

Compact Binary Waveform Center-of-Mass Corrections

Charles J. Woodford,¹ Michael Boyle,² and Harald P. Pfeiffer^{1,3}

¹Canadian Institute for Theoretical Astrophysics, University of Toronto, 60 Saint George Street, Toronto, Ontario M5S 3H8, Canada

²Cornell Center for Astrophysics and Planetary Science, Cornell University, Ithaca, New York 14853, USA

³Albert-Einstein-Institut, Max-Planck-Institut für Gravitationsphysik, D-14476 Potsdam-Golm, Germany

(Dated: April 11, 2019)

We present a detailed study of the center-of-mass (COM) motion seen in simulations produced by the Simulating eXtreme Spacetimes (SXS) collaboration. We investigate potential physical sources for the large COM motion in binary black hole (BBH) simulations and find that a significant fraction of the COM motion cannot be explained physically, thus concluding that it is largely a gauge effect. These large COM displacements cause mode mixing in the gravitational waveform, most easily recognized as amplitude oscillations caused by the dominant $(2, \pm 2)$ modes mixing into subdominant modes. This mixing *does not* diminish with increasing distance from the source; it is present even in *asymptotic* waveforms, regardless of the method of data extraction. We describe the current COM-correction method used by the SXS collaboration, which is based on counteracting the motion of the COM as measured by the trajectories of the apparent horizons in the simulations, and investigate potential methods to improve that correction to the waveform. We also present a complementary method for computing an optimal COM correction or evaluating any other COM transformation based solely on the asymptotic waveform data.

I. INTRODUCTION

Binary black hole (BBH) systems have been studied for decades, beginning with analytic work and branching out into numerical relativity. With the introduction of gravitational-wave detectors, particularly LIGO, the pursuit of BBH gravitational waveforms has intensified in an attempt to create and fill vast waveform template banks. Gravitational waveforms created through numerical relativity are generally the most accurate waveforms available, and are used for not only gravitational-wave detection and parameter estimation, but also to compare and improve semi-analytic and analytic models of BBHs [1–3].

While numerical-relativity waveforms are the most accurate BBH waveforms, there are concerns regarding their validity, accuracy, and reproducibility. There have been numerous discussions on how to measure the accuracy of a numerical relativity simulation and some sources of error in the simulations have been investigated, including numerical truncation errors, error due to extraction at finite radius or imperfect extrapolation to infinite radius, and errors between simulations of different lengths that otherwise have identical parameters [4–6].

Though not strictly a source of *error* like those named above, there are also consequences due to the gauge freedom of general relativity that may be confused for errors if not properly understood, and will effectively become sources of error if ignored [7–9]. Gauge freedom in general relativity affects all numerical relativity simulations, and thus far no numerical relativity results in the literature have been in a completely specified gauge due to this inherent gauge freedom. While a full accounting of the effects of general gauge freedom is beyond the scope of this work, we will address the translation and boost degrees of freedom. Reference [9] identified these transformations as important for counteracting the observed motion of the center of mass in simulations produced by the Simulating eXtreme Spacetimes (SXS) collaboration. Here, we expand on that analysis, using the recently updated catalog

of 2018 SXS simulations [10, 11], and investigating possible improvements to the correction method.

A translation $\vec{\alpha}$ and a boost $\vec{\beta}$ will transform the waveform h , measured at some point distant from the source, as

$$h(t) \rightarrow h\left(t + (\vec{\alpha} + \vec{\beta}t) \cdot \hat{n}\right) + \mathcal{O}\left(|\vec{\beta}|h\right) + \mathcal{O}\left(|\vec{\alpha} + \vec{\beta}t|^2 \partial_t^2 h\right), \quad (1a)$$

$$\approx h(t) + \partial_t h(t) (\vec{\alpha} + \vec{\beta}t) \cdot \hat{n}, \quad (1b)$$

where \hat{n} is the direction to the observer from the source [9]. Note that this is independent of the distance to the source; even the asymptotic waveform will exhibit this dependence regardless of any extrapolation, Cauchy-characteristic extraction, or similar techniques that may be applied to the data. We can understand this intuitively by thinking about a sphere surrounding the source. If we displace the source away from the center of the sphere, an emitted signal will arrive at the part of the sphere closest to the source before it will arrive at the opposite side of the sphere. The difference in arrival times is independent of the radius of the sphere; it only depends on the size of the displacement. The additional term in Eq. (1b) introduces an angular dependence that is not generally included in waveform models.

Figure 1 demonstrates the effects of COM motion for two systems from the SXS catalog. The most striking example is the upper pair of panels, which show data from a non-precessing system with mass ratio 1.23. On physical grounds, there is nothing to suggest modulations in the mode amplitudes on the orbital timescale; this is a relatively symmetric system with very low eccentricity. The dominant physical behavior on the orbital timescale is simply rotation, which should have no effect on the amplitudes of the modes, along with a secular increase toward merger. Nonetheless, the raw waveform data from the simulation (thin dark lines in the upper left panel) shows very clear amplitude modulations of the sub-dominant modes on the orbital timescale. These modulations—like the COM trajectories seen in the upper right panel—show no signs




FIG. 1. Center-of-mass motion and its effect on waveforms. These plots show data for two systems from the SXS catalog. The upper panels correspond to the non-precessing system SXS:BBH:0314 [12], which has a mass ratio of 1.23, with spins of 0.31 and -0.46 aligned with the orbital angular momentum. The lower panels correspond to the precessing system SXS:BBH:0622 [13], which has a mass ratio of 1.2, with randomly oriented spins of magnitude 0.85. The panels on the right side show the COM trajectories in the simulation coordinates calculated from the apparent horizons of each black hole for a variety of resolutions. In addition to the roughly circular motion, offsets and drifts are also apparent in each case and are the portions of the motion that we remove using inertial transformations. Note that no convergence is evident, reinforcing the idea that the motion is effectively random. The panels on the left show the dominant mode amplitudes of each system, both before and after COM correction—the thin darker lines being the raw waveform data, and the thicker transparent lines being the corrected data. In the non-precessing case (upper panel), we see modulations in the raw data that are not expected on physical grounds; even the relatively small COM motion gives rise to clearly visible effects. No such modulations are visible in the corrected data. In the precessing case (lower panel), modulations are present in both the raw and corrected data, caused by mode mixing due to the precession of the system itself. It is not obvious from this plot alone that the correction makes any improvement to the data. In Sec. IV, we define a quantitative measure of the waveform that very clearly distinguishes the corrected data as a significant improvement.

of convergence with increasing numerical resolution in the simulation, even though the initial data for each resolution is created from identical high-resolution initial data.

As we discuss below, the COM motions found in the SXS catalog are effectively random and apparently independent of any physical parameters of the systems. Therefore, they comprise an essentially random source of unmodeled and unphysical contributions to waveforms from numerical relativity. In particular, they are not systematic; the modulations found in waveforms for one set of physical parameters will be uncorrelated with the modulations in waveforms even for nearly identical physical parameters. Clearly, expecting waveform models such as EOB, phenomenological, and surrogate models to accurately represent these features across a range of physical parameters is tantamount to expecting them to fit large, discontinuous, random signals.

However, by simply compensating for the inertial part of the measured COM motion, the modulations can be almost completely eliminated (thick transparent lines in the upper left panel). It is notable that the COM only drifts by roughly $0.1M$ during almost the entire inspiral for the system shown in the upper panels of Fig. 1, but still has such a drastic effect on the waveform’s modes. Even though this is only a gauge choice—which we have been trained to consider irrelevant in principle—in practice, gauge choices must be made consistently and systematically for the waveforms to be really useful. In this sense, we might suggest that the COM-centered gauge is really an *optimal* choice.

A more difficult comparison is for the precessing system shown in the lower panels of Fig. 1. The precession already mixes the modes drastically, leading to a complicated waveform with pronounced amplitude modulations, even after COM

correction. Clearly the COM correction *changed* the data, but it is not obvious that we can say it was a change *for the better*—at least from looking at this plot alone. To make the comparison more quantitative, we introduce a new measure of a waveform’s “simplicity” in Sec. IV. Essentially, this quantity measures the residual when the waveform is modeled by simple linear-in-time amplitudes in the corotating frame [14]. The value of this residual is 117 times smaller for the COM-corrected data than for the raw data in this precessing system, showing that the corrected waveform is clearly and objectively better in this sense at least.

To address the miscalculation of the COM and its correction, this paper is organized as follows:

- In Sec. II, we discuss the current definition of the COM and the consequences this definition and its use have on SXS gravitational waveforms.
- In Sec. III, we discuss the current method for correcting waveform data and selecting an optimal gauge. Any correlations found between simulation parameters and the COM correction factors are discussed.
- In Sec. IV, we discuss a quantitative method in evaluating the “correctness” of the gauge a waveform is currently in. We compare the waveform data in its original, unoptimized gauge to the COM-corrected gauge described in Sec. III.
- In Sec. V, we discuss how we may improve the definition of the COM to find a better correction, potentially leading to a further optimized choice of gauge. This section also investigates alternative definitions of the COM with a focus on potential physical causes of COM motion like that seen in Fig. 1, including post-Newtonian definitions and considerations of linear-momentum recoil.
- Finally, we present our findings and results in Sec. VI.

II. THE NEED FOR COM CORRECTIONS

The primary concern with BBH simulations is the validity of their gravitational waveforms. Above all, the output from a BBH simulation should result in a reliable, reproducible waveform that can then be released for public usage. In the case of the SXS collaboration, many of the waveforms produced are also compressed into a catalog that is released to LIGO for data analysis and waveform comparisons with their gravitational wave detector data.

Gravitational waveforms can be defined using the gravitational-wave strain h , or the Weyl component Ψ_4 . In regards to detecting gravitational waves, h and Ψ_4 contain the same information, and the analysis and corrections applied in this work may be applied to either with the same results. For simplicity, we will focus on h .

Waveforms from SXS are represented by mode weights, or amplitudes, for spin-weighted spherical harmonics (SW-SHs). The gravitational wave strain may be represented by

the transverse-traceless projection of the metric perturbation caused by the gravitational waves at time t and location (θ, ϕ) relative to the binary, and can be combined into a single complex quantity, given by

$$h(t, \theta, \phi) := h_+(t, \theta, \phi) - ih_\times(t, \theta, \phi). \quad (2)$$

For each slice in time, the combined perturbation $h(t, \theta, \phi)$ is measured on the coordinate sphere. The angular dependence of this measurement can then be expanded in SWSHs. The quantity $h(t, \theta, \phi)$ has a spin weight of -2 [15–17], and may be represented as

$$h(t, \theta, \phi) = \sum_{l,m} h^{l,m}(t) {}_{-2}Y_{l,m}(\theta, \phi), \quad (3)$$

where $h^{l,m}(t)$ are referred to as modes or mode weights, and are much more convenient when analyzing BBH than the total perturbation in any particular direction. Usually the $(l, m) = (2, \pm 2)$ modes, $h_{2,\pm 2}$, dominate the waveform. SWSHs are further discussed in Appendix A.

