Rates and populations of compact binary mergers

Caltech LIGO SURF 2019, Second Interim Report

Phoebe McClincy (phoebemcclincy@gmail.com)
Department of Astronomy and Astrophysics, Pennsylvania State University

Mentors: Alan Weinstein, Jonah Kanner, Liting Xiao
LIGO Laboratory, California Institute of Technology
(Dated: 2 August 2019)

The LIGO and Virgo detectors have been observing the cosmos in search of gravitational waves (GW) since 2000. All three detectors were upgraded to Advanced versions, which for LIGO began observing in 2015 and for Virgo in 2017. In Advanced LIGO's first (12 September 2015 to 19 January 2016) and second (30 November 2016 to 25 August 2017) observing runs (O1 and O2, respectively), the detectors found 10 GW signals from binary black hole (BBH) mergers, and 1 from a binary neutron star (BNS) merger, all with high significance, or low probability of being due to instrumental noise fluctuations. Already in the first several months of O3, which began in April 2019, dozens of candidates have been seen with such high significance. The two aforementioned categories, along with neutron star/black hole mergers (NSBH), are collectively known as compact binary coalescence (CBC). In the coming years, as the detectors' sensitivities are improved, we expect to accumulate tens, hundreds, or thousands of CBC events. From such large samples, we expect to be able to infer the underlying population of CBC systems as a function of their masses, component black hole spins, and redshift. This, in turn, will allow us to better understand the astrophysical processes governing the formation, evolution, and final fate of such systems, as tracers of the most massive stars. In this project, we aim to develop tools and techniques to accomplish this through detailed simulation and Bayesian inference. *We report seven weeks of progress during this project.*

I. INTRODUCTION

Since Advanced LIGO began collecting data in 2015 (and Advanced Virgo in 2017), the collaboration has made many detections of CBC events. As the detectors improve, the number of events that we are able to recover will increase drastically. With a population size akin to the one we expect to see, we are able to identify and measure the event rate as a function of mass, spin, redshift. We may then use this information to infer the mechanisms of formation and evolution that govern such binaries.

Since we expect to accumulate large amounts of events in the future, it is now the optimal time to develop and debug the tools necessary to use when such a time comes. Therefore, in this project, we concentrate on developing the tools and techniques to do so.

II. MOTIVATIONS

A. Measuring binaries with gravitational waves

Gravitational radiation comes from rapidly-changing gravitational fields, and is observed in the form of a wave which propagates at the speed of light. When a GW propagates, it distorts spacetime, causing a change in the distance between objects. We thus observe a change in the length of the arms of our Michelson laser interferometer. This quantity is called strain:

$$h = \frac{\Delta L}{L}, \tag{1}$$

where L is the original length of the detector arm, and ΔL is the change in that length as induced by a GW of strain h.

General relativity allows us to compute the strain of a GW if we know the binary's parameters. Strain depends primarily on the intrinsic parameters, mass and spin, which directly affect the shape of the observed waveform. The other parameters, which are all extrinsic (observer-dependent), govern only the strength of the signal. We can thus analyze the evolution of the waveform's shape by using Bayesian inference techniques to determine the chirp mass $\mathcal{M}$, as well as the symmetric mass ratio η. The component masses are implicit in η [1]:

$$\eta = \frac{m_1 m_2}{M_{tot}^2} = \left(\frac{\mathcal{M}}{M_{tot}}\right)^{\frac{5}{3}}. \tag{2}$$

The ultimate goal of this project is to learn about the universe's more massive stars, which lead to the formation of BBHs. The component masses of a BBH are related to the masses of its progenitor stars. More massive stars, such as those that form BBHs, have lower metallicity and were formed early in the universe's evolution (before heavier elements existed). Younger, higher-metallicity stars have masses too small to form BBH; as

such, it is important to determine the underlying distribution of the high-mass progenitor stars so that we may be informed about their evolution into BBHs.

At present, there exist well-supported models of the star formation rate as a function of redshift (z), such as the Salpeter Initial Mass Function [2] and the Madau-Dickinson model [3]. These models are based on observations from electromagnetic radiation [4, 5]. For BHs, we postulate a relation between BH event rate density $\mathcal{R}$, z, and mass. We describe $\mathcal{R}$, or the number of events per unit comoving volume per unit time, as

$$\mathcal{R}(m_1, m_2, z) = \frac{dN}{dV_c dt_s}. \tag{3}$$

We aim to determine the dependence of $\mathcal{R}$ on z, m_1, and m_2. Such a measurement will allow us to infer the underlying mass distribution and track the evolution of BBHs in cosmic spacetime.

