Searching for gravitational waves from the coalescence of high mass black hole binaries

LIGO SURF Progress Report 1

Johnathon Lowery
Advisors: Surabhi Sachdev, Tjonnie Li, Kent Blackburn and Alan Weinstein
Indiana University, Bloomington, IN 47405
(Dated: July 6, 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. Output from the detectors is analyzed using an advanced form of matched filtering which helps to pull out small signals from a background of noise. One promising class of sources of gravitational waves is binary black holes. This progress report details steps taken toward optimizing the analysis pipelines for Advanced LIGO in the context of binary black hole detection.

I. STATUS REPORT

In these first few weeks I have been developing intuition about different aspects of gravitational wave detection including the inspiral process and detector response. These investigations have allowed to me to get to know my mentors better and have given me opportunities to engage with the other SURF students working in my group. I have also read a good deal about the LIGO analysis pipelines in preparation for making modifications to different pipeline parameters. In the next month I hope to continue learning about astrophysical sources of gravitational waves by attending the Caltech Gravitational Wave Astrophysics School and to begin work on optimizing the `gstlal` analysis pipeline by changing various tunable parameters and running with signal injections.

What follows is a working draft of my final report which contains an updated summary of what I have learned and worked on so far.

II. BACKGROUND

A. Gravitational Waves and Their Sources

Gravitational waves (GWs) were first predicted by Einstein in a 1916 paper, 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 [1]. 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 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.

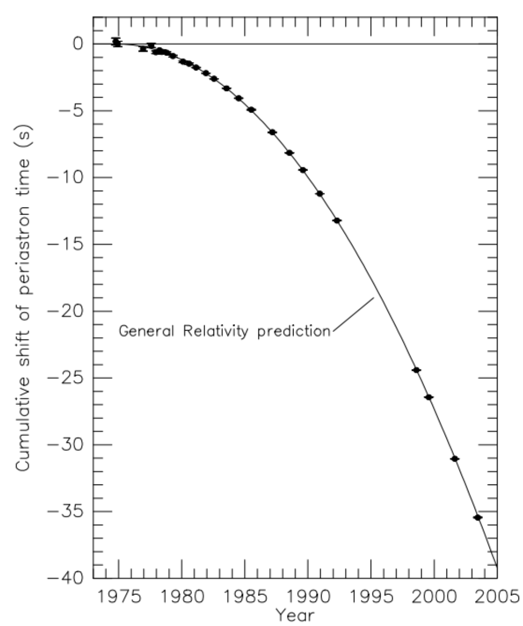

FIG. 1. Accumulated shift in the orbital phase relative to an assumed orbit with constant period caused by energy loss to gravitational waves. The straight line represents the prediction with no losses and the curve is the prediction from general relativity. Plot taken from [2].

The generation of gravitational waves requires a quadrupole source as opposed 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 [3].

Binary systems evolve through 3 phases: an inspiral phase driven by the release 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 [4]. During this entire process, called compact binary coalescence (CBC), the binary system is emitting gravitational waves. Each phase of the binary coalescence represents a distinct regime of the gravitational wave signal and must be computed independently. As an example, consider Fig. 3 which shows a gravitational waveform for the inspiral of a pair of 50 $M_{\odot}$ black holes at a distance of 1 Mpc.

When spin is introduced, the waveforms generated by binary systems become even more complicated. In the simpler case of aligned spins (spins pointing in the same direction as the orbital angular momentum), 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) [5]. Table I lists these parameters and gives descriptions. The plots in Fig. 2 show how varying different parameters of the binary system impacts the waveform.

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 disturbances in space-time caused by 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 [3].

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 [6]. 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 [7]. An example waveform in the frequency domain for a binary system is shown against the early Advanced LIGO noise curve in Fig. 5. It is expected that Advanced LIGO will detect 40 neutron star mergers per year and between 30 and 100 black hole mergers [8].