While the $h_{2,\pm 2}$ modes are dominant, it is important to consider the behavior of the other modes present in the waveform. The other modes may not be used directly for BBH detection, as they are much smaller in magnitude compared to the $h_{2,\pm 2}$ modes, but can be important for parameter estimation. They are also useful for verifying the reliability and potentially the accuracy of the waveform. If the shape, variability, magnitude, or any other characteristic of the higher order, or subdominant, modes are found to not suitably match with theory, then this could indicate a possible flaw in the simulation.

One clear issue is the coordinate system, or gauge choice, for the simulation, as spherical harmonics and hence SWSHs depend on the coordinates. The center chosen for the simulation is the COM of the system, as calculated and set in the initial data. It is expected that the COM will move slightly throughout the simulation; however, large movements are not expected and suggest a flaw in the choice of gauge. If the COM moves significantly, there is mode mixing [9]. The dominant effect [9] is leaking of the $h_{2,\pm 2}$ modes of BBH waveforms into the higher modes, and this leakage can be at least partially removed through COM drift corrections, as described in Sec. III.

Mode mixing is manifested in the waveforms as oscillating amplitudes, which can clearly be seen in the left column panels of Fig. 1, especially for the $(2, 1)$, $(3, 1)$, and $(3, 3)$ modes in the top panel for SXS:BBH:0314. Precessing simulations, like SXS:BBH:0622, are expected to have some amplitude modulations purely due to the orientation of the system. The worse the COM calculation is for a simulation, the more altered the SWSH representations are, and the worse the mode mixing becomes.

It is easily seen in Fig. 1 that the applied COM correction removes what we will find to be unphysical waveform amplitude modulations for non-precessing systems, however for precessing systems it is not so clear or obvious. Therefore, by-eye and qualitative measures for evaluating the “correctness” of the COM correction and gauge, especially for precessing systems, requires a quantitative method. This is discussed in Sec. IV.

The current definition used during BBH simulations for the COM is the usual Newtonian definition:

$$\vec{x}_{\text{COM}} = \frac{m_a}{M} \vec{x}_a + \frac{m_b}{M} \vec{x}_b, \quad (4)$$

where $M = m_a + m_b$ is the total mass of the system, m_a, m_b are the masses of the primary and secondary black holes respectively, and \vec{x}_a, \vec{x}_b are the coordinates of the centers for the primary and secondary black hole respectively. This is a Newtonian expression for the COM, and from output of the simulations like in Fig. 1, we know is not a perfect description of the optimal COM. The tracking of the COM throughout the simulations can be seen for SXS:BBH:0314 and SXS:BBH:0622 in the right column panels of Fig. 1.

The COM motion is an effect of the initial data. One aspect of the initial data construction method proposed in Ref. [18] is the elimination of ADM linear momentum in the initial data for precessing systems, namely enforcing $\vec{P}_{ADM} = 0$. The work done in Ref. [18] proposed a new, and now adopted, method for calculating and constructing the initial data for BBH simulations. The improved method for calculating initial data has far-reaching effects in SpEC simulations and most of the simulations in the SXS simulation catalog were completed using this relatively new method. This had the effect of reducing specific components of mode mixing as seen in the gravitational waveforms, showcased in Fig. 1.

As is further discussed in Sec. VB, linear-momentum recoil is an expected physical contribution to the motion of the COM. However, unphysical contributions to the linear momentum in the initial data of simulations introduce unphysical motion in the COM, essentially imparting spurious linear-momentum kicks. By controlling the linear momentum and removing it, this effect from the initial data is removed. However, even for simulations with initial data constructed using the method described in Ref. [18], significant translations and boosts, and the resulting mode mixing, are still present in the gravitational waveforms. This warrants further investigation into the COM motion and the application of a COM correction.

It had been suggested in Refs. [9] and [15] that much of the COM motion depicted in the right column panels of Fig. 1, and seen in all SXS simulations, was largely unphysical and could be removed from the data. The COM correction used to remedy the unphysical COM motion is discussed in the following section. Additionally, there are alternative definitions of the COM and physical effects that are expected to cause the COM to move, or imply that the COM is not moving at all. The more obvious of these physical effects are post-Newtonian (PN) corrections for the COM which may include effects explaining the COM motion, and linear-momentum recoil from the system. PN and linear momentum contributions are examined in Secs. VA and VB.

III. COM CORRECTION METHOD

Previous work [9, 15] suggests that the COM motion is largely a result of gauge choice. Therefore, understanding the

COM correction begins with understanding the permissible gauge transformations. More specifically, we are interested in the gauge transformations that will affect the waveform measured by distant observers. Because a gravitational-wave detector will typically be very far from the source, only the asymptotic behavior of the waves is generally considered relevant—specifically at future null infinity, \mathcal{I}^+ .

While the asymptotic gauge of waveforms from numerical relativity has not been extensively investigated, it is certainly fair to say that no results in the literature thus far have been in a completely specified gauge. Even the strongest claims of “gauge-invariant” asymptotic waveforms [19] are only invariant modulo the *infinite-dimensional* Bondi-Metzner-Sachs (BMS) gauge group [20, 21], which is the asymptotic symmetry group corresponding to the Bondi gauge condition. An important feature of Bondi gauge is that the gravitational waves measured by *any* distant inertial observer (at least over a duration short compared to the distance to the source) are approximately given by the asymptotic metric perturbation at fixed spatial coordinates as a function of retarded time in some member of this gauge class—and conversely, any such function corresponds to a signal that could be measured by some distant inertial observer [22]. Essentially, we might think of Bondi gauge to be “as simple as possible, but not simpler” for the purposes of gravitational-wave detection. And because the BMS group alters the waveform while preserving Bondi gauge, we consider it to be the fundamental symmetry group relevant to gravitational-wave modeling.¹

Because BMS transformations preserve the inertial property of observers, we cannot expect to counteract all of the COM motion seen in Fig. 1—particularly the cyclical behavior. However, in addition to the cyclical behavior, these coordinate tracks begin with some overall displacement from the origin, and then drift away from that initial location over the entire course of the inspiral. Thus, we expect that a space translation and a boost are needed to negate the effects of some of the COM motion. In particular, we will choose the translation $\vec{\alpha}$ and boost $\vec{\beta}$ to minimize the average distance between the measured COM and the origin.²

¹ Other possible gauge choices exist. For example, Newman-Unti gauge [23] is closely related to Bondi gauge, and is invariant under the same asymptotic symmetry algebra [24]. More generally, it is not even clear that waveforms from numerical relativity are actually expressed in either of these well-defined gauge classes, in which case more general gauge transformations may be of interest. Ultimately, the gauge freedoms relevant to counteracting COM motion are simply space translations and boosts. As long as these transformations are allowed, this discussion of COM motion remains relevant. Previous work [25] suggests that SXS waveforms are consistent with waveforms in Bondi gauge, though further research is warranted.

² This measure will be invariant under time translation and rotation, and they are generally dealt with separately during gravitational-wave analysis, so we simply ignore those degrees of gauge freedom. Furthermore, it is not at all clear how a higher-order supertranslation should affect the coordinates close to the center of a simulation, and so we leave discussion of more general supertranslations to future work.




FIG. 2. Magnitudes of COM offsets and drifts for all simulations in the SXS catalog. The top row shows values for non-precessing systems (i.e., non-spinning, spin-aligned, and spin-antialigned) and the bottom row shows values for precessing systems. The horizontal axis for each plot is the magnitude of the COM value shown ($|\vec{\alpha}|$, $|\vec{\beta}|$, or $|\vec{\delta}| = |\vec{\alpha} + \vec{\beta}t_{\text{merger}}|$, where t_{merger} is the first reported instance of a common apparent horizon found between the two BHs) and the vertical axis is the number of simulations that have COM values of that bin magnitude. Note that typical values of $|\vec{\beta}|$ are quite small, but accumulate over the course of a simulation to cause a large overall displacement by merger. Blue indicates runs using the initial-data method described in Ref. [18]; orange indicates runs using the previous initial-data method. These results suggest that this procedure improves the initial location of the center of mass, but does little to improve its drift.

A. Choosing the translation and boost

We follow Appendix E of Ref. [9] in choosing the translation $\vec{\alpha}$ and boost $\vec{\beta}$ to minimize the average distance between the COM measured in the raw data and the origin of the corrected frame. That is, we choose $\vec{\alpha}$ and $\vec{\beta}$ to minimize the function

$$\Xi(\vec{\alpha}, \vec{\beta}) = \int_{t_i}^{t_f} |\vec{x}_{\text{COM}} - (\vec{\alpha} + \vec{\beta}t)|^2 dt. \quad (5)$$

It is not hard to find the minimum of this quantity analytically. We define two moments of the COM position

$$\langle \vec{x}_{\text{COM}} \rangle = \frac{1}{t_f - t_i} \int_{t_i}^{t_f} \vec{x}_{\text{COM}}(t) dt, \quad (6a)$$

$$\langle t \vec{x}_{\text{COM}} \rangle = \frac{1}{t_f - t_i} \int_{t_i}^{t_f} t \vec{x}_{\text{COM}}(t) dt. \quad (6b)$$

Then, the minimum of Eq. (5) is achieved with

$$\vec{\alpha} = \frac{4(t_f^2 + t_f t_i + t_i^2) \langle \vec{x}_{\text{COM}} \rangle - 6(t_f + t_i) \langle t \vec{x}_{\text{COM}} \rangle}{(t_f - t_i)^2}, \quad (7a)$$

$$\vec{\beta} = \frac{12 \langle t \vec{x}_{\text{COM}} \rangle - 6(t_f + t_i) \langle \vec{x}_{\text{COM}} \rangle}{(t_f - t_i)^2}. \quad (7b)$$

We then apply this transformation to the asymptotic waveform using the method described in Ref. [9]. Note that this rests on implicit assumptions about how directly comparable the coordinates of the apparent-horizon data and the asymptotic coordinates are. For example, this assumes that the time coordinate of the apparent-horizon data and the asymptotic retarded-time coordinate are equal. While there is no rigorous motivation for this assumption, the results of Sec. IV bear out its approximate validity.

Using this minimization method, the COM offsets for every public waveform in the SXS catalog have been corrected in the waveform data. The first instance of the COM corrections to waveforms in the SXS public waveform catalog was in January of 2017. COM corrected waveform data is recommended over non-corrected data in all cases, and corresponding files are listed in the SXS public waveform catalog as files ending in `CoM.h5`. An overview of the COM correction values is shown in Fig. 2.

It is clear from the upper-left panel of Fig. 1 that COM removal is “helpful” in the sense that it reduces the amplitude oscillations, which are not expected on physical grounds.




