

Testing General Relativity with Binary Black Hole Mergers

Radha Mastandrea
MIT

Alan J. Weinstein
Caltech

(Dated: September 21, 2017)

The recent detections of gravitational waves (GWs) by the Laser Interferometer Gravitational-Wave Observatory (LIGO) have provided researchers with the first opportunities to test general relativity (GR) in the strong-field and highly-dynamical limit. Qualitative tests of the agreement between LIGO’s GW observations and classical GR have already been done; we have carried out more quantitative tests in terms of controlled, parameterized deviations from GR. In this project, we simulate a number of binary black hole (BBH) merger waveforms with known amplitude and phase deviations from those predicted by GR that are governed by the real and imaginary parts, respectively, of a complex parameter λ . We use Bayesian analysis to recover the deviation. We then provide an estimate of the number of GW detections from BBH mergers that are necessary to establish a given deviation from classical GR, notably finding that under 80 events are necessary to determine λ to a precision of 0.025 (a fractional precision of 5% for $\lambda = 0.5$).

I. INTRODUCTION

Einstein’s theory of general relativity (GR), introduced to the scientific community in 1915, remained largely untested in the strong-field limit for over a century. However, in what is becoming a familiar story to scientists, on September 14th of 2015, the LIGO Scientific Collaboration made their first detection of gravitational waves [1, 2], marking the opening of a “unique observational window on the universe”[3]. LIGO’s detections simultaneously gave scientists a confirmation of Einstein’s theory of GR, which predicted the existence of gravitational waves (GWs) emitted from compact binary coalescences (CBCs), and the potential to disprove GR, since GR had never before been tested in the strong-field highly-dynamical limit.

Since these initial discoveries, many tests of GR have been made on these GW observations. As an initial confirmation of GR, tests that were conducted on GW150914 found no statistically significant evidence that the detected waveform was emitted by a system that could not be described by Einstein’s field equations [4]. Tests on the GW150914 waveform involved comparing the BBH mass and spin values recovered from the inspiral and post-inspiral phases of the waveform independently of each other (“IMR” tests), and checking the post-newtonian coefficients derived from the waveform against those predicted by GR [4]. Another group elaborated on mechanisms of IMR testing [5], outlining how tight constraints on deviations from GR can be generated with a large number of mid-to-low signal-to-noise ratio (SNR) events, which they point out as being the expected output from LIGO in future observing runs (as opposed to a few high-SNR events). However, as the authors of [4] note, such tests on GR only look for deviations from Einstein’s postulated theory, and do not aim to “constrain parameters that might arise from specific alternative theories.”

This is not to say that no research has been done on specific alternate theories of gravity. Sources such as [6] and [7] survey a number of hypotheses of beyond-GR theories or parameterizations that could be tested with the advent of gravitational wave astronomy, but do not actually provide quantitative details as to how the waveforms produced by CBCs can be analyzed to test these theories. Other researchers have more definitively explored specific alternative theories of gravity. One group explores the theories of gravity that allow for alternate polarizations of gravitational waves, both noting that detecting these extra polarizations would directly disprove Einstein’s GR and exploring a mechanism for detecting such alternative polarizations [8–10].

In this project, we directly and quantitatively test a specific alternative parameterization of deviations from GR in the strong-gravitational-field regime. In section II, we introduce the waveform modulation for the particular parameterization used in this analysis. In section III, we set up the Bayesian parameter estimation, motivating the use of a likelihood function based heavily on the matched-filter analysis [11]. In section IV, we describe the results of testing the modulated theory of GR on two sets of simulated CBC data classified by their SNRs. In section V, we present data on how number of BBH events affects the accuracy and precision of recovery of the modulated theory. In section VI, we outline conclusions of our analyses and discuss the next steps to be taken.

II. WAVEFORM MODULATIONS

Any BBH signal, whether generated using numerical relativity or observed in the strain data at a LIGO detector, can be written in the frequency domain as

$$\tilde{h}_{GR}(f) = A(f)e^{i\phi(f)}, \quad (1)$$

where $A(f)$ represents the amplitude of the waveform and $e^{i\phi(f)}$ represents the complex phase. Since GR is well-tested in the weak-field, quasi-static regime, deviations from GR, if they exist, will manifest themselves in the strong-field limit where $\frac{v}{c}$ becomes an appreciable fraction of 1. Further, $\frac{v}{c}$ is the expansion parameter in the post-Newtonian approximation to GR, making that parameter acceptable to use in testing general relativity [12]. So, we will model deviations from GR as having a strong dependence on this quantity in the form of the multiplicative factor: $e^{\lambda \frac{v(f)^2}{c^2}}$.