Although there are six parameters used to describe spin, only two combinations of them dominate in the phase and amplitude evolution of a gravitational waveform: $\chi_{\rm eff}$ and $\chi_{\rm p}$. $\chi_{\rm eff}$ is the mass-weighted combination of the component spins along the binary's orbital angular momentum vector [6]. $\chi_{\rm p}$ refers to the components of spin perpendicular to the orbital angular momentum which cause the binary's orbit to precess. These quantities carry information about the mechanisms by which a binary was formed, and are described [7] by

$$\chi_{\rm eff} = \frac{a_1 m_1 \cos\theta_1 + a_2 m_2 \cos\theta_2}{m_1 + m_2}, \tag{4}$$

and

$$\chi_{\rm p} = \max\big(a_1 \sin\theta_1, (\frac{4m_1 + 3m_2}{4m_2 + 3m_1})(\frac{m_1}{m_2}) a_2 \sin\theta_2\big), \tag{5}$$

where θ_1 and θ_2 are the angles between the angular momentum vectors of each component and the binary's total angular momentum vector.

Because spin is a higher-order effect in gravitational waveform evolution, it is more difficult to measure, whereas mass is of first-order and its range of values can be more accurately constrained. As such, we will primarily focus on mass for now.

B. Channels of formation

If we are able to determine the underlying distribution of masses, we are then able to gain valuable insights into possible methods by which a binary in question was formed. There are many proposed formation channels, but two (shown in Fig. 1) that are of particular interest to us [8]. The first main channel is common evolution. In one common evolution sub-channel marked by a common envelope phase, the BBH evolves via traditional collapse of both components from a progenitor stellar binary, within a single cloud of gas. In another sub-channel, the BBH forms via chemically homogeneous evolution, in which the orbit of the binary components does not expand traditionally during main sequence helium production, but rather remains compact [9].

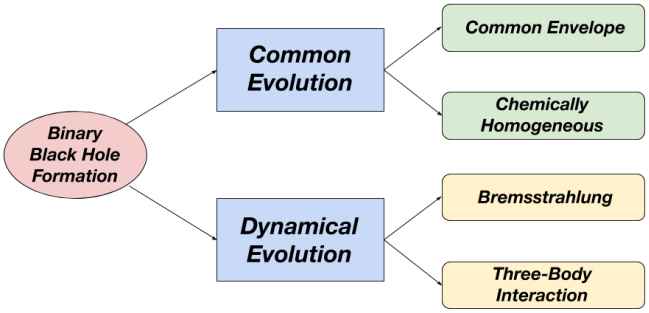

FIG. 1. Formation channels of interest for binary black holes. The first such channel is common evolution, in which binaries are either formed from common envelope evolution or chemically-homogeneous evolution. The second such channel is dynamical formation, in which black holes formed in separate environments interact, in either two-body or three-body interactions.

The second main channel is dynamical formation, which entails the interaction of components formed independently of each other. In one such case, a binary system interacts with a third body in a dense stellar cluster, which results in the ejection of the binary's less massive component and the capture of the third, more massive body into the binary. In another dynamical formation case, one single body captures a second body via the gravitational Bremsstrahlung radiation caused by the acceleration of the second body through the first body's gravitational field. We aim to be able to distinguish between these formation channels, primarily from mass, then eventually from spin.

C. Populations, mass distributions, and mass gaps

Within the mass distribution of CBC events (shown in Fig. 2), there exist three proposed regions of scarcity (mass gaps): one below 1 $M_\odot$, one between ∼2-5 $M_\odot$, and another between ∼50-150 $M_\odot$. The first is proposed to exist due to the small likelihood that traditional stellar collapse mechanisms would produce any compact bodies in this region [10]. The second mass gap is thought to be caused by a disparity in NS and BH masses; NS masses gather between 1-2 $M_\odot$, whereas BH masses tend to begin around 5 $M_\odot$ [11]. The latter mass gap is proposed to exist due to pulsational pair-instability supernovae [12], in which progenitor binary stars with component masses between 100-150 $M_\odot$ eject a significant amount of their mass upon going supernova. This theoretically causes the

subsequent BBH component masses to settle around ∼40 $M_\odot$ [13]. We can thus infer truths about this particular mass gap by measuring the underlying mass distribution of BBH.