Current limits on the rate of binary black hole mergers are set by combined searches from the LIGO and 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 $M_{\odot}$ [9], systems with total mass between 25 and 100 $M_{\odot}$ [10], and systems with total mass between 100 and 450 $M_{\odot}$ [11]. No detections have been made yet, but the upgraded Advanced LIGO detectors may find the first definitive gravitational wave signal from these sources.

III. THE `gstlal` PIPELINE

When the LIGO detectors are in operation, there is a constant stream of data that must be analyzed to look for gravitational wave signals. 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 `gstlal` currently has latency of the order of 30 seconds [12]. This fast separation of signal and noise is accomplished through an advanced form of matched filtering [13]. In what follows, matched filtering will first be discussed followed by a description of some of the many additional techniques LIGO must use.

A. Matched Filtering

Basic matched filtering correlates a template signal (which is known a priori) with data to detect a potential signal. This means that if the gravitational wave signal

TABLE I. 15 of the parameters specifying a compact binary system. The sky position angles, orientation relative to the line of sight, 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 [5].

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}

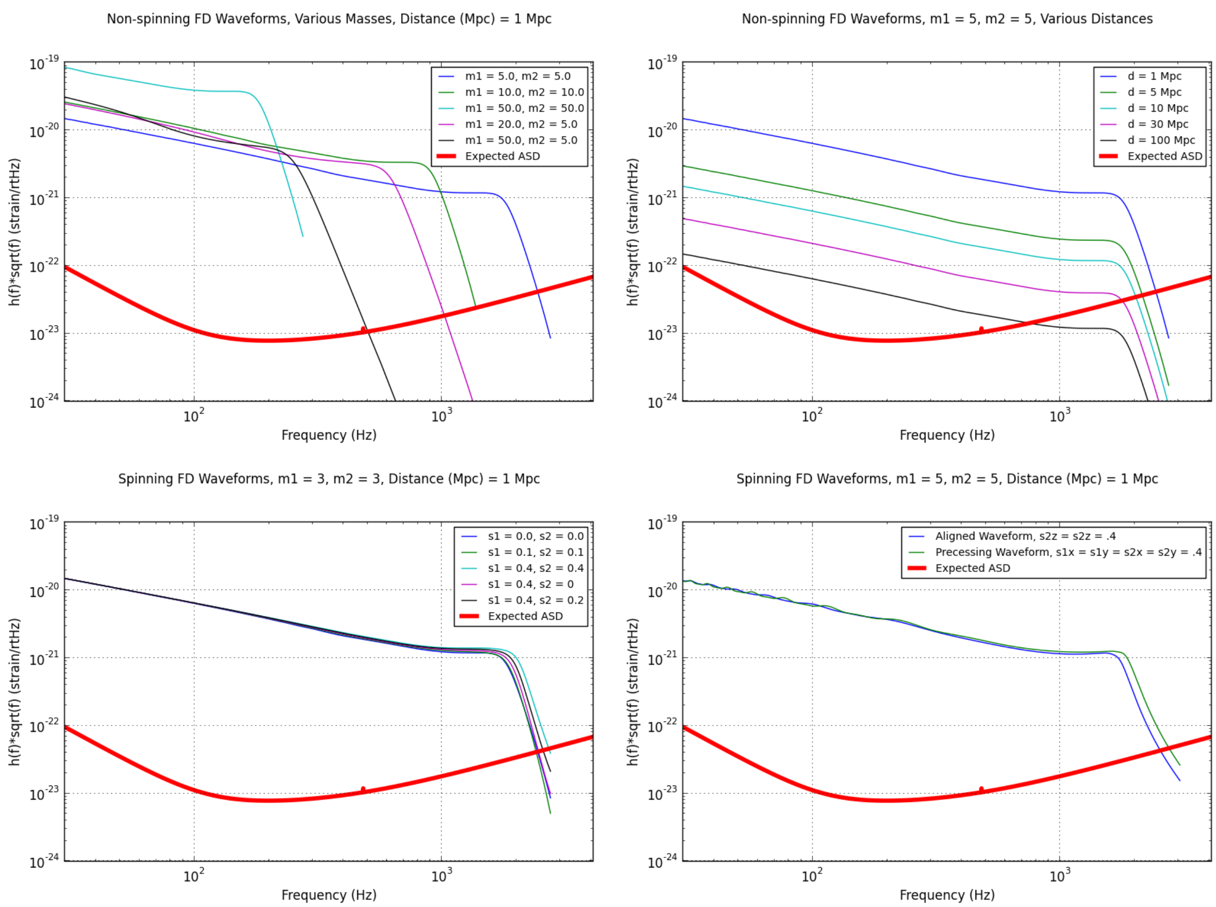