Note that $\frac{v(f)^2}{c^2} \Rightarrow \frac{4\pi^2 r(f)^2}{P^2 c^2}$ (applying quasi-circular motion) $\Rightarrow \frac{GM_{tot}}{r(f)c^2}$ (using Kepler's 3rd law, where $\frac{1}{P} = f$, the orbital frequency), implying that $e^{\lambda \frac{v(f)^2}{c^2}}$ is equivalent to $e^{\lambda \frac{GM_{tot}}{r(f)c^2}}$. While this relation was derived classically, Newtonian physics provides a reasonable approximation to a CBC system up until the merger; additionally, $\frac{GM_{tot}}{r(f)c^2}$

is a quantity measuring compactness of strong-field relativistic objects [12, 13], so it is valid to use it in a GR-based model.

In this model, λ is a complex number, where $\Re(\lambda)$ corresponds to an amplitude modulation, and $\Im(\lambda)$ corresponds to a phase modulation.

We can see this by applying the modulation to the GR waveform in Eq. (1):

$$\tilde{h}_{non-GR}(f, \lambda) = e^{\lambda \frac{GM_{tot}}{r(f)c^2}} A(f)e^{i\phi(f)}, \quad (2)$$

which is equivalent to

$$\tilde{h}_{non-GR}(f, \lambda) = e^{\Re(\lambda) \frac{GM_{tot}}{r(f)c^2}} A(f)e^{i(\Im(\lambda) + \phi(f))}. \quad (3)$$

In our analyses, we will use LIGO's IMRPhenomPv2 waveform model (as outlined in [14]) to generate CBC waveforms in the frequency domain. Figure 1 shows some sample waveforms along with their modulated versions for two illustrative values of λ .




FIG. 1. Original and modulated waveforms for a BBH system with $M_1 = 63.6$, $M_2 = 46.7$, and $d = 328$ Mpc. Top: an amplitude modulation with $\lambda = .5$; Bottom: a phase modulation with $\lambda = 4i$. Left: waveform in the time domain; Middle: phase evolution $\phi(t)$; Right: waveform in the frequency domain, overlaid on the strain-equivalent noise amplitude spectral density of the Advanced LIGO design [15].

III. MATCHED FILTER AND BAYESIAN ANALYSIS

We will now motivate the specific posterior and prior functions used in our Bayesian analysis framework. In

deriving our posterior, we began with the equation for a matched filter, which can be expressed as

$$z(t) = 4 \int_0^\infty \frac{\tilde{s}^*(f)\tilde{h}(f)e^{2\pi ift}}{S_n(f)} df \quad (4)$$

[11, 16], where $z(t)$ is a complex function of time, $S_n(f)$ is the power spectral density for LIGO detector noise (in this project, taken from the LIGO Livingston detector a few seconds away from the GW150914 signal [17]), $\tilde{s}(f)$ is the data (signal + noise) in the frequency domain, and $\tilde{h}(f)$ is the template in the frequency domain [11, 16]. Further, the template normalization can be expressed as

$$\sigma^2 = 4 \int_0^\infty \frac{\tilde{h}^*(f)\tilde{h}(f)}{S_n(f)} df \quad (5)$$

[11, 16]. In practice, both integrals have a lower bound of 20Hz, below which LIGO noise becomes several orders of magnitude larger than any potential signal that could arrive at the detector, and an upper bound of 2048Hz, the Nyquist frequency.

Then the SNR, ρ , is written as

$$\rho(t) = \frac{|z(t)|}{\sigma^2} \quad (6)$$

[11]. For our analyses, we wanted to test how well a certain value of λ could be recovered with Bayesian parameter estimation. Therefore, our likelihood function is

$$\prod_t^n \max_t(\rho_n(t)). \quad (7)$$

Here, $\rho_n(t)$ is being maximized over time, and is evaluated with

$$\tilde{s}(f) = \tilde{h}_{non-GR}(f, \lambda_f)$$

(where λ_f is the parameter whose value we want to recover), and

$$\tilde{h}(f) = \tilde{h}_{non-GR}(f, \lambda_v)$$

(where λ_v is a variable parameter and is sampled using MCMC in an attempt to find the value of λ_v that best maximizes the SNR). Note that n is just an index for the number of events used in the analyses, and that both λ_f and λ_v are the same for all n .

