LASER INTERFEROMETER GRAVITATIONAL WAVE OBSERVATORY
- LIGO -
CALIFORNIA INSTITUTE OF TECHNOLOGY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Technical Note **LIGO-T2400223–v1** 2024/08/02

Second Interim SURF Report

Erin Coleman

Distribution of this document:

AIC, ISC

California Institute of Technology
LIGO Project, MS 18-34
Pasadena, CA 91125
Phone (626) 395-2129
Fax (626) 304-9834
E-mail: info@ligo.caltech.edu

Massachusetts Institute of Technology
LIGO Project, Room NW22-295
Cambridge, MA 02139
Phone (617) 253-4824
Fax (617) 253-7014
E-mail: info@ligo.mit.edu

LIGO Hanford Observatory
Route 10, Mile Marker 2
Richland, WA 99352
Phone (509) 372-8106
Fax (509) 372-8137
E-mail: info@ligo.caltech.edu

LIGO Livingston Observatory
19100 LIGO Lane
Livingston, LA 70754
Phone (225) 686-3100
Fax (225) 686-7189
E-mail: info@ligo.caltech.edu

http://www.ligo.caltech.edu/

1 Introduction

Gravitational waves (GWs) are perturbations in spacetime caused by compact objects moving at relativistic speeds. They are now routinely observed with interferometry by GW observatories around the world [1]. One major source of GWs is binary black hole (BH) mergers. The final stage of a BH merger is known as the ringdown.

The ringdown occurs because a BH merger results in a single perturbed BH. This perturbed BH radiates gravitational energy to settle into a more stable state. The ringdown can be decomposed into quasinormal modes (QNMs) indexed with indices (ℓ, m, n). Each QNM has angular dependence given by spin-weighted spheroidal harmonics and a complex frequency $\omega_{\ell mn} = 2\pi f_{\ell mn} - i/\tau_{\ell mn}$. These modes are called QNMs because, unlike normal modes, they decay over time. The real component of $\omega_{\ell mn}$ indicates the oscillatory frequency of that particular mode, and the imaginary component indicates how quickly it decays [2, 3, 4].

Within general relativity GWs have two possible polarizations, denoted by h_+ and $h_\times$. It is convenient to describe the overall strain in both polarizations as a complex number $h = h_+ - ih_\times$. Setting the ringdown start time $t_0 = 0$, the overall ringdown waveform at time $t > 0$ is given by the equation

$$h(t, \theta, \varphi) = h_+(t, \theta, \varphi) - ih_\times(t, \theta, \varphi) = \sum_{\ell=2}^{\infty} \sum_{m=-\ell}^{\ell} \sum_{n=0}^{\infty} {}_{-2}S_{\ell mn}(\theta, \varphi) C_{\ell mn} e^{-i\omega_{\ell mn} t} \tag{1}$$

where ${}_{-2}S_{\ell mn}$ are the spheroidal harmonic functions and $C_{\ell mn} = A_{\ell mn} e^{i\phi_{\ell mn}}$ is a complex amplitude. In the coordinate system typically used by numerical relativity (NR) simulations, the binary BHs are initially along the x axis, orbiting one another in the x-y plane, and the z-axis points in the direction of the initial orbital angular momentum vector. The angles θ and φ represent the polar and azimuthal angles in this frame, as in the typical formulation of spherical coordinates. This project will focus on the amplitudes $C_{\ell mn}$.

However, in NR, h is written as a sum of modes

$$h = \sum_{\ell=2}^{\infty} \sum_{m=-\ell}^{\ell} h_{\ell m}(t) {}_{-2}Y_{\ell m}(\theta, \varphi) \tag{2}$$

over the spin-weighted spherical harmonics ${}_{-2}Y_{\ell m}$ instead of the spheroidal harmonics ${}_{-2}S_{\ell mn}$. In this basis, every spherical harmonic mode has contributions from every QNM of the same m. This is an effect known as mode mixing.

If a QNM can be recovered from a real GW signal, the resultant final BH mass M_f and dimensionless spin χ_f can be calculated. The No-Hair Theorem, which is a significant theorem in general relativity, states that these properties are the only properties of an astrophysical BH [4]. Thus, measuring the amplitudes of different QNMs in a GW from a binary BH merger is an important test of general relativity. This technique is known as black hole spectroscopy [5]. We can measure the frequencies of all QNMs received, and determine whether they are consistent with a single value of χ_f and M_f. We can study the ringdown of NR waveforms to inform our analysis of real GW data.

Recent studies have focused on fitting the amplitudes of the overtones of the $(2, 2)$ mode, which dominates the waveform. The overtones are tones with $n \geq 1$. Giesler et al. found that overtones up to $n = 7$ could be fit to the waveform, stretching back to the time of peak strain [3]. However, the same study found that the overtone amplitudes did not decay as expected when fitting the ringdown at different times. A more careful approach is required. A more recent study by Clarke et al. fit QNM overtone amplitudes using Bayesian analysis to classify tone amplitudes as stably recovered, unstably recovered, or unresolved depending on the consistency of the amplitude and phase of the recovered tone with subsequent fits [2].

2 Progress

In the first part of this project, I studied least-squares fits of the ringdown of binary black hole (BBH) mergers. In particular, we investigated the effects of fitting different numbers of overtones of the fundamental $(2, 2, 0)$ mode. We have run fits on the dominant $\ell = m = 2$ spherical harmonic mode of a set of CCE simulated waveforms for BBH mergers with spins aligned along the z axis. We have worked with the aligned-spin CCE waveforms available in the SXS catalog [6, 7]. These were mapped to the remnant ("superrest") frame via the `scri` [8] Python package.

2.1 Mismatch studies

We are studying the "usefulness" of adding additional overtones to a fit. One measure of usefulness is the additional signal time which can be used in a fit. Fitting additional overtones allows the signal to be a good fit earlier. This is desirable because the amplitudes of QNMs decay exponentially, so earlier times have a higher signal to noise ratio.

