

Transformations of asymptotic gravitational-wave data

Michael Boyle

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

(Received 17 September 2015; published 18 April 2016)

Gravitational-wave data is gauge dependent. While we can restrict the class of gauges in which such data may be expressed, there will still be an infinite-dimensional group of transformations allowed while remaining in this class, and almost as many different—though physically equivalent—waveforms as there are transformations. This paper presents a method for calculating the effects of the most important transformation group, the Bondi-Metzner-Sachs (BMS) group, consisting of rotations, boosts, and supertranslations (which include time and space translations as special cases). To a reasonable approximation, these transformations result in simple coupling between the modes in a spin-weighted spherical-harmonic decomposition of the waveform. It is shown that waveforms from simulated compact binaries in the publicly available SXS waveform catalog contain unmodeled effects due to displacement and drift of the center of mass, accounting for mode mixing at typical levels of 1%. However, these effects can be mitigated by measuring the average motion of the system’s center of mass for a portion of the inspiral, and applying the opposite transformation to the waveform data. More generally, controlling the BMS transformations will be necessary to eliminate the gauge ambiguity inherent in gravitational-wave data for both numerical and analytical waveforms. Open-source code implementing BMS transformations of waveforms is supplied along with this paper in the supplemental materials.

DOI: [10.1103/PhysRevD.93.084031](https://doi.org/10.1103/PhysRevD.93.084031)**I. INTRODUCTION**

As the era of gravitational-wave astronomy approaches, models of gravitational waveforms from physical systems become crucial to the extraction of scientific results from the data. The basic goal of this effort is to make the claim that a waveform measured in a detector corresponds to some particular physical model. But a treacherous gulf lies between any waveform and its corresponding physical model, abounding in subtle and delicate challenges—most least of which is the gauge flexibility of general relativity. This paper describes the gauge transformations most relevant to studies of gravitational waves and shows how to calculate their effects on waveforms. We will see that, in order to obtain accurate waveform models, we must account for gauge effects.

The literature on gravitational-wave analysis almost universally allows for two standard gauge ambiguities: time translations and phase rotations. For example, the standard technique of matched filtering involves optimizing the match over the time and phase of the signal [1,2]. Similarly, comparisons between numerical evolutions, between numerical and analytical waveforms, and between different approximate analytical waveforms have generally allowed for time and phase offsets [3,4]. These transformations alter the waveforms, but in well-behaved ways which can be expressed fairly simply as functions of the transformations. More recently, the harder problem of analyzing precessing systems has required generalizing phase rotations to include the full three-dimensional rotation group, which induces slightly more complicated—though

still well understood—transformations of the waveforms [5–8].

Because of the essential diffeomorphism invariance of general relativity, it might seem that the natural end point of this progression would include all possible gauge transformations. This would be problematic, to say the least, because accounting for the effects of arbitrary diffeomorphisms on a waveform would be intractable. Fortunately, by making certain standard approximations, we can avoid accounting for the *complete* diffeomorphism freedom, and restrict to a smaller gauge group. The end result will be somewhat larger than the familiar Poincaré group—in fact infinitely so, at least in principle—yet entirely tractable and far smaller than the diffeomorphism group.

To see how this is possible, we must first note that near-field effects in the waveforms (effects appearing at second order in the distance between the emitter and observer) should be quite small in data collected in the vicinity of Earth, because even the leading-order waveform will be hard to detect. Thus, the model waveforms only need to capture asymptotic features of the radiation far from the source. In particular, we assume that the model spacetime is asymptotically flat, and calculate the asymptotic waveform in the limit of future null infinity, \mathcal{I}^+ , which is described below. Though it is not believed that our Universe is asymptotically flat, this is a useful construction approximating an isolated source when the intervening curvature is typically small. The gravitational-wave signal observed by a detector in the vicinity of Earth will then be very well approximated by the waves along some geodesic of \mathcal{I}^+ —up to a scaling related to distance from the source. The

benefit of assuming such an asymptotic structure is that it allows us to impose certain conditions on the gauge, the most common of which is called Bondi gauge [9–14].

Essentially, Bondi gauge consists of a special class of coordinates that manifest the asymptotic behavior of the spacetime such that the metric and its derivatives, when expressed in these coordinates, approach those of Minkowski spacetime at large radii. The allowed gauge transformations are symmetry transformations of this metric, which form a group known as the Bondi-Metzner-Sachs (BMS) group [9,10,12,15–18]. This group simply extends the Poincaré group with generalized translations. Bondi coordinates exist in a neighborhood of \mathcal{I}^+ for any asymptotically flat system [12], and any two Bondi coordinate systems are related by some BMS transformation [10]. This means that the BMS group encompasses all possible gauge transformations we need to be concerned with when discussing the limits of an asymptotically flat spacetime.

Bondi gauge also has a particularly nice feature related to the inertial observers in a neighborhood of \mathcal{I}^+ . At very large radii, curves of constant spatial coordinates parametrized by the retarded-time coordinate are nearly time-like geodesics—becoming more exactly geodesic at larger radii. Thus, if we extract a quantity on \mathcal{I}^+ along a simple curve of constant spatial coordinate, we approximate the signal an inertial observer measures as a function of proper time (up to the usual amplitude scaling with radius). Moreover, in the approximately flat asymptotic region, any two inertial observers are related by an element of the Poincaré group. But that is a subgroup of the BMS group, so we can use the BMS group to easily generate all possible signals that might be measured by any inertial observer.

Taken together, these facts mean that Bondi gauge is not only sufficiently general to describe any signal observed at great distance from a source in the asymptotically flat approximation, but is also a convenient choice that allows us to construct waveforms using simple curves and BMS transformations. We will therefore assume that any waveform is expressed in Bondi gauge, and narrow our focus to the BMS group. These concepts are reviewed pedagogically in Sec. II; the eager reader may prefer to skim that section, and simply refer to Eq. (8) for the key expressions describing the BMS transformations.

Having understood the BMS group itself, we will then need to understand its effects on waveforms, which can be separated into two parts. First is the effect at a single spacetime event. The transformation changes the differential structure in a neighborhood of that event, and since the gravitational field is fundamentally a measure of that differential structure, it should come as no surprise that the waveform will change under a transformation. To make these ideas more precise, however, we will need a careful treatment of asymptotic flatness. Section III will review a convenient formalism for developing asymptotic flatness,

then use that formalism to calculate the transformation properties of a waveform at a point.

Of course, a waveform is not simply measured at a single spacetime event: a gravitational-wave detector will measure it along some worldline, whereas model waveforms are typically expressed over an extended portion of the (future) celestial sphere of the source, as a function of time. In practical terms, this means expressing the waveform as a function of some coordinate system, and that coordinate system also changes under a BMS transformation. So the second part of a BMS transformation involves rewriting the waveform as a function of these new coordinates. This is a fairly simple bookkeeping exercise in principle, but involves numerous delicate manipulations and various minor subtleties for practical implementation, as discussed in Sec. IV.

Portions of the BMS group have been discussed previously in the context of transforming gravitational waveforms produced by numerical simulations. Gualtieri *et al.* [19] considered rotations and boosts, neglecting quantities of order v^2/c^2 . Kelly and Baker [20] looked at the effect of supertranslations on ringdown modes, to first order in the time derivative of the waveform. This paper, however, presents an exact algorithm for the full BMS group. The practical implementation of the algorithm is only limited by numerical precision and the accuracy of interpolation of the input waveform as a function of time.

Section V will give a brief overview of the size of these effects, for various types of BMS transformations and simple waveforms. Basic analytical arguments will show that the leading-order coupling due to a supertranslation will be proportional to the size of the translation and the dominant frequency of the coupled mode (or generally, the mode’s logarithmic derivative); for boosts the leading-order coupling will be proportional to the speed of the boost. In both cases, the constants of proportionality are typically of order 1, though there are various geometric factors involved.

In Sec. VI, mode coupling will be demonstrated for a full waveform from a numerical simulation of a binary black-hole system in the public waveform catalog maintained by the SXS collaboration [21,22]. This example system is chosen for its seeming symmetry, being equal mass and nonprecessing, though with a spin on one black hole aligned with the orbital angular velocity. Indeed, simple coordinate-based measures suggest that the center of mass only strays from the coordinate origin by only about $0.1M$ over the course of the simulation (where M is the total mass of the system). Nonetheless, we will find that near merger more than 1% of each mode—most notably the dominant (2,2) mode—will mix into other modes. In fact, in the raw data the (3,3) and (3,1) modes are completely dominated by power leaking in from the (2,2) mode. Overall, the third-largest mode in the data, (2,1), is comprised of leaked power by up to 30%. These couplings give rise to curious features in the smaller modes that are not present in the

post-Newtonian model of this system, for example. The mode couplings, and resulting curious features, can be dramatically decreased by measuring the motion of the center of mass (in simulation coordinates), and applying the opposite transformation to the waveform.

The example system was chosen so that we will be able to see clearly that the unexpected features are removed. More complicated systems—in particular, precessing systems—will have more complicated waveforms, but also larger anomalous motion of the center of mass. A survey of the entire SXS catalog suggests that the center of mass in more complicated simulations will drift from the origin by larger amounts, up to $8M$ for the most extreme system. This implies correspondingly larger mode couplings for these systems. Ossokine *et al.* [23] showed that it is possible to greatly reduce the size of these displacements by adjusting the initial data, improving the outlook for future simulations. Nonetheless, the current SXS waveform catalog must still be adjusted, and any recoil that develops during future evolutions will need to be accounted for. Moreover, boosts and translations only account for six of the infinitely many degrees of freedom in the BMS group; the general supertranslations in particular are still uncontrolled, even after eliminating the drift of the center of mass.

Given the amount of work that needs to be done to account for gauge effects in numerical waveforms, it is reasonable to wonder how this affects searches for gravitational waves in detector data. To understand this issue we will need to know more about BMS transformations of waveforms, and so we delay the full discussion until the end of the paper, Sec. VII. The upshot is that, while we must take BMS transformations into account when constructing waveforms, we do not need to search over all waveforms generated by the BMS group. For an isolated observer, the gauge ambiguities reduce to time translation and Lorentz rotations, which are already known.

Open-source code, in the form of a Python module SCRI, is supplied in the supplemental materials along with this paper [24]. It implements the BMS transformations of the most common gravitational waveforms, including the Newman-Penrose quantity ψ_4 , the Bondi news function, the shear spin coefficient σ , and the transverse-traceless metric perturbation h —as well as the remaining Newman-Penrose quantities ψ_0 through ψ_3 . Several Appendices describe details about various constructions and calculations from geometric algebra [25–27] used in this paper and in the SCRI module, including the method of implementing a boost of the sphere described in Sec. II A. A final Appendix details the crude method of measuring and removing the center-of-mass drift found in the numerical data, which is used in Sec. VI.

A. A note on conventions

More extensive descriptions of the conventions used in this paper are given in the Appendices, but a few basic

comments are appropriate here. We will assume that all transformations are proper and orthochronous; both spatial orientation and the direction of time are preserved. In general, a strictly improper transformation can be written as the product of a proper transformation and a parity operator. The effects of parity operations on modes of a spin-weighted spherical-harmonic decomposition are described in Appendix B of Ref. [8]. Similarly, anachronous transformations can be written as the product of an orthochronous transformation and the time-reversal operator. The effect of the time-reversal operator is to simply negate the time coordinate, and in some cases to change the sign and labeling of waveform quantities. Since these can be dealt with separately in ways that are already understood, we dispense with them entirely, and will not bother to repeat below that all transformations in this paper are proper and orthochronous.

Points on the sphere will be labeled interchangeably by the usual spherical coordinates (θ, ϕ) , by the standard stereographic coordinate ζ , or by the unit spatial vector \mathbf{r} pointing in that direction from the origin when the sphere is considered as being embedded in Euclidean 3-space and centered on the origin. While stereographic coordinates are the preferred representation throughout much of the literature, and do occasionally simplify theoretical calculations, they are unsuited for practical computations because of their infinite range and the nature of the point at infinity. Perhaps surprisingly, (θ, ϕ) presents a more useful parametrization for practical applications. Despite the coordinate singularities in its representation of the sphere S^2 , it actually provides a *nonsingular* parametrization of a portion of the rotation group that covers S^2 , which is more relevant for dealing with spin-weighted functions. As a result, software packages such as SPINFAST [28,29] that implement numerical routines involving spin-weighted spherical harmonics use the (θ, ϕ) representation. Consequently, the same representation is used by the SCRI package accompanying this paper.

II. THE BMS GROUP

We introduce the Bondi-Metzner-Sachs (BMS) group in a simple and familiar setting. This will provide a common basis and useful motivation for the coming sections. Though more of the formalism of asymptotically flat spacetime will be needed below, for our purposes in this section, it will be sufficient to consider the standard compactification of Minkowski space. In particular, this compactification will provide all the understanding needed for more general asymptotically flat spacetimes. We will use null rays to relate coordinates of timelike geodesics at finite radius to coordinates at future null infinity. By relating timelike geodesics to each other, we will then be able to describe the effect of a BMS transformation on the coordinates of future null infinity.




FIG. 1. Extending local coordinates to \mathcal{I}^+ . This conformal diagram shows the worldlines of a pair of inertial emitters, with emitter \mathcal{B} moving at speed $0.5c$ relative to emitter \mathcal{A} , and a distant observer \mathcal{O} stationary with respect to \mathcal{A} . The origins of the emitters' coordinate systems coincide at $\tau_{\mathcal{A}} = \tau_{\mathcal{B}} = 0$. We construct the null cone \mathcal{N} emanating from that event, which allows us to extend coordinates to \mathcal{I}^+ . The intersection of \mathcal{N} with \mathcal{I}^+ is a sphere; all points on that sphere are assigned time coordinate $u_{\mathcal{A}} = 0$ by emitter \mathcal{A} and $u_{\mathcal{B}} = 0$ by emitter \mathcal{B} . Each point is also labeled by the direction of the null generator extending from the given emitter to that point. Note that a rotation obviously does not affect the set of points comprising \mathcal{N} , though the labeling of points will change—except for the points along the axis of rotation. Similarly, a boost leaves the null cone invariant, but will change the labeling of any point not along the boost velocity vector, as discussed in Sec. II A. The observer \mathcal{O} can also be assigned coordinates based on the null rays emitted by \mathcal{A} and, in the limit of very large separation, any field it observes will approach the field observed on \mathcal{I}^+ , up to a scaling based on radius. This is the basic motivation for using asymptotically flat spacetimes to model radiation.