We chose to use the Jeffrey's uninformative prior, which takes the form $p(\lambda) = \frac{1}{|\lambda|}$ [18]. Code to implement Bayesian parameter estimation was implemented using the emcee [19] module for python.

IV. SIMULATED MERGERS

The Bayesian analysis was run on 2 sets of 20 (each) simulated mergers: the ‘‘High-SNR’’ and ‘‘Low-SNR’’ sets. These sets were generated with different methods and consequently have different functions. The High-SNR events were used as tests of the Bayesian estimation code, while the Low-SNR events serve as a more realistic means of testing GR.

A. High-SNR events

In generating this set of events, we elected to use artificially small masses and small distances to create artificially high SNRs. Working with low-mass systems, as opposed to working with stellar black hole mass systems at low distances (which would also lead to high-SNR events), allowed us to work with longer duration waveforms that extended over a larger frequency range. This had the benefit of allowing us to test noisy λ recovery across a larger portion of the frequency band of the LIGO Livingston detector.

We first generated 20 random values of M_{tot} , drawing from the Salpeter Initial Mass Function (IMF) with $M_{min} = 1M_{sol}$ [20]. We then generated 20 random values of $\eta = \frac{m_1 m_2}{(m_2 + m_1)^2}$ (the symmetric mass ratio), drawing from a half-normal distribution with $\mu = .25, \sigma = .05$ with the restriction that $0 \leq \eta \leq .25$. This allowed us to generate 20 (m_1, m_2) pairs. Figure 2 shows a graph of this distribution for the High-SNR events. Note that this range of merger masses is expected to be populated by binary neutron star (BNS), not BBH, mergers. For BNS mergers, tidal distortion and disruption of the neutron stars near merger will produce waveform distortions that will make it impossible to test GR in that regime; however, for purposes of this study, we neglect any tidal effects. Note that, as of this writing, LIGO has only detected BBH events with significantly higher masses than used for the study described in this section.

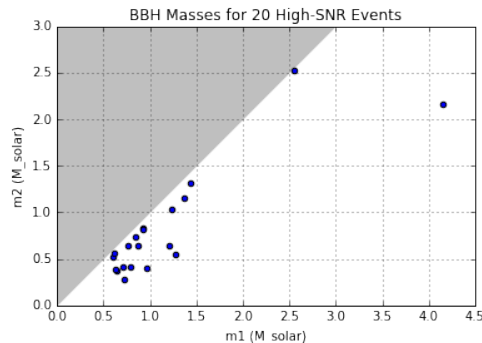


FIG. 2. m_1 vs m_2 for the High-SNR set of 20 events. This graph indicates that the High-SNR events are more likely to be BNS mergers than BBH mergers.

To generate distances, we drew random distances from a distance probability distribution function $p(r)dr \propto r^2 dr$, with the restriction that $0 \leq r \leq 1\text{Mpc}$. Figure 3 shows a graph of this distribution. Note that the dependence on r^2 is true for CBCs that exist in the universe in the limit of constant merger rate per unit volume (which does not actually hold at such small distance scales), but since LIGO’s sensitivity decreases with distance, the CBC events that LIGO can detect does not follow this distribution.

Other BBH parameters, such as right ascension, dec-

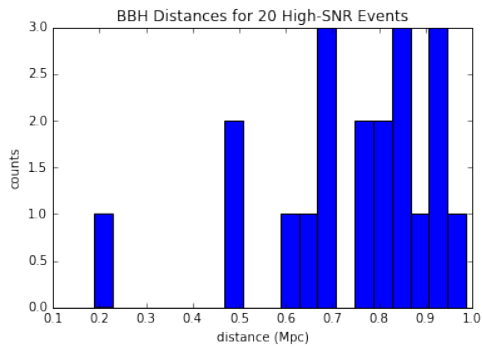


FIG. 3. Distance distribution. Note that the dependence on r^2 is true not for CBC mergers we can detect but for CBC mergers that exist; in reality, LIGO’s sensitivity decreases with distance, so we would see a falloff of detections at larger values of r .




FIG. 4. SNR distribution. Since these events were generated under extremely artificial conditions, the SNRs are 1-2 orders of magnitude larger than what we expect LIGO to detect.

lination, and inclination angle, were sampled randomly from unbiased probability distributions, with the excep-

tions of spins in the x and y directions, which were set to be 0. A graph of the SNRs for these events is shown in Figure 4. These SNRs are higher than anything LIGO is expected to detect by more than an order of magnitude; such high SNRs allow this set of CBC events to be used as a diagnostic for the Bayesian analysis code. If λ is unrecoverable or biased for these High-SNR events, it will be unrecoverable for more realistic events.