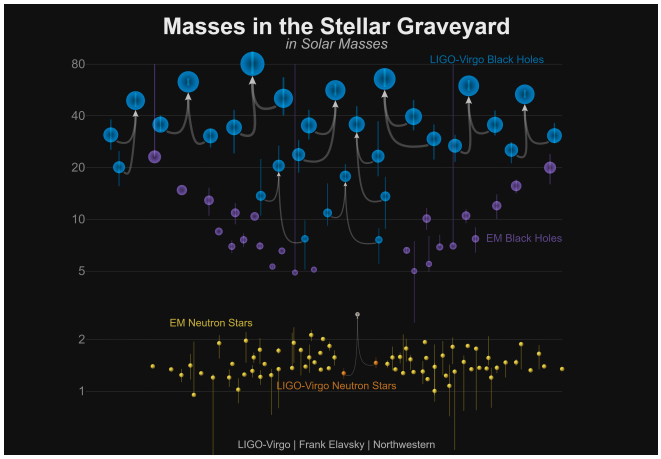

FIG. 2. Distribution of NS and BH masses. Yellow represents NS detected via electromagnetic radiation, and orange represents NS detected via Advanced LIGO. Purple represents BHs detected via electromagnetic radiation, and blue represents BHs detected via Advanced LIGO. There are regions of scarcity below 1 $M_\odot$ and between ∼2-5 $M_\odot$, as well as a possible mass gap between ∼50-150 $M_\odot$. The x-axis does not have significance in this figure.

III. PROJECT

At present, LIGO has catalogued over 30 CBC events, most of which have been detected solely in the third observing run. As our detectors improve, the volume of spacetime in which LIGO is able to observe increases. This quantity is known as the sensitive spacetime volume

$$\langle VT \rangle = \frac{4}{3}\pi D_{avg}^3 T, \tag{6}$$

where D_{avg} represents LIGO's sensitive distance and T represents the observation time of the LIGO detectors [10]. This represents the sensitive spacetime volume in Euclidean space, and this equation no longer holds in the presence of cosmological effects. It is important to note that D_{avg} is a strong function of mass; systems with a larger overall mass produce louder GW, and result in a larger D_{avg}. Because $\langle VT \rangle$ is proportional to the sensitivity of our detectors, we are able to observe larger regions of spacetime as we produce higher-sensitivity detectors.

As larger regions of spacetime are observed, LIGO is expected to recover events in larger and larger numbers. A larger sample size of BBHs offers the unique opportunity to reveal the underlying naturally-occurring relationship between merger rate and BBH masses, spins, and redshift. By studying populations of BBHs, we may better understand the relationship between the progenitor star initial mass function and the mass function that governs BBH. We also stand to uncover information about models that describe formation channels for BBHs, and which formation channel is most prevalent in nature. We plan to carry out detailed simulation and Bayesian inference to do so.

The actual number of events that we have observed (or our collected data) is described as N. The true number of events that occur in nature is described as

$$\hat{N}_{true} = \int \frac{dN}{dm_1 dm_2 dz dt_s} dm_1 dm_2 dz dt_s. \tag{7}$$

The component masses m_1 and m_2, as well as source time t_s and z, can be written as a series of parameters called $\vec{\theta}$. We assume that $\hat{N}_{true}$ depends on m_1 and m_2, the distribution of which we will describe by hyperparameters α and β, respectively. We also assume that $\mathcal{R}$ has a dependence on z, which is described by another hyperparameter γ. Collectively, we denote these hyperparamters $\vec{\lambda}$. Spin is also believed to be a dependent of $\hat{N}_{true}$; however, we will not address spin at this point in the project. We can thus rewrite Eqn. (7) as

$$\hat{N}_{true} = \int \frac{dN(\lambda)}{d\vec{\theta}} d\vec{\theta}, \tag{8}$$

where

$$\frac{dN(\lambda)}{d\vec{\theta}} = \mathcal{R}(1+z)^\gamma f(m_1|\alpha) f(m_2|\beta) \frac{dV_c}{dz} \frac{dt_d}{dt_s} \frac{1}{T_d}, \tag{9}$$

with t_d being the time as measured at the detector (which has been dilated with respect to the time at the source due to cosmological expansion), T_d the observation time of the detector, and V_c the comoving volume, whose relation to z is determined in accordance with the ΛCDM cosmological model [14]. From this, we can construct an expression for the expected amount of events that we will observe:

$$\hat{N}_{det} = \int \frac{dN(\vec{\lambda})}{d\vec{\theta}} \mathcal{E}(\vec{\theta}) d\vec{\theta}, \tag{10}$$

where $\mathcal{E}(\vec{\theta})$ represents the efficiency of detection. Upon collecting N, we can construct a Poissonian probability distribution for N such that

$$P(N|\hat{N}_{det}, \vec{\lambda}) = \frac{\hat{N}_{det}^N e^{-\hat{N}_{det}}}{N!}. \tag{11}$$