Let $t_0^{N,\mathcal{M}}$ be the earliest time at which a fit with N overtones in addition to the fundamental mode is a good fit, as quantified by the mismatch. To determine $t_0^{N,\mathcal{M}}$ for a particular signal, we generated a mismatch curve for the signal. The mismatch $\mathcal{M}$ between two signals h_1 and h_2 is defined to be

$$\mathcal{M} = 1 - \frac{\langle h_1 | h_2 \rangle}{\sqrt{\langle h_1 | h_1 \rangle \langle h_2 | h_2 \rangle}}$$

where the inner product $\langle h_1 | h_2 \rangle$ is defined as

$$\langle h_1 | h_2 \rangle = \mathrm{Re}\left[\int_{t_0}^{T} h_1(t) h_2^*(t) \mathrm{d}t \right] \tag{3}$$

and t_0 and T are the start and end times of the two signals. In the mismatch curve, we plotted $\mathcal{M}$ vs t_0 with $T = t_0 + 100M$ for $N = 0$ through $N = 21$. Beyond $N = 21$, it becomes difficult to calculate QNM frequencies, but as will be detailed later, 21 overtones is sufficient. We used the knee-finding Python package `kneed` [9, 10] to find the knee of the mismatch curve for each fit. This is the value of t_0 at which the mismatch curve begins to increase rapidly, which means that it is the earliest time at which a fit with N overtones is a good fit. Thus, each knee is at $t_0^{N,\mathcal{M}}$. Figure 1 shows the mismatch curve for CCE waveform `SXS_BBH_ExtCCE_0001` (CCE:01). Note that although $t_0^{N,\mathcal{M}}$ is the earliest time at which a fit with N overtones is a good fit, that fit is not necessarily physical. For early enough values of t, the GW signal is from the BH merger, not the ringdown, so ringdown models should not be applied.

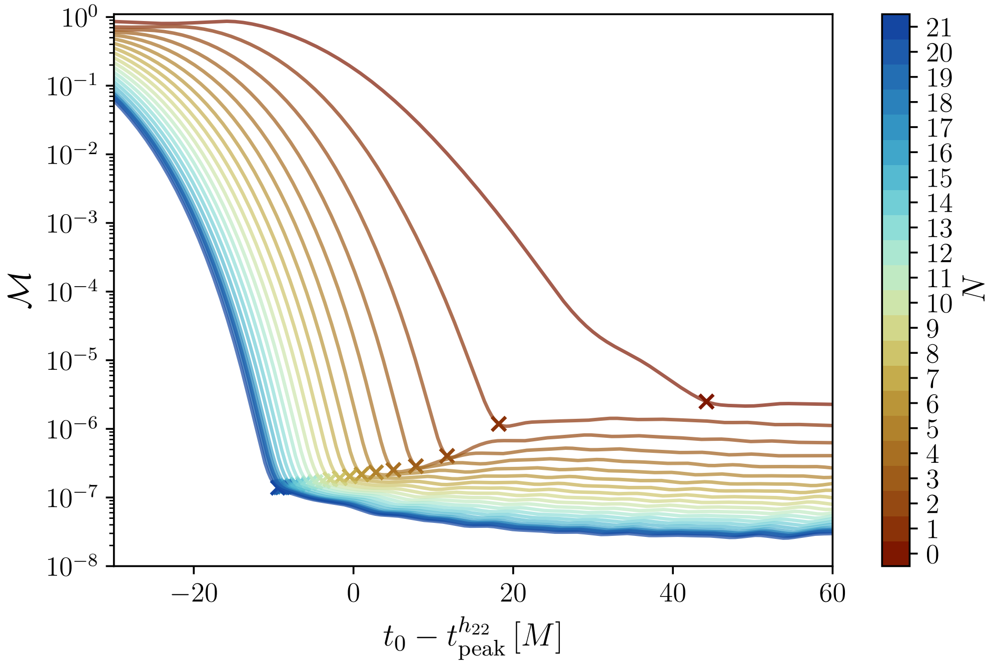

Figure 1: The mismatch curve for CCE:01 for fits of 0 through 21 overtones of the $(2,2)$ mode. Crosses mark the location of $t_0^{N,\mathcal{M}}$ for each N value.

We have defined one measure of usefulness to be the difference $\Delta t_0^{N,\mathcal{M}} = t_0^{N,\mathcal{M}} - t_0^{N-1,\mathcal{M}}$. Note that $\Delta t_0^{N,\mathcal{M}}$ is not defined for $N < 1$. This is a measure of how much additional signal can be fitted by adding the Nth overtone. It should always be positive, because more overtones will fit a signal better, and allow the signal to be fitted at earlier times. Figure 2 shows $\Delta t_0^{N,\mathcal{M}}$ as a function of N for ten CCE waveforms. For all waveforms studied, $\Delta t_0^{N,\mathcal{M}}$ decreases rapidly toward zero—there is a limit to how early overtones can fit CCE waveforms. However, at $N = 18$, every waveform studied had $\Delta t_0^{N,\mathcal{M}} < 0$. This was unexpected, because every additional overtone should extend the fit to earlier times. This suggests that at $N = 18$ overtones, we begin to run up against the limit of the NR simulation precision, and that the maximum usable number of overtones is 17.

We are interested in the location of the "knees" in Figure 1. We do not yet understand why the mismatch levels out at a value of approximately $\mathcal{M} = 10^{-6}$ for times after $t_0^{N,\mathcal{M}}$. To try to investigate this phenomenon further, we created an injection waveform.

Because the ringdown is made up of spheroidal harmonics, each spherical harmonic mode has contributions from all other QNMs of the same m. The QNM with the largest contribution to the $(2,2,0)$ mode is the $(3,2,0)$ mode. We fitted CCE:01 with the QNMs $(2,2,0)$ through $(2,2,17)$ and $(3,2,0)$ at $t_0 = t_0^{17,\mathcal{M}}$. Then, we created a pure waveform using the fitted amplitudes of each mode, with the sum

$$h_{\mathrm{I}}(t) = \left(\sum_{n=0}^{17} A_{22n} e^{-i\omega_{22n}(t-t_0)}\right) + A_{320} e^{-i\omega_{320}(t-t_0)}.$$

We then fitted this waveform to a model containing up to $N = 20$ overtones of the $(2,2,0)$ QNM, without including the $(3,2,0)$ QNM in the model.