B. High-SNR Results

We ran the Bayesian analysis on these High-SNR events, first holding $\Im(\lambda)$ constant while varying $\Re(\lambda)$, then holding $\Re(\lambda)$ constant while varying $\Im(\lambda)$.

Figure 5 shows the residuals for λ , which are equal to $recovered(\lambda) - actual(\lambda)$, along with 2σ errorbars ($1\sigma = 1$ standard deviation). The top left figure shows residuals for $\Re(\lambda)$ with varying $\Re(\lambda)$; the top right shows for $\Im(\lambda)$ with varying $\Im(\lambda)$; the bottom left shows residuals for $\Im(\lambda)$ with varying $\Re(\lambda)$ (for these runs, we expect $\Im(\lambda)$ to be 0, since we’re holding $\Im(\lambda)$ constant at 0); the bottom right shows residuals for $\Re(\lambda)$ with varying $\Im(\lambda)$ (for these runs, we also expect $\Re(\lambda)$ to be 0).

The recovered residuals for $\Re(\lambda)$ are always consistent with 0 to within 1σ both for varying $\Re(\lambda)$ and $\Im(\lambda)$. In fact, the distance between the residuals and 0 is often much smaller than 1σ . The residuals for the recovered values of $\Im(\lambda)$ have slightly larger errors, but are still consistent with 0 to within 2σ . There is evidence of a small negative bias in the residuals for $\Im(\lambda)$ except near 0; the origin of this bias has not yet been identified.

Overall, these results are promising, as they show that the tested values of λ can be precisely and accurately recovered to within 2σ . This implies that the code itself does what we expect it to do, so we can have confidence when we run it on lower-SNR CBC events.

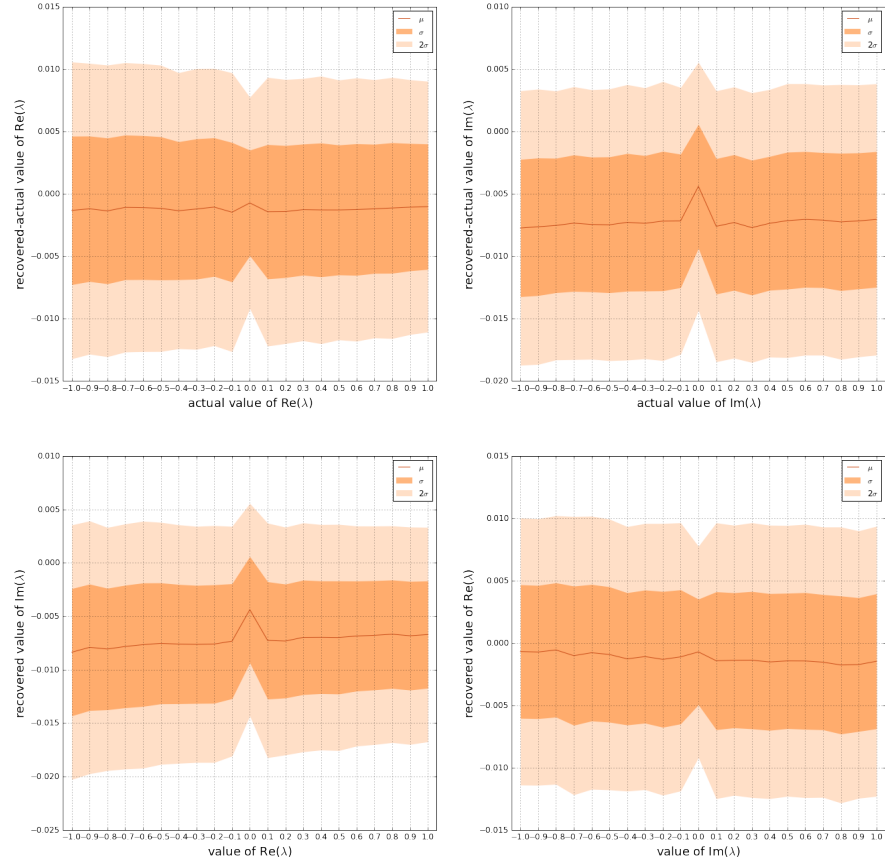


FIG. 5. Residuals for varying values of λ with High-SNR events, plotted with error bars of 1 and 2 standard deviations. Top: residuals of $\Re(\lambda)$ (left) and $\Im(\lambda)$ (right) when their actual values vary from -1 to 1; Bottom: residuals of $\Im(\lambda)$ (left) and $\Re(\lambda)$ (right) when their actual values are 0.