The conformal diagram displayed in Fig. 1 provides the standard picture [17,30,31]. Here, i^- and i^+ are past and future timelike infinity; i^0 is spacelike infinity; \mathcal{I}^- is past null infinity; and \mathcal{I}^+ is future null infinity. We will be concerned almost exclusively with \mathcal{I}^+ , as that is the asymptotic limit of outgoing gravitational radiation. Also shown in the diagram is a pair of emitters, \mathcal{A} and \mathcal{B} , traveling along timelike geodesics. We can extend coordinates defined in a neighborhood of an inertial emitter to coordinates throughout the spacetime and to \mathcal{I}^+ using null generators. For example, suppose \mathcal{A} emits a null ray at proper time $\tau_{\mathcal{A}} = 0$ in a direction given in local coordinates by the angular coordinates (θ, ϕ) , or equivalently the stereographic coordinate ζ . Any point at finite distance along that ray can be assigned coordinates (u, r, ζ) , where $u = \tau_{\mathcal{A}}$ is the retarded time and r is an affine parameter along the geodesic—in Minkowski spacetime, we can think

of this as the distance between the emitter and that point as measured in the frame of the emitter. The future limit of the null ray will represent a unique point on \mathcal{I}^+ ; we typically assign that point the coordinates (u, ζ) , dropping r because it will, of course, be infinite. Continuing in this way for all directions, emitter \mathcal{A} can provide coordinates for the entire null cone \mathcal{N} and, in particular, the sphere \mathcal{S}^+ given by the intersection of \mathcal{N} with \mathcal{I}^+ .

Of course, coordinates can equivalently be constructed in the same way by emitter \mathcal{B} . The set of points \mathcal{S}^+ will naturally be the same in both cases; any relative rotation of the two emitters will simply take one null ray into another, and a boost leaves null rays invariant. But the coordinates labeling each null ray—hence the coordinates labeling each point on \mathcal{S}^+ —will be different for the two systems whenever \mathcal{A} and \mathcal{B} are related by any Lorentz transformation. We discuss the effect of these transformations in Sec. II A.

But the Lorentz transformations only relate a subset of possible emitters, and hence a subset of possible coordinate systems on \mathcal{I}^+ . In Minkowski spacetime, we are familiar with translations as the remaining freedom relating coordinate systems. However, a translation at finite radius has somewhat surprising effects on the coordinates at \mathcal{I}^+ . This is explained more fully in Sec. II B. The conclusion will be that a translation is equivalent to an offset of the retarded time u that depends on direction, though in a simple way.

The surprising result of early studies [9,10] was that the familiar Lorentz and translation groups are not sufficient for describing all of the asymptotic symmetries of asymptotically flat spacetimes. It turns out that a focus on null cones in Minkowski spacetime is too restrictive. Since we are only prescribing the asymptotic behavior, we can disregard all but a neighborhood of \mathcal{I}^+ —which means that our null “cones” need no longer look like cones, in the sense that the generators need not meet at a point. In general, then, the simple angular dependence of the offset of the retarded time u induced by translations must be generalized to an arbitrary (smooth) function of the angles. These transformations are referred to as *supertranslations*. Together with the Lorentz group,¹ these form the complete BMS group, as discussed further in Sec. II C.

A. Rotations and boosts

We first confine ourselves to a single null cone, and the corresponding sphere \mathcal{S}^+ on \mathcal{I}^+ , by considering emitters with coordinate systems having identical origins but which are related by elements of the Lorentz subgroup: rotations and boosts. The situation is depicted in Fig. 1, where two

¹To be precise, the supertranslations form a normal subgroup \mathbf{T} of the BMS group, and the factor group of the BMS group by \mathbf{T} is precisely the Lorentz group $\text{SO}^+(3, 1)$. However, the Lorentz group is not a *normal* subgroup. Thus, the BMS group is the semidirect product $\mathbf{T} \rtimes \text{SO}^+(3, 1)$.

emitters give off null rays at the same spacetime event. It is, of course, one of the fundamental conclusions of basic special relativity that boosts take null rays to null rays. Thus, the collection of all null rays (the null cone) originating at a particular spacetime event will be invariant under boosts. However, the coordinates assigned to the direction of a given null ray within that collection will change under a boost. In the same way, a rotation maps the null cone onto itself, while simply changing the coordinates of individual rays. Our objective in this section, then, is to find how directions in one coordinate system map to directions in another system under rotations and boosts.

We begin with the simpler case, in which the coordinate systems of \mathcal{A} and \mathcal{B} are simply related by some known rotation, with no relative boost. For a scalar field, if we suppose that the field is known in frame \mathcal{A} , we can find the value of the field at any point \mathbf{r}_B in frame \mathcal{B} by simply rotating that point back to \mathbf{r}_A in \mathcal{A} and evaluating the field there. However, for spin-weighted fields, there is an additional complication. A spin-weighted field at a point is defined with respect to the basis of the tangent space to the sphere at that point—usually represented by a complex tangent vector \mathbf{m} . But if \mathcal{A} and \mathcal{B} are rotated relative to one another, the tangent vector \mathbf{m}_A at \mathbf{r}_A will also be rotated relative to the tangent vector \mathbf{m}_B at \mathbf{r}_B by some angle λ , referred to as the “spin phase.” (This factor λ is defined more precisely in Appendix C.) This situation is depicted in Fig. 2, where a standard grid is shown in frame \mathcal{B} on the left-hand side, and a grid representing the same physical points in the coordinates of \mathcal{A} is shown in the center, along with the spin phase.

A simple way of dealing with the complication of the spin phase is described in Appendix B of Ref. [8]. The

essential idea is to evaluate spin-weighted fields directly in terms of a rotation operator. Thus, if the field value is needed at (θ, ϕ) in \mathcal{B} , this is represented by a rotation operator $R_{\theta, \phi}$, described in Appendix B (of this paper). Now, if frame \mathcal{B} is obtained from frame \mathcal{A} by a frame-rotation R_f , then the value of the field can be found by evaluating the field in \mathcal{A} at $R_f R_{\theta, \phi}$. The spin phase is automatically accounted for.

Similarly, we can find the value of the field in \mathcal{B} if it is related to \mathcal{A} by a pure boost. We assume that \mathcal{B} moves with respect to \mathcal{A} with three-velocity \mathbf{v} , and use the conventions established in the Appendices to directly compare components in the two frames. Suppose that \mathcal{A} measures an angle Θ_A between \mathbf{v} and the spatial component, \mathbf{r} of some null direction. That is, we have $\cos \Theta_A = \mathbf{v} \cdot \mathbf{r} / |\mathbf{v}| |\mathbf{r}|$. Similarly \mathcal{B} measures an angle Θ_B between \mathbf{v} and the spatial component of that same null direction. Note that the spatial subspaces will, of course, generally be different for the two frames, except along the axis containing \mathbf{v} . Nonetheless, we can relate the angles measured in the two frames, as shown in Appendix C, by the formula

$$\tan \frac{\Theta_B}{2} = e^\varphi \tan \frac{\Theta_A}{2}, \tag{1}$$

where $\varphi = \operatorname{arctanh} |\mathbf{v}|$ is the usual rapidity parameter.

We can use this equation to transform a physical scalar field measured in one frame into the other frame. Suppose that this physical field is known on \mathcal{S}^+ as a function of the null direction measured by \mathcal{A} , and we wish to know the value of the field in some null direction \mathbf{r}_B as measured by \mathcal{B} , noting that the angle between \mathbf{v} and \mathbf{r}_B is Θ_B . We first take the direction \mathbf{r}'_B in the frame of \mathcal{A} having the same




FIG. 2. Transformation of a grid under Lorentz rotations. To decompose a spin-weighted field into spin-weighted spherical harmonics, the values of the field are needed on the colatitude-longitude grid of frame \mathcal{B} seen on the left-hand side. If the field is known in frame \mathcal{A} , and \mathcal{B} is simply rotated relative to \mathcal{A} , then the appropriate values can be found by evaluating the field in \mathcal{A} on the grid shown in the center. The points on the two grids represent the same physical points. Similarly if \mathcal{B} is boosted relative to \mathcal{A} (in this example with velocity $0.5c$ to the right and out of the page), the values can be found by evaluating the field in \mathcal{A} on the grid shown on the right-hand side. However, for spin-weighted fields, the locations of the points alone are not sufficient; we also need to know the relative alignment angle between the tangent basis constructed by \mathcal{A} and the tangent basis constructed by \mathcal{B} at each point. This angle is the spin phase λ described in the text, represented here by the size and color of the marker at each grid point. The transformed grid positions and λ values are calculated using Eq. (2).

components with respect to the basis of \mathcal{A} as \mathbf{r}_B has with respect to the basis of \mathcal{B} , even though this is a different frame. We then rotate this vector in the \mathbf{v} - \mathbf{r}'_B plane until we arrive at a new vector \mathbf{r}_A that makes an angle with \mathbf{v} of Θ_A , satisfying Eq. (1). The physical field measured at \mathcal{I}^+ by \mathcal{A} in this direction is the same as the physical field as measured by \mathcal{B} , and is thus the result we sought.

Again, there are complications involved with spin-weighted fields. However, as shown in Appendix C, these complications are automatically dealt with when using the rotation-operator approach described above. The scenario is illustrated in Fig. 2, where the grid in \mathcal{B} is shown on the left-hand side, and the same grid of physical points is shown in the coordinate system of \mathcal{A} on the right-hand side. The basic idea is the same: the grid points are simply moved around the sphere and associated with some spin phase λ . Of course, in this case, the points are moved in different ways, and λ is a different function of position. Nonetheless, it is still beneficial to evaluate the field in \mathcal{A} directly in terms of the rotation operator. Here, however, rather than the constant frame-rotation operator R_f accounting for the difference between frames \mathcal{A} and \mathcal{B} , we need to use a position-dependent rotation operator B' . Using the notation of quaternions, we can write this operator as

$$B'(\theta, \phi) = \exp \left[\frac{\Theta_B - \Theta_A}{2} \frac{\mathbf{r}'_B \times \mathbf{v}}{|\mathbf{r}'_B \times \mathbf{v}|} \right], \quad (2)$$

where \mathbf{r}'_B is the unit spatial vector in the (θ, ϕ) direction of frame \mathcal{A} , Θ_B is the angle between that vector and \mathbf{v} , and Θ_A is related to it by Eq. (1). This operator represents a rotation through $\Theta_B - \Theta_A$ about the $\mathbf{r}'_B \times \mathbf{v}$ axis. In this case, the field can be evaluated in \mathcal{A} from $B'R_{\theta, \phi}$. More generally, an arbitrary element of the Lorentz group can be written as the product of a frame rotation and a boost, $B'R_f$, then the field can be evaluated from $B'R_f R_{\theta, \phi}$.

It is worth exhibiting the effect of a Lorentz transformation in terms of stereographic coordinates. Though they are ill suited to actual computations involving data, there are certain advantages to using the stereographic formalism for theoretical calculations. In particular, if the sphere S^+ is parametrized by the stereographic coordinate ζ_A in frame \mathcal{A} and by ζ_B in frame \mathcal{B} , then under a general Lorentz transformation the two are related by²

$$\zeta_B = \frac{a\zeta_A + b}{c\zeta_A + d}, \quad (3)$$

²The stereographic coordinates are usually thought of as elements of the complex plane augmented by adjoining the point at infinity, also known as the Riemann sphere. In this form, the transformation shown here is usually known as a Möbius transformation—an element of the Möbius group, which is isomorphic to the group of conformal transformations of the sphere, the projective special linear group $\text{PSL}(2, \mathbb{C})$, and the proper orthochronous Lorentz group $\text{SO}^+(3, 1) \cong \text{SO}^+(1, 3)$.

where (a, b, c, d) is a collection of complex coefficients satisfying $ad - bc = 1$. Because of its compactness, the representation in stereographic coordinates is useful for descriptions, and occasionally for deriving results. We will encounter this formalism again in Sec. III, though the stereographic coordinates themselves will not appear in the final results. For all other purposes, quaternions and related formalism will be used because of their computational superiority.

Finally, we also note that this transformation of a spin-weighted field under Lorentz transformations is only part of the story. More generally, different fields will mix with each other because \mathbf{m} remains neither tangent to the sphere nor even purely spatial under a Lorentz transformation. For example, to calculate the transformation of the Newman-Penrose quantity ψ_3 on \mathcal{I}^+ we will also need a contribution from ψ_4 . That simple behavior is only a result of the peeling theorem; at finite radii all Newman-Penrose quantities could mix with each other under a Lorentz transformation. This will be discussed further in Sec. III B. Throughout the remainder of this section, however, we will be able to focus solely on the movement of points at which a field is evaluated.

B. Translations and supertranslations

Now, having understood the transformations that preserve the light cone, we can move on to more general transformations—though still considering only inertial emitters in Minkowski space. In particular, we have translations of both time and space. Generalizing these, we will be led to the encompassing notion of supertranslations.

It is instructive to begin with the simple case of time translations. As noted in the introduction to this section, every point on a null cone originating at emitter \mathcal{A} , at a proper time of τ_A , is assigned the same retarded time $u_A = \tau_A$; similarly $u_B = \tau_B$. Now, if the emitters' time scales are related by a simple time translation such that $\tau_B = \tau_A - \delta t$, we clearly have the simple relation between retarded-time coordinates $u_B = u_A - \delta t$. This is depicted in Fig. 3. The notable feature of this transformation is that it is isotropic; the change in the retarded-time coordinate does not depend on the direction. This seemingly trivial observation is important because it is not true of space translations, and generalizing this notion will be key to understanding the broader class of supertranslations.

We can now consider space translations as depicted in Fig. 4. Emitter \mathcal{B} is simply displaced from \mathcal{A} by a spatial vector $\delta \mathbf{x}$ but the two are stationary with respect to each other. The null cone \mathcal{N}_B emanates from the origin of \mathcal{B} , $u_B = 0$, and intersects \mathcal{I}^+ at two points on this diagram. Those same points of \mathcal{I}^+ are on null rays from two separate null cones of \mathcal{A} —one in the $-\mathbf{x}$ direction with retarded time $u_{A1} = -|\delta \mathbf{x}|$, the other in the \mathbf{x} direction emitted at $u_{A2} = |\delta \mathbf{x}|$. Of course, these two points correspond to the two points of the sphere S^0 because the three spatial




FIG. 3. Effect of time translation on coordinates of \mathcal{I}^+ . Here, we see two different local coordinate systems extended to \mathcal{I}^+ : \mathcal{A} and \mathcal{B} represent the same emitter with the same spatial coordinates, but different origins for the time coordinate. The two null cones correspond to the two origins of the time coordinate. We see that the time translation $\tau_B = \tau_A - \delta\tau$ affects the time coordinates on \mathcal{I}^+ isotropically—in fact, the transformation of the retarded-time coordinate is simply $u_B = u_A - \delta\tau$.