However, we are looking to infer what the true, naturally-occurring number of events is; thus, we step through this process in the reverse.

To do this, we are currently randomly generating a population of masses and proposing possible values for the hyperparameters enveloped in $\vec{\lambda}$. We will develop methodology for measuring $\vec{\lambda}$ using Bayesian inference to compare how close our experimentally-recovered values are to our proposed values. In particular, we will use the dynamic nested sampler dynesty [15] to carry this process out. We will do this in three major steps:

(1) Simulate a dataset consisting of many observed binary systems, following the distribution outlined in Eqn. (9) with an arbitrary choice of $\vec{\lambda}$.

(2) Use Bayesian inference techniques to recover the most accurate posterior probability distribution for the underlying $\vec{\lambda}$.

(3) Reconstruct the distribution for $\hat{N}_{true}$.

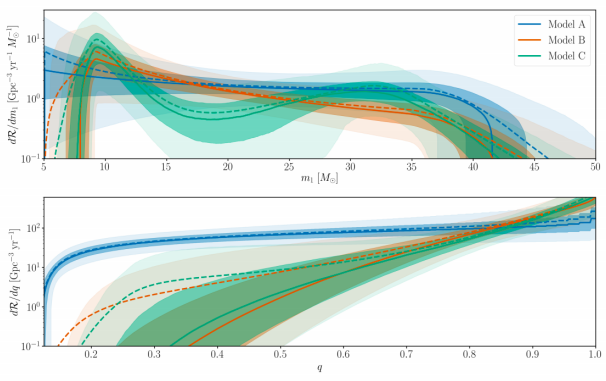

FIG. 3. Differential merger rate distribution for BBHs as a function of primary mass and mass ratio (q) for proposed Models A, B, and C in [16]. At lower masses, Model C follows a power law distribution, and at higher masses, it follows a Gaussian distribution. This distribution is based on data from O1 and O2 only (10 BBH mergers); we plan to begin our testing using this model.

This has been attempted by other groups. In [16], three models were presented for the BBH primary mass distribution, denoted Models A, B, and C. Model A fixes m_{min} to be 5 $M_\odot$, and allows m_{max} to vary. Model B allows both mass limits to vary. Model C allows multiple functions to describe the distribution; a second component of Gaussian nature appears due to the pair instability in massive progenitor stars. As such, for Model C, a power law distribution fits at lower masses, and a Gaussian distribution fits at higher masses. In this case, α and β generally refer to the power law indices, and γ describes the Gaussian component. We plan to use Model C for the analysis done in this project.

IV. PROGRESS

Thus far, I have worked on generating my own distributions and conducting parameter estimation on them. My warm-up assignment was to generate a distribution of masses that followed a power law (in particular, the Salpeter Initial Mass Function (IMF), $N = M^{-2.35}$). I did this using 50 randomly-generated points governed by Gaussian distributions with standard deviations of 0.002, shown in Fig. 4.

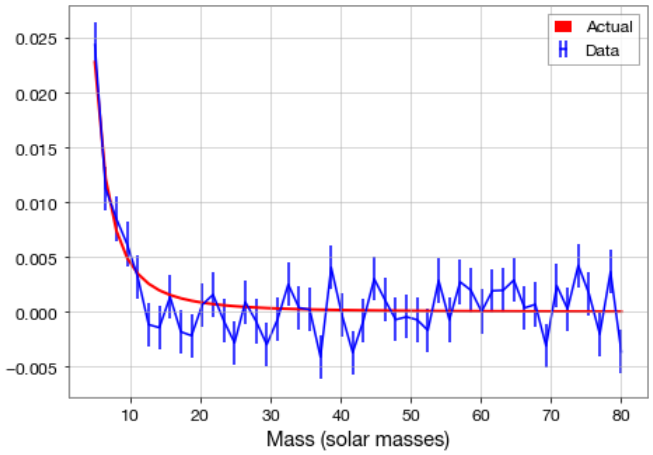

FIG. 4. Results from the initial random generation of masses following the Salpeter IMF; the red curve is the actual distribution, whereas the blue curve represents the 50 generated points and their Poisson errors.