C. Low-SNR events

After confirming that the Bayesian parameter estimation worked for High-SNR events, we generated a set of more realistic BBH events that are similar to the ones LIGO has detected in its first two observing runs. In creating this set, we first generated 20 values of M_{tot} by randomly drawing values from the Salpeter IMF with $M_{min}=10M_{sol}$. We then drew 20 random values of η (with the same distribution as before) to generate 20 random pairs (m_1, m_2) . The distribution is shown in Figure 6.




FIG. 6. m_2 vs m_1 for the Low-SNR events.

For each pair, we calculated the horizon distance h_d (using a base SNR of 15 with random Gaussian variations to incorporate noise effects), then drew a random distance d from the probability distribution function $p(r)dr \propto r^2dr$, with the restriction that $0 \leq r \leq h_d$. The distance distribution is shown in Figure 7. Note the falloff at higher distances, which is indicative of the fact that the actual LIGO detectors lose sensitivity at larger distances.

The uncertainties are markedly larger than for the High-SNR events, although this is to be expected; they are consistent with the scaling of residuals $\propto \frac{1}{SNR}$. The residuals for the recovered values of $\Im(\lambda)$ are always consistent with 0 to within 1σ (an interesting reversal from the High-SNR events results). The residuals for the recovered values of $\Re(\lambda)$ again have slightly larger errors, but are still consistent with 0 to within 2σ .

These results show evidence of a slight bias, as was seen in the High-SNR tests; the bias is positive this time, and the cause is still unknown. However, the results they are promising in that they imply that the tested values of λ can be accurately recovered within 2σ even for mid-to-low range SNR events.




FIG. 7. Distance distribution for Low-SNR events.

The other BBH parameters were chosen as before. The SNR distribution is shown in Figure 8.




FIG. 8. SNR distribution for Low-SNR events

D. Low-SNR Results

Figure 9 shows the residuals for λ along with 2σ errorbars. The graph layout is the same as in Figure 5; the top left figure shows residuals for $\Re(\lambda)$ with varying $\Re(\lambda)$; the top right shows residuals for $\Im(\lambda)$ with varying $\Im(\lambda)$; the bottom left shows residuals for $\Im(\lambda)$ with varying $\Re(\lambda)$ (for these runs, we expect $\Im(\lambda)$ to be 0); the bottom right shows residuals for $\Re(\lambda)$ with varying $\Im(\lambda)$ (for these runs, we also expect $\Re(\lambda)$ to be 0).




FIG. 9. Residuals for varying values of λ with Low-SNR events, plotted with error bars of 1 and 2 standard deviations. Top: residuals of $\Re(\lambda)$ (left) and $\Im(\lambda)$ (right) when their actual values vary from -1 to 1; Bottom: residuals of $\Im(\lambda)$ (left) and $\Re(\lambda)$ (right) when their actual values are 0.

V. COMBINING MULTIPLE EVENTS

Using the same techniques as for the Low-SNR events, we generated 100 mid-to-low range BBH events which have mass, distance, and SNR distributions as shown in Figures 6-8. We then repeated the series of recovering- λ runs as in the previous section, but for a varying number of events. For each given set of runs with n events, we combined the values residuals over all varied values of

λ , then created a probability distribution function. As a clarifying example: the curve in the top left corner of Figure 10 with the label “30 events” is a probability distribution function on the values of $recovered(\lambda) - actual(\lambda)$ for $\Re(\lambda)$, over all the variations of $\Re(\lambda)$ from 0 to 1.

Close-ups of the graphs in Figure 10 are shown in Figure 11, along with a vertical line showing the target value of the residuals.