FIG. 4. Effect of space translation on coordinates of \mathcal{I}^+ . Here, \mathcal{A} and \mathcal{B} represent emitters displaced relative to each other. A single null cone emanates from \mathcal{B} and intersects \mathcal{I}^+ in two points. The same points of \mathcal{I}^+ are found on two separate null cones emitted by \mathcal{A} . Thus, a space translation has a nonisotropic effect on the retarded time coordinates of \mathcal{I}^+ . More generally, allowing for all three spatial dimensions, the effect of a translation $\delta\mathbf{x}$ will transform the retarded-time coordinate in a direction \mathbf{r} as $u_B = u_A + \delta\mathbf{x} \cdot \mathbf{r}$.

dimensions have been collapsed to one in this simple diagram. More generally, for any point on the sphere S^2 the relationship between the retarded time coordinates is $u_B = u_A + \delta\mathbf{x} \cdot \mathbf{r}$.

We can combine these two transformation laws into a single law for general spacetime translations:

$$u_B = u_A - \sum_{\ell \in \{0,1\}} \sum_{m=-\ell}^{\ell} \alpha^{\ell,m} Y_{\ell,m}(\theta, \phi), \quad (4)$$

where

$$\alpha^{0,0} = \sqrt{4\pi} \delta t, \quad (5a)$$

$$\alpha^{1,-1} = -\sqrt{\frac{2\pi}{3}} (\delta x + i\delta y), \quad (5b)$$

$$\alpha^{1,0} = -\sqrt{\frac{4\pi}{3}} \delta z, \quad (5c)$$

$$\alpha^{1,1} = -\sqrt{\frac{2\pi}{3}} (-\delta x + i\delta y), \quad (5d)$$

using $\delta\mathbf{x} = (\delta x, \delta y, \delta z)$. Note that the sum over ℓ is restricted to $\{0, 1\}$ here. This suggests the final generalization we need to arrive at the BMS group: expanding the range of the sum over ℓ to all positive integers, while retaining the condition that $\alpha^{\ell,m} = (-1)^m \bar{\alpha}^{\ell,-m}$ to ensure that the retarded-time coordinate remains real. More precisely, we construct a transformation of the coordinates such that

$$u' = u - \alpha, \quad (6)$$

where α is any real-valued function on the sphere. To simplify later analyses, we can also add the conditions that α be square integrable and twice differentiable. This transformation—which encompasses spacetime translations—is referred to as a *supertranslation*. It can be shown that supertranslations are asymptotic symmetries of asymptotically flat spacetimes [9,10], and thus are indeed members of the BMS group.

One way of thinking about supertranslations is to imagine a network of observers located on a sphere surrounding the source. Ideally, we could combine the signals detected by these observers, but to do so we would need some idea of how their time coordinates compared to each other; we would need to have some synchronization between their clocks. But if we now move the network to \mathcal{I}^+ , such a synchronization becomes impossible. We could supply a separate time offset to each observer without changing the physics. Roughly speaking, a supertranslation is just the limit of this direction-dependent time translation

where there is a different observer in every possible direction.

Supertranslations present an interesting departure from the other, more basic, types of transformations constituting the familiar Poincaré group. If u is constructed as given above by light cones emitting from an inertial world line \mathcal{A} , then we know (by construction) that the null rays generating a surface of constant u meet in a common point—the vertex of the null cone. On the other hand, if the function α has any $\ell > 1$ components, the null rays generating a surface of constant u' , as given by Eq. (6), *do not* meet in a common point. This is why the notation changed in Eq. (6), dropping the subscripts denoting the emitter, because in general we do not require the retarded time to be constructed by an emitter.

As another, possibly more enlightening, consideration of this peculiar nature of supertranslations, we can imagine light cones originating at an emitter in an asymptotically flat spacetime containing some nontrivial geometry. In the example shown in Fig. 5, we see a simple cartoon of a




FIG. 5. Null rays in complicated spacetimes. When the interior of the spacetime is not Minkowski, we cannot expect to construct retarded time coordinates globally based on null cones. The interior of this diagram is a rough cartoon, in which the shaded region represents the space between two merging black holes. Some of the null generators from any emitter in this spacetime must pass through this region, and may be affected in erratic ways. Given null rays near \mathcal{I}^+ , we cannot say whether or not they originated at the same point. Clearly, then, it is too much to demand that general asymptotically flat spacetimes must have their coordinates given by a construction like the one given for Minkowski space. Instead, we simply place requirements on the compactified spacetime in a neighborhood of \mathcal{I}^+ . Conversely, it is too much to ask that a “nice” coordinate system on \mathcal{I}^+ correspond to null cones that meet at one spacetime event in general. This motivates our intuitive acceptance of supertranslations.

merging binary. The emitter \mathcal{A} gives off two null cones, \mathcal{N}_1 followed by \mathcal{N}_2 . The rays given off to the right intersect \mathcal{I}^+ as we would expect, \mathcal{N}_1 followed by \mathcal{N}_2 . The rays given off to the left, however, behave more erratically. Here, the first null ray interacts strongly with the black holes and is delayed, arriving at \mathcal{I}^+ *after* the null ray that was emitted later. Obviously, coordinates constructed from null cones of \mathcal{A} will be “bad” coordinates, with singularities resulting from caustics of the null rays.

So for general asymptotically flat spacetimes, it is simply a bad idea to expect that the retarded time coordinate should be constructible from null rays emitted from a timelike worldline. Instead, we should only expect to have “good” or “nice” coordinates in a neighborhood of \mathcal{I}^+ . In fact, the motivation for the original paper by Newman and Penrose that introduced the δ operator [16] was to impose a condition on u in a neighborhood of \mathcal{I}^+ to fix the $\ell > 1$ supertranslation freedom. This is also (at least partially) the motivation for the “good cut” construction [18], the “nice section” construction [32], and the “regularized null cone cut” construction [33].

C. The complete BMS group

One final element is needed to complete the construction of the BMS group. In Sec. II A, we assumed that the origins of the two emitters coincided, but only looked at the effect of a boost on the null cone emitted at that common origin. Obviously, at later times, the null cones by which those emitters extend their local coordinates to \mathcal{I}^+ will not originate at the same event; there will also be some translation involved. Thus, we expect that the simple formula from the previous section $u' = u - \alpha$ must be modified in some way by the boost.

The situation is easily described pictorially, as in Fig. 6. The origins of proper time for the observers coincide where their paths cross. At some later time τ_B , emitter \mathcal{B} constructs a null cone \mathcal{N}_B . A simple exercise in special relativity shows that \mathcal{A} must emit a null ray in direction \mathbf{r} at time $\tau_A = \gamma(1 - \mathbf{v} \cdot \mathbf{r})\tau_B$ in order to reach the same point of \mathcal{I}^+ as the null ray emitted in that direction by \mathcal{B} . Because we will encounter this factor frequently, we define

$$k := \frac{1}{\gamma(1 - \mathbf{v} \cdot \mathbf{r})}. \quad (7)$$

As shown in Appendix D, this k factor is also the conformal factor of a boost appropriate to the spherical metric. In this case, where the spatial origins coincide at $\tau_A = \tau_B = 0$, we have the transformation law $u' = ku$.

Finally, we can combine this with the supertranslation of Eq. (6) and the angular effects of the Lorentz transformation given by Eq. (3) to find the general BMS transformation of coordinates on \mathcal{I}^+ , representing an initial supertranslation, followed by a Lorentz transformation:




FIG. 6. Effect of boost on coordinates of \mathcal{I}^+ with $u \neq 0$. Here, \mathcal{A} and \mathcal{B} represent emitters related by a simple boost. Though their origins coincide, at some later time τ_B emitter \mathcal{B} constructs the null cone \mathcal{N}_B . It is not hard to solve for the time $\tau_{A1} = \gamma(1 - v)\tau_B$ at which emitter \mathcal{A} must construct a null cone to overlap with the right-going null ray of \mathcal{N}_B . We can similarly solve for the time $\tau_{A2} = \gamma(1 + v)\tau_B$ at which emitter \mathcal{A} must construct a null cone to overlap with the left-going null ray of \mathcal{N}_B .

$$u' = k(u - \alpha) \quad (8a)$$

$$\zeta' = \frac{a\zeta + b}{c\zeta + d}. \quad (8b)$$

Again, (a, b, c, d) is a collection of complex coefficients satisfying $ad - bc = 1$, representing the Lorentz transformation, and α is an arbitrary real-valued square-integrable and twice-differentiable function on the sphere. The implementation of these transformations to be described below will use Eqs. (1) and (2) to represent Lorentz transformations rather than the stereographic coordinates shown here.

It is also important to note that the transformation is constant; a, b, c, d, k , and α are all independent of time. This may seem to give us a static transformation—though we know that a boost should, in some sense, result in a time-dependent translation. To simplify matters, we assume no rotation and $\alpha = 0$, leaving only a boost. The transformation law for time in this case might be rewritten as $\gamma u' = u + \gamma \mathbf{v} \cdot \mathbf{r}u'$. We can interpret this as a rescaling of the time coordinate, in agreement with the standard time dilation, along with a translation by $\gamma u' \mathbf{v}$, much as we might expect. Interestingly, it is awkward to express this translation as being proportional to u [e.g., by expanding the factor of k as spherical harmonics in Eq. (8a)], because this would imply that a boost gives

rise to a time-dependent *supertranslation*. This suggests a minor subtlety of nomenclature when defining the supertranslation, due to noncommutativity of the boost and supertranslation.

In this section, we have built up the BMS group through heuristic arguments in order to come to an intuitive and pedagogical understanding of how coordinates change under a BMS transformation—though of course, the same result is also obtained through more rigorous methods [9,10,15]. In particular, Sec. IV C of Ref. [15] describes the associated Lie algebra \mathfrak{bms} . In short, the rotations and boosts correspond to the standard generators of infinitesimal (Lorentz) rotations in a plane of Minkowski space, while the generators of supertranslations are given by the basis $Y_{\ell,m} \partial / \partial u$. We will not find this infinitesimal presentation directly useful, however, because it is not easily applicable to finite transformations. Moreover, we will see in Sec. III B that these operators only account for the change in how coordinates label points, but not changes in the waveforms themselves.

III. ASYMPTOTIC FLATNESS AND TRANSFORMATIONS AT A POINT

Now, having seen the effects of the BMS transformation on *coordinates* on \mathcal{I}^+ , we need to understand the effects on *waveforms* measured at \mathcal{I}^+ . We begin, in this section, by examining the effect on the waveform at a single point, where the transformation leaves that point fixed. This will be extended in Sec. IV by allowing the point to vary, which will involve the relatively simple task of evaluating the known function at different points—in practice, requiring mostly interpolation and other bookkeeping.

Though we will not yet vary the coordinates of our selected point, the coordinates of *neighboring* points will change. For our purposes, a waveform measures some piece of the differential structure of spacetime. But waveforms are not true scalars, in the sense that they are not invariant under coordinate transformations—in fact, they are inherently defined with respect to coordinates. More precisely, the tetrad with respect to which they are constructed is *defined* in terms of coordinates. It is, of course, possible to perform a coordinate transformation while leaving the tetrad fixed. But this is not relevant; waveforms expressed in different coordinate systems use different tetrads. Therefore, a BMS transformation that changes the coordinates of *nearby* points should also change the waveform *at* the given point.

To make these ideas precise, we need to be more specific about our representation of \mathcal{I}^+ , and the spacetime in a neighborhood of \mathcal{I}^+ . It will then be a relatively simple matter to calculate the transformations of standard curvature quantities. The reader who is willing to take these results on faith may simply refer to Eqs. (17), (19), (21), and (23), and otherwise skip this section.

A. Asymptotically flat spacetime

Numerous formulations describe the basic idea of asymptotic flatness, most prominently developed by Penrose [30]. For definiteness, we will follow the development by Moreschi [12,34]. The essential idea is to begin with a physical spacetime (M, Γ_{ab}) , and identify it with a portion of a model spacetime (m, γ_{ab}) representing the asymptotic completion of the physical spacetime. Here and in the following, to simplify notation, quantities in the physical spacetime will be represented by uppercase characters, while quantities in the asymptotic spacetime will be represented by lowercase.³

We begin with the physical spacetime (M, Γ_{ab}) , which has Weyl spinor (the spinor form of the standard Weyl tensor) Ψ_{ABCD} . We impose the assumption of (future) asymptotic flatness by requiring the existence of another spacetime (m, γ_{ab}) , with boundary \mathcal{I}^+ such that as topological spaces $\mathcal{I}^+ = S^2 \times \mathbb{R}$ and $M = m \setminus \mathcal{I}^+$. In particular, for any point $P \in M$ we have a point identified as $P \in m$, so that we can interchangeably describe any function at a point not on \mathcal{I}^+ as being defined either on M or m . We further assume the existence of a real-valued function ω that is continuous on m and smooth on M , and satisfies the following conditions:

- (1) $\omega|_M > 0$.
- (2) $\omega|_{\mathcal{I}^+} = 0$.
- (3) $d\omega|_{\mathcal{I}^+} \neq 0$.

Given this function, the spacetimes are also required to obey the following conditions:

- (4) $\gamma_{ab}|_M = \Gamma_{ab}\omega^2|_M$.
- (5) At every point of \mathcal{I}^+ , there ends a future-directed null geodesic of m .
- (6) In some neighborhood of \mathcal{I}^+ , there exist quantities \hat{r} and \tilde{r} on m such that the Riemann tensor of (M, Γ_{ab}) satisfies

$$R_{abc}{}^d = f(\omega)\hat{r}_{abc}{}^d + \tilde{r}_{abc}{}^d, \quad (9)$$

where

- (a) $df/d\omega > 0$,
- (b) $\lim_{\omega \rightarrow 0} f = 0$,
- (c) \hat{r} is regular at \mathcal{I}^+ , and
- (d) \tilde{r} goes to zero faster than f as $\omega \rightarrow 0$.