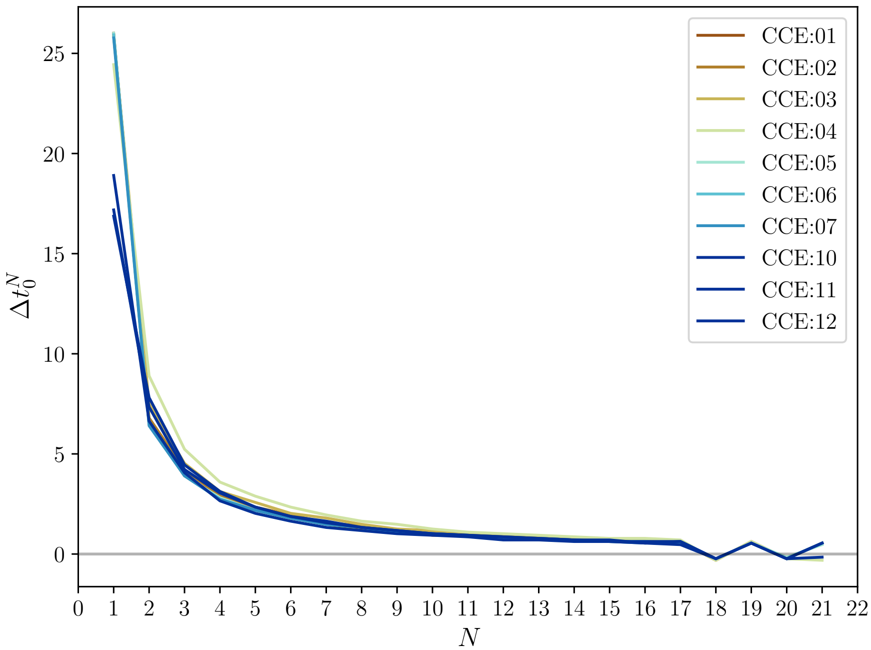

Figure 2: $\Delta t_0^{N,\mathcal{M}}$ vs N for ten different CCE waveforms.

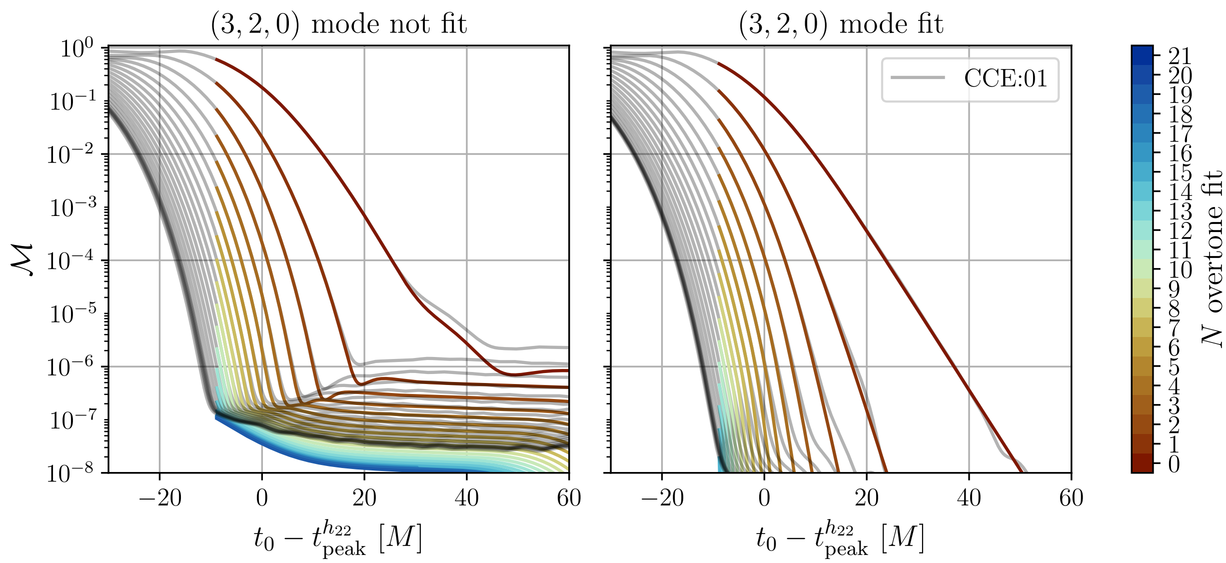

Figure 3: Plots of the mismatch between the injection waveform for fits of 0 through 21 overtones of the $(2, 2, 0)$ QNM. The gray lines show the mismatch curves for CCE:01 and the colored lines show mismatch curves for an injection waveform. The left panel shows fits where the $(3, 2, 0)$ QNM was not included, and the right panel shows fits where the $(3, 2, 0)$ mode was included.

As shown in Figure 3, we were able to reproduce the knee behavior seen in Figure 1. When we fitted the injection to a model which *did* contain the $(3, 2, 0)$ mode, the mismatch between the fit and the signal did not have a knee where it ceased to decrease. It decreased until it reached the numerical precision of the program. This is a an important step in understanding the behavior of overtones: we know that the $(3, 2, 0)$ mode is influential in the mismatch leveling off after a knee. Although we have not yet discovered the exact way in which the $(3, 2, 0)$ QNM determines the location of knees, we know that mode is very influential.

Previous research on fitting overtones of the ringdown fit a maximum of seven overtones [3]. However, there is no particular reason to stop at seven overtones. We have shown that, for these waveforms, the maximum number of overtones that can be fit is 17. The mismatch is a mathematical metric for evaluating how well a model fits the simulated data, but does not describe whether a particular model is in physical agreement with the simulation. The following subsection will introduce a more physical metric for goodness of fit and compare the times $t_0^{N,\mathcal{M}}$ generated by both methods.

2.2 Remnant mass and spin studies

Giesler et al. introduced a metric ϵ which measures the physical agreement between a fit and a NR simulation [3]. The No-Hair Theorem states that the mass M_f and dimensionless spin χ_f of the final BH completely determine the QNM frequencies. Thus, we can vary the QNM frequencies of a model by varying the model's mass and spin. For a model with N overtones and start time t_0, we minimize the mismatch $\mathcal{M}$ between the model and the signal by varying M_f and χ_f. In this work, we use `minimize` from `scipy.optimize` to perform this calculation. Then, δM_f and $\delta\chi_f$ are the differences between the simulation final mass and spin and the final mass and spin which minimize $\mathcal{M}$.

Then,

$$\epsilon = \sqrt{(\delta M_f/M)^2 + (\delta\chi_f)^2}.$$

We plot ϵ vs t_0 for CCE:01 in Figure 4.

Just as models with more overtones have a low mismatch at earlier times, models with more overtones have a low ϵ at earlier times. Let $t_0^{N,\epsilon}$ be the earliest time at which $\epsilon(t_0)$ reaches a local minimum. As can be seen in Figure 4, the curve has several local minima, but we choose the earliest one. This metric differs from $t_0^{N,\mathcal{M}}$ because it determines the earliest time at which a model physically matches a NR simulation.

