 102004 (2007).

[18] I. W. Harry, B. Allen, and B. S. Sathyaprakash, Phys. Rev. D **80**, 104014 (2009).

[19] S. E. Field et al., Phys.Rev.Lett. **106**, 221102 (2011).

[20] K. Cannon et al., Phys.Rev. **D82**, 044025 (2010).

[21] K. Cannon et al., Astrophys. J. **748**, 136 (2012).

[22] I. W. Harry and S. Fairhurst, Phys.Rev. **D83**, 084002 (2011).

[23] C. Hanna, *Searching For Gravitational Waves From Binary Systems In Non-stationary Data*, PhD thesis, LSU, 2008.

[24] B. Allen, Phys.Rev. **D71**, 062001 (2005).

[25] J. Abadie et al., (2012).

[26] B. J. Owen, Phys. Rev. D **53**, 6749 (1996).

[27] S. Babak, R. Balasubramanian, D. Churches, T. Cokelaer, and B. S. Sathyaprakash, Classical and Quantum Gravity **23**, 5477 (2006).

[28] R. Prix, Classical and Quantum Gravity **24**, 481 (2007).

[29] P. Ajith et al., Phys. Rev. D **89**, 084041 (2014).

[30] S. Babak, Classical and Quantum Gravity **25**, 195011 (2008).

[31] B. J. Owen and B. S. Sathyaprakash, Phys. Rev. D **60**, 022002 (1999).