The last condition is to be understood componentwise, with respect to an orthogonal tetrad of (m, γ_{ab}) that is regular at \mathcal{I}^+ , like the one constructed below in Eq. (12). We also define a spinor on M by $\psi_{ABCD} := \omega^{-1}\Psi_{ABCD}$, which we can extend continuously to \mathcal{I}^+ . Note, however, that \hat{r} and ψ need not be the Riemann tensor and Weyl spinor of (m, γ_{ab}) .

It is possible [10,12,15] to choose coordinates (u, θ, ϕ) on \mathcal{I}^+ , where u labels a slice of \mathcal{I}^+ with topology S^2 and (θ, ϕ) are the standard coordinates of the unit sphere. The

³Indices, of course, will not be included in this distinction. Instead, lowercase indices will denote tensor indices, while uppercase indices will denote spinor indices, as usual.

latter are frequently expressed—at least for theoretical work—as the usual stereographic coordinate ζ and its complex conjugate $\bar{\zeta}$ [35]. These coordinates can be extended into a neighborhood of \mathcal{I}^+ by taking ω as an additional coordinate along future-directed null geodesics, where (u, θ, ϕ) labels the geodesics. Moreschi [12] showed that, up to irrelevant gauge freedom, the ω function is related to the luminosity distance r_L by

$$\omega = \frac{1}{r_L} + \mathcal{O}\left(\frac{1}{r_L^3}\right). \quad (10)$$

These coordinates are essentially what are known as Bondi coordinates, and allow the metric to be put in a particularly simple form [9,10]. This form of the metric is asymptotically invariant under BMS transformations.

An orthonormal spin dyad [11,35–38] (O^A, I^A) and its asymptotic counterpart (o^A, i^A) can also be defined related to these coordinates, such that we have orthonormal tetrads

$$l^a := \sigma_{AA'}^a o^A o^{A'}, \quad L^a := \sigma_{AA'}^a O^A O^{A'}, \quad (11a)$$

$$m^a := \sigma_{AA'}^a o^A i^{A'}, \quad M^a := \sigma_{AA'}^a O^A I^{A'}, \quad (11b)$$

$$\bar{m}^a := \sigma_{AA'}^a i^A o^{A'}, \quad \bar{M}^a := \sigma_{AA'}^a I^A O^{A'}, \quad (11c)$$

$$n^a := \sigma_{AA'}^a i^A i^{A'}, \quad N^a := \sigma_{AA'}^a I^A I^{A'}, \quad (11d)$$

where $\sigma_{AA'}^a$ are the Infeld-van der Waerden symbols, and at leading order in ω we have

$$l_a = (du)_a \simeq L_a, \quad (12a)$$

$$m_a = -\frac{\sqrt{2}}{1 + \zeta\bar{\zeta}}(d\bar{\zeta})_a \simeq \omega M_a, \quad (12b)$$

$$\bar{m}_a = -\frac{\sqrt{2}}{1 + \zeta\bar{\zeta}}(d\zeta)_a \simeq \omega \bar{M}_a, \quad (12c)$$

$$n_a = -(d\omega)_a \simeq \omega^2 N_a. \quad (12d)$$

We also denote by

$$\mathfrak{D} \simeq \omega \delta \quad (13)$$

the spin-raising differential operator introduced (at finite radius as δ) by Geroch, Held, and Penrose [37].⁴

This completes the basic framework we use to describe asymptotically flat spacetimes, allowing us to understand the asymptotic behavior of the physical fields. Next, we

⁴Note that the operator δ_{NP} originally introduced by Newman and Penrose [16] is generally different from the operator $\delta_{\text{GHP}} \equiv \mathfrak{D}$ introduced by Geroch, Held, and Penrose, in that only the latter has well-defined transformation behavior under boosts. There is also a discrepancy in the normalization such that $\delta_{\text{NP}} = \sqrt{2}\delta_{\text{GHP}}$ for scalar functions.

will show how a BMS transformation alters this framework, and use that result to find the changes in curvature quantities expressed within this framework.

B. Transformations

Equation (8) describes the general BMS transformation. However, this transformation is only defined *on* \mathcal{I}^+ . Because the curvature quantities we are interested in measure the differential structure of spacetime, understanding those quantities requires understanding the transformation *in a neighborhood of* \mathcal{I}^+ . Moreschi [34] found the general transformation to first order in ω that preserves the leading-order Bondi form of the metric:

$$\check{u} = k(u - \alpha) - \omega \frac{\check{\delta}u' \bar{\delta}u'}{k}, \quad (14a)$$

$$\check{\zeta} = \frac{a\zeta + b}{c\zeta + d} - \omega \frac{\check{\delta}u' \bar{\delta}\zeta' + \check{\delta}\zeta' \bar{\delta}u'}{k}, \quad (14b)$$

$$\check{\omega} = k\omega. \quad (14c)$$

Here, u' and ζ' are the leading-order terms in their respective equations—also given by the standard BMS transformation of Eq. (8). This transformation is defined in a neighborhood of \mathcal{I}^+ , so we can evaluate the differentials in Eq. (12) and take the limit as $\omega \rightarrow 0$, to find the transformation laws for the tetrad and infer the effects on the spinor basis:

$$o'^A = \frac{e^{i\lambda/2}}{\sqrt{k}} \left(o^A - \frac{\check{\delta}u'}{k} t^A \right), \quad (15a)$$

$$t'^A = \frac{e^{-i\lambda/2}}{\sqrt{k}} t^A. \quad (15b)$$

Here, λ is the spin phase described in Sec. II A and Appendix C. Because the curvature quantities are defined with respect to these spinors and their spatial dependence, this is enough to calculate the transformation laws of the curvature quantities.

The first and simplest set of curvature quantities we will need is the collection of Newman-Penrose scalars. The following are the definitions of these scalars on \mathcal{I}^+ , along with their leading-order relationship to the corresponding finite-radius scalars:

$$\psi_0 := \psi_{ABCD} o^A o^B o^C o^D \simeq \omega^{-5} \Psi_0, \quad (16a)$$

$$\psi_1 := \psi_{ABCD} o^A o^B o^C t^D \simeq \omega^{-4} \Psi_1, \quad (16b)$$

$$\psi_2 := \psi_{ABCD} o^A o^B t^C t^D \simeq \omega^{-3} \Psi_2, \quad (16c)$$

$$\psi_3 := \psi_{ABCD} o^A t^B t^C t^D \simeq \omega^{-2} \Psi_3, \quad (16d)$$

$$\psi_4 := \psi_{ABCD} t^A t^B t^C t^D \simeq \omega^{-1} \Psi_4. \quad (16e)$$

Because ψ_{ABCD} is a geometric object, it does not transform under a change of coordinates, so the transformation law for these scalars is given simply by replacing the spinors o^A and t^A with their transformed values, which leads to a simple hierarchy with a basic combinatorial pattern:

$$\psi'_0 = \frac{e^{2i\lambda}}{k^3} \left[\psi_0 - 4 \frac{\check{\delta}u'}{k} \psi_1 + 6 \left(\frac{\check{\delta}u'}{k} \right)^2 \psi_2 - 4 \left(\frac{\check{\delta}u'}{k} \right)^3 \psi_3 + \left(\frac{\check{\delta}u'}{k} \right)^4 \psi_4 \right], \quad (17a)$$

$$\psi'_1 = \frac{e^{i\lambda}}{k^3} \left[\psi_1 - 3 \frac{\check{\delta}u'}{k} \psi_2 + 3 \left(\frac{\check{\delta}u'}{k} \right)^2 \psi_3 - \left(\frac{\check{\delta}u'}{k} \right)^3 \psi_4 \right], \quad (17b)$$

$$\psi'_2 = \frac{1}{k^3} \left[\psi_2 - 2 \frac{\check{\delta}u'}{k} \psi_3 + \left(\frac{\check{\delta}u'}{k} \right)^2 \psi_4 \right], \quad (17c)$$

$$\psi'_3 = \frac{e^{-i\lambda}}{k^3} \left[\psi_3 - \frac{\check{\delta}u'}{k} \psi_4 \right], \quad (17d)$$

$$\psi'_4 = \frac{e^{-2i\lambda}}{k^3} [\psi_4]. \quad (17e)$$

The simplicity of this result is surprising because the Newman-Penrose quantities Ψ_n represent the components of a tensor, so at finite radius all of these components would mix with each other. However, in the limit as \mathcal{I}^+ is approached, the “peeling-off property” [10,15,30] of asymptotically flat spacetimes comes into play, as seen in the right-hand column of Eq. (17), so that the pattern emerges with lower-index quantities (e.g., ψ_0) being irrelevant to the transformed values of higher-index quantities (e.g., ψ'_4). On the other hand, given this peeling behavior, it may also seem surprising that effects from the higher-index quantities do not overwhelm the lower-index scalars. The reason is that at finite radii $\check{\delta}$ is replaced by $\omega^{-1}\check{\delta}$, and since each higher-index scalar appears in the expressions for lower-index scalars accompanied by powers of $\check{\delta}$, the resulting factors of ω^{-1} are exactly enough to cancel the dominance of the higher-index scalars—to leading order in ω .

The most interesting remaining quantity is the spin coefficient representing the shear:

$$\sigma := o^A o^B \bar{t}^{B'} \nabla_{BB'} o_A. \quad (18)$$

Because of the derivative, this is somewhat more difficult to evaluate than the Newman-Penrose scalars. However, after the suitable limit has been taken, we arrive at the simple formula [9,10,12]

$$\sigma' = \frac{e^{2i\lambda}}{k} [\sigma - \delta^2 \alpha]. \quad (19)$$

This is consistent with Eq. (17e) and the asymptotic relation

$$\psi_4 = -\frac{\partial^2}{\partial u^2} \bar{\sigma} \quad (20)$$

because $\partial/\partial u' = \frac{1}{k} \partial/\partial u$, and the BMS transformation is constant, so that the λ , k , and α functions are independent of u . The shear is also related to the more commonly used [39] strain of the transverse-traceless metric perturbation⁵ $h = \bar{\sigma}$, which implies the transformation law

$$h' = \frac{e^{-2i\lambda}}{k} [h - \bar{\delta}^2 \alpha]. \quad (21)$$

Similarly, the Bondi news function [9,10,40–42] satisfies

$$N \simeq n = -\frac{\partial}{\partial u} \bar{\sigma}, \quad (22)$$

which implies the transformation law

$$n' = \frac{e^{-2i\lambda}}{k^2} n. \quad (23)$$

We note, however, that these relationships between ψ_4 , σ , h , and n are only valid asymptotically, and only in Bondi coordinates; more generally, the relationships would be more complicated.

The expressions given here for the transformations of the waveform quantities are fairly simple, and can all be constructed given the waveforms and a choice of transformation—as described by the functions k , λ , and α . However, these expressions hide a complication: all of the quantities involved are functions of position. To actually implement a BMS transformation, we need to know how to express these functions in terms of the coordinates, both old and new. This requires combining the ideas of the present section with those of the previous section.

IV. IMPLEMENTATION OF BMS TRANSFORMATIONS OF WAVEFORMS

We assume that the field is known in some observer's frame \mathcal{O} , as a function of that observer's Bondi coordinates throughout some portion of \mathcal{I}^+ . A second observer \mathcal{O}' is related to the first by some known BMS transformation as in Eq. (8). In particular the frame of \mathcal{O}' can be obtained from \mathcal{O} by an initial supertranslation α , followed by a frame

rotation R_f , followed by a boost of velocity \mathbf{v} . Our objective will be to find the field as decomposed into modes of a spin-weighted spherical-harmonic expansion, at a series of discrete retarded times $\{u'_i\}$.

The first step is to find the mode weights of all quantities we will need in frame \mathcal{O} . At the most basic level, we need the modes of the waveform in question. We denote this waveform by f , which may represent h , ψ_4 , or any of the other quantities discussed in Sec. III B. A waveform f of spin weight s will be decomposed into modes of the spin-weighted spherical harmonic expansion as

$$f(u, \theta, \phi) = \sum_{\ell, m} f^{\ell, m}(u)_s Y_{\ell, m}(\theta, \phi), \quad (24)$$

where the relevant data are the modes $f^{\ell, m}(u)$. In practice, the sum over ℓ extends up to some maximum integer ℓ_{\max} .

If f represents a Newman-Penrose quantity ψ_n with $n < 4$, we also need all the higher-index Newman-Penrose quantities as shown in Eq. (17), as well as the quantity $\delta u'/k$. For σ and h , we will need $\delta^2 \alpha$ (or its complex conjugate). Given an arbitrary function w of spin weight s , we can calculate the modes of the differentiated quantity δw in terms of the modes of the original function as⁶

$$(\delta w)^{\ell, m} = \begin{cases} 0 & \ell < |s + 1|, \\ \sqrt{\frac{(\ell - s)(\ell + s + 1)}{2}} w^{\ell, m} & \text{otherwise.} \end{cases} \quad (25)$$

Similarly, we have

$$(\bar{\delta} w)^{\ell, m} = \begin{cases} 0 & \ell < |s - 1|, \\ -\sqrt{\frac{(\ell + s)(\ell - s + 1)}{2}} w^{\ell, m} & \text{otherwise.} \end{cases} \quad (26)$$

We will assume that the supertranslation α is given directly in terms of its modes. This makes it trivial to compute either $\delta^2 \alpha$ or $\bar{\delta}^2 \alpha$, noting that α has spin weight $s = 0$, while $\delta \alpha$ has spin weight $s = 1$ and $\bar{\delta} \alpha$ has spin weight $s = -1$. On the other hand, to compute $\delta u'/k = u \delta k/k - \delta \alpha$, we need to know k in terms of its spherical-harmonic modes. This could be done analytically with exact expressions involving Wigner's \mathfrak{D} and $3-j$ functions. For practical purposes, a more efficient approach is to evaluate k as given in Eq. (7) on a series of grid points, and feed the results into software that computes the modes, as discussed below.

Now, given the modes of the various fields at some discrete set of times $\{u_i\}$, we need to be able to interpolate

⁵It is worth pointing out that this relation is only true for the asymptotic fields. Trivially, we have $h \simeq rH$, whereas $\sigma \simeq r^2 \Sigma$. That is, the two finite-radius fields behave differently in the limit $r \rightarrow \infty$. But terms at higher relative order in $1/r$ may differ more substantially.