Now, let $\Delta t_0^{N,\epsilon} = t_0^{N,\epsilon} - t_0^{N-1,\epsilon}$. This is another measure of "usefulness" for each additional overtone. The greater the value of $\Delta t_0^{N,\epsilon}$ is, the more value (e.g. earlier fitting time and higher SNR) is gained by adding the Nth overtone. We calculate $\Delta t_0^{N,\epsilon} = t_0^{N,\epsilon} - t_0^{N-1,\epsilon}$ in the same way as $\Delta t_0^{N,\mathcal{M}}$. Indeed, $\Delta t_0^{N,\mathcal{M}}$ and $\Delta t_0^{N,\epsilon}$. are highly correlated, especially for higher numbers of overtones (See Figure 5).

The plot shows a general downward trend in $\Delta t_0^{N,\epsilon}$, similar to Figure 2, although some CCE waveforms significantly deviate from this trend.

So far, we have defined two definitions of usefulness: $\Delta t_0^{N,\mathcal{M}}$ and $\Delta t_0^{N,\epsilon}$. Both measures show diminishing returns for each additional overtone fitted, although the decline in $\Delta t_0^{N,\epsilon}$ is less

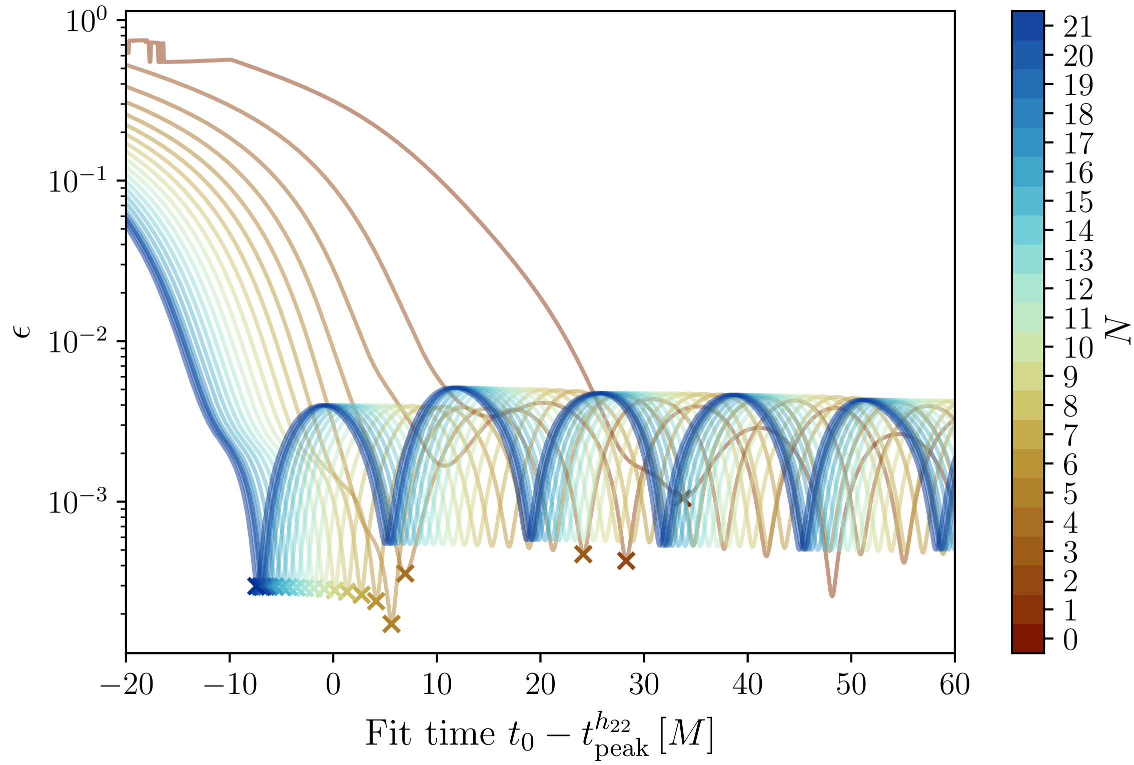

Figure 4: A plot showing the value of ϵ as a function of t_0 for CCE:01. Crosses mark the location of $t_0^{N,\epsilon}$.

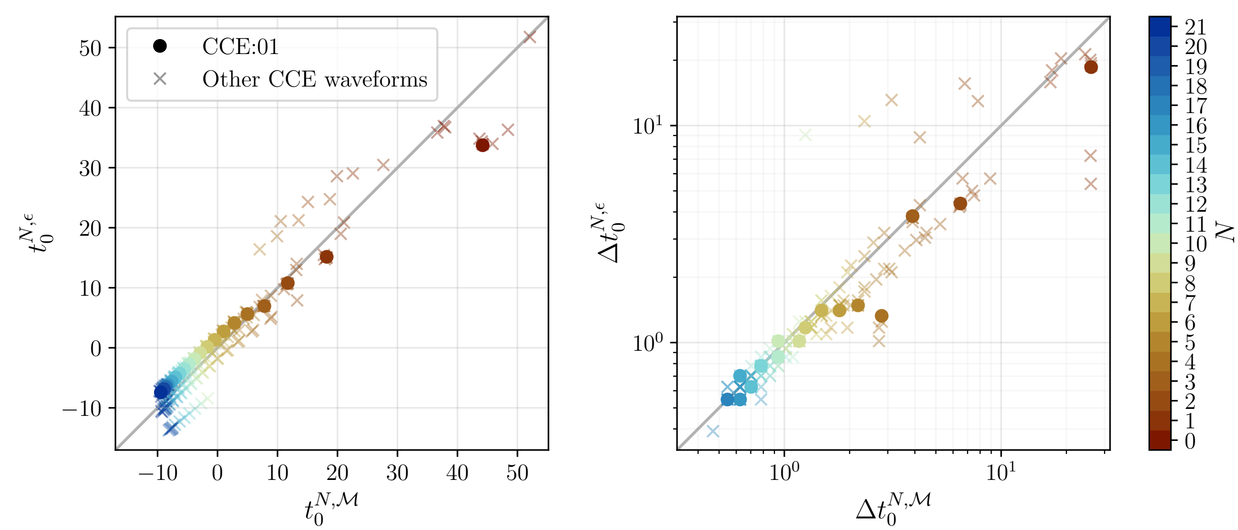

Figure 5: A plot showing $t_0^{N,\mathcal{M}}$ vs $t_0^{N,\epsilon}$ and $\Delta t_0^{N,\mathcal{M}}$ vs $\Delta t_0^{N,\epsilon}$ for the ten CCE waveforms used in this analysis for fits with $N \in [0, 7]$ overtones. Each point represents one fit, with color indicating the number of overtones used. The two times $t_0^{N,\mathcal{M}}$ and $t_0^{N,\epsilon}$ are correlated, especially for higher values of N.