FIG. 3. Comparing the size of COM corrections in the SXS catalog, $|\vec{\mu}_{00}|$, to how much those corrections change under small variations in the endpoints of integration used to compute the COM correction. The vertical axis shows the largest change in the COM correction if we shift t_i and/or t_f later by half an orbit. The systems with the largest COM corrections—where these corrections are presumably the most important—change by small fractional amounts. On the other hand, there are several systems in which the COM correction changes by more than the original correction; those systems also have some of the smaller COM corrections in the catalog. The median percentage change is 4% of the original correction, and even the largest individual change is smaller than the median value of $|\vec{\mu}_{00}|$.

Unfortunately, this by-eye analysis is not quantitative, and it is not clear how it would apply to a precessing system, as seen in the lower-left panel of the same figure. We discuss a better measure of how the waveform quality is impacted by COM corrections in Sec. IV.

B. Choosing the integration region

The determination of $\vec{\alpha}$ and $\vec{\beta}$ is made over a subset of the total simulation time, from t_i to t_f [see Eq. (5)]. Choosing different values of t_i and t_f may affect the resulting $\vec{\alpha}$ and $\vec{\beta}$ values. For the corrections performed on the SXS catalog, a standard subset of the simulation time was chosen. All waveforms had their COM-correction values calculated from $t_i = t_{\text{relax}}$, the “relaxation” time after which the initial transients have dissipated, to 10% before the end of the inspiral: $t_f = 0.9t_{\text{merger}}$. These time bounds were chosen to avoid including periods of junk radiation as well as the merger and ringdown stages.

However, changing t_i , t_f by small amounts could change the COM correction values. As there is epicyclical motion of the COM (as seen in Fig. 1, for example), changing the beginning or ending time may cause the resulting $\vec{\alpha}$ and $\vec{\beta}$ to change, depending on where t_i and t_f fall on an epicycle. For example, if t_i and t_f are separated by an integer number of epicycles, then we might expect any effect from the epicycles on the calculation of $\vec{\alpha}$ and $\vec{\beta}$ to cancel out. However, if t_i and t_f are separated by

a non-integer number of epicycles, especially by a half-integer number of epicycles, the epicycles may induce significant bias in $\vec{\alpha}$ and $\vec{\beta}$. The overall number of epicycles included in the calculation of $\vec{\alpha}$ and $\vec{\beta}$ may also affect how sensitive they are to this bias.

Here, we compare the size of the COM correction using the standard prescription to the size of the correction when t_i and/or t_f are changed by half an orbit. This will give us some idea of the stability of the COM-correction procedure. However, it must be noted that, at a larger scale, the choices of t_i and t_f are quite arbitrary. For some purposes, it may be preferable to choose those values to range over only the first half of the inspiral, or even just the ringdown stage. The values used in the SXS catalog were chosen for robustness and easy reproducibility.

To simplify the comparison, we describe the COM motion using the quantity $\vec{\mu}$, which gives the most distant position of the corrected origin of coordinates throughout the inspiral, relative to the origin used in the simulation. Specifically, we can define $\vec{\mu}$ according to

$$\vec{\mu} = \begin{cases} \vec{\alpha} & \text{if } |\vec{\alpha}| > |\vec{\alpha} + \vec{\beta}t_{\text{merger}}|, \\ \vec{\alpha} + \vec{\beta}t_{\text{merger}} & \text{if } |\vec{\alpha}| \leq |\vec{\alpha} + \vec{\beta}t_{\text{merger}}|. \end{cases} \quad (8)$$

For 96% of the simulations in the SXS catalog, we find that $\vec{\mu} = \vec{\alpha} + \vec{\beta}t_{\text{merger}}$. The 4% of simulations with $\vec{\mu} = \vec{\alpha}$ have no apparent correlations with system parameters, and are effectively random. We also introduce subscripts, so that $\vec{\mu}_{00}$ is the result of this calculation when using the original values of t_i and t_f ; $\vec{\mu}_{10}$ is the result when moving t_i later by half an orbit; $\vec{\mu}_{01}$ is the result when moving t_f later by half an orbit; and $\vec{\mu}_{11}$ is the result when moving both t_i and t_f later by half an orbit.

In Fig. 3, we examine $\max_{jk} |\vec{\mu}_{00} - \vec{\mu}_{jk}|$ as a measure of how robust the COM corrections are with respect to these small adjustments in the choices of t_i and t_f . In the great majority of systems the COM changes by less than $10^{-2}M$. This is, for example, just one tenth the size of the displacements seen in the upper panels of Fig. 1. The median change is $3.1 \times 10^{-3}M$, and in all cases is smaller than the median value of $|\vec{\mu}_{00}|$ itself, which is $6.9 \times 10^{-2}M$. The systems with the largest COM corrections in the SXS catalog change by fractions of a percent, suggesting that the results are certainly stable in the cases where applying a COM correction is most important. There are several cases where the fractional change is greater than 100%, though these are systems with relatively small values of $|\vec{\mu}_{00}|$. The median fractional change is 4.3%. It is also notable that the data points separate roughly into three groups. The group in the lower left corner of Fig. 3 is comprised exclusively of equal-mass non-spinning simulations with various eccentricities, though several of these are also found in the central group. The central group is where all equal-mass simulations with equal but nonzero spins are found, which includes ten systems with significant precession. Every other type of system is in or near the largest group, on the upper right.




FIG. 4. COM correction values and relevant simulation parameters. χ_{eff} is the effective spin, χ_p is the effective precessing spin, q is the mass ratio, e is the eccentricity, Orbits represents the total number of orbits the simulation had at t_{merger} , and $\vec{\alpha}$, $\vec{\beta}$, $\vec{\delta} = \vec{\alpha} + \vec{\beta}_{t_{\text{merger}}}$ are the COM correction values representing the spatial translation, boost, and total COM displacement respectively. Red represents spin-aligned simulations, and teal represents precessing simulations.

C. Correlations between COM correction values and physical parameters

Along with having COM corrected the waveforms, we have also performed an analysis of the values of the boosts and translations needed by each simulation in the SXS public waveform catalog.

No obvious correlations can be seen in Fig. 2 between spin aligned and precessing systems. We also show a more in-depth correlation plot in Fig. 4, taking more of the simulation parameters into consideration. It can be seen that typically precessing simulations may have larger overall COM displacement, $\vec{\delta} = \vec{\alpha} + \vec{\beta}_{t_{\text{merger}}}$, and that larger boost values $\vec{\beta}$ correspond with larger overall displacement values $\vec{\delta}$ for both spin aligned and precessing systems.

Outside of the correlations between the boost $\vec{\beta}$ and total displacement $\vec{\delta}$ of the COM, there does not appear to be any other strong correlations present for the current SXS simulation catalog. It was expected that precessing, high mass ratio, and eccentric systems should have vastly different COM correction values than spin aligned, low mass ratio, and more circular systems, however no such correlations are present with this data set. Weak correlations may be present between mass ratio q and boost $\vec{\beta}$ and total displacement $\vec{\delta}$ for spin-aligned systems, such that systems with larger mass ratios have larger $\vec{\beta}$ and $\vec{\delta}$, but this is not clear.

For Fig. 4, we use the eccentricity e , number of orbits, and mass ratio q reported by SpEC at the end of the simulation. We calculate the effective spin χ_{eff} and the effective precessing

spin χ_p from the definitions given in Eq. (6) and Eq. (9) in the GW150914 LIGO parameter-estimation paper [26]:

$$\chi_{\text{eff}} = \frac{c}{GM} \left(\frac{\vec{S}_a}{m_a} + \frac{\vec{S}_b}{m_b} \right) \cdot \frac{\vec{L}}{|\vec{L}|} = \frac{\vec{\chi}_a m_a + \vec{\chi}_b m_b}{M} \cdot \frac{\vec{L}}{|\vec{L}|}, \quad (9)$$

$$\chi_p = \frac{c}{B_1 G m_a^2} \max(B_1 |\vec{S}_{a,\perp}|, B_2 |\vec{S}_{b,\perp}|), \quad (10)$$

where in simulation coordinates, $c = G = 1$, $M = m_a + m_b$ is the total mass of the system, $\vec{S}_i = \vec{\chi}_i m_i^2$ is the spin of the black hole being considered and $\vec{\chi}_i$ is the dimensionless spin, $B_1 = 2 + 3m_b/2m_a$, $B_2 = 2 + 3m_a/2m_b$, and the subscript \perp indicates the quantity perpendicular to the angular momentum of the system, e.g., $\vec{S}_{a,\perp} = \vec{S}_a - \frac{\vec{S}_a \cdot \vec{L}}{L^2} \vec{L}$. Note that χ_p gives a measure of how much a system is precessing during a simulation.

IV. QUANTIFYING COM CORRECTION USING WAVEFORMS ALONE

Any discussion of COM based on the positions of the individual black holes will suffer from the same fundamental ambiguity: reliance on coordinates—specifically in the highly dynamical region between the two black holes—that are subject to unknown gauge ambiguities. The only region of the space-time where the gauge freedom is limited in any useful sense is the asymptotic region, in which we assume the only freedom is given by the BMS group (described in Sec. III). While there are many suggestions in the literature [27–36] for using asymptotic information to specify the asymptotic gauge more narrowly, they all require more information than is available from most catalogs of numerical-relativity waveforms—such as additional Newman-Penrose quantities or more precise characterization of the asymptotic behavior of the various fields.

Here, we present a simplistic but effective measure of COM effects that can be applied exclusively to asymptotic waveform data h or Ψ_4 . The basic idea is that we expect to be able to decompose a waveform measured in COM-centered coordinates into modes that are, at least for small portions of the inspiral, given by a slowly changing complex amplitude times a complex phase that varies proportionally with the orbital phase. When the waveform is decomposed in off-center coordinates, those well-behaved modes mix, so that the amplitude and phase do not behave as expected. Therefore, we will attempt to model a given waveform in a sort of piecewise fashion that *assumes* the expected behavior, and simply measure the residual between the model and the waveform itself. For a given transformation applied to the waveform, we will minimize the residual by adjusting the parameters to the model while keeping the waveform fixed. The smaller the minimized residual, the more accurately the waveform with that transformation can be modeled in this simple way, and the more nearly we expect that the waveform is decomposed in COM-centered coordinates. Roughly speaking, we can think of this as a measure of the “simplicity” of the waveform, which is not only in line with our basic expectations for waveforms in the appropriate coordinates,

but also a useful measure of how accurately simple waveform models (EOB, surrogate, etc.) will be able to capture features in the numerical waveforms. This criterion is obviously totally distinct from any criteria involving the BH positions, but is important precisely because it provides a complementary way of looking at the data. Finding agreement between the results of this method and another will lend support to the idea that the other method is suitable.