I then used `dynesty` for parameter estimation, and generated posterior distributions for both hyperparameters of the power law, the amplitude and power law slope. Overall, the true values and the data values were consistent, as shown in Fig. 5.

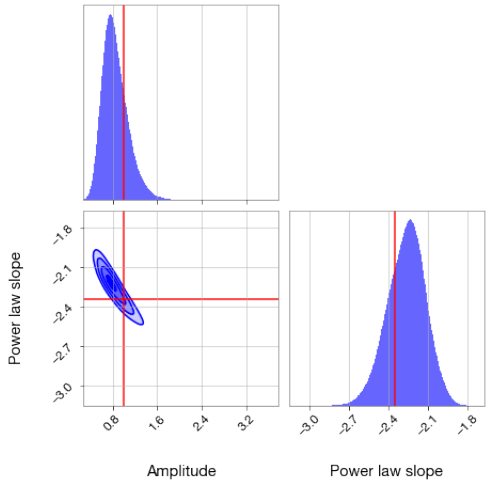

FIG. 5. Posterior distributions for amplitude and power law slope as generated by dynesty, along with the parameters' correlation.

I then repeated this process using a different method called the Inverse Transform Method, which entails using the inverse cumulative density function (CDF) to sample the original distribution. This is done by generating numbers selected from a uniform distribution between 0

and 1, and transforming them into points on the original distribution using the inverse CDF.

I have recently been working on writing scripts for use in this project using Python 3.7 code. There is a series of four scripts that I have written to carry out the work for this project.

The first such script focuses on the generation of one "injection." An injection is a waveform that we have created from predetermined parameters (such as mass, spin, etc.). To do this, I have used LIGO-created software called `PyCBC`. This code allows one to use already-defined programs and functions to generate waveforms, in both the time and frequency domains. As such, I used my script to first generate a waveform in the time domain.

I then applied more `PyCBC` code to translate the ideal waveform to what would be seen in the detector's frame. Because CBC systems are usually very far away in the universe and are relativistic in nature, the waveform that is emitted from the CBC system is not what is observed on Earth. We observe the now redshifted and time-dilated event occurring at lower frequencies and higher masses, and must subsequently correct for this.

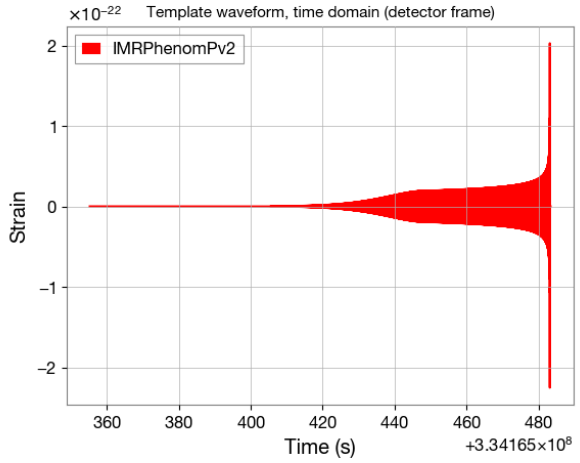

FIG. 6. A gravitational wave pictured in the time domain. This signal has been tapered at the beginning. Inspiral occurs until approximately 485s, at which point merger causes the signal to sharply spike, and ringdown follows.

Because we are concerned with populations of BBH, we must repeat this process for many iterations. I created a second script that does just this. The second script calls the "make injection" function I defined in the first script repeatedly. It is also here that the parameters for each waveform are randomly generated from their respective distributions. At present, we are generating masses from the same power law distribution as before (the Salpeter IMF). The spin, right ascension, polarization angle, coalescence phase, inclination, and time are all drawn from uniform distributions. Declination (δ) is determined by drawing from a distribution uniform in $sin(\delta)$. Distance is determined by drawing from a distribution uniform