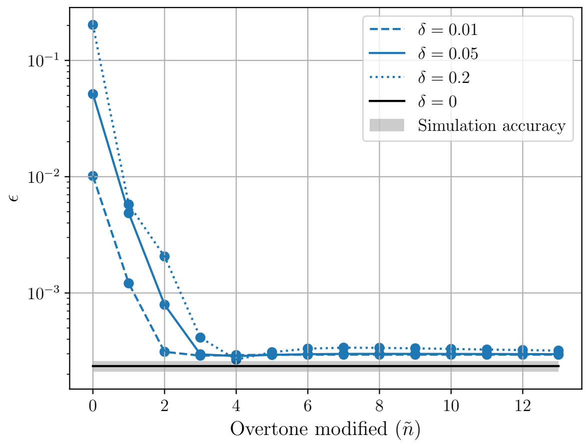

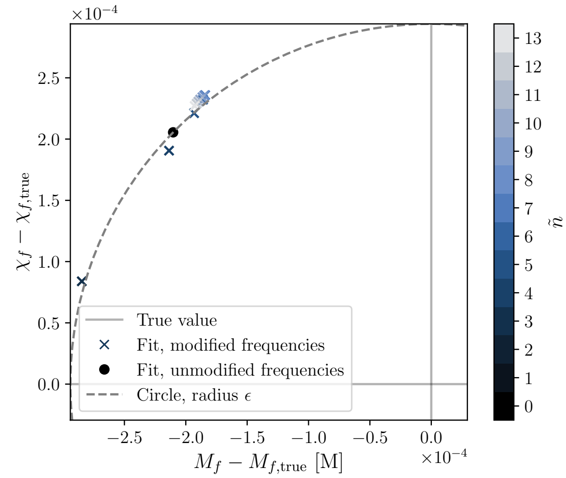

Figure 6: *Left:* A plot showing the deviation ϵ from the true M_f and χ_f for different size frequency modifications δ for CCE:01 in fits with $N = 13$. The behavior of the other CCE waveforms used in this study is similar.*Right:* A plot showing the difference between best-fit M_f and χ_f values and the true M_f and χ_f for fits with 13 overtones, where each cross represents the best fit mass-spin pair for a different overtone was modified. A circle showing all points with ϵ equivalent to the unmodified fit with 13 overtones is provided for reference. Crosses inside the circle indicate that a particular modified fit provided more accurate mass and spin than the unmodified fit.

uniform.

Another method of defining the usefulness of an overtone is by how much it influences the fit overall. A major goal of gravitational wave science is to test general relativity (GR), so the impact of deviations from GR in each overtone is an important metric. This was calculated by introducing a variation in frequency δ. Within GR, the frequency $\omega_{\ell,m,n}$ of the (ℓ, m, n) QNM is completely determined by the mass and dimensionless spin of the final black hole. To introduce deviations form GR in one mode,the QNM frequency for a particular mode $(2, 2, \tilde{n})$ were modified by a factor of $1 + \delta$, such that $\tilde{\omega}_{\ell,m,\tilde{n}} = \omega^{\mathrm{GR}}_{\ell,m,\tilde{n}}$. Then a least-squares fit was performed on the phases and amplitudes of the set of modes with frequencies $\{\omega_{2,2,0}, \omega_{2,2,1} \ldots, \tilde{\omega}_{2,2,\tilde{n}}, \ldots, \omega_{2,2,N}\}$. ϵ was calculated for this fit, and plotted vs $\tilde{n}$ in Fig. 6 We found that each overtone has a smaller effect on the best-fit mass-spin of the system.

This result also provides insight into a result from Giesler et al. (2019). That study found that when all overtones are modified, as the number of overtones increases, the value of ϵ decreases, suggesting that each overtone added does indeed improve the match between signal and model [3]. We repeated this analysis, but found that the decrease in ϵ for higher values of N is primarily due to adding more degrees of freedom to the fit. The main contribution to an increased ϵ is from the modified $(2, 2, 1)$ mode. This effect can be seen in Fig. 7. The NR waveforms used in this analysis were created using numerical solutions to the equations of general relativity, so the $\delta = 0$ case should yield the best results. However, we find that unmodified modes perform worse than modified modes for low values of δ for $N > 4$. Modifying all overtones. Thus, caution must be used when applying this type of analysis.