A. Defining the method

We now describe this method more precisely. The initial inputs are some translation $\vec{\alpha}$ and boost $\vec{\beta}$ that we wish to evaluate. We transform the waveform by those inputs and denote the result $\mathcal{T}_{\vec{\alpha},\vec{\beta}}[h]$. We then transform to a “corotating frame”, which is a time-dependent frame chosen so that the waveform in that frame is varying as slowly as possible [14]. Only the angular velocity of this frame, $\vec{\Omega}$, is determined by the condition that the waveform vary slowly; it is integrated in time to obtain one such frame [37], but the result is only unique up to an overall rotation. We choose that overall rotation so that the \vec{z}' axis of the final corotating frame is aligned as nearly as possible throughout the inspiral portion of the waveform with the dominant eigenvector [38, 39] of the matrix

$$\langle L_a L_b \rangle := \int_{S^2} L_a \{ \mathcal{T}_{\vec{\alpha},\vec{\beta}}[h] \}^* L_b \{ \mathcal{T}_{\vec{\alpha},\vec{\beta}}[h] \} dA, \quad (11)$$

where L_a is the usual angular-momentum operator. This still leaves the frame defined only up to an overall rotation about \vec{z}' , but such a rotation will have no effect on our results. The transformed waveform in this corotating frame will be denoted $\mathcal{R} \{ \mathcal{T}_{\vec{\alpha},\vec{\beta}}[h] \} (t, \theta, \phi)$, though we will usually suppress the parameters, and may decompose the angular dependence in terms of SWSH mode weights as usual. This is the quantity we will be attempting to model.

For the model itself, we first break the inspiral up into smaller spans of time; we will be modeling the waveforms using simple linear-in-time approximations, so we cannot expect to accurately reproduce the nonlinear evolution over a very long portion of the inspiral using just one such model. The relevant measurement of the waveform’s dynamical behavior is the angular velocity of the corotating frame. More specifically, we define $\Omega_{z'} = \vec{\Omega} \cdot \vec{z}'$, and use that to determine a phase³ $\Phi_{z'} = \int \Omega_{z'} dt$. An obvious span of time would be a single cycle of this phase, which would include enough data so that the fit would actually reflect the behavior of the waveform, but not so much that we would expect a poor fit due to evolution on the inspiral timescale. However, we will essentially be fitting

³ This phase is loosely related to the orbital phase of the binary. The angular velocity $\vec{\Omega}$, however, is defined solely with respect to the waveform at \mathcal{S}^+ , and entirely without reference to any quantities at finite distance in the system. Nonetheless, for reasonably well-behaved coordinate systems, we would expect it to agree roughly with the orbital phase deduced from the trajectories of the black holes, especially during the early inspiral regime.

oscillatory terms with linear models. In the simple case of fitting a line to a basic sine function, it is not hard to see that the optimal line has the expected slope of zero—independent of the phase of the sine function—when the fit region is such that the argument of the sine function goes through a phase change of $\varphi \approx 8.9868$ [or other solutions of $\varphi = 2 \tan(\varphi/2)$]. Therefore, we select each span of time so that it extends over a phase Φ_z of approximately φ , thereby determining the difference in time between $t_{i,1}$ and $t_{i,2}$ so that they satisfy

$$\Phi_z(t_{i,2}) - \Phi_z(t_{i,1}) = \varphi. \quad (12)$$

We find that this choice does drastically reduce the oscillations in the optimal fit parameters as we shift the fitting window. While the individual time spans extend over this range, we find that remaining effects from oscillation are minimized by selecting successive time spans to be separated by *half* of a period—so that Φ_z changes by exactly π between $t_{i,1}$ and $t_{i+1,1}$:

$$\Phi_z(t_{i+1,1}) - \Phi_z(t_{i,1}) = \pi. \quad (13)$$

So that the model may be reasonably accurate, without encountering excessive numerical noise or excessively dynamical behavior at merger, we limit the region over which we choose these time spans to be the central 80% of time between the “relaxation time” listed in the waveform metadata and the time of maximum signal power in the waveform. This establishes $t_{0,1}$, and all successive times can be computed from that using Eqs. (12) and (13).

Now, we model the waveform “piecewise” on these spans of time, though the pieces are overlapping. The advantage of transforming the waveform as described above is that each mode separates [15] into two parts that are symmetric and antisymmetric under reflection along the z axis. The symmetric part varies on an inspiral timescale because the primary rotational behavior has been factored out by transforming to the corotating frame; the antisymmetric part is mostly due to spin-orbit coupling and therefore varies most rapidly by a complex phase with frequency equal to the rotational frequency of the frame itself, though possibly with opposite sign. We model these two parts separately as simple linear-in-time complex quantities, with an additional phase-evolution term for the antisymmetric parts. For each time span i , we write

$$\begin{aligned} \mu_i(t, \theta, \phi) = & \sum_{l,m} \left[s_i^{l,m} + \dot{s}_i^{l,m}(t - t_{i,c}) \right. \\ & \left. + \left(a_i^{l,m} + \dot{a}_i^{l,m}(t - t_{i,c}) \right) e^{i\sigma(m,l)\Phi_z(t)} \right] {}_{-2}Y_{l,m}(\theta, \phi). \end{aligned} \quad (14)$$

Here, each of the coefficients $s_i^{l,m}$, $\dot{s}_i^{l,m}$, $a_i^{l,m}$, and $\dot{a}_i^{l,m}$ is a complex constant, we use $t_{i,c} = (t_{i,1} + t_{i,2})/2$ to mitigate degeneracy between the constant-in-time and linear-in-time terms, and the function σ is given by

$$\sigma(m, l) = \begin{cases} 1 & |m| < l, \\ -1 & |m| = l. \end{cases} \quad (15)$$

These signs are chosen because they represent the dominant behavior of the corresponding terms in the data. Note that the

symmetry properties imply that once the quantities $s_i^{l,m}$, etc., are chosen for positive m , they are automatically known for negative m from the relations

$$s_i^{l,m} = (-1)^l \bar{s}_i^{l,-m}, \quad \dot{s}_i^{l,m} = (-1)^l \bar{\dot{s}}_i^{l,-m}, \quad (16a)$$

$$a_i^{l,m} = (-1)^{l+1} \bar{a}_i^{l,-m}, \quad \dot{a}_i^{l,m} = (-1)^{l+1} \bar{\dot{a}}_i^{l,-m}. \quad (16b)$$

Because the $m = 0$ modes of the SXS waveforms we use are not considered reliable [6, 11], we simply ignore those modes in both the model and the data. That is, the sum in Eq. (14) does not include any $m = 0$ modes. If the sum over modes extends from $l = 2$ to some maximum $l = L$, the total number of (real) degrees of freedom in this model is $4L(L + 1) - 8$ for each span of time. While the data we use contains up to $l = 8$, the highest-order modes contribute little to the result, and drastically increase the number of degrees of freedom in the problem (and therefore the time taken to optimize the model). Therefore, we use only up to $l = 6$ in constructing the model and evaluating the residual, reducing the degrees of freedom from 280 to 160 per time span. Finally, because this is still such a large number of degrees of freedom, we limit the evaluation to only the first two and last two time spans; we find that including the rest has no significant effect on the result, but vastly increases the amount of processing time required. This leaves us with a manageable 640 degrees of freedom in this model.

Now, using this model, we define the objective function

$$\begin{aligned} \Upsilon(\vec{\alpha}, \vec{\beta}) = & \min_{s, \dot{s}, a, \dot{a}} \sum_i \int_{t_{i,1}}^{t_{i,2}} \int_{S^2} \left| \mathcal{R}\{\mathcal{T}_{\vec{\alpha}, \vec{\beta}}[h]\} - \mu_i \right|^2 dA dt \\ = & \min_{s, \dot{s}, a, \dot{a}} \sum_i \int_{t_{i,1}}^{t_{i,2}} \sum_{l,m} \left| \mathcal{R}\{\mathcal{T}_{\vec{\alpha}, \vec{\beta}}[h]\}^{l,m} - \mu_i^{l,m} \right|^2 dt. \end{aligned} \quad (17)$$

We will use this function in two ways: first, to simply evaluate Υ for given values of $(\vec{\alpha}, \vec{\beta})$, where those values are obtained from the methods described in other sections; second, to minimize Υ over possible values of $(\vec{\alpha}, \vec{\beta})$ to find the optimum COM correction.

B. Results for the standard COM-correction method

We can now compare the value of Υ defined in Eq. (17) for all the waveforms discussed in the previous sections. First, we compare its value Υ_{raw} in the raw data to its value Υ_{OPN} using $\vec{\alpha}$ and $\vec{\beta}$ as given by Eqs. (7), where \vec{x}_{COM} is given by the Newtonian (OPN) formula. The latter corresponds to the technique actually used in the current SXS data, for waveforms found in the SXS simulation catalog with file names ending in CoM. The results of this comparison are shown in Fig. 5. The vast majority of systems improve significantly by this measure. The notable feature is that even though the naive OPN COM trajectory is so fundamentally different from Υ , this plot suggests that they agree in the sense that the OPN correction improves the waveforms for all but three systems—and even




FIG. 5. Relative difference between Υ_{raw} evaluated on the raw waveform data and Υ_{0PN} evaluated using the values $\vec{\alpha}$ and $\vec{\beta}$ given by the simplest Newtonian (OPN) approximation of Eqs. (7)—the same COM correction used in the current SXS catalog. In the vast majority of cases, the value of Υ decreases substantially (though it actually *increases* very slightly in three cases with significant eccentricity). This suggests that even though Υ and the coordinate-based COM are such entirely different measures and based on completely different data, they agree that the changes introduced by naive COM corrections are generally improvements.

for those three the change is very small.⁴

We can also actively optimize Υ over the values of $\vec{\alpha}$ and $\vec{\beta}$. The results are shown in the lower-right panel of Fig. 6. Naturally, Υ improves in every case because it is specifically being optimized. In Fig. 6, we see the pattern that the vast majority of systems are changing by small fractions. In this case, there are just three systems in which Υ changes at the percent level. These are some of the same systems that changed the most in going from the raw data to the OPN-corrected data. These particular systems also happen to be extremely long, with significant overall accelerations during the inspiral. This suggests that the corrections will be sensitive to the precise span of times over which the corrections are being made, which may explain why they continue to change so much by optimization. However, as discussed in Sec. III B, changing the beginning and ending fractions do not significantly change $\vec{\alpha}$ and $\vec{\beta}$. Nevertheless, the overall scale of the changes seen in this plot suggests once again that the naive OPN COM correction is achieving near-optimal results in the vast majority of cases.

⁴ These three systems are unusual, in that they are quite short (having 13 to 15 orbits before merger, compared to an average of 22), and have eccentricities (0.215 and higher) that place them among the 12 most eccentric in the SXS catalog. Furthermore, the magnitude of the change in Υ for each of them is very small—in the lowest percentile for the entire catalog—which suggests that the negative results may be consistent with numerical error, and in any case are not cause for much concern.