⁶These equations differ from the similar Eq. (3.22) of Newman and Penrose [16] and Eq. (2.7) of Goldberg *et al.* [43] by the factor $1/\sqrt{2}$ here. As noted previously, this is because the operator here is—up to the factor given in Eq. (13)—identical to the one given by Geroch, Held, and Penrose [37], which intentionally introduced the $1/\sqrt{2}$ factor.

as a function of time, because slices of constant u will not typically correspond to slices of constant u' . Of course, because of the direction dependence of these time slices, interpolation of the mode weights themselves is not possible in general. Instead, we must transform the modes into a series of points in physical space, interpolate the values of the field at each spatial point to the appropriate time, and then transform back to modes. The current state-of-the-art numerical code for transforming between physical space and modes is the SPINFAST package [28,29]. The points in physical space used by this package form an equiangular grid in colatitude-longitude:

$$\theta'_j = \frac{\pi j}{N_\theta - 1} \quad \text{for } j \in \{0, 1, \dots, N_\theta - 1\}, \quad (27a)$$

and

$$\phi'_k = \frac{2\pi k}{N_\phi} \quad \text{for } k \in \{0, 1, \dots, N_\phi - 1\}. \quad (27b)$$

Note that the poles, $\theta'_j = 0$ and π are each covered by N_ϕ pairs of (θ, ϕ) values, but each such pair represents a different alignment of the tangent basis at that point. For the sake of accuracy, it is best to choose $N_\theta > 2\ell_{\max}$ and $N_\phi > 2\ell_{\max}$ [29]. In practice, it seems to be sufficient to simply choose $N_\theta = N_\phi = 2\ell_{\max} + 1$. Of course, this grid is given in the frame of \mathcal{O}' ; since the waveform is given in the frame of \mathcal{O} , we need to know the points in that frame corresponding to the points $\{(\theta'_j, \phi'_k)\}$. Moreover, a spin-weighted field in \mathcal{O}' is defined with respect to the tangent vectors to the sphere, canonically defined in terms of the (θ', ϕ') coordinates. Thus, we also need to know what these tangent vectors correspond to in the basis of \mathcal{O} .

Adapting the discussion of Sec. II A, we begin by defining the rotor (in quaternion notation),

$$R'_{j,k} := e^{\phi'_k z/2} e^{\theta'_j y/2}, \quad (28)$$

where \mathbf{x} , \mathbf{y} , and \mathbf{z} are the orthonormal basis vectors of \mathcal{O} . Note the mixing of coordinates from \mathcal{O}' with basis elements of \mathcal{O} . We then define the unit vector

$$\mathbf{r}'_{j,k} = R'_{j,k} \mathbf{z} R'^{-1}_{j,k}, \quad (29)$$

which points in the direction (θ'_j, ϕ'_k) as measured by \mathcal{O} , and define the angle

$$\Theta'_{j,k} := \arccos \frac{\mathbf{v} \cdot \mathbf{r}'_{j,k}}{|\mathbf{v}| |\mathbf{r}'_{j,k}|}. \quad (30)$$

The equivalent angle in the unprimed frame is

$$\Theta_{j,k} = 2 \arctan \left[e^{-\phi} \tan \frac{\Theta'_{j,k}}{2} \right], \quad (31)$$

where $\phi = \text{artanh}|\mathbf{v}|$ is the rapidity. Then, we can define

$$B'_{j,k} := \exp \left[\frac{\Theta'_{j,k} - \Theta_{j,k}}{2} \frac{\mathbf{r}'_{j,k} \times \mathbf{v}}{|\mathbf{r}'_{j,k} \times \mathbf{v}|} \right], \quad (32)$$

unless $\Theta' = \Theta = 0$ or π , in which case we simply have $B'_{j,k} = 1$. Finally, we arrive at the required rotor

$$R_{j,k} := B'_{j,k} R_{\text{f}} R'_{j,k}. \quad (33)$$

The physical point labeled by (θ'_j, ϕ'_k) in \mathcal{O}' is given by

$$\mathbf{r}_{j,k} = R_{j,k} \mathbf{z} R_{j,k}^{-1} \quad (34)$$

in \mathcal{O} , while the complex tangent vector $\mathbf{m}'_{j,k}$ at that point in \mathcal{O}' corresponds to the vector

$$\mathbf{m}_{j,k} = R_{j,k} \frac{\mathbf{x} + i\mathbf{y}}{\sqrt{2}} R_{j,k}^{-1} \quad (35)$$

in \mathcal{O} . The spin phase is determined by the relative rotation between $\mathbf{m}_{j,k}$ as given here and the natural canonical \mathbf{m} vector given at the same point by \mathcal{O} . This was depicted in Fig. 2, and is explained in more detail in Appendix C.

More directly, we can evaluate any field, along with its appropriate spin-phase factor, by evaluating the mode-weighted spin-weighted spherical harmonics directly as functions of $R_{j,k}$. As detailed in Appendix B of Ref. [8], this is made possible by redefining the spin-weighted spherical harmonics to be functions of a single unit-quaternion argument in terms of Wigner's \mathfrak{D} matrices as⁷

$${}_s Y_{\ell,m}(R) := \sqrt{\frac{2\ell+1}{4\pi}} \mathfrak{D}_{-s,m}^{(\ell)}(R). \quad (36)$$

Thus, for example, when transforming the strain h , part of the right-hand side of Eq. (21) can be calculated very simply as

⁷This relationship was originally noted by Goldberg *et al.* [43], though they essentially restricted the possible rotations to rotors of the form $R = e^{\phi \mathbf{z}/2} e^{\theta \mathbf{y}/2}$. The problem with such a limited interpretation is that the spin-weighted spherical harmonics so defined *do not* transform among themselves under rotations, and are incapable of expressing the correct spin-phase behavior. By expanding the meaning of the spherical harmonics in this way we eliminate those problems, while maintaining agreement with the original definition and standard usage.

$$e^{-2i\lambda} [h - \bar{\delta}^2 \alpha] |_{\theta_j, \phi'_k} \\ = \sum_{\ell, m} [h^{\ell, m} - (\bar{\delta}^2 \alpha)^{\ell, m}] {}_{-2}Y_{\ell, m}(R_{j, k}). \quad (37)$$

Note that no additional manipulation is required to find the spin-phase factor $e^{-2i\lambda}$; it is implicitly calculated by ${}_{-2}Y_{\ell, m}(R_{j, k})$. There is, however, the remaining factor of $1/k$ to calculate. Including this factor is best done by evaluating this factor as [compare Eq. (7)]

$$\frac{1}{k} = \gamma(1 - \mathbf{v} \cdot \mathbf{r}_{j, k}), \quad (38)$$

then multiplying this result by the result of Eq. (37). In a similar way, other waveforms can be computed as necessary by pointwise combination of the relevant quantities given in the transformation laws of Sec. III B.

Proceeding in this way for all values of the discrete indices, we obtain the waveform values $f'(u'_{i, j, k}, \theta'_j, \phi'_k)$, where

$$u'_{i, j, k} = k(\theta'_j, \phi'_k)[u_i - \alpha(\theta'_j, \phi'_k)]. \quad (39)$$

Next, we simply need to interpolate these values in each direction to a corresponding set of times $\{u'_i\}$ representing some target time slices of observer \mathcal{O}' . There is a minor ambiguity here, in that this set of times is somewhat arbitrary. In practice, the input data may be sampled unevenly in time—for example, to provide better resolution of the merger-ringdown portion of a waveform, while reducing the amount of data representing the slow inspiral. It would presumably be best to retain this sampling in the transformed data set. To a reasonable approximation, this can be done by assigning

$$u'_i = \frac{1}{\gamma}(u_i - \alpha^{0,0}/\sqrt{4\pi}), \quad (40)$$

which is the value of u' for which the average value of u over the sphere (on the slice of constant u') is precisely u_i . To clarify the notation, $\{u'_{i, j, k}\}$ is the set of time coordinates already present in the data, whereas $\{u'_i\}$ is the set of times to which we might wish to interpolate.

However, we must deal with a subtlety first. In some directions, interpolation to some of the values of u'_i given by Eq. (40) would require data at times earlier than u_0 or later than u_{N_u-1} . This is because we have simply used the average value to derive Eq. (40), while neglecting the direction dependence. To avoid extrapolation, then, we must restrict the set $\{u'_i\}$ to the range of times $u'_{\min} \leq u' \leq u'_{\max}$, where

$$u'_{\min} = \max_{j, k} u'_{0, j, k}, \quad (41a)$$

$$u'_{\max} = \min_{j, k} u'_{N_u-1, j, k}. \quad (41b)$$

We denote the resulting subset by $\{u'_i\}$, which is the final set of times to which we will interpolate the data. The index i' is used to indicate that it comes from a slightly different indexing set than the index i used for the input data.

Though the construction of $\{u'_i\}$ suggested here is by no means unique, we will always be limited to using a *proper* subset of the input data, whenever the boost and super-translation components with $\ell > 0$ are nontrivial, because some of the input time steps will correspond to slices of u' for which the input data represent an incomplete sphere, and thus insufficient data for computing spin-weighted spherical-harmonic modes. Nonetheless, this choice of $\{u'_i\}$ is well defined and easy to implement, it roughly preserves the sampling of the input data, and it uses the input data to nearly the fullest possible extent.

Finally, for each value of (i', j, k) , we interpolate the waveform values $f'(u'_{i, j, k}, \theta'_j, \phi'_k)$ in time to $f'(u'_i, \theta'_j, \phi'_k)$. For each i' , we then feed these values into a software package like SPINFAST to obtain the modes as measured by observer \mathcal{O}' , thus arriving at our goal: the set of modes $f'^{\ell, m}$ for each time u'_i . The entire transformation is implemented in the Python module SCRI, which is included in the supplemental materials provided with this paper [24].

V. EFFECTS OF TRANSFORMATIONS ON WAVEFORM MODES

It will be instructive to observe the effect of typical transformations on waveform modes. Because of the peculiar nature of \mathcal{I}^+ and the highly nonlinear behavior of waveforms under these transformations, we will not be able to rely on any intuition for transformations of multipole moments that we may have gained in studying electromagnetism, for example. However, we can take advantage of the fact that mode decompositions are linear, so that it is sufficient to observe the transformation of a single mode at a time. In particular, we will define input waveforms having a single nonzero mode. For further simplicity, that mode will behave as a pure phase rotation at constant angular velocity. We can then transform this model waveform, and see how the power in the chosen mode leaks out into other modes. Because rotations are already well understood—in fact, they behave identically to the more familiar rotations of spin-zero spherical harmonics—we will focus here only on translations and boosts.

To be precise, let us choose the nonzero mode $(\ell_{\text{nz}}, m_{\text{nz}})$ and define our model waveform by its modes as

$$\psi_4^{\ell, m}(u) = \delta_{\ell_{\text{nz}}, m_{\text{nz}}}^{\ell, m} e^{i\omega u}. \quad (42)$$

Here, ω represents an angular velocity. For purposes of illustration, let us choose $\omega = 0.3 \frac{1}{M}$, which is a typical

value for the $(\ell, m) = (2, 2)$ mode of comparable-mass binaries just before merger, where M is the total mass of the system.

As a first example, we see the effect of translations in Fig. 7. The four cases shown here correspond to $(\ell_{\text{nz}}, m_{\text{nz}}) = (2, 2)$ or $(\ell_{\text{nz}}, m_{\text{nz}}) = (4, 2)$, and translations of $\alpha = 0.1M \sin \theta \cos \phi$ or $\alpha = 0.1M \cos \theta$. These are displacements of $0.1M$ in the x and z directions, respectively, corresponding to typical displacements found in the publicly available catalog of waveforms from the SXS collaboration [21], as will be discussed further in Sec. VI.

The z displacement evidently has a very simple effect on the modes; power is transferred to all other ℓ modes with $m = m_{\text{nz}}$, where the transferred power goes roughly as $\epsilon^{|\ell - \ell_{\text{nz}}|}$ for some parameter $\epsilon \approx 0.01$. This simplicity is a result of the fact that the waveform of Eq. (42) is effectively rotating about the z axis, so a translation along that direction preserves a great deal of the symmetry of the system. A similar but far more complicated pattern can be seen in the x translations, where now the power is transferred into essentially all modes. While there is a similar dependence in ℓ —where the coupling seems to get smaller