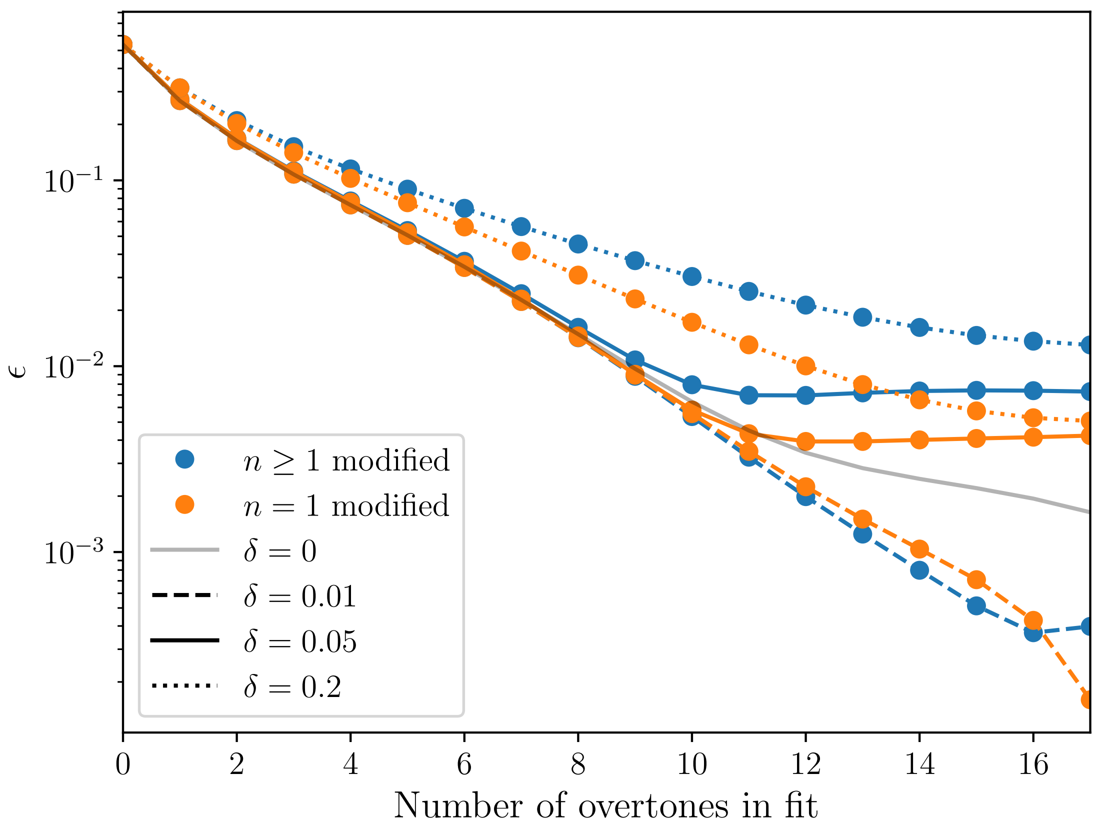

Figure 7: ϵ for fits to CCE:01 with varying numbers of overtones and modifications to the GR frequency of one or more overtones. In blue are fits which modify the frequency of only the $(2, 2, 1)$ mode, and in orange are fits which modify the frequencies of all overtones of the $(2, 2, 0)$ mode.

2.3 Amplitude studies

We are also interested in the stability of the amplitudes of overtones for different fitting times. When a fit is performed at time t_0, it returns the amplitude of each mode at t_0. To compare amplitudes, we adjust the amplitude to be the amplitude at some time $t_{\rm ref}$. This is possible because each QNM decays at a known rate. Throughout this work, we set $t_{\rm ref} = 0 = t_{\rm peak}^{h_{22}}$.

Then, we can rescale $A_{\ell mn}(t_0)$ to the reference time $t_{\rm ref}$ using the relationship

$$A_{\ell mn}(t_{\rm ref}) = A_{\ell mn}(t_0)\exp\left(\frac{t_0 - t_{\rm ref}}{\tau_{\ell mn}(M_f,\chi_f)}\right)$$

where $\tau_{\ell mn}(M_f,\chi_f)$ is the decay time of the mode [11]. $\tau_{\ell mn}$ is related to the complex frequency of the (ℓ, m, n) QNM $\omega_{\ell mn}$ by the equation

$$\tau_{\ell mn} = -\frac{1}{\mathrm{Im}[\omega_{\ell mn}]}.$$

Now that we can compare the amplitudes of QNMs fitted at different times, we measure the stability of the fitted amplitudes. Let $A_{\ell mn}^N$ be the amplitude of the (ℓ, m, n) mode in a fit with N overtones, fitted at $t_0 = t_0^{N,\mathcal{M}}$. We define

$$\Delta A_{\ell mn}^N = \frac{A_{\ell mn}^N - A_{\ell mn}^n}{A_{\ell mn}^n}.$$

The amplitude of each QNM increases rapidly with n, but $\Delta A_{\ell mn}^N$ normalizes the change in amplitude with respect to the amplitude of the nth overtone in a fit with $N = n$ overtones. By definition, $\Delta A_{\ell mn}^n = 0$. We plot $\Delta A_{\ell mn}^N$ vs N for up to seven overtones in Figure 8. It is possible to plot $\Delta A_{\ell mn}^N$ for larger numbers of overtones, but the plot becomes very large.

Figure 8 shows that the amplitudes of higher overtones are more unstable than the amplitude of lower overtones. In general, this plot also suggests that the amplitude of the nth overtone decreases as more overtones are added, although the magnitude of each decrease gets smaller.

However, least-squares fitting can overfit to numerical noise in the data. In particular, least-squares fits of a QNM after it has decayed below the noise floor of a simulation attempt to fit that QNM to the noise, rather than recognizing that it has an amplitude consistent with zero. To avoid this issue, we began using `numpyro` and `jax` to create posterior distributions for QNM amplitudes. Since the $(3, 2, 0)$ mode dominates the noise, as seen in Fig. 3, we want to use the decaying amplitude of the $(3, 2, 0)$ mode as the noise measurement for a Bayesian analysis of the data.

So far, we have worked to get the `NUTS` sampler working properly. We have tested the sampler's limits using zero-noise injections injected waveforms, and found that the sampler struggles for signals with a large number of overtones. We created injections with a varying number of overtones, then sampled the posterior distributions for each mode's amplitude and phase using a fixed signal noise floor. We considered the samples for a parameter to be consistent with the injected value if the 90% credible range included the parameter's injected value. We tested a range of noise floor values σ and N values, and the results for the number