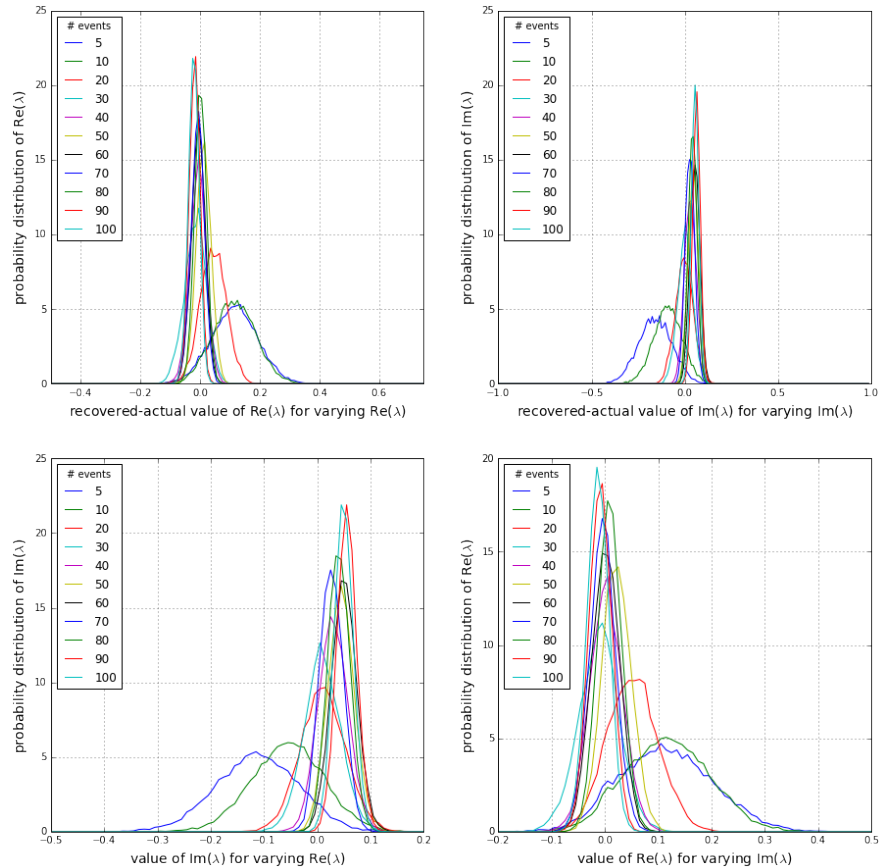


FIG. 10. Probability distribution functions for residuals for different numbers of events analyzed. Top: the pdf on the residuals of $\Re(\lambda)$ (left) and $\Im(\lambda)$ (right) when their actual values vary from 0 to 1. Bottom: the pdf on the residuals of $\Im(\lambda)$ (left) and $\Re(\lambda)$ (right) when their actual values are 0 and $\Re(\lambda)$ (left) and $\Im(\lambda)$ (right) vary, respectively.

Figures 10 and 11 show that the residuals for $\Re(\lambda)$, both for varying $\Re(\lambda)$ and $\Im(\lambda)$, generally improve with increasing number of events, although the probability distribution functions for 90 and 100 events have slightly worse residuals than those for 50-70 events. In contrast, the residuals for $\Im(\lambda)$, both for varying $\Re(\lambda)$ and $\Im(\lambda)$, are consistently too large; they are in fact closest to 0 for 20-30 event runs. This offset may be indicative of a correlation between $\Re(\lambda)$ and $\Im(\lambda)$, but it is not statistically significant even for higher- n runs.

Finally, Figure 12 plots the means of the residuals as a function of the number of events used in the analysis. As we might expect, the standard deviation of the residuals decreases with increasing number of events. The means

of the residuals approach 0 (as we would like) for recovered values of $\Re(\lambda)$ both for varying $\Re(\lambda)$ and $\Im(\lambda)$. However, the recovered residuals for $\Im(\lambda)$ are less satisfactory, since they overestimate 0 by as much as 1σ . The origin of this discrepancy has not yet been identified. However, these residuals are consistent with 0 within 2σ .

The same data is displayed in Tables 1 and 2.

Based on Table II, if we want to recover a purely imaginary λ , we can get to a precision of 0.025 with only 70 events. We can do a little better with a recovery of a purely real λ , which requires at least 50 BBH events to get a recovery with precision of 0.025. Table I tells us that the residuals for these detections will be consistent with 0 to within 1σ for a purely real λ , and consistent with