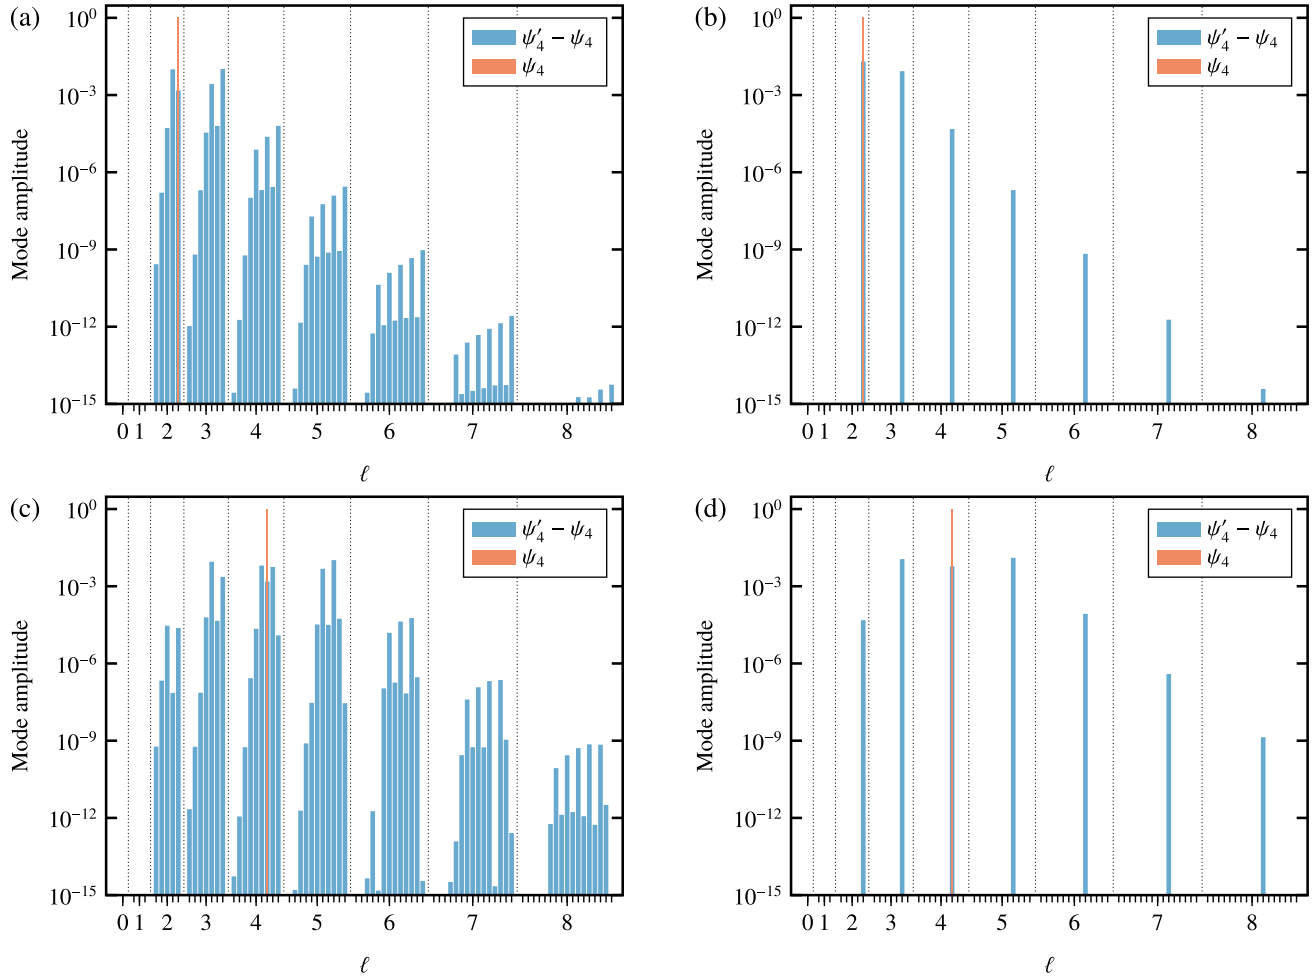


FIG. 7. Mode transformations under translation. These plots show the changes to the amplitudes of the waveform modes when the system is translated. The modes are grouped by ℓ value, with individual m values increasing from $-\ell$ on the left to ℓ on the right in each group. In each case, the initial waveform ψ_4 is made up of a single mode, as in Eq. (42). In the two upper panels, (a) and (b), the nonzero mode is $(\ell_{\text{nz}}, m_{\text{nz}}) = (2, 2)$; in the two lower panels, (c) and (d), the nonzero mode is $(\ell_{\text{nz}}, m_{\text{nz}}) = (4, 2)$. These waveforms are then transformed to ψ'_4 by a translation of magnitude $0.1M$. The panels on the left, (a) and (c), depict a translation in the x direction (the same translation in the y direction would look almost identical here); the panels on the right, (b) and (d), depict a translation in the z direction. We see that a translation in the x direction tends to move power into modes with a wide variety of (ℓ, m) values, whereas a translation in the z direction only moves power into modes with the same m values as the original waveform. This is a result of the fact that the simulated waveform is effectively rotating about the z axis, so a z translation preserves a certain amount of symmetry, whereas the x translation violates that symmetry. As explained in the text, the power leakage is roughly given by powers of the product of displacement and frequency, which is roughly 0.03 in this case. The frequency was chosen to be typical of frequencies seen just prior to the merger stage of comparable-mass binaries. Earlier during the inspiral portion, the frequencies will be an order of magnitude smaller, and the size of these effects correspondingly smaller.

exponentially with $|\ell - \ell_{\text{nz}}|$ —there is a more complicated dependence on m .

These patterns can be understood by looking at the effect of a translation on the time coordinate. Kelly and Baker [20] pointed out that the effect on ψ_4 of a supertranslation α (without any accompanying boost or rotation) can typically be approximated by the first few terms of the Taylor-series expansion⁸

$$\psi'_4(u', \theta', \phi') = \sum_{j=0}^{\infty} \frac{1}{j!} \left(-\alpha(\theta, \phi) \frac{\partial}{\partial u} \right)^j \psi_4(u, \theta, \phi), \quad (43a)$$

where $(\theta, \phi) = (\theta', \phi')$. With Eq. (42), this specializes to

$$\psi'_4(u', \theta', \phi') = \sum_{j=0}^{\infty} \frac{1}{j!} (-i\omega\alpha(\theta, \phi))^j {}_{-2}Y_{\ell_{\text{nz}}, m_{\text{nz}}}(\theta, \phi) e^{i\omega u}. \quad (43b)$$

In each case shown in Fig. 7, α is an $\ell = 1$ function, and thus couples with ${}_{-2}Y_{\ell_{\text{nz}}, m_{\text{nz}}}$ to progressively higher orders—and hence at larger “distances” from $(\ell_{\text{nz}}, m_{\text{nz}})$, in some sense—with increasing values of the summation index j . On the other hand, these couplings also include progressively higher powers of ω times the amplitude of α , roughly 0.03, and thus progressively smaller amplitudes. Besides the factor of $1/j!$, there are further geometric factors involved in the normalization of the spin-weighted spherical harmonics, which means that the ratios of power in the various modes do not follow a particularly simple pattern, but it is clear that these considerations lead to the correct qualitative behavior and—when accounting for the factorial and geometric factors—the correct quantitative behavior.

The nonzero input waveform mode in each of these cases has amplitude 1. Of course, the effect of mode mixing is linear, so the plots in Fig. 7 should essentially be read as fractional coupling between the modes. For example, in

⁸The term $-\alpha\partial/\partial u$ in Eq. (43a) is the generator of the supertranslation α , as described at the end of Sec. II C and in Ref. [15]. Thus, this equation is simply the exponentiation of that element of the Lie algebra \mathfrak{bms} , which gives us the corresponding element of the Lie group BMS. It must be noted, however, that such exponentiation is not typically a sufficient method for transforming waveform data. For example, Eq. (17e) shows that we generally also have a factor $e^{-2i\lambda}/k^3$ in the transformation law. In our particular case, this factor happens to be 1, which is why exponentiation works. More generally, the generator of a boost would not supply the correct factor of k . But even for supertranslations, exponentiation would fail to correctly transform other quantities. For example, in transforming h [Eq. (21)], the term $-\bar{\delta}^2\alpha$ would not appear. A more extreme example is provided by ψ_0 [Eq. (17a)]; action of the \mathfrak{bms} generators would fail to supply the terms ψ_1 through ψ_4 . Moreover, because of the infinite nature of this expansion, it may be useful for gaining qualitative insight into the approximate coupling between modes, but it is not useful for accurate implementation of these transformations.

Fig. 7(a), we see that a little more than 1% of the power in the (2,2) mode is mixed into the (2,1) and (3,3) modes. But in many cases, the physical (2,2) mode is strongly dominant over either of these modes, so that the expected ratio of amplitudes would be less than 1%. In such cases, the measured (2,1) and (3,3) modes would actually be primarily made up of power leaking from the (2,2) mode. While we might typically expect the frequency of the (2,1) mode in real binary systems to be roughly 1/2 that of the (2,2) mode, and the frequency of the (3,3) mode to be roughly 3/2 that of the (2,2) mode, the frequency of the mixed component would be nearly the same as that of the (2,2) mode. Taken together, these features can provide a signature of mixing due to transformations.

As we have seen, translations comprise a special case of supertranslations having $\ell = 1$. Similar behavior results from supertranslations with $\ell > 1$, except that the coupling between modes is more extensive. For example, if the original waveform has nonzero mode $(\ell_{\text{nz}}, m_{\text{nz}}) = (2, 2)$, a supertranslation with nonzero $(\ell, m) = (2, 0)$ component couples power at a roughly equal level into both the (3,2) and (4,2) modes of the transformed waveform. It should also be noted that supertranslations with $\ell > 1$ can directly alter the value of, for example, the strain waveform h through the term $-\bar{\delta}^2\alpha$ in Eq. (21). The operator $\bar{\delta}^2$ eliminates the modes of α with $\ell \leq 1$, as it must for a field of spin weight $s = -2$.

We can make a similar comparison of waveform modes before and after a boost. Figure 8 shows essentially the same thing as Fig. 7, except that instead of translations, the waveforms have been subjected to boosts. The speed of the boost is $\beta = 0.01c$ in each case, directed in either the x or z direction. The most obvious feature here is the remarkable similarity between Figs. 7 and 8. The coupling due to translation falls off more quickly with increasing distance from the dominant mode, but the general patterns are very similar. The time has been chosen as $u = u' = 0$, so that only the boost itself factors into this transformation. This means that the translation induced by the boost, as discussed near the end of Sec. II C, is zero. In this case, only the movement of the points around the sphere—as depicted in Fig. 2—comes into the transformation. In particular, the transformation law given by Eq. (17e) is

$$\psi'_4(0, \theta', \phi') = e^{-2i\lambda}\gamma^3(1 - \mathbf{v} \cdot \mathbf{r})^3 \psi_4(0, \theta, \phi). \quad (44)$$

The spin factor $e^{-2i\lambda}$ is primarily just transforming the tangent bases to avoid singularities in the basis vector fields, and has no strong effect on the modes in this case. The value of γ^3 is approximately 1.00015, which accounts for the change to the (2,2) mode, to reasonable accuracy. The third factor multiplies the waveform by roughly $1 - 3\beta \sin \theta \cos \phi$ for the boost in the x direction, and roughly $1 - 3\beta \cos \theta$ for the boost in the z direction. But this factor can only explain coupling between modes with $\Delta\ell = \pm 3$, whereas we

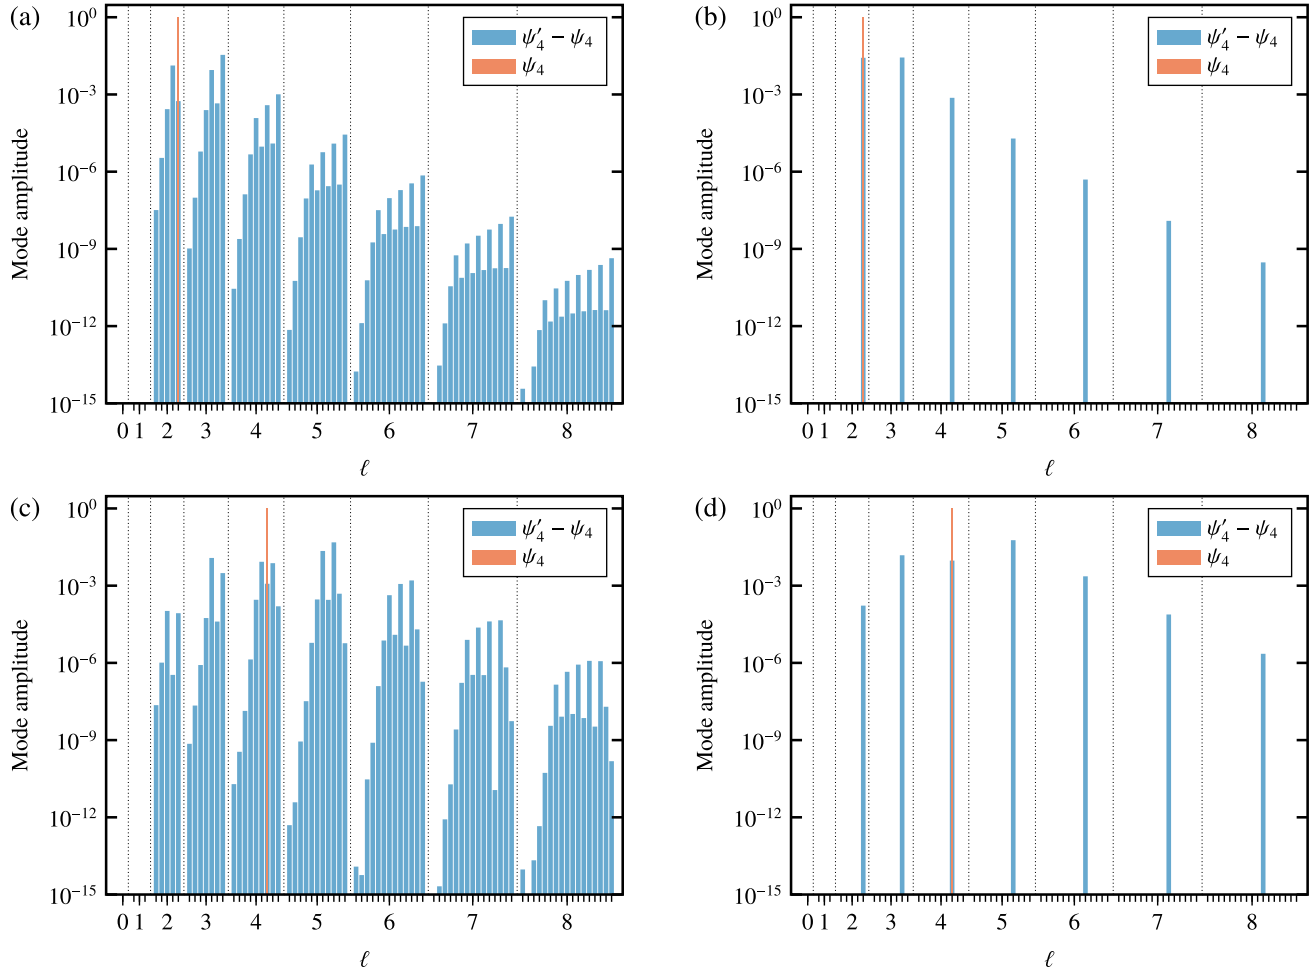


FIG. 8. Mode transformations under boost. These plots show the changes to the amplitudes of the waveform modes when the system is boosted. These are similar to the plots in Fig. 7, except that there is no translation, only a boost. Again, the initial waveform in the upper panels has nonzero mode $(\ell_{\text{nz}}, m_{\text{nz}}) = (2, 2)$; in the two lower panels, the nonzero mode is $(\ell_{\text{nz}}, m_{\text{nz}}) = (4, 2)$. In this case, the two panels on the left depict a boost of speed $0.01c$ in the x direction, while the two panels on the right depict a boost of speed $0.01c$ in the z direction. The mode amplitudes change as a function of time in this case, because the boost is essentially a time-dependent supertranslation, as described at the end of Sec. II C. The quantities shown here correspond to $u = u' = 0$, so that the translation induced by the boost is actually zero. Despite the difference in their construction, these results are qualitatively very similar to those of Fig. 7. However, the boost used to produce these figures is orders of magnitude greater than any found in the numerical waveforms considered below. The basic conclusion is that the boost matters in those cases only to the extent that even a small boost can induce a significant translation over the course of a long simulation.

clearly see more extensive coupling in Fig. 8. In fact, if we expand the boost rotor of Eq. (2) in powers of β , and use that to expand the argument of $\psi_4(0, \theta', \phi')$ in a Taylor series, we find another factor multiplying $\psi_4(0, \theta, \phi)$:

$$\begin{aligned}
 & 1 - |\mathbf{v} \times \mathbf{r}| + \frac{1}{2} |\mathbf{v} \times \mathbf{r}|^2 \mathbf{v} \cdot \mathbf{r} + \dots \\
 & = \begin{cases} 1 - \beta \sin \phi + \frac{1}{4} \beta^2 \sin 2\phi \sin \theta + \dots & x \text{ boost,} \\ 1 - \beta \sin \theta + \frac{1}{4} \beta^2 \sin 2\theta + \dots & z \text{ boost.} \end{cases} \quad (45)
 \end{aligned}$$

Because of the geometry, the largest couplings from this factor are typically several times smaller than the couplings from the $(1 - \mathbf{v} \cdot \mathbf{r})^3$ factor. That is, the largest peaks in Fig. 8

will be dominated by the 3β term, but smaller peaks with $|\delta\ell| > 3$ will be dominated (and in fact made possible) by the more complicated factor of Eq. (45).