FIG. 2. Frequency domain waveforms for compact binary systems generated by varying the binary companions' masses (top left), distances from the detectors (top right), and spins in the aligned (z) direction (bottom left). The bottom right plot shows the effect of adding non-zero spins in the x- and y-directions to create a precessing waveform.

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 where the form of the signal is known 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 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 a Gaussian probability distribution that is independent of time. Thus the detector output takes the form

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

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

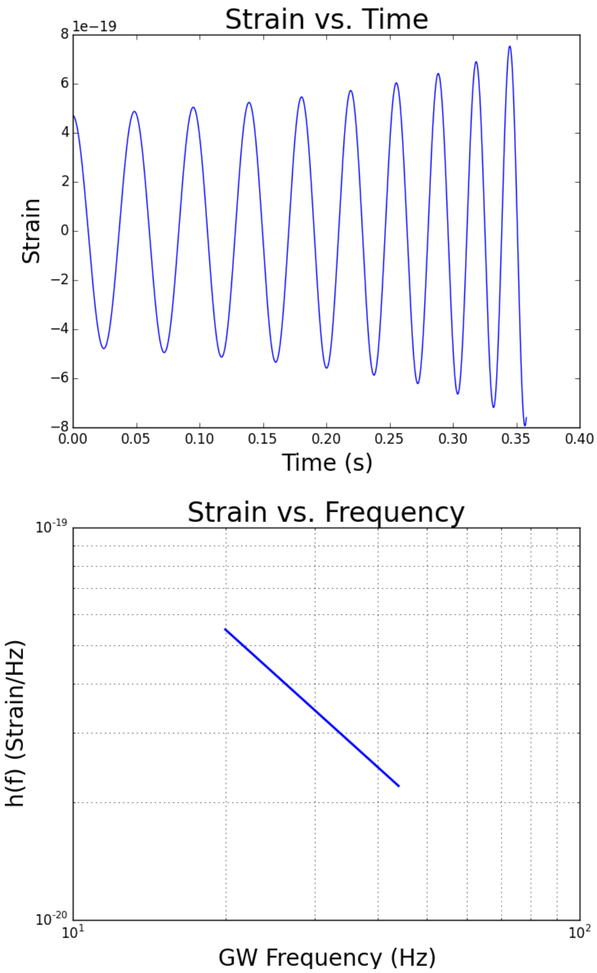

FIG. 3. A time domain waveform (top) and a frequency domain waveform (bottom) generated using the quadrupole and stationary phase approximations for the inspiral of a pair of non-spinning 50 $M_\odot$ black holes at a distance of 1 Mpc.

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. BH/BH mergers can be detected at such large distances that cosmological effects are important and so the detection distance is instead given as a redshift, z [4].

	NS/NS	NS/BH	BH/BH
$\mathcal{D}$	300 Mpc	650 Mpc	$z = 0.4$
$\mathcal{R}$, yr^{-1}	1-800	$\lesssim$ 1-1500	$\lesssim$ 30-4000

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

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

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 [6].