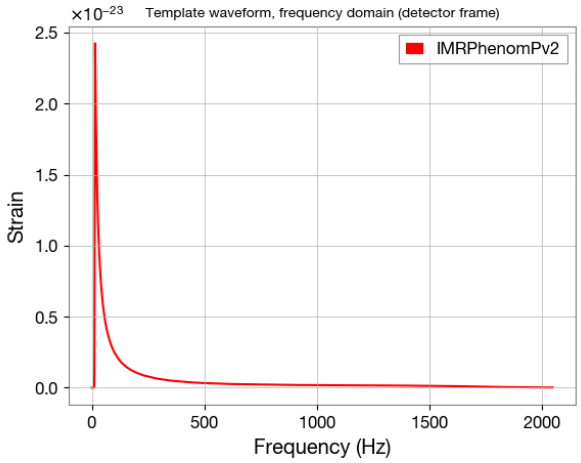

FIG. 7. A gravitational wave pictured in the frequency domain. Here, we see the chirp at a seemingly low frequency. This may be due to the masses of the component BHs, as higher masses result in lower merger frequencies.

in comoving volume; here, we must be sure to determine what the maximum distance (rather, redshift) will be. These parameters are generated for each individual waveform. The second script stores these parameters in a file.

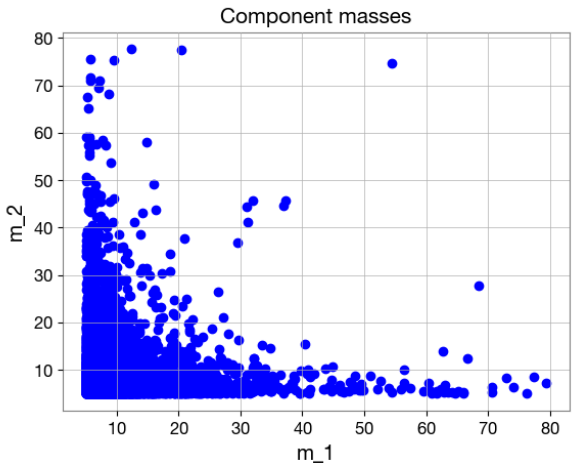

FIG. 8. Primary mass vs. secondary mass for 10,000 events, drawn from a power law distribution (the Salpeter IMF). The code that generates these masses will later be altered to always use the larger of the two masses as the primary mass.

It is important to consider whether or not we would be able to recover these simulated waveforms in application. The efficiency referenced in Eqn.(10) plays a large part in whether or not we are able to recover a signal. To test this, I have written a third script (yet to be finalized) that recovers the optimal signal-to-noise ratio (SNR) from each dataset. This can be done two ways; the first is to run a matched-filter-based search on each

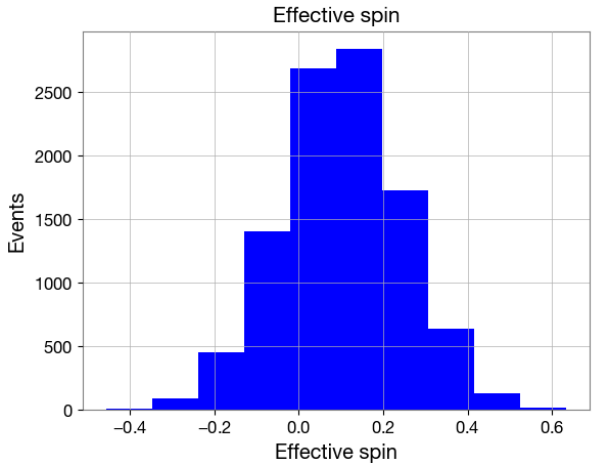

FIG. 9. Histogram of effective spin for 10,000 events, as calculated from Eqn.(4). The z-component of each BH's spin is drawn from a Gaussian distribution centered at 0.1 with standard deviation 0.2.

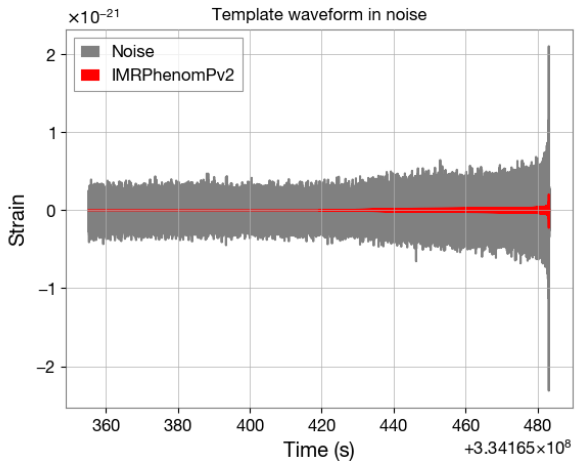

FIG. 10. A gravitational wave in the time domain within noise. This noise will be simulated by drawing a randomly-generated number from a Gaussian distribution centered at 0 with a standard deviation of 1, and adding it to the SNR.