In the plots of Fig. 8, $3\beta = 0.03$, which is also the approximate scale of the effects of the translation for the plots of Fig. 7. This explains why the magnitude of the coupling is so similar in the two cases, at least for the dominant coupling terms. Of course, there is no $1/j!$ term for the boost couplings, as in Eq. (43). This explains why the couplings in Fig. 7 should fall off so much faster than those in Fig. 8.

In fact, this numerical equality between the couplings for translations and boosts is the reason $\beta = 0.01c$ was chosen

for these examples, to ease comparison between the effects of a translation and of a boost. But we must note that this value was chosen entirely for the purpose of illustration; it is an order of magnitude larger than the largest speed found in the SXS catalog discussed below, and several orders of magnitude larger than typical speeds. This might appear to suggest that the effect of the boost itself is entirely negligible for those simulations. However, we have thus far only described the transformation due to a boost on the $u = u' = 0$ slice. At any later time u' , an additional coupling is present, which is essentially identical to a translation by $\gamma u'v$, as we can see from the arguments toward the end of Sec. II C. Even a very small boost can build up to a significant translation over the course of a long simulation. In fact, we will find that near merger, boosts play a significantly more important role than translations in the SXS catalog.

VI. REMOVING DRIFT FROM NUMERICAL WAVEFORMS

To demonstrate one way in which BMS transformations are important at a practical level, we examine the publicly available catalog of simulations from the SXS collaboration [21,22]. First, we will illustrate a particular system to see unexpected effects in its waveform modes, and see how these effects can be reduced by applying a spatial translation and a boost derived from the simulation data. Then, we will briefly examine the size of the translation and boost for other simulations in the catalog.

The first simulation we consider is labeled in the waveform catalog as SXS:BBH:0004, and represents an (approximately) equal-mass system in which one black hole has dimensionless spin $S_1/M_1^2 = 0.5$ along the $-z$ axis, while the other black hole has no spin. This system is interesting because it is not precessing, and so retains enough symmetry to allow us to unambiguously identify some curious features. But it is nonetheless not perfectly symmetric, and thus exhibits those nontrivial features.

We can see the first example of nontrivial features in this system by simply plotting the center of mass. Using the Christodoulou masses and coordinate positions of the black holes, we form the usual center of mass.⁹ The result is plotted in Fig. 9. Because the system is not symmetric, we expect to see some asymmetry in the emission of gravitational waves in the orbital (x - y) plane [8,44], and thus some force in this plane. But that force should have roughly constant magnitude on the orbital time scale, and should simply rotate with the system. So we expect the center of mass to be pushed around in a circle. This is essentially what we find in the data. The center of mass starts nearly at the origin, so this circle is initially not centered on the

⁹These quantities are all stored in the waveform catalog in files named HORIZONS.H5.




FIG. 9. Center of mass motion. This plot shows the coordinates in the x - y plane of the center of mass for system SXS:BBH:0004, throughout the inspiral of the system. Motion in the z direction is far smaller. The coordinates are given in units of $M = M_1 + M_2$, the total Christodoulou mass [45] of the system. The raw data from the simulation results in the curve labeled “Original.” This is the motion of the center of mass, as seen in the same frame in which the waveforms are measured. There is a small initial offset, as well as a strong drift velocity. We can also apply a spatial translation and boost to the system, in which case the center of mass appears to rotate more simply around the origin, as seen in the curve labeled “Transformed.”

origin. But there is another strong effect: an overall drift. Evidently, this drift is due to residual linear momentum in the initial data. For future evolutions, Ref. [23] introduced a method to eliminate such residual momentum from the initial data. However, for the present waveform catalog and any future simulations in which a large initial translation is present, or a significant recoil develops during the inspiral, we must transform the data to eliminate the offset and drift.

The approach taken here is crude, but will serve the purpose of illustration. By minimizing the average distance between the center of mass and the origin, we can find the optimal translation and boost, as described in Appendix E. For this system, the results are

$$\delta\mathbf{x} = (-9.1 \times 10^{-3}, 7.8 \times 10^{-3}, -4.0 \times 10^{-9}), \quad (46a)$$

$$\mathbf{v} = (9.4 \times 10^{-6}, -5.3 \times 10^{-6}, 2.6 \times 10^{-12}). \quad (46b)$$

Over the course of this $\sim 11000M$ simulation, the small boost grows into a larger translation than the initial offset $\delta\mathbf{x}$. Applying this transformation to the center of mass we see a much cleaner-looking curve, essentially orbiting the origin, in Fig. 9. Although the center of mass measured in this way is based on coordinates, and thus susceptible to all the vagaries of gauge in the most extreme regions of the simulated spacetime, we will nonetheless find that the same transformation applied to the waveform removes

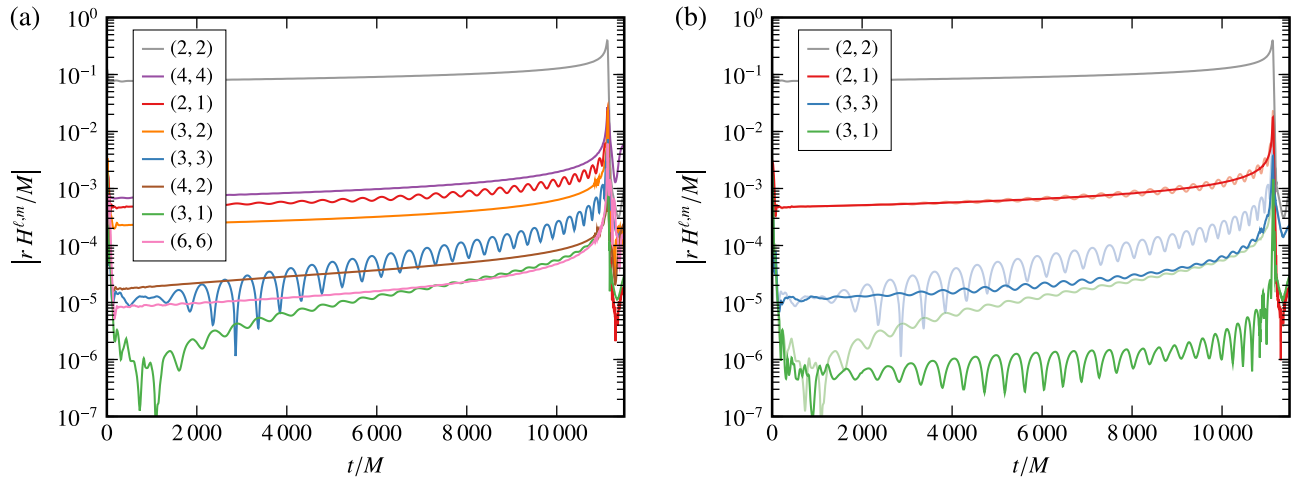


FIG. 10. Original and corrected waveform data. These plots show the waveform data for SXS:BBH:0004, and the effects of the transformation. The plot on the left, (a), shows the raw waveform obtained from the SXS catalog, using extrapolation with polynomials of order $N = 3$ [46,47]. The eight most significant modes are shown. Of these, the (2,1), (3,3), and (3,1) modes exhibit substantial oscillations. Oscillations are somewhat unexpected because analytical models of this system contain no such features. However, these are also the modes that couple most strongly to the dominant (2,2) mode under a translation in the x - y plane, as shown in Fig. 7(a). The plot on the right, (b), shows the (2,2) mode and the three oscillating modes (removing the other modes for clarity) after the waveform has been transformed by the translation and boost given in Eq. (46). For comparison, the original modes are also shown in the corresponding colors, with lower opacity. The effect of the transformation is largest near merger, when the translation induced by the boost is largest and the frequency is highest. Despite the crude way in which the transformation parameters were determined, the transformation itself eliminates the oscillations of the (2,1) mode, while reducing the overall amplitudes of the (3,3) and (3,1) modes by as much as a factor of 20.

features that we would not expect based on naive analytical models.

Figure 10(a) shows the largest modes¹⁰ in the waveform. This is the original data taken from the SXS catalog. The (2,2) mode is entirely dominant, as expected. The post-Newtonian model [44] of this system predicts smooth, monotonic waveform amplitudes during the inspiral. Yet the (2,1), (3,3), and (3,1) modes exhibit distinctive oscillations that are not visible in the other modes. These modes also have the largest couplings with the (2,2) mode under translation in the x - y plane, as seen in Fig. 7(a), and the oscillations are at the same frequency as the (2,2) mode. These facts strongly suggest that the oscillations are caused by mode coupling due to the motion of the center of mass. In fact, we can even predict the size of these couplings. The original system ends up translated from the origin by about $0.1M$ at merger, and the frequency of the (2,2) mode just prior to merger exceeds $0.3/M$. These were the parameters used to construct Fig. 7(a), which means that the couplings shown in that plot should be roughly the couplings found in this waveform near merger. Specifically, we expect to find

¹⁰This is as measured shortly before merger. We ignore modes with negative m values because, for this system, they have essentially the same magnitudes as their counterparts with positive m values. Also, the (2,0) mode comes out above the (4,2) mode when extrapolating with $N = 3$ polynomials, but is evidently not trustworthy [46], so we ignore it.

mode couplings in this waveform starting at just over 1% of the magnitude of the (2,2) mode near merger.

Figure 10(b) shows the dominant mode—which is not visibly changed at this scale—and the oscillating modes after the transformation of Eq. (46) has been applied to the waveform. For comparison, the original modes are shown in the same colors with lower opacity. In each case, the effect of the transformation is smallest at the beginning of the simulation, when the offset and frequency are smallest; conversely, it is largest at merger, when the offset and frequency are largest. This is just as we would expect, given the arguments of Sec. V. Moreover, we can look at the changes to these modes as a fraction of the (2,2) amplitude, and find that they do agree nicely with Fig. 7(a): the (2,1) and (3,3) modes change by just over 1.3% at merger, and the (3,1) mode by around 0.35%. Finally, we can subtract the transformed waveform from the original data and measure the frequency of the difference; for each of these three modes, we find that it matches the frequency of the (2,2) mode, rather than the frequency of the transformed mode. These facts all suggest that the changes to these three modes are primarily undoing leakage of the (2,2) mode.

The oscillations have been essentially removed from the (2,1) mode. This mode is the third-largest overall, after the (2,2) and (4,4) modes. Yet its amplitude is altered in this transformation-induced coupling by several percent throughout the inspiral, growing to 30% at merger (relative to the transformed values). The oscillations of the much

smaller (3,3) mode are reduced substantially, though not entirely eliminated. This is not surprising, given that the change in the waveform is so large, and the method of choosing the transformation of Eq. (46) so crude. In fact small changes in the parameters used to choose the transformation (t_i and t_f in Appendix E) lead to significant changes in the smoothness of this transformed mode, suggesting that there may be better choices. However, the remarkable feature of this transformation is the size of the change, which ranges up to 350% late in the inspiral (again, relative to the transformed values). Even more extreme is the change in the (3,1) mode, which reaches typical values over 2000% toward the end of the inspiral.

We can conclude from this that throughout almost the entire simulation, the (3,3) and (3,1) modes as given in the original data from the SXS catalog are entirely dominated by coupling from the (2,2) mode, while the (2,1) mode is strongly affected—though not completely dominated. This means that any attempt to use these modes without accounting for the effect of the residual velocity in the initial data will be prone to errors.

The mode couplings we have seen here are all caused by a very small residual velocity, which leads to an anomalous translation of just $0.1M$ around merger. It may be surprising that such large effects can follow from such a seemingly small cause. But it is more surprising that this anomalous translation is typical of the simulations in the SXS catalog. Figure 11 shows the initial and final displacements of the center of mass for every system in the SXS catalog. We can see that the size of the transformation in SXS:BBH:0004 is fairly typical of systems in the catalog. In fact, the translation

near merger [Fig. 11(b)] for this system (0.108) is slightly above the median (0.070), and just half the mean (0.216).

Closer inspection of the data show that all the systems with very small translations [less than about 10^{-2} in Fig. 11(b)] are symmetric, with equal masses and spins, and the spins are all aligned with the orbital axis. If the masses or spin magnitudes are not equal, there is generally a larger translation. Still larger translations are typically found in systems for which one or both black holes have spin components in the orbital plane. On the other hand, simulations that run for longer have more opportunity to develop a large translation; the very largest values result from very long simulations, rather than extraordinarily asymmetric physics.

In this section, we have found that applying translations and boosts determined from the orbital trajectories of a simulation in the SXS catalog can have a very large effect on the distribution of power in the modes, and can diminish unmodeled features in the waveform. Moreover, we have seen that this particular system is fairly typical of systems in the SXS catalog, with numerous systems expected to exhibit significantly larger mode couplings. The mode transformations described in this paper can be expected to substantially improve the agreement between analytical waveform models and numerical waveforms in these cases.