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

where the tilde denotes the Fourier transform given by

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

Now, 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

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

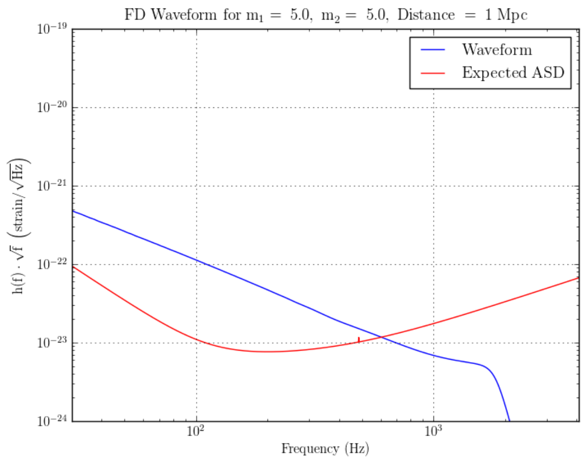

FIG. 5. An example frequency domain gravitational waveform (blue) generated for a binary system with two 5 solar mass companions at a distance of 1 Mpc. The expected noise amplitude spectral density (ASD) for early Advanced LIGO is shown in red.

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 in the red curve in Fig. 5. 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 function $a(t)$ 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

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

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

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

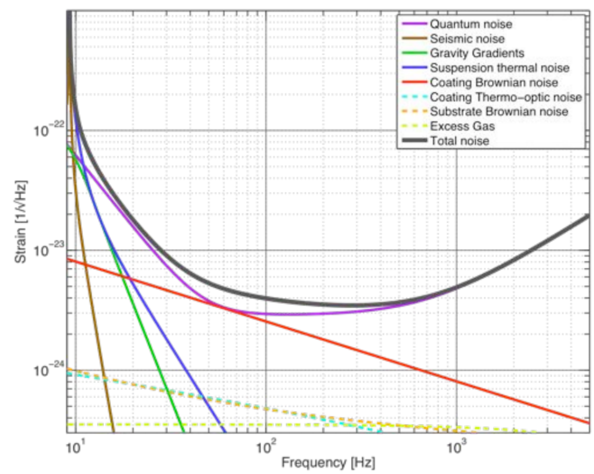

FIG. 6. The Advanced LIGO noise ASD with the different noise contributions shown. Plot taken from [14].

The signal to noise ratio 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 contain non-stationary, non-Gaussian noise which cannot be accounted for in the standard matched filter method. Below I will discuss some of the techniques used by `gstlal` to ovecome 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 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. An example of a template bank spanning the parameter space of masses and the effective spin parameter χ_{eff} is shown in Fig. 7.

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.

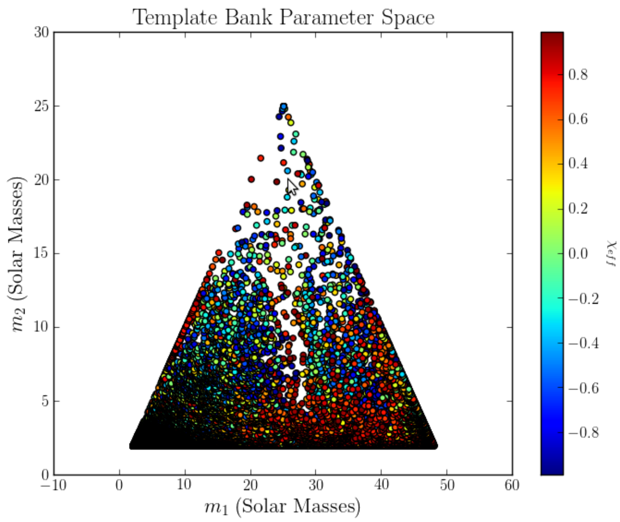

FIG. 7. An example of the parameter values used in a template bank where the color of the dots indicates the value of the effective spin parameter χ_{eff}.

The procedure of constructing an optimal template bank has been extensively studied and several methods have emerged. In the case of searches over just the mass parameters of the system (not as relevant now that spin is considered important in these searches), it has been shown that hexagonal template placement is optimal [15]. 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 those around them [16]. This method has the advantage that is scalable. Yet another method uses a modification of the Gram-Schmidt process to construct an optimal basis of templates [17]. Below I will discuss some methods used by the Advanced LIGO pipelines to compress the template banks to make the search procedure faster.

C. 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 [18]. In what follows, some of the details of this technique along with the computational benefits are discussed.

To begin, we construct the $m \times n$ 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.