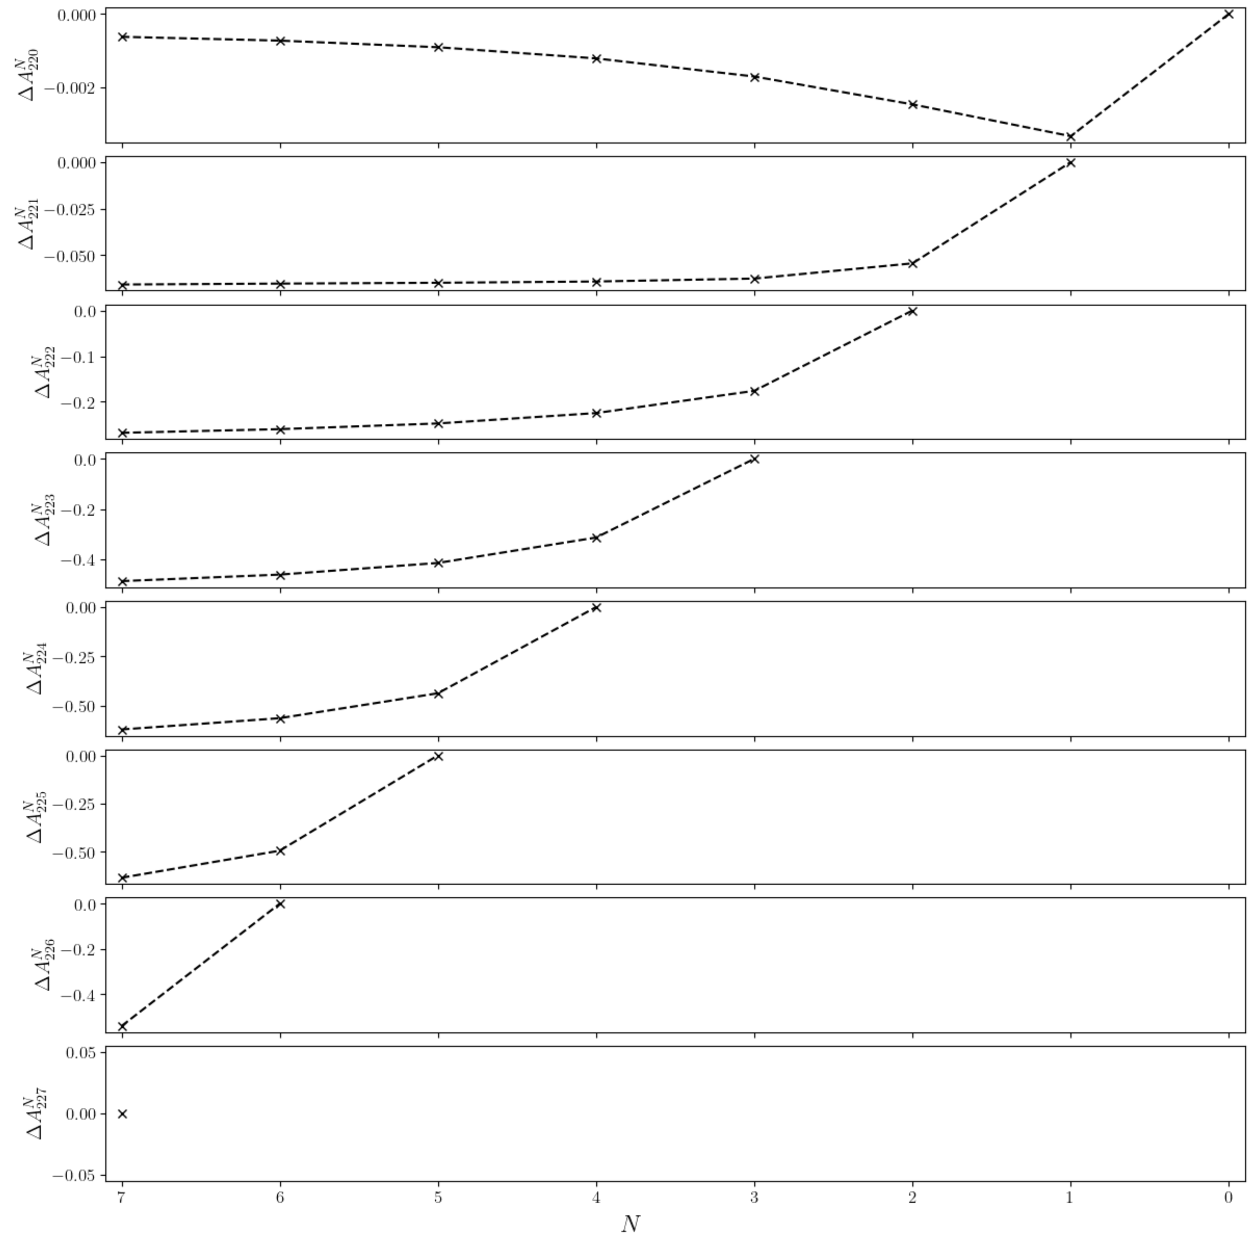

Figure 8: A plot showing the normalized change in amplitude of each QNM for a fit performed at $t_0^{N,\mathcal{M}}$, adjusted to $t_{\rm ref} = 0$, as a function of N.