FIG. 6. Comparison between the value Υ_{opt} for which $\vec{\alpha}$ and $\vec{\beta}$ are optimized, and Υ_{0PN} evaluated using the values $\vec{\alpha}$ and $\vec{\beta}$ given by the simplest Newtonian (OPN) approximation—the same COM correction used in the current SXS catalog. The vertical axis shows the relative improvement in going from the Newtonian correction to the optimized correction. Optimization improves the results for the great majority of systems by less than 1%. The three exceptions to this rule are particularly long systems.

V. IMPROVING THE COM CORRECTION

A. Post-Newtonian COM definition

To characterize the motion seen in the COM in the raw simulation data, an obvious first step to finding a more accurate definition of the COM during the simulation is to try low orders of post-Newtonian (PN) corrections. Note that the COM should be near the origin of the coordinate system, with minimal motion around the origin from linear-momentum recoil, as discussed in Sec. V B. 1PN order corrections are analytically trivial, and are zero for circular systems. However, SXS simulations are not perfectly circular and so we investigate 1PN and 2PN order corrections to the COM. We implement the PN corrections given by Eq (4.5) in [40]. This formalism goes up to 3.5 PN for the COM vector in time. The 2PN version of the expression that we used can be found in Appendix B.

Using this formalism, the effects of the correction on the COM are small but measurable at 1PN and significantly different at 2PN for many systems. We are not concerned with the coordinates of the COM itself changing, but of the COM correction values changing, as discussed in Sec. III. In general, we assume that the COM drifts in a linear fashion and can be corrected with a translation $\vec{\alpha}$ and boost $\vec{\beta}$. Our results from the 1PN and 2PN analysis are shown in Fig. 7, which shows the relative difference between $\vec{\alpha}$, $\vec{\beta}$, and total COM displacement $\vec{\delta} = \vec{\alpha} + \vec{\beta}_{t_{\text{merger}}}$ for the 1PN (top panels) correction to the COM and the 2PN (bottom panels) correction to the COM. The 1PN corrections show small changes for most simulations. However, the 2PN correction shows more sizable changes in the COM




FIG. 7. Differences in the maximum displacements between the COM correction computed using the Newtonian (0PN) COM formula and the 1PN or 2PN COM formulas. The quantity $\vec{\mu}$, defined in Eq. (8), is the largest displacement between the origin of coordinates in the simulation and the corrected origin. The post-Newtonian corrections change the COM correction values by roughly 1% in the majority of cases. Systems with larger changes are consistent with systems that are sensitive to small changes in the endpoints of integration used to find the COM correction, as seen in Fig. 3.

correction values, which may indicate that including at least up to 2PN corrections to the COM will give better accuracy either to the COM itself or to the correction factors.

To see any potential correlations with large 1PN or 2PN corrections and simulation parameters, we compared the relative difference in the COM corrections to the eccentricity e of the system. As shown in Fig. 9, no correlations between the magnitude of the 1PN and 2PN corrections to the COM and the eccentricity are apparent, despite the definition of the 1PN contribution to the COM being dependent on e .

Using the method described in Sec. IV, comparisons of Υ , as defined in Eq. 17, between 1PN, 2PN, and the original COM correction method (dubbed 0PN), can be found in Fig. 8. The striking feature of these plots is that a significant number of systems actually have *larger* values of Υ when including either of these corrections. While it is reassuring that the majority of systems in each case only exhibit quite small changes—changes of order 10^{-4} or less in a quantity that already improved significantly from the raw data—the 1PN and 2PN corrections plots include a large group of systems that change at the percent level. These systems also happen to be the same systems that changed most drastically in going from the raw data to the 0PN-corrected data (found near the upper-right corner of the plots), and are particularly biased towards increasing values of Υ . That is to say, it appears that the 1PN and 2PN corrections do worst for the most extreme systems. This shouldn’t come as a great surprise, since those systems tend to be the ones with the most extreme mass ratios and precession, so that post-Newtonian

analysis is also expected to be at its least accurate.

B. Linear-momentum recoil

Any binary with asymmetric components will emit net linear momentum in the form of gravitational waves, which will cause a recoil of the binary itself. As the system rotates, the direction of recoil will also rotate, pushing the COM roughly in a circle [42]. In principle, this effect could cause the epicyclic motion apparent in the COM trajectories, which is further characterized in Sec. V D. To see if recoil is responsible, we use methods described in Refs. [15, 41, 42] to investigate the size of the linear-momentum recoil implied by the gravitational-wave emission in these systems.

As shown in the right-column panels of Fig. 1, the COM motion follows an overall linear track with additional epicyclic motion. The linear motion of the COM is well understood, and discussed in Sec. III. For this analysis, we assume that the linear part of the COM motion may be removed from the data without loss of information, leaving the epicyclic behavior about the coordinate origin.

Fitchett [42] used low-order approximations to the waves emitted by a BBH to calculate the radius of the COM motion induced by recoil, finding

$$r_{\text{recoil}} = a \left(\frac{29\sqrt{2}}{105} \right) \left(\frac{a}{R_s} \right)^{-7/2} f(m_1/m_2) \quad (18)$$

for a system comprised of non-spinning BHs in a circular Keplerian orbit. Here, a is the semi-major axis of the orbit and R_s is the Schwarzschild radius for the total mass of the BBH. f is a polynomial function in the BBH masses, where $f(m_1/m_2) = (1 + m_2/m_1)^{-2}(1 + m_1/m_2)^{-3}(1 - m_1/m_2)$ and $|f| < 0.02$ for all mass ratios. As mentioned in Ref. [42], the radius estimate given is accurate for order-of-magnitude comparisons only, as it is derived assuming a Keplerian orbit and using a “quasi-Newtonian” approach.

Subtracting only the linear motion out of the COM motion and comparing the radius of the measured motion to that of Eq. (18) immediately shows that the measured COM motion—specifically the epicyclic motion—is significantly larger than what can be explained by linear-momentum reaction for all simulations. Fig. 10 shows a comparison across the SXS public waveform catalog for the ratio between the measured COM radius about the coordinate origin and the estimated COM radius given by Eq. (18). The measured COM radius is calculated by averaging the distance of the COM away from the line $\vec{\alpha} + \vec{\beta}t$ between times t_i and t_f after the linear motion has been removed:

$$r_{\text{measured}} = \frac{1}{t_f - t_i} \int_{t_i}^{t_f} |\vec{x}_{\text{COM}} - (\vec{\alpha} + \vec{\beta}t)| dt. \quad (19)$$

The results show that the measured radius is typically hundreds of times larger than the radius implied by Fitchett’s analysis.

It is also possible to go beyond Fitchett’s analysis by using the *measured* gravitational waves to compute the linear momentum flux, and compare the acceleration that would cause to the




FIG. 8. Comparison of post-Newtonian contributions for determining the center of mass. These plots show the difference between the value of Υ [Eq. (17)] resulting from the naive 0PN method based on the coordinate trajectories of the apparent horizons and the value of Υ when incorporating 1PN and 2PN effects, as described in the text. The horizontal axes show the relative magnitude of the change when going from the raw data to the corrected waveform. In most cases, the values actually become significantly *larger* when going from the 0PN value to other values. Those systems are shown as crosses, while systems with smaller values are shown as circles.




FIG. 9. Differences between the Newtonian COM and the 1PN corrected COM (top row) and 2PN corrected COM (bottom row) correction values versus the eccentricity of the simulation at relaxation time. No correlations are evident between either the 1PN or 2PN correction and eccentricity values.

measured acceleration of the COM in the simulation. The acceleration of the COM, \vec{a}_{COM} was calculated by taking two time derivatives of the coordinate position of the COM after removing the linear motion:

$$\vec{a}_{\text{COM}} = \frac{d^2}{dt^2} (\vec{x}_{\text{COM}} - (\vec{\alpha} + \vec{\beta}t)). \quad (20)$$

The linear-momentum flux may be calculated from the gravita-

tional radiation as

$$\frac{d\vec{p}}{d\Omega dt} = \frac{R^2}{16\pi} \left| \frac{dh}{dt} \right|^2 \hat{r}, \quad (21)$$

where \hat{r} is the direction from the source to the point in question, R is the distance from the source to the observation sphere, and Ω represents all angles on the sphere. Integrating over all angles to find the total linear momentum flux $d\vec{p}/dt = \dot{\vec{p}}$, we need to expand h into SWSHs as done in Eq. (3). We can also decompose \hat{r} accordingly as

$$\hat{r} = \sqrt{\frac{2\pi}{3}} (Y_{1,-1} - Y_{1,1}, iY_{1,-1} + iY_{1,1}, \sqrt{2}Y_{1,0}). \quad (22)$$

Then integrating over all angles gives

$$\dot{p}_j = \frac{R^2}{16\pi} \sum_{l,l',m,m'} \hat{r}_j^{l,m'-m} \bar{h}^{l,m} \bar{h}^{l',m'} (-1)^{m'} \sqrt{\frac{3(2l+1)(2l'+1)}{4\pi}} \begin{pmatrix} l & l' & 1 \\ m & -m' & m' - m \end{pmatrix} \begin{pmatrix} l & l' & 1 \\ 2 & -2 & 0 \end{pmatrix}, \quad (23)$$

where the last two factors are Wigner 3- j symbols. Explicit expressions for the calculation of the linear momentum flux are given in Appendix C. Our calculation of the linear momentum flux can be found in the open-source package `spherical_functions` [43].

Of course, there are multiple methods for calculating the linear momentum for a BBH system. We compared our method of using the $h^{l,m}$ modes directly with the method proposed in Ref. [41], which derives the linear momentum flux from the energy emitted by the system. We found that the two derivations give the same results to within numerical accuracy.




FIG. 10. Comparing measured COM motion to motion caused by emission of linear momentum carried away by gravitational waves. The panel on the left shows the average ratio of the measured radius of COM motion, given by the time-averaged magnitude of the COM epicycles and stated in Eq. (19), to the radius expected from leading-order calculations given by Eq. (18). The center and right panels show the average ratio of the measured COM acceleration, given by the second time derivative of the COM coordinate positions and stated in Eq. (20), to the acceleration due to asymmetric momentum flux carried by the measured gravitational waves—for near-equal mass ratios and larger mass ratios, respectively. Blue indicates runs using the initial-data method described in Ref. [18], orange indicates runs using the previous initial-data method. In every case, the measured motion is at least an order of magnitude larger than the motion expected from gravitational-wave recoil.

We compared the linear-momentum flux divided by the total mass of the system with the measured acceleration of the COM and find that the values for most of the SXS public waveform catalog do not agree. We further found that the COM acceleration for non-equal mass systems is consistently larger than the acceleration found through the linear momentum flux, confirming that most of the COM motion is not due to linear-momentum recoil.