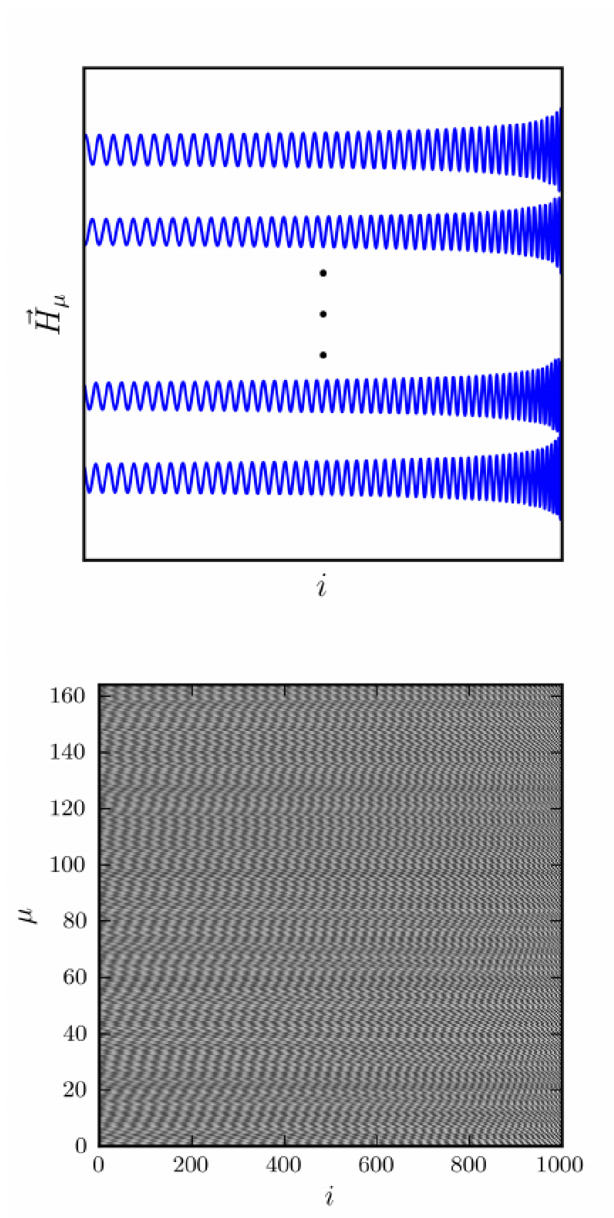

FIG. 8. Top: a series of templates in the time domain. Bottom: a template matrix $\mathbf{H}$ whose rows are templates and the columns are time samples. Fig. taken from [18].

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:

$$\mathbf{H} = \mathbf{U\Sigma V^T} \tag{12}$$

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

$$H_{\mu j} = \sum_{\nu=1}^{N} v_{\mu\nu}\sigma_\nu u_{\nu j}. \tag{13}$$

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 $\mathbf{H}$. As a common convention, the σ_i are listed in descending order so that choosing the largest N' of the σ_i is equivalent to choosing the first N'. This truncation produces an approximate reconstruction of $\mathbf{H}$ given by

$$H_{\mu j} \approx H'_{\mu j} = \sum_{\nu=1}^{N'} v_{\mu j}\sigma_\nu u_{\nu j} \tag{14}$$

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

$$\left\langle \frac{\delta\rho}{\rho} \right\rangle = \frac{1}{2N} \sum_{\nu=N'+1}^{N} \sigma_\nu^2. \tag{15}$$

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 [18], 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.

D. Multivariate Filtering

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. This predictable pattern in the signal can be taken advantage of by using different sampling rates. Dividing the signal into different pieces that are sampled at different rates is called multivariate filtering (or multibanding) and is a technique used to efficiently analyze LIGO data [19].

There is a theorem in signal processing called the Nyquist-Shannon Theorem which 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:

$$h[k] = \sum_{s=0}^{S-1} \begin{cases} h^s[k] & \text{if } t^s \le k/f^0 < t^{s+1} \\ 0 & \text{otherwise} \end{cases} \tag{16}$$

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:

$$h^s[k] = \begin{cases} h[k\frac{f^0}{f^s}] & \text{if } t^s \le k/f^s < t^{s+1} \\ 0 & \text{otherwise} \end{cases} . \tag{17}$$

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.