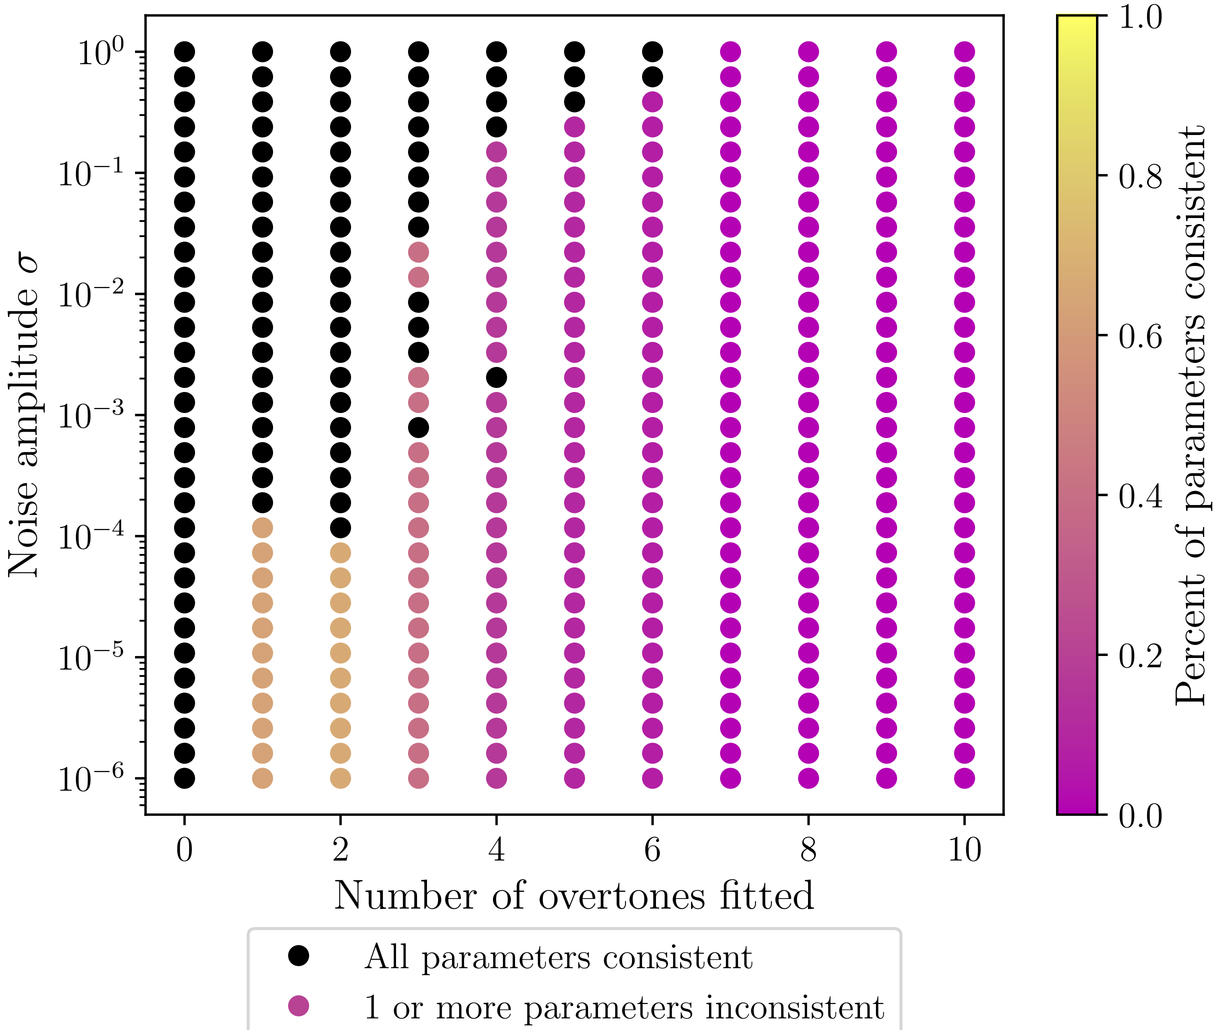

Figure 9: The consistency of different combinations of noise floor values σ and number of QNMs fit.

of consistent parameters are shown in Fig. 9. Each fit with N overtones was performed at $t_0 = t_0^{N,\mathcal{M}}$.

We find that the minimum sigma value required to consistently recover all parameters increases exponentially with the number of overtones. We were not able to consistently recover all parameters for more than six overtones at any σ tested. As n increases, the amplitude of the $(2,2,n)$ mode increases, but each QNM is out of phase with the previous one, keeping the sum amplitude relatively small. This interference effect means that the amplitudes and phases of higher overtones are highly correlated. Samplers work best when variables are not correlated, so using other metrics (such as the difference between two highly correlated amplitudes) in the fit is likely to improve results. We have also found that the posteriors for some parameters are strongly influenced by the random number generator seed. This needs further investigation and more samples to ensure that the sampler is behaving consistently.

Within the current constraints on noise threshold and number of overtones, the sampler has produced excellent results. At late times, when an overtone is expected to have decayed below the noise floor, the posterior shows that the amplitude of that overtone is most likely to be near-zero.

3 Future Work

For the rest of the summer, I hope to gain a better understanding of the Bayesian sampler and its behavior. We plan to try to reduce correlation in fit parameters to increase the number of overtones which can be fit reliably and test the consistency of amplitudes across fits with different numbers of overtones. I would like to be able to fit real data, using the amplitude of the $(3,2,0)$ mode as a function of time as the noise floor, since our previous findings show that the $(3,2,0)$ mode dominates the mismatch in fits to the $(2,2,0)$ mode and its overtones.

4 Acknowledgements

This work was supported by the National Science Foundation Research Experience for Undergraduates (NSF REU) program, the LIGO Laboratory Summer Undergraduate Research Fellowship program (NSF LIGO), and the California Institute of Technology Student-Faculty Programs.

The scientific colour map `roma` [12] is used in this report to prevent visual distortion of the data and exclusion of readers with color-vision deficiencies [13].

References

[1] R. Abbott et al. "GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo during the Second Part of the Third Observing Run". In: *Phys. Rev. X* 13.4 (2023), p. 041039. DOI: 10.1103/PhysRevX.13.041039. arXiv: 2111.03606 [gr-qc].

[2] Teagan A. Clarke et al. “Toward a self-consistent framework for measuring black hole ringdowns”. In: *Phys. Rev. D* 109.12 (2024), p. 124030. DOI: 10.1103/PhysRevD.109.124030. arXiv: 2402.02819 [gr-qc].

[3] Matthew Giesler et al. “Black Hole Ringdown: The Importance of Overtones”. In: *Phys. Rev. X* 9.4 (2019), p. 041060. DOI: 10.1103/PhysRevX.9.041060. arXiv: 1903.08284 [gr-qc].

[4] Maximiliano Isi et al. “Testing the no-hair theorem with GW150914”. In: *Phys. Rev. Lett.* 123.11 (2019), p. 111102. DOI: 10.1103/PhysRevLett.123.111102. arXiv: 1905.00869 [gr-qc].

[5] Gregorio Carullo, Walter Del Pozzo, and John Veitch. “Observational black hole spectroscopy: A time-domain multimode analysis of GW150914”. In: *Physical Review D* 99.12 (June 2019). ISSN: 2470-0029. DOI: 10.1103/physrevd.99.123029. URL: http://dx.doi.org/10.1103/PhysRevD.99.123029.

[6] Michael Boyle et al. “The SXS Collaboration catalog of binary black hole simulations”. In: *Class. Quant. Grav.* 36.19 (2019), p. 195006. DOI: 10.1088/1361-6382/ab34e2. arXiv: 1904.04831 [gr-qc].

[7] SXS Collaboration. *The Ext-CCE Waveform Database*. https://data.black-holes.org/waveforms/extcce_catalog.html.

[8] Michael Boyle et al. *scri*. Version v2022.8.10. Feb. 2024. DOI: 10.5281/zenodo.10709136. URL: https://doi.org/10.5281/zenodo.10709136.

[9] Ville Satopaa et al. “Finding a “Kneedle” in a Haystack: Detecting Knee Points in System Behavior”. In: *2011 31st International Conference on Distributed Computing Systems Workshops*. 2011, pp. 166–171. DOI: 10.1109/ICDCSW.2011.20.

[10] Kevin Arvai. *kneed*. Version v0.8.5. July 2023. DOI: 10.5281/zenodo.8127224. URL: https://doi.org/10.5281/zenodo.8127224.

[11] Eliot Finch. “Black-hole Ringdown: Quasinormal Modes in Numerical-relativity Simulations and Gravitational-wave Observations”. PhD thesis. University of Birmingham, 2023.

[12] Fabio Crameri. *Scientific colour maps*. Version 8.0.1. Oct. 2023. DOI: 10.5281/zenodo.8409685. URL: https://doi.org/10.5281/zenodo.8409685.

[13] Fabio Crameri, Grace E. Shephard, and Philip J. Heron. “The misuse of colour in science communication”. In: *Nature Communications* 11.1 (Oct. 2020), p. 5444. ISSN: 2041-1723. DOI: 10.1038/s41467-020-19160-7. URL: https://doi.org/10.1038/s41467-020-19160-7.