An overview of the ratios between \vec{a}_{COM} and $\dot{\vec{p}}/M$ can be seen in Fig. 10. Note that linear-momentum recoil in equal mass or near-equal mass BBH systems is expected to be very small, and so these systems are isolated as a separate case in the middle panel of Fig. 10.

C. Causes of unphysical COM motion

If the COM motion seen in the SXS BBH simulations cannot be explained by physical processes such as inclusion of PN terms or linear-momentum recoil, then why is it there? There are two potential causes for the appearance of unphysical or erroneous COM motion: (i) the presence of uncontrolled residual linear momentum in the initial data, and (ii) the emission of unresolved junk radiation at the beginning of the BBH simulation causing effectively random and unphysical coordinate kicks to the system. The presence of uncontrolled residual linear momentum was addressed and partially rectified in Ref. [18], leading to a new method of creating initial data for the BBH simulations. However, this method was not used for all systems in the SXS catalog, and does not completely resolve the issues of spurious translations and boosts even when it is applied.

Another factor that seems to cause COM motion starting from early in the simulations appears to be junk radiation, which

is an inherent part of BBH simulations. It is the radiation emitted when a BBH relaxes from its initial-data “snapshot”, which is only an approximation to the true state of a long binary inspiral at the time the simulation starts. Junk radiation is physical, in the sense that if the entire spacetime were actually in the configuration given by the initial data, it would indeed emit this radiation as the system evolved. However, it is not *astrophysical*, in the sense that no real system in the universe is expected to contain this type of radiation. Junk radiation contains high frequencies that are largely unresolved in BBH simulations because of limits of computational power and time. As the resolution increases in the simulations, more of the junk radiation is accurately treated. This is a potential cause for the difference in initial kicks of the COM, as seen in Fig 2. Even in our higher-resolution simulations, not all of the junk radiation is accounted for. This failure to resolve all of the junk radiation possibly contributes to the COM kicks observed in SpEC BBH simulations.

Fortunately, kicks from the emission of unresolved junk radiation can be corrected using the gauge transformations discussed in the previous section. The large epicyclic motion in the COM cannot be accounted for using a BMS gauge transformation like the linear motion of the COM. The cause for such large, seemingly unphysical epicyclic motion is unknown, and is left for future work.

D. Epicycle Quantification

As seen in the right column panels of Fig. 1, the COM motion has both a linear and epicyclic component. The linear component of the COM motion has been discussed and is already considered in the current COM correction technique.

The size of the epicyclic motion in the COM cannot be solely




FIG. 11. Change in the size of the COM correction, $\vec{\mu}$ as defined in Eq. (8), when removing epicycles before fitting for the COM correction, as described in Eq. (26). These changes are comparable to, but almost always smaller than, the changes due to variations in the endpoints of integration as seen in Fig. 3; they are also smaller than most of the post-Newtonian corrections seen in Fig. 7, except for systems changing here by more than about 10%.

explained by linear-momentum recoil, as discussed in Sec. V B. The leftmost plot in Fig. 10 illustrates that the expected radius of the epicycles from linear-momentum recoil, even on an approximate basis, is orders of magnitude smaller than what is actually seen, given by r_{measured} defined in Eq. 19. The actual size of r_{measured} is fairly consistent across the SXS simulation catalog regardless of simulation parameters or initial data construction, and tends to be between 0.01 and 0.1 with an average value of 0.026 in simulation units. There is not a significant change in the distribution or magnitude of r_{measured} between the 0PN, 1PN, or 2PN representations of the COM.

We assume that the epicycle motion seen in the COM after the translation and boost are applied is from the calculated COM, \vec{x}_{COM} , being displaced from the optimal COM by a small amount. “Displaced” here means that we assume \vec{x}_{COM} is displaced from the optimal COM along the separation vector $\vec{r}_{ab} = \vec{x}_a - \vec{x}_b$, and not in any other direction.

Regardless of the origin of these epicycles, removing them to calculate the linear COM correction should improve the quality of the waveforms. Their removal should reduce the error associated with the averaging done to calculate the translation and boost correction values. As discussed regarding the choice of beginning and ending times in Sec. III B, the presence of large epicycles has the potential to affect the reproducibility and reliability of the calculation of \vec{a} and $\vec{\beta}$. Removing the epicycles, or at the very least minimizing them, to calculate a more accurate COM correction should further reduce mode mixing. Of course, epicycle motion cannot be completely subtracted from the waveform itself as this would require an acceleration correction, which is not an allowed BMS transformation.

To accurately describe the epicycles, we need to define the co-rotating coordinate frame. For our simulations, we have three unit vectors that describe the rotating coordinate frame:

$$\begin{aligned}\hat{n} &= \frac{\vec{x}_a - \vec{x}_b}{|\vec{x}_a - \vec{x}_b|} = \frac{\vec{r}_{ab}}{|\vec{r}_{ab}|}, \\ \hat{k} &= \frac{\vec{r}_{ab} \times \vec{r}_{ab}}{|\vec{r}_{ab} \times \vec{r}_{ab}|} = \frac{\vec{\omega}}{|\vec{\omega}|}, \\ \hat{\lambda} &= -\hat{n} \times \hat{k},\end{aligned}\quad (24)$$

where \hat{n} points along the separation vector \vec{r}_{ab} , \hat{k} points along the orbital angular velocity $\vec{\omega}$, and $\hat{\lambda}$ points along the direction of rotation.

This leads us to a potential method for epicycle removal. First, we calculate the estimated spatial translation \vec{a} and boost $\vec{\beta}$ from the original COM \vec{c} using the current averaging method

$$\vec{c}_1 = \vec{x}_{\text{COM}} - (\vec{a} + \vec{\beta}t). \quad (25)$$

Note that $\langle |\vec{c}_1| \rangle = r_{\text{measured}}$ as defined in Eq. (19), using the angle-bracket notation to imply averaging over time, as for the moments of the COM position in Eqs. (6). We then take the new COM values \vec{c}_1 and calculate the corresponding co-rotating coordinate frame unit vectors given in Eq. (24). These unit vectors are then applied to the original COM \vec{x}_{COM} as

$$\vec{c}_2 = \vec{x}_{\text{COM}} - \Delta_n \hat{n} - \Delta_\lambda \hat{\lambda} - \Delta_k \hat{k}, \quad (26)$$

where $\Delta_n = \langle \vec{c}_1(t) \cdot \hat{n}(t) \rangle$, $\Delta_\lambda = \langle \vec{c}_1(t) \cdot \hat{\lambda}(t) \rangle$, $\Delta_k = \langle \vec{c}_1(t) \cdot \hat{k}(t) \rangle$, are the time averaged projections of the linearly corrected COM onto the rotational coordinate system unit vectors. The epicycle corrected COM \vec{c}_2 is then fed back into the averaging method to get the final spatial translation \vec{a}_{epi} and boost $\vec{\beta}_{\text{epi}}$ values, which can then be applied to the waveform data.

Applying this method to the BBH simulations in the SXS public simulation catalog, we can calculate \vec{a}_{epi} and $\vec{\beta}_{\text{epi}}$ COM correction values. Figure 11 compares the usual “0PN” COM correction values with the epicycle-removed values. The values plotted involve $\vec{\mu}$, defined in Eq. (8), which is the largest displacement between the origin of coordinates in the simulation and the corrected origin. We see that epicycle removal changes the COM correction values at a scale comparable to the changes caused by varying the endpoints of integration used to determine the COM, as seen in Fig. 3. The changes due to epicycle removal are generally somewhat smaller than changes due to varying endpoints. If we remove the epicycles *before* applying those variations, the changes seen in Fig. 3 are reduced by a typical factor of two—though there is no apparent effect on roughly 10% of systems. Systems that change by more than 10% in this figure are also typically changing by more than the post-Newtonian changes shown in Fig. 7.

Using the method outlined in Sec. IV, we also find that we cannot reliably conclude that the epicycle correction actually makes a significant improvement in the waveforms. The results of the Υ comparisons between the original COM correction method outlined in Sec. III and the epicycle removal method




FIG. 12. This plot shows the difference between the value of Υ [Eq. (17)] resulting from the naive OPN method based on the coordinate trajectories of the apparent horizons and the value resulting from the epicycle removal method described in Sec. V D. The horizontal axis shows the relative magnitude of the change when going from the raw data to the corrected waveform, while the vertical axis shows the change when incorporating the epicycle corrections. In most cases, the values of Υ actually become significantly *larger* when going from the OPN value to the epicycle corrected value. Those systems are shown as crosses, while systems with smaller values are shown as circles.

in this section can be found in Fig. 12. This plot shows that approximately 30% of simulations improve using the epicycle method, and approximately 70% get worse. This is not enough of a benefit to warrant the use of the epicycle removal in all simulations, and implies that the epicycle removal, at this stage of BBH simulations, is an unnecessary step in calculating and applying the COM correction, a somewhat disappointing conclusion.

As mentioned in Sec. III B, epicycles are also a potential source of instability regarding choice of beginning and ending times t_i , t_f . Initial investigation into how the epicycle removal method affects changes in the COM correction values due to differing t_i , t_f implies that the epicycle removal method does not diminish changes in the COM correction values. This is also an unintuitive and disappointing result, and may imply that other methods are required to calculate the COM correction values $\vec{\alpha}$ and $\vec{\beta}$ after epicycle removal or a different method for epicycle removal entirely. We leave such an investigation to future work.

E. Position of the COM

During the analysis of the epicycles present in the COM, we investigated the position of the COM relative to the two black holes. We still assumed, that \vec{x}_{COM} would lie along the separation vector between the two black holes. As seen in Fig. 13, this is not the case. The COM deviates significantly from the separation vector between the two black holes and

typically lies in the rotation direction $\pm\hat{\lambda}$. The projection in the $\pm\hat{\lambda}$ direction, Δ_λ , is typically not 0, and averages at $-0.44r_{\text{measured}}$ when considering all simulations in the SXS Catalog. On average, the projection of \vec{c}_1 into the $\pm\hat{n}$ direction, Δ_n , is smaller than Δ_λ , with an average ratio Δ_λ/Δ_n for spin-aligned systems of -1.43 and -2.48 for precessing systems. The projection in the $\pm\hat{k}$ direction, Δ_k , is significantly smaller than Δ_λ , with an average ratio Δ_λ/Δ_k of -3.22×10^3 for spin-aligned systems and -2.78 for precessing systems.

Having typically most of the COM correction vector \vec{c}_1 in the direction of $\pm\hat{\lambda}$ potentially indicates unaccounted spin-orbit effects on the COM, unknown effects from unresolved junk radiation, or additional gauge effects that cannot be compensated for with BMS transformations. It is also suspicious that much of the COM correction vector lies out of the orbital plane in the $\pm\hat{k}$ direction. Both attributes of \vec{c}_1 warrant further investigation, and is left to future work.