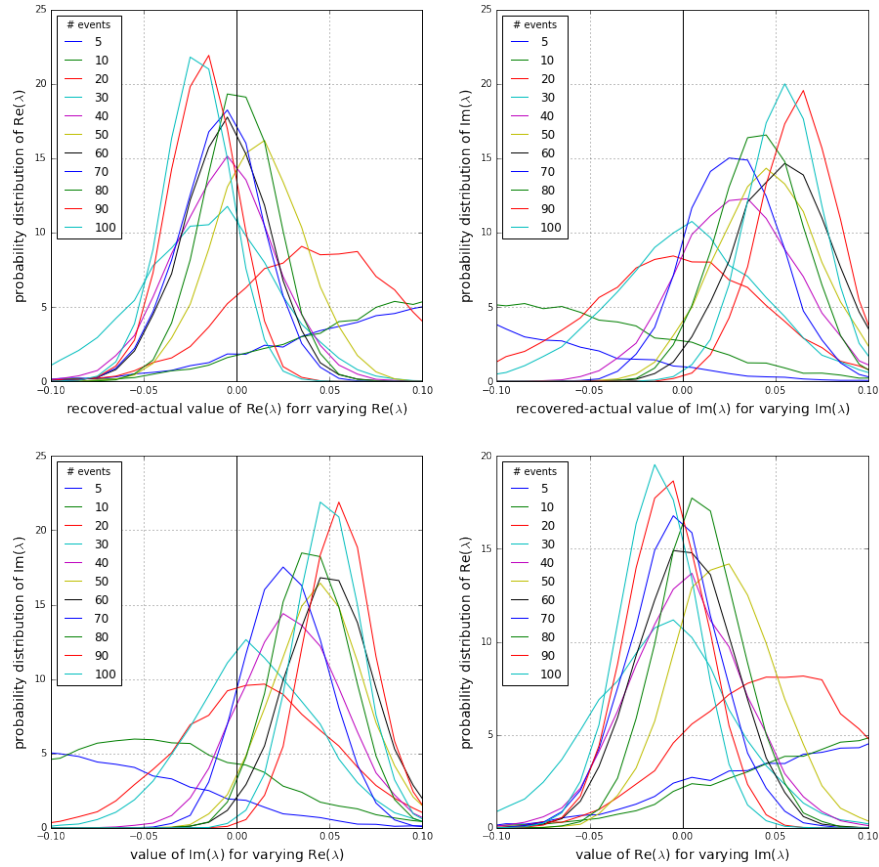


FIG. 11. Probability distribution functions for residuals for different numbers of events analyzed. Top: the pdf on the residuals of $\Re(\lambda)$ (left) and $\Im(\lambda)$ (right) when their actual values vary from 0 to 1. Bottom: the pdf on the residuals of $\Im(\lambda)$ (left) and $\Re(\lambda)$ (right) when their actual values are 0 and $\Re(\lambda)$ (left) and $\Im(\lambda)$ (right) vary, respectively. Equivalent to Figure 10, but zoomed in near 0.

events	$(\Re(\lambda), \Re(\lambda))$	$(\Re(\lambda), \Im(\lambda))$	$(\Im(\lambda), \Re(\lambda))$	$(\Im(\lambda), \Im(\lambda))$
5	0.115	-0.113	0.114	-0.159
10	0.114	-0.0558	0.117	-0.0904
20	0.0438	0.0106	0.0515	-0.00314
30	-0.0164	0.00940	0.0112	0.00248
40	-0.00735	0.0298	-0.00212	0.0300
50	0.0116	0.0445	0.0220	0.0458
60	-0.00643	0.0499	0.00133	0.0538
70	-0.00844	0.0266	0.00480	0.0273
80	0.00118	0.0404	0.00887	0.0416
90	-0.0185	0.0553	-0.00870	0.0619
100	-0.0214	0.0491	-0.0130	0.0545

TABLE I. means of the residuals; in the format (residuals on λ , varying λ)

events	$(\Re(\lambda), \Re(\lambda))$	$(\Re(\lambda), \Im(\lambda))$	$(\Im(\lambda), \Re(\lambda))$	$(\Im(\lambda), \Im(\lambda))$
5	0.0785	0.0771	0.0876	0.0902
10	0.0722	0.0675	0.0787	0.0777
20	0.0436	0.0422	0.0475	0.0480
30	0.0364	0.0341	0.0382	0.0405
40	0.0279	0.0278	0.0309	0.0317
50	0.0248	0.0249	0.0274	0.0280
60	0.0231	0.0235	0.0258	0.0267
70	0.0223	0.0218	0.0244	0.0250
80	0.0201	0.0205	0.0226	0.0232
90	0.0183	0.0184	0.0210	0.0212
100	0.0180	0.0179	0.0204	0.0207

TABLE II. standard deviations of the residuals; in the format (residuals on λ , varying λ)

0 to just over 1σ for a purely imaginary λ , so the recoveries are both accurate and precise. Notably, while the precision decreases with increasing number of events, the recovered residuals become less consistent with 0 (though always consistent to within 2σ), a result which will re-

quire further investigation.

Detecting 50-100 events is by no means an unachievable feat: it's likely that this number of BBH events will have been detected before the end of the decade, with the existing Advanced LIGO detectors operating at design sensitivity.