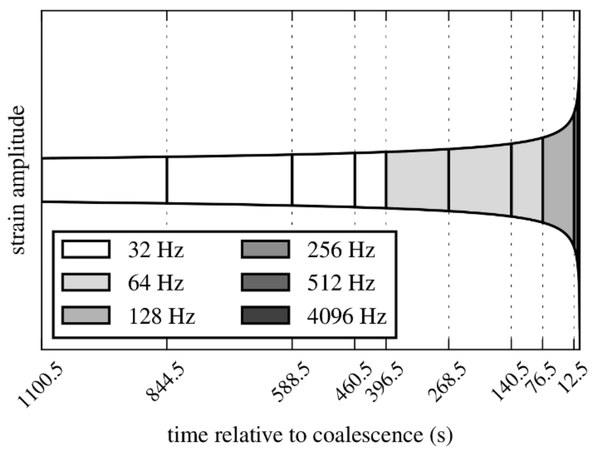

FIG. 9. A time domain chirp waveform with time slices chosen. The sampling frequency of each time slice is indicated by the shading. Fig. taken from [19].

E. 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. Below I will describe two χ^2 tests 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

$$(a(f), b(f))_p = \int_{-\Delta f_p \cup \Delta f_p} \frac{a^*(f)b(f)}{S_n(f)} df. \qquad (18)$$

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

$$(\tilde{h}, \tilde{h})_p = \frac{1}{p} \qquad (19)$$

Now, if we consider the signal to take the simple form

$$s(t) = \frac{D}{d} h(t - t_0) \qquad (20)$$

where D is the distance to the source, d is the effective distance, and 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

$$\rho = (s|s)^{1/2} = \frac{D}{d}(h \cdot h)^{1/2} = \frac{D}{d} \qquad (21)$$

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

$$\langle \rho_j \rangle = \frac{1}{p}\frac{D}{d} \quad \langle \rho_j^2 \rangle = \frac{1}{p} + \frac{1}{p^2}\left(\frac{D}{d}\right)^2 \qquad (22)$$

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

$$\Delta\rho_j \equiv \rho_j - \frac{\rho}{p}. \qquad (23)$$

Then the χ^2 is defined to be [20]:

$$\chi^2 \equiv p \sum_{j=0}^{p} (\Delta\rho_j)^2 \qquad (24)$$

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.

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 [5].

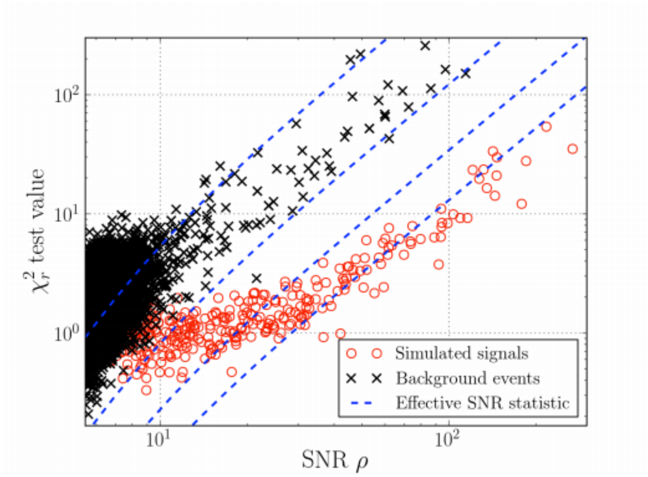

FIG. 10. A plot showing signal and background events with different SNR and χ^2 values. It is clear that a simple SNR threshold is not enough to separate signal from background. Fig. taken from [5].