dataset using `PyCBC` code to recover the SNR time series, and respectively the peak SNR. The second is to calculate the optimal SNR for each dataset using

$$\rho^2{}_{opt} = \int \frac{\hat{h}^*(f)h(f)}{\mathcal{S}(f)}df, \tag{12}$$

where $\hat{h}^*(f)$ and $h(f)$ both represent the waveform (since we are working with the optimal SNR), and $\mathcal{S}(f)$ is the PSD (power spectral density, generated from Advanced LIGO's colored noise curve). We will use the second method. These mathematics must be done in the frequency domain, so it is essential to first use `PyCBC` code to convert both the template and signal to frequency series, as the PSD is already in the frequency domain. To determine if we will recover or miss an injection, we will compute both the individual detector SNRs and the network SNR. We will account for noise by drawing a random number from a Gaussian centered at 0 with a standard deviation of 1; we will then add this number to the SNR value. For an injection, if any detector's SNR is greater than 4, it is factored into the network SNR; otherwise, it is not. Then, if the injection has a network SNR above the threshold of 9, it is recovered; otherwise, the injection is missed. To differentiate between the two, recovered injections will be assigned a value of 1, and missed injections will be assigned a value of 0. We can then see how the recovered and missed distributions compare.

A fourth script (also yet to be finalized) will then read the injections file to retrieve the mass distribution. It will run parameter estimation on the distribution using `dynesty` as the preferred sampler, and will recover the hyperparameters of the distribution (in the initial case of a power law, the amplitude and power law slope). We will then produce, for each distribution, a corner plot detailing the hyperparameters and their respective correlations with the expected values. We will also be able to compare the distribution that was generated initially to the distribution that was recovered.

In the past month, I have made a great deal of progress on the aforementioned scripts. I started with a blank Python notebook and have been able to successfully write programs that carry out these processes.

In regards to data, at this time we are only able to view the parameters that have been randomly generated, such as mass, spin, etc., and are therefore able to view the waveforms that correspond to each parameter set. However, the most interesting data will come in the form of a finished fourth script that runs parameter estimation on the generated distributions. Here, we will be able to see how well, in the future, we will be able to infer the underlying distributions and their connection to redshift and event rate. At present, we are only working with a power law distribution for the masses in order to test the scripts. Later, we will analyze the more complex Model C from [16], which involves a power law component at low masses, but evolves into a Gaussian distribution at higher masses. These changes and observations will be made in the final three weeks of this project.

V. CHALLENGES AND FUTURE PROSPECTS

In the beginning, the main sources of my challenges were figuring out how to use `dynesty` and using the Inverse Transform Method. At first, I had trouble understanding exactly what `dynesty` was used for and how it tied into the work I would be doing. During my time working on this project, I have been able to use `dynesty`

to complete the warm-up exercise and have learned quite a bit about Bayesian inference, which has aided in my understanding of `dynesty`'s use. However, I am still learning more about Bayesian inference and `dynesty`'s applications, so although it remains challenging, I hope that as I continue to learn even more, it will become easier to use – this is my experience with the code I have written thus far. In regards to the Inverse Transform Method, I carried out the process both analytically and numerically. I spent a while working on it analytically before realizing that I was working with the wrong function, which was a seemingly trivial and thus frustrating mistake. However, once I was using the correct function, I was able to complete the exercise with no problems.

Overall, my largest challenge has been successfully executing the scripts. I have experience with Python, but not in the regard of writing from scratch. For prior research, I generally only needed to change minor things about already-written scripts to do my work, and the single programming class I have taken did not teach anything past extremely basic concepts. I subsequently struggled quite a bit to write my own scripts efficiently, but I absolutely feel that I learned a lot from this experience and have become a much better programmer as a result (which was my foremost goal for the summer).

The most prominent challenge with the code has been identifying what within the first and third scripts causes a certain recurring error. In order to do the aforementioned mathematical analysis of a waveform to find the SNR, the three components (template, data, PSD) must be the same length. My code generated these components, but with different lengths for all three. It was a challenge to figure out how to fix this, but it was ultimately resolved for the first script by making all the components the length of the template. We have run into a similar issue for the third script. However, this error is more interesting to fix; the length of the data and PSD is exactly twice the length of the template, and thus the program will not run. We believe this particular error may be due to a mistake in specifying the time step or frequency step of the template, and are working on resolving this issue presently.