Correlations between Δ_n , Δ_λ , and Δ_k with pertinent simulation parameters are shown in Fig. 13. This plot shows the correlations between Δ_n , Δ_λ , Δ_k , χ_{eff} , χ_p , and q .

A few notable correlations are apparent in Fig. 13, the most obvious being the correlation between Δ_n and Δ_λ . Most simulations tend to have negative Δ_λ values that grow larger in magnitude with increasing Δ_n , however there is also a cluster of aligned-spin simulations with Δ_n , Δ_λ values close to zero. Additionally, Δ_λ becomes more negative with increasing q and there are some weak correlations between both COM position offsets Δ_n , Δ_λ with χ_{eff} , but not with χ_p . Δ_k does not appear to have any strong correlations. It is apparent that spin aligned simulations tend to have small or 0 Δ_k values, which is unsurprising as Δ_k represents the projection of the COM out of the orbital plane. It can also be seen in Fig. 13 that larger Δ_k values cluster around $\Delta_n = 0$ and $\Delta_\lambda < 0$, which may only be due to Δ_n values being symmetric around 0 and Δ_λ values being mostly negative. No apparent correlations are present for other simulation parameters.

VI. CONCLUSIONS

In this work, we have investigated the effects of COM motion waveforms, removed unphysical COM motion through allowed gauge transformations to the waveforms, and have investigated methods for improving the COM correction. Having unphysical motion in the COM causes mode mixing in the gravitational waveforms, and thus a power loss from the dominant $(2, \pm 2)$ modes to the less-dominant, higher-order modes—which is typically visible as amplitude modulation in the higher-order modes.

We found that the COM motion observed in the SXS simulations cannot be entirely accounted for by PN corrections or linear-momentum recoil. We also found that the motion of the COM does not lie along the separation vector between the black holes and is offset from the estimated COM within and out of the orbital plane, which is not expected on physical grounds.

The current method for correcting the COM motion uses




FIG. 13. Comparisons of effective spin χ_{eff} as defined in Eq. 9, effective precessing spin χ_p as defined in Eq. 10, mass ratio q , and the time averaged projections of \vec{c}_1 onto \hat{n} , $\hat{\lambda}$, and \hat{k} ; Δ_n , Δ_λ , Δ_k . Red represents spin-aligned simulations, and teal represents precessing simulations.

allowed BMS transformations, namely a spatial translation and boost that counteracts the linear motion from the COM and removes a large amount of the mode mixing from the waveform. The translation and boost are calculated for all simulations at all resolutions, as the COM motion is not consistent between different resolutions of the same system.

We attempted to improve the COM correction by developing a method to remove the large epicycles from the COM motion before calculating the BMS translation and boost. We found that the resulting changes to the translation and boost values were not significant and did not improve the waveforms compared to the originally calculated values.

Lastly, we introduced a complementary method to quantify the effect of the COM correction on the waveforms. We used this method to determine that PN corrections and the epicycle removal technique did not improve the COM correction transformations, and thus would not further improve the waveforms

or accurately describe the COM physically.

Future work includes investigating spin-orbit effects on the COM and the peculiarity of the COM position being offset from the separation vector of the BBH. Further investigation is required specifically on the unaccounted for size of the radius of the epicycles seen in the COM motion, which may be due to unknown spin-orbit or unresolved junk radiation effects, and may be corrected with additional gauge transformations that minimize the epicycles.

ACKNOWLEDGMENTS

It is our pleasure to thank Saul Teukolsky for useful conversations and guidance on this project. We also thank Leo Stein for helpful insight and suggestions regarding notation. This project was supported in part by the Sherman Fairchild Foundation; by NSF Grant PHY-1606654 at Cornell; and by

NSERC of Canada grant PGSD3-504366-2017 as well as the Canada Research Chairs Program.

The computations described in this paper were performed on the Wheeler cluster at Caltech, which is supported by the Sherman Fairchild Foundation and by Caltech; and on the GPC and Gravity clusters at the SciNet HPC Consortium, funded by: the Canada Foundation for Innovation under the auspices of Compute Canada; the Government of Ontario; Ontario Research Fund–Research Excellence; and the University of Toronto.

Appendix A: Spin-Weighted Spherical Harmonics

Spin-weighted spherical harmonics (SWSHs) are typically used to generalize the well-known standard spherical harmonics. Specifically, SWSHs provide a decomposition of general Spin-weighted spherical functions (SWSFs) into a sum of SWSHs. SWSFs themselves provide a vital way to study waves radiating from bounded regions, and so have an obvious and important application in gravitational-wave astronomy, which is the focus of this work. SWSFs play two key roles in this field: (1) describing the magnitude of the wave given any direction of emission or observation, and (2) providing polarization information. There are a number of subtleties in defining SWSFs and hence SWSHs, including dependencies on the chosen coordinate system and the explicit definition of SWSFs.

The spin weight of a function is defined by how it transforms under rotation of the spacelike vectors $\mathfrak{R}(m^\mu)$, $\mathfrak{I}(m^\mu)$ where m^μ is a complex null vector tangent to the coordinate sphere S^2 . The rotation of these spacelike vectors is given by

$$(m^\mu)' = e^{i\psi} m^\mu. \quad (\text{A1})$$

A function η is then said to have a spin weight s if it transforms as

$$\eta' = e^{s i \psi} \eta. \quad (\text{A2})$$

In the case of gravitational waves, the metric perturbation h has a spin weight of -2 [16, 17] and this decomposition has been used in numerical relativity extensively.

The classic definition of SWSHs [17] writes the functions in terms of spherical coordinates for S^2 , giving them as explicit formulas using polar and azimuthal angles (θ, ϕ) and using two integer variables that define the order of spherical harmonic to be used. The convention used here and in SXS is $Y_{l,m}$.

SWSHs are thus classically defined as

$${}_s Y_{l,m} = \begin{cases} \left[\frac{(l-s)!}{(l+s)!} \right]^{1/2} \delta^s Y_{l,m}, & 0 \leq s \leq l, \\ (-1)^s \left[\frac{(l+s)!}{(l-s)!} \right]^{1/2} \bar{\delta}^{-s} Y_{l,m}, & -l \leq s \leq 0, \end{cases} \quad (\text{A3})$$

where δ is effectively a covariant differentiation operator in the

surface of the sphere. δ is defined [17] as

$$\delta \eta = -(\sin \theta)^s \left\{ \frac{\partial}{\partial \theta} + \frac{i}{\sin \theta} \frac{\partial}{\partial \phi} \right\} \{ (\sin \theta)^s \eta \} \quad (\text{A4})$$

when operating on some function η that has a spin weight s .

The above classic method inherits an unfortunate dependency on the chosen coordinates. In particular, SWSHs *cannot* be written as functions on the sphere S^2 ; they can only be written as functions on *coordinates of* S^2 . As such, SWSHs as defined in Eq. (A3) do not transform among themselves under rotation of the sphere (or, equivalently, rotation of coordinates of the sphere). That is, a SWSH in a given coordinate system cannot generally be expressed as a linear combination of SWSHs in another coordinate system. A more correct method for defining SWSFs that does not inherit these coordinate-system dependencies is to represent them as functions from $\text{Spin}(3) \approx \text{SU}(2)$, which in turn maps onto S^2 [44]. By forming a representation of $\text{Spin}(3)$, SWSHs defined in this way do transform among themselves and still agree with the classic definition. SWSHs may then be defined as

$${}_s Y_{l,m}(\mathbf{R}) := (-1)^s \sqrt{\frac{2l+1}{4\pi}} \mathfrak{D}_{m,-s}^{(l)}(\mathbf{R}), \quad (\text{A5})$$

where \mathfrak{D} is a Wigner matrix, which are representations of the spin group, and \mathbf{R} is the $\text{Spin}(3)$ argument. Taking \mathbf{R} to be in the unit-quaternion representation of $\text{Spin}(3)$, the \mathfrak{D} matrices may be expressed as

$$\mathfrak{D}_{m',m}^{(l)}(\mathbf{R}) = \sqrt{\frac{(l+m)!(l-m)!}{(l+m')!(l-m')!}} \sum_{\rho} \binom{l+m'}{\rho} \binom{l-m'}{l-\rho-m} (-1)^{\rho} R_s^{l+m'-\rho} \bar{R}_s^{l-\rho-m} R_a^{\rho-m'+m} \bar{R}_a^{\rho}, \quad (\text{A6})$$

where R_s and R_a are the geometric projections of the quaternion into symmetric and antisymmetric parts, and are complex combinations of components of the quaternion. The definition given in Eq. (A5) is consistent with the definition of SWSHs typically used within SXS and is the assumed formulation for this work.

For further information on SWSFs and SWSHs, a comprehensive in-depth discussion of the history, details, and additional formulations of SWSHs can be found in [44].

Appendix B: Post-Newtonian Correction to the COM

As discussed in Sec. V A, we used the 1PN and 2PN corrections to the COM as outlined in Ref. [40]. The COM up to 2PN order is given in Eq. (4.5) of Ref. [40] as

$$\begin{aligned}
G^i &= m_a y_a^i + \frac{1}{c^2} \left[y_a^i \left(-\frac{Gm_a m_b}{2r_{ab}} + \frac{m_a v_a^2}{2} \right) \right] \\
&+ \frac{1}{c^4} \left[v_a^i Gm_a m_b \left(-\frac{7}{4}(n_{ab} v_a) - \frac{7}{4}(n_{ab} v_b) \right) + y_a^i \left(-\frac{5G^2 m_a^2 m_b}{4r_{ab}^2} + \frac{7G^2 m_a m_b^2}{4r_{ab}^2} + \frac{3m_a v_a^4}{8} \right. \right. \\
&\left. \left. + \frac{Gm_a m_b}{r_{ab}} \left(-\frac{1}{8}(n_{ab} v_a)^2 - \frac{1}{4}(n_{ab} v_a)(n_{ab} v_b) + \frac{1}{8}(n_{ab} v_b)^2 + \frac{19}{8}v_a^2 - \frac{7}{4}(v_a v_b) - \frac{7}{8}v_b^2 \right) \right) \right] + a \iff b
\end{aligned} \tag{B1}$$

where the superscript i designates the component being summed (x, y, z), subscripts a, b designate which object is being considered, \vec{y} is the position of the body being considered, $r_{ab} = |\vec{y}_a - \vec{y}_b|$ is the distance between body a and b , \vec{v} is the velocity of the body being considered and likewise v is the magnitude of the velocity, and $\vec{n}_{ab} = \vec{r}_{ab}/r_{ab}$. Parentheses here represent the scalar product of the interior values, e.g., $(n_{ab} v_b) = \vec{n}_{ab} \cdot \vec{v}_b$.