FIG. 12. Graph of the means of the residuals as a function of the number of events, plotted with 1σ errorbars.

VI. CONCLUSIONS AND FUTURE WORK

In this report, we have shown that for a modulated theory of GR characterized by a complex parameter λ that independently introduces both amplitude and frequency modulations, nonzero values of λ can be accurately and precisely recovered using noisy data (where the noise is

LIGO detector noise) with mid-to-low-range SNRs. We have also given a rough estimate of the number of events required to bound λ to a given degree of precision, notably showing that under 80 events are necessary to determine λ to a precision of 0.025 (a fractional precision of 5% for $\lambda = 0.5$). Attaining this number of events is certainly feasible within the next decade. It is only a matter of years before stringent limits can be put on the value of λ , or, perhaps, Einstein's GR can be decisively ruled out.

Future work involves investigating the cause of the bias in residuals seen in the high- n events. This may be resolved by modifying the λ recovery code such that the Bayesian parameter estimation samples over the other 15 dimensions of parameter space that characterize a BBH merger (i.e. masses, spins, orbital orientation, sky location). While this analysis has shown that λ can be recovered well without a full Bayesian analysis, it is likely that an alternate theory of gravity would lead to slightly different recovered masses, distances, or possibly other parameters which should be explored for completion.

VII. ACKNOWLEDGEMENTS

The authors gratefully acknowledge the LIGO SURF program, Caltech SURF, and the National Science Foundation (NSF) for funding, as well as the many researchers involved in the LIGO SURF program for providing advice and support without which this research would not have been possible.

-
- [1] B. P. Abbott et al., *GW151226: Observation of Gravitational Waves from a 22-Solar-Mass Binary Black Hole Coalescence*. Phys. Rev. Lett. 116, 241103 (2016).
 - [2] B. P. Abbott et al., *Binary Black Hole Mergers in the first Advanced LIGO Observing Run*. Phys. Rev. X 6, 041015 (2016).
 - [3] Vassiliki Kalogera and Albert Lazzarini, *LIGO and the opening of a unique observational window on the universe*. PNAS March 21, 2017 vol. 114 no. 12 3017-3025.
 - [4] B. P. Abbott et al., *Tests of general relativity with GW150914*. Phys. Rev. Lett. 116, 221101 (2016).
 - [5] Abhirup Ghosh et al., *Testing general relativity using gravitational wave signals from the inspiral, merger and ringdown of binary black holes*. arXiv:1704.06784
 - [6] Nicolas Yunes and Xavier Siemens, *Gravitational-Wave Tests of General Relativity with Ground-Based Detectors and Pulsar-Timing Arrays*. arXiv:1304.3473
 - [7] Clifford M. Will, *The Confrontation between General Relativity and Experiment*. Living Reviews in Relativity, Vol.17 No.1 (2014).
 - [8] Maximiliano Isi et al., *Probing dynamical gravity with the polarization of continuous gravitational waves*. Phys. Rev. D 96, 042001 (2017)
 - [9] Maximiliano Isi et al., *Detecting Beyond-Einstein Polarizations of Continuous Gravitational Waves*. Phys. Rev. D 91, 082002 (2015)
 - [10] Thomas Callister et al., *Tests of General Relativity with the Stochastic Gravitational-Wave Background*. LIGO DCC P1700059 arXiv:1704.08373
 - [11] Bruce Allen et al., *FINCHIRP: an algorithm for detection of gravitational waves from inspiraling compact binaries*. Phys Rev D. 85, 122006 (2011).
 - [12] Blanchet, L. Living Rev. Relativ. (2014) 17: 2. doi:10.12942/lrr-2014-2
 - [13] Charalampos Markakis et al., *Neutron star equation of state via gravitational wave observations*. J. Phys.: Conf. Ser. 189 012024 (2009).
 - [14] Rory Smith et al., *Fast and Accurate Inference on Gravitational Waves from Precessing Compact Binaries*. Phys. Rev. D 94, 044031 (2016).
 - [15] <https://dcc.ligo.org/LIGO-P1200087/public>
 - [16] The LIGO Scientific Collaboration, *Tuning matched filtering searches for compact binary systems*. LIGO DCC LIGO-T070109-01-Z
 - [17] <https://losc.ligo.org/events/GW150914/>
 - [18] <https://www2.stat.duke.edu/courses/Fall11/sta114/jeffreys.pdf>

- [19] Daniel Foreman-Mackey et al., *emcee: The MCMC Hammer*. arXiv:1202.3665
- [20] https://en.wikipedia.org/wiki/Initial_mass_function