For the remainder of this project, my goals are to finish the third and fourth scripts using the power law distribution to complete initial testing, and later produce a final parameter estimation analysis of Model C from [16]. Ultimately, by the end of this project, I would like to be able to visualize how the Model C mass distribution is connected to redshift/event rate.

This is a similar but different goal from what I had in mind at the beginning of this project. I think writing the scripts to do the analysis has taken significantly longer than was planned. To adjust for this, I believe I may possibly only have time to analyze Model C. Depending on how quickly I finish with Model C, I may still be able to analyze other models, but my primary goal is to do an in-depth analysis of only Model C.

[1] B. S. Sathyaprakash and B. F. Schutz, "Physics, astrophysics and cosmology with gravitational waves," *Living reviews in relativity*, vol. 12, no. 1, p. 2, 2009.

[2] G. Chabrier, "The initial mass function: from salpeter 1955 to 2005," in *The Initial Mass Function 50 Years Later*, pp. 41–50, Springer, 2005.

[3] P. Madau and M. Dickinson, "Cosmic star-formation history," *Annual Review of Astronomy and Astrophysics*, vol. 52, pp. 415–486, 2014.

[4] A. E. Bauer, N. Drory, G. Hill, and G. Feulner, "Specific star formation rates to redshift 1.5," *The Astrophysical Journal Letters*, vol. 621, no. 2, p. L89, 2005.

[5] F. Mannucci, G. Cresci, R. Maiolino, A. Marconi, and A. Gnerucci, "A fundamental relation between mass, star formation rate and metallicity in local and high-redshift galaxies," *Monthly Notices of the Royal Astronomical Society*, vol. 408, no. 4, pp. 2115–2127, 2010.

[6] K. K. Ng, S. Vitale, A. Zimmerman, K. Chatziioannou, D. Gerosa, and C.-J. Haster, "Gravitational-wave astrophysics with effective-spin measurements: Asymmetries and selection biases," *Physical Review D*, vol. 98, no. 8, p. 083007, 2018.

[7] C. Talbot and E. Thrane, "Determining the population properties of spinning black holes," *Physical Review D*, vol. 96, no. 2, p. 023012, 2017.

[8] C. L. Rodriguez, M. Zevin, C. Pankow, V. Kalogera, and F. A. Rasio, "Illuminating black hole binary formation channels with spins in Advanced LIGO," *The Astrophysical Journal Letters*, vol. 832, p. L2, 2016.

[9] S. de Mink, M. Cantiello, N. Langer, and O. Pols, "Chemically homogeneous evolution in massive binaries," in *AIP Conference Proceedings*, vol. 1314, pp. 291–296, AIP, 2010.

[10] R. Magee, A.-S. Deutsch, P. McClincy, C. Hanna, C. Horst, D. Meacher, C. Messick, S. Shandera, and M. Wade, "Methods for the detection of gravitational waves from subsolar mass ultracompact binaries," *Physical Review D*, vol. 98, no. 10, p. 103024, 2018.

[11] F. Özel, D. Psaltis, R. Narayan, and J. E. McClintock, "The black hole mass distribution in the galaxy," *The Astrophysical Journal*, vol. 725, no. 2, p. 1918, 2010.

[12] C. Fryer, S. Woosley, and A. Heger, "Pair-instability supernovae, gravity waves, and gamma-ray transients," *The Astrophysical Journal*, vol. 550, p. 372, 2001.

[13] C. Talbot and E. Thrane, "Measuring the binary black hole mass spectrum with an astrophysically motivated parameterization," *The Astrophysical Journal*, vol. 856, p. 173, 2018.

[14] D. W. Hogg, "Distance measures in cosmology," *arXiv preprint astro-ph/9905116*, 1999.

[15] J. S. Speagle, "dynesty: A dynamic nested sampling package for estimating bayesian posteriors and evidences," *arXiv preprint arXiv:1904.02180*, 2019.

[16] B. Abbott, R. Abbott, T. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, R. Adhikari, V. Adya, C. Affeldt, *et al.*, "Binary black hole population prop-

erties inferred from the first and second observing runs of Advanced LIGO and Advanced Virgo," *arXiv preprint arXiv:1811.12940*, 2019.