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 $\langle h, h\rangle = 1$. Then, the SNR time series is:

$$\rho(\tau) = \langle n, he^{2\pi i f\tau}\rangle + A\langle h, he^{2\pi i f\tau}\rangle \tag{25}$$
$$= \langle n, he^{2\pi i f\tau}\rangle + \alpha(\tau) \tag{26}$$

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 ρ_{max}, $\rho(\tau)$, and $\alpha(\tau)$ are easily computable from the templates and data while the quantity $\langle n, he^{2\pi i f\tau}\rangle$ will be Gaussian distributed when the noise is Gaussian. Thus, it is possible to compute a χ^2 of the form

$$\chi^2 = \int_0^{T_{max}} |\rho(\tau) - \rho_{max}\alpha(\tau)|^2 d\tau. \tag{27}$$

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 $N = T_{max}/\Delta t$).

F. 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.

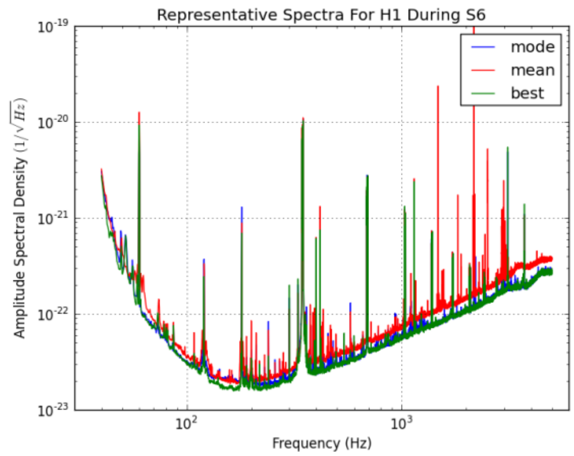

FIG. 11. Noise ASD curves for the H1 detector during S6 (May 9, 2010 - November 4, 2010). The ASD for the mode (blue), mean (red), and best (green) operation are shown. Plot taken from [21].

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 [5]. 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.

G. Tuning the Pipeline

Another method for handling noise is to tune the matching algorithms used in the analysis pipelines. A simple example of a parameter that can be varied is the threshold signal-to-noise ratio used in determining if a signal is significant. A more involved example is the type of χ^2 (either a traditional χ^2, an autocorrelation χ^2, or one of many others that have been developed [22, 23]) used in the statistical tests performed on the data to determine significance.

[1] J. H. T. Jr., Binary pulsars and relativistic gravity, Nobel Prize lecture by Joseph Taylor., 2014.
[2] D. R. Lorimer, Living Reviews in Relativity **8** (2005).
[3] B. Sathyaprakash and B. F. Schutz, Living Reviews in Relativity **12** (2009).
[4] C. Cutler and K. S. Thorne, page 72 (2013).
[5] 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.
[6] G. M. Harry and the LIGO Scientific Collaboration, Classical and Quantum Gravity **27**, 084006 (2010).
[7] J. Abadie et al., Phys.Rev. **D83**, 122005 (2011).
[8] R. K. Kopparapu et al., The Astrophysical Journal **675**, 1459 (2008).
[9] J. Abadie et al., Phys.Rev. **D85**, 082002 (2012).
[10] J. Aasi et al., Phys.Rev. **D87**, 022002 (2013).
[11] J. Abadie et al., Phys.Rev. **D85**, 102004 (2012).
[12] K. C. Chad Hanna, gstlal_inspiral: Code performance and scaling, Technical report, LIGO Scientific Collaboration, 2015.
[13] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. Creighton, Phys.Rev. **D85**, 122006 (2012).
[14] J. Aasi et al., Class.Quant.Grav. **32**, 074001 (2015).
[15] T. Cokelaer, Phys.Rev. **D76**, 102004 (2007).
[16] I. W. Harry, B. Allen, and B. Sathyaprakash, Phys.Rev. **D80**, 104014 (2009).
[17] S. E. Field et al., Phys.Rev.Lett. **106**, 221102 (2011).
[18] K. Cannon et al., Phys.Rev. **D82**, 044025 (2010).
[19] K. Cannon et al., Astrophys. J. **748**, 136 (2012).
[20] B. Allen, Phys.Rev. **D71**, 062001 (2005).
[21] J. Abadie et al., (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.