Note that this representation of the COM position does not include an overall division by the total mass of the system, and so our calculations deviate from Eq. (B1) only by including an

overall denominator $M = m_a + m_b$.

Appendix C: Linear Momentum Flux from $h^{l,m}$ modes

As mentioned in Sec. VB, our calculation of the linear momentum flux from the simulations is based on the formalism outlined in [15], which uses the SWSH structure of the gravitational strain h . Starting from Eq. (23), which gives the general form the linear momentum flux in $h^{l,m}$ modes, we can evaluate the components of the linear momentum flux:

$$\begin{aligned}
\dot{p}_x &= \frac{R^2}{16\pi} \left[\sum_{l,l',m,m-1} -(-1)^m \dot{h}^{l,m} \bar{h}^{l',m-1} \sqrt{\frac{(2l+1)(2l'+1)}{2}} \begin{pmatrix} l & l' & 1 \\ m & 1-m & -1 \end{pmatrix} \begin{pmatrix} l & l' & 1 \\ 2 & -2 & 0 \end{pmatrix} \right. \\
&\quad \left. - \sum_{l,l',m,m+1} \dot{h}^{l,m} - (-1)^m \bar{h}^{l',m+1} \sqrt{\frac{(2l+1)(2l'+1)}{2}} \begin{pmatrix} l & l' & 1 \\ m & -1-m & 1 \end{pmatrix} \begin{pmatrix} l & l' & 1 \\ 2 & -2 & 0 \end{pmatrix} \right]
\end{aligned} \tag{C1}$$

$$\begin{aligned}
\dot{p}_y &= \frac{R^2}{16\pi} \left[\sum_{l,l',m,m-1} -i(-1)^m \dot{h}^{l,m} \bar{h}^{l',m-1} \sqrt{\frac{(2l+1)(2l'+1)}{2}} \begin{pmatrix} l & l' & 1 \\ m & 1-m & -1 \end{pmatrix} \begin{pmatrix} l & l' & 1 \\ 2 & -2 & 0 \end{pmatrix} \right. \\
&\quad \left. + \sum_{l,l',m,m+1} \dot{h}^{l,m} - i(-1)^m \bar{h}^{l',m+1} \sqrt{\frac{(2l+1)(2l'+1)}{2}} \begin{pmatrix} l & l' & 1 \\ m & -1-m & 1 \end{pmatrix} \begin{pmatrix} l & l' & 1 \\ 2 & -2 & 0 \end{pmatrix} \right]
\end{aligned} \tag{C2}$$

$$\dot{p}_z = \frac{R^2}{16\pi} \sum_{l,l',m} \dot{h}^{l,m} \bar{h}^{l',m} (-1)^m \sqrt{(2l+1)(2l'+1)} \begin{pmatrix} l & l' & 1 \\ m & -m & 0 \end{pmatrix} \begin{pmatrix} l & l' & 1 \\ 2 & -2 & 0 \end{pmatrix}, \tag{C3}$$

where x, y, z refer to the simulation coordinates, with the orbit typically lying in the $x - y$ plane, R is the distance to the observation sphere, and \dot{h}, \bar{h} are the time derivative and its

conjugate of the mode amplitudes. The matrices at the ends of the summations are Wigner 3- j symbols.

- [1] G. Lovelace, C. O. Lousto, J. Healy, M. A. Scheel, A. G. a, R. O'Shaughnessy, M. Boyle, M. Campanelli, D. A. Hemberger, L. E. K. r, H. P. Pfeiffer, B. Szilagyi, S. A. Teukolsky, and Y. Zlochower, *Class. Quant. Grav.* **33**, 244002 (2016), arXiv:1607.05377 [gr-qc].
- [2] J. Healy *et al.*, *Phys. Rev.* **D97**, 064027 (2018), arXiv:1712.05836 [gr-qc].

- [3] P. Kumar *et al.*, *Phys. Rev.* **D89**, 042002 (2014), arXiv:1310.7949 [gr-qc].
- [4] J. Blackman, S. E. Field, M. A. Scheel, C. R. Galley, D. A. Hemberger, P. Schmidt, and R. Smith, *Phys. Rev.* **D95**, 104023 (2017), arXiv:1701.00550 [gr-qc].
- [5] T. Chu, H. Fong, P. Kumar, H. P. Pfeiffer, M. Boyle, D. A. Hemberger, L. E. Kidder, M. A. Scheel, and B. Szilagyi, *Class.*

- Quant. Grav. **33**, 165001 (2016), arXiv:1512.06800 [gr-qc].
- [6] M. Boyle and A. H. Mroué, *Phys. Rev. D* **80**, 124045 (2009).
- [7] L. Gualtieri, E. Berti, V. Cardoso, and U. Sperhake, *Phys. Rev. D* **78**, 044024 (2008).
- [8] B. J. Kelly and J. G. Baker, *Phys. Rev. D* **87**, 084004 (2013).
- [9] M. Boyle, *Phys. Rev. D* **93**, 084031 (2016), arXiv:1509.00862 [gr-qc].
- [10] <https://www.black-holes.org/waveforms>.
- [11] M. Boyle, D. Hemberger, G. Lovelace, S. Ossokine, H. P. Pfeiffer, M. A. Scheel, L. C. Stein, C. J. Woodford, A. B. Zimmerman, N. Afshari, K. Barkett, J. Blackman, K. Chatziioannou, T. Chu, N. Demos, S. Field, H. Fong, E. Foley, A. Garcia, M. Giesler, F. Hebert, I. Hinder, D. A. B. Iozzo, R. Katebi, H. Khan, L. E. Kidder, P. Kumar, K. Kuper, H. Lim, T. Ramirez, S. Rodriguez, P. Schmidt, B. Szilagy, S. A. Teukolsky, V. Varma, and M. Walker, “The SXS collaboration catalog of binary black hole simulations,” Not yet published.
- [12] SXS Collaboration, “Binary black-hole simulation SXS:BBH:0314,” (2019), mass ratio 1.23, with aligned spins.
- [13] SXS Collaboration, “Binary black-hole simulation SXS:BBH:0622,” (2019), mass ratio 1.2, precessing spins of magnitude 0.85 in random directions.
- [14] M. Boyle, *Phys. Rev. D* **87**, 104006 (2013).
- [15] M. Boyle, L. E. Kidder, S. Ossokine, and H. P. Pfeiffer, “Gravitational-wave modes from precessing black-hole binaries,” (2014), arXiv:1409.4431 [gr-qc].
- [16] D. Brown, S. Fairhurst, B. Krishnan, R. A. Mercer, R. K. Koppa-rapu, L. Santamaria, and J. T. Whelan, (2007), arXiv:0709.0093 [gr-qc].
- [17] E. T. Newman and R. Penrose, *J. Math. Phys.* **7**, 863 (1966).
- [18] S. Ossokine, F. Foucart, H. P. Pfeiffer, M. Boyle, and B. Szilagy, *Class. Quant. Grav.* **32**, 245010 (2015), arXiv:1506.01689 [gr-qc].
- [19] C. J. Handmer, B. Szilagy, and J. Winicour, *Class. Quant. Grav.* **32**, 235018 (2015), arXiv:1502.06987 [gr-qc].
- [20] H. Bondi, M. G. J. van der Burg, and A. W. K. Metzner, *Proc. Roy. Soc. Lond.* **A269**, 21 (1962).
- [21] R. K. Sachs, *Proc. Roy. Soc. Lond.* **A270**, 103 (1962).
- [22] A. Strominger, *Lectures on the infrared structure of gravity and gauge theory* (Princeton University Press, Princeton, New Jersey, 2018) pp. 157–159.
- [23] E. T. Newman and T. W. J. Unti, *Journal of Mathematical Physics* **3**, 891 (1962).
- [24] G. Barnich and P.-H. Lambert, *Adv. Math. Phys.* **2012**, 197385 (2012), arXiv:1102.0589 [gr-qc].
- [25] N. W. Taylor, M. Boyle, C. Reisswig, M. A. Scheel, T. Chu, L. E. Kidder, and B. Szilagy, *Phys. Rev. D* **88**, 124010 (2013).
- [26] B. P. Abbott *et al.* (Virgo, LIGO Scientific), *Phys. Rev. Lett.* **116**, 241102 (2016), arXiv:1602.03840 [gr-qc].
- [27] E. T. Newman and R. Penrose, *Journal of Mathematical Physics* **7**, 863 (1966).
- [28] T. M. Adamo and E. T. Newman, *Classical and Quantum Gravity* **27**, 245004 (2010), arXiv:1007.4215 [gr-qc].
- [29] T. M. Adamo, E. T. Newman, and C. Kozameh, *Living Reviews in Relativity* **15**, 1 (2012).
- [30] E. Deadman and J. M. Stewart, *Classical and Quantum Gravity* **26**, 065008 (2009), arXiv:0902.0481 [gr-qc].
- [31] C. Kozameh, E. T. Newman, and G. Silva-Ortigoza, *Classical and Quantum Gravity* **25**, 145001 (2008), arXiv:0802.3314 [gr-qc].
- [32] L. B. Szabados, *Living Reviews in Relativity* **7**, 4 (2004).
- [33] A. D. Helfer, *General Relativity and Gravitation* **39**, 2125 (2007), arXiv:0709.1078 [gr-qc].
- [34] T. Dray and M. Streubel, *Classical and Quantum Gravity* **1**, 15 (1984).
- [35] A. Ashtekar and J. Winicour, *Journal of Mathematical Physics* **23**, 2410 (1982).
- [36] É. É. Flanagan, D. A. Nichols, L. C. Stein, and J. Vines, *Phys. Rev. D* **93**, 104007 (2016).
- [37] M. Boyle, *Advances in Applied Clifford Algebras* **27**, 2345 (2017).
- [38] R. O’Shaughnessy, B. Vaishnav, J. Healy, Z. Meeks, and D. Shoemaker, *Phys. Rev. D* **84**, 124002 (2011).
- [39] M. Boyle, R. Owen, and H. P. Pfeiffer, *Phys. Rev. D* **84**, 124011 (2011).
- [40] V. C. de Andrade, L. Blanchet, and G. Faye, *Class. Quant. Grav.* **18**, 753 (2001), arXiv:gr-qc/0011063 [gr-qc].
- [41] D. Gerosa, F. Hébert, and L. C. Stein, *Phys. Rev. D* **97**, 104049 (2018), arXiv:1802.04276 [gr-qc].
- [42] M. J. Fitchett, *Mon. Not. R. Astron. Soc.* **203**, 1049 (1983), <http://oup.prod.sis.lan/mnras/article-pdf/203/4/1049/18223796/mnras203-1049.pdf>.
- [43] https://github.com/moble/spherical_functions.
- [44] M. Boyle, *J. Math. Phys.* **57**, 092504 (2016), arXiv:1604.08140 [gr-qc].