VII. EFFECTS ON DATA ANALYSIS FOR GRAVITATIONAL-WAVE DETECTORS

The previous section showed that seemingly small transformations can have pronounced effects on waveform

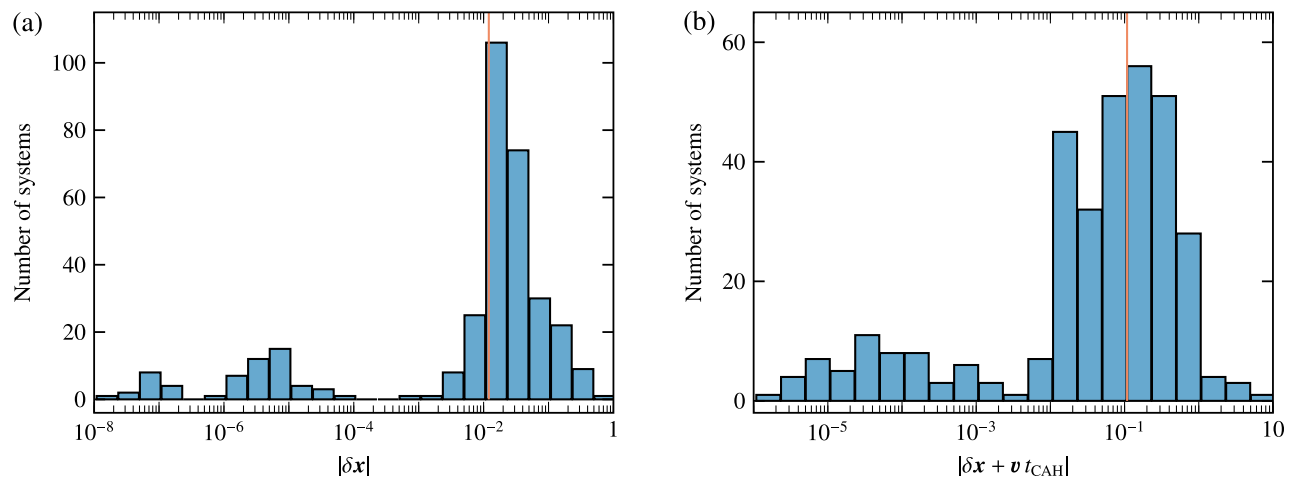


FIG. 11. Survey of translations and boosts in the SXS catalog. These plots show the number of systems with a given initial displacement, (a), and a given displacement at merger, (b). In the latter plot, t_{CAH} is the time at which a common apparent horizon is found, which is a typical definition of the merger time. The data include every system in the SXS catalog, but only the highest-resolution instance of each system. The vertical red line in each plot shows the value for the system SXS:BBH:0004, which is the one described in Sec. VI and Figs. 9 and 10. Systems with small values in each plot are typically simple systems with little or no spin, or high symmetry; larger values generally indicate asymmetries in the masses, unequal spins, and especially strongly precessing systems. The displacement at merger is dominated in most cases by the translation due to boost, rather than initial displacement—though the boost actually *reduces* the displacement from its initial value in roughly one fifth of the systems in the catalog.

modes. Having understood the nature and origin of these effects, we can now address the issue of what must be done about them. This section briefly discusses the impact these transformations have on two aspects of detections of gravitational waves: the production of waveform models, and the construction of template banks for searches in detector data.

The first step in detecting a gravitational wave is to devise a model of a waveform we might expect to find in the data. From astrophysical arguments, the most reliable candidates for detection are mergers of compact binary systems. Because of the nonlinear nature of mergers, they can only be modeled accurately by computers. On the other hand, using a computer to generate the entire signal is simply unrealistic for most of the expected systems [48]. Thus, at some level, waveform modeling must combine numerical and analytical results. But because the waveforms come from different approaches, we should expect to find differences in their gauges—as amply demonstrated by the boosts and translations discovered in the SXS data, which will not naturally appear in any analytical model.

These gauge differences will have real impacts on any model that uses numerical waveforms. For example, when “calibrating” effective-one-body models [49–51], the analytical waveform must be aligned to the numerical waveform. If the numerical waveform has spurious features, the waveforms will appear to align poorly, so the calibration will be less than optimal and result in inaccurate waveforms. Other phenomenological waveform models [52,53] and surrogate models [54,55] would experience the same biases, trying to fit simple formulas to waveforms with effectively random gauge effects. Similarly, when constructing hybrid waveform models [48], the hybrids will be imperfect or even discontinuous in the region where one switches from analytical to numerical data.

As mentioned in Sec. I, some of these gauge freedoms—time translations and rotations—are entirely familiar, and routinely dealt with simply by applying a gauge transformation to one waveform to minimize some measure of the difference between the waveforms. In principle at least, this approach could also be extended to the full BMS group, though the supertranslations would obviously be represented only up to some finite spherical-harmonic order, and the numerical implementation may be delicate. It may also be feasible to resolve the gauge ambiguities using any of various methods presented in the literature [33,56–58], though it is not clear that such an approach would be numerically feasible. In any case, a simplistic approach like the one found in Sec. VI is presumably a helpful first step.

Now, assuming that we have a waveform model for a particular astrophysical system produced with the appropriate care for gauge ambiguities, searches of detector data require a family of template waveforms—specific instances of waveforms from the broader class making up the model.

As noted in Sec. I, the signal measured along any simple curve on \mathcal{I}^+ of constant spatial coordinates is a good approximation for the signal measured by some inertial observer, because the metric in Bondi gauge is *manifestly* asymptotically flat. Here, we consider the signal to be measured as a function of the retarded-time coordinate u along some direction (θ, ϕ) , in which case the limiting process as $r \rightarrow \infty$ is well defined and the signals at finite radius and on \mathcal{I}^+ can be compared meaningfully. Of course, since BMS transformations preserve Bondi gauge, we can apply any BMS transformation to generate another curve on \mathcal{I}^+ , and another corresponding waveform. We might worry that we would need a separate template for each member of the BMS group—or at least for some discrete sampling of the BMS group. But given its infinite dimensionality, this could still be a very large or even impossible task.

Fortunately, the situation is not quite so dire. We only need to generate templates for elements of the BMS group that produce *detectably distinct* waveforms. But there are degeneracies among the templates created in this way, particularly among the supertranslations. Given that our detector will lie along a single direction from the source, the supertranslation α will be evaluated along a particular direction. Ignoring the Lorentz transformation for the moment, the retarded time transforms as $u' = u - \alpha(\theta, \phi)$. As far as its effect on the time variable is concerned, all those infinitely many degrees of freedom in α reduce to a single number. In principle, the angular dependence of α does lead to a transformation of the quantity h measured by a gravitational-wave detector, as shown in Eq. (21). However, the term $\bar{\delta}^2\alpha$ is constant in time, and so is not detectable. Thus, for detections along a single line of sight from the source, the entire supertranslation sector of the BMS group is reduced to a single time offset.

Factoring the supertranslations out of the BMS group leaves us with the familiar Lorentz group of rotations and boosts. The rotations determine the sky position of the detector relative to the source—or equivalently the orientation of the source relative to the detector.¹¹ We can further separate rotations into a rotation along the line of sight between detector and source, and 2 other degrees of freedom that we might call latitude and longitude. This first rotation is directly degenerate with the spin phase λ described in Sec. II A. But both this and the time offset are “extrinsic” parameters, already dealt with in searches by

¹¹There is still another rotation that must be accounted for in data analysis, related to the orientation of the detector relative to the source—or equivalently the sky position of the source relative to the detector. This second rotation is partially degenerate with the first, in that it will also affect the spin phase. The other 2 degrees of freedom in this rotation determine the detector’s sensitivity to the signal via the “antenna pattern.” However, since this rotation relates the detector’s orientation to the coordinates we have already been dealing with, it does not fit comfortably within the scope of this discussion.

simply finding the element of a discrete Fourier transform with the largest magnitude [1,59,60]. The remaining rotational degrees of freedom are described in more detail elsewhere [61–63]. In brief, it appears that accounting for them could provide benefits for localization and parameter estimation, but could actually be counterproductive for detection. The impact of the boost degrees of freedom is likely to be much smaller, and indistinguishable from an error in the total mass of the system.

To summarize this section, let us reiterate how it is that supertranslations are so important for waveform modeling, but not important for detection. Supertranslations are important to waveform models for two reasons: (1) the models must be able to describe the waveform in any direction from the source; and (2) at some point we generally need to compare or combine two different models, so the gauge freedom must be accounted for. On the other hand, a detector lies along a single direction from the source, which means that all the degrees of freedom in the supertranslation are degenerate. If we had a network of detectors located in significantly different directions from a source, and we wished to combine their information, we would need some control over their relative time offsets which could be considered equivalent to supertranslation degrees of freedom. This is not expected to be a pressing concern in the near future.

VIII. CONCLUSIONS

There is no such thing as a gauge-invariant gravitational waveform. It is possible to find a gauge-*fixed* waveform—for example, one measured in the standard Bondi gauge. However, even this is not a particular coordinate system, but an infinite-dimensional class of equally acceptable systems. We can transform between members of the class using any element of the infinite-dimensional Bondi-Metzner-Sachs group, which shows that the class is very large indeed. Moreover, we have seen that such transformations can affect the waveform dramatically, even when the transformation seems to be small. This means that any comparison between waveforms—whether numerical, analytical, or even experimentally measured waveforms—will be affected by the gauge in which the waveform is expressed (or equivalently, the frame in which the waveform is measured). There is no obvious preferred frame. Instead, all we can (and should) do is to insist that the waveforms are at least in the *same* frame. Doing so requires understanding the BMS group, and how its elements transform waveforms.

This paper has explored the BMS group, and illustrated some of its impact on gravitational-wave analysis. We began with a thorough and pedagogical introduction to the group itself, to provide a common starting point to be used in the remainder of the paper. We then examined asymptotically flat spacetime, and found how the BMS group transforms various types of waveforms. We then used these

insights to see how such transformations can be implemented in practice. This is applied in the Python package SCRI accompanying this paper on its arXiv page. The following section then showed how these transformations should affect the spin-weighted spherical-harmonic modes of a waveform with simplified numerical models, and we found good agreement with analytical approximations for the leading-order couplings. Anomalous translations and boosts were found in the publicly available SXS catalog. A particular example was used to show that the original data contains large effects from these anomalies, including modes that are several to dozens of times larger than they would be expected to be. These modes can be transformed to simplify their structure, and bring them more closely in line with what is expected from analytical models. However, more complicated systems will have even larger mode couplings. The size of the coupling is expected to scale roughly linearly with the size of the translation involved—since the direct contribution of the boost is relatively small compared to the influence of the translation it gives rise to—and some simulations in the SXS catalog have translations almost 100 times greater than the example system. Finally, we discussed the effect of the BMS gauge freedom on data analysis for gravitational-wave detectors, showing that it must be accounted for when creating model waveforms, but the supertranslations do not complicate searches.

The waveforms found in the SXS catalog are not wrong, *per se*; but they contain effects that may not be expected. For example, they will not be consistent with the usual post-Newtonian waveforms; using the raw waveforms to construct hybrids with PN waveforms would result in mismatches between the modes. Using raw waveforms to calibrate effective-one-body waveforms [49–51], surrogate models [54,55], or other phenomenological waveform models [52,53] would degrade the quality of the numerous fits inherent to the calibration process, by subjecting them to effectively random noise in the input. A broader and deeper survey of the effects of these transformations on waveforms in the SXS catalog will be the subject of an upcoming paper [64].

Essentially, we have a more general form of the familiar alignment problem in which arbitrary time and phase offsets need to be removed. Those simple alignments are just special cases of the one described here, restricted to the subgroup of BMS transformations consisting of time translations and rotations about the z axis. This more general alignment problem will necessitate using more general elements of the BMS group. With the algorithm presented in this paper, we can begin to investigate ways to achieve such alignment. Previous investigations have suggested ways of using asymptotic data to determine the center of mass, and more generally resolve the supertranslation ambiguity [33,56–58]. While these are promising theoretical developments, additional work will be

needed to make these methods practicable—towards which the present work is a crucial first step.

ACKNOWLEDGMENTS

It is my pleasure to thank David Nichols and Leo Stein for helpful comments on the paper draft; and Scott Field, Éanna Flanagan, Dan Hemberger, Larry Kidder, Richard O’Shaughnessy, Sergei Ossokine, Harald Pfeiffer, Mark Scheel, Patricia Schmidt, and Saul Teukolsky for useful conversations. This project was supported in part by the Sherman Fairchild Foundation and by NSF Grants No. PHY-1306125 and No. AST-1333129. The computations discussed in this paper were performed on the ZWICKY cluster hosted at Caltech by the Center for Advanced Computing Research, which was funded by the Sherman Fairchild Foundation and the NSF MRI-R² program; 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: CONVENTIONS

We start with some fiducial frame $(t, \mathbf{x}, \mathbf{y}, \mathbf{z})$, and some corresponding observer \mathcal{O} . A spacetime event is a point p , and is represented by some vector corresponding to the displacement from the origin of \mathcal{O} to that point. The point p can be given coordinates (p_t, p_x, p_y, p_z) such that its corresponding vector is $p_t \mathbf{t} + p_x \mathbf{x} + p_y \mathbf{y} + p_z \mathbf{z}$.

Another observer $\hat{\mathcal{O}}$ moves at velocity \mathbf{v} relative to \mathcal{O} , which means that the location of the spatial origin of $\hat{\mathcal{O}}$ relative to the (absolute) origin of \mathcal{O} is of the form $\eta(\mathbf{t} + \mathbf{v})$ for some η , where we assume the speed of light is $c = 1$. We also define the following shorthand notations:

$$\beta := |\mathbf{v}|, \quad (\text{A1a})$$

$$\varphi := \operatorname{artanh} \beta, \quad (\text{A1b})$$

$$\gamma := \frac{1}{\sqrt{1 - \beta^2}}, \quad (\text{A1c})$$

$$\boldsymbol{\varphi} := \varphi \frac{\mathbf{v}}{\beta}. \quad (\text{A1d})$$

It is worth noting the convenient identities,

$$\gamma \equiv \cosh \varphi, \quad (\text{A1e})$$

$$\beta \gamma \equiv \sinh \varphi, \quad (\text{A1f})$$

$$\gamma(1 + \beta) \equiv \cosh \varphi + \sinh \varphi \equiv e^\varphi, \quad (\text{A1g})$$

$$\gamma(1 - \beta) \equiv \cosh \varphi - \sinh \varphi \equiv e^{-\varphi}, \quad (\text{A1h})$$

$$\frac{1}{2} \ln \frac{1 + \beta}{1 - \beta} \equiv \varphi. \quad (\text{A1i})$$

The frame $(\hat{t}, \hat{x}, \hat{y}, \hat{z})$ of observer $\hat{\mathcal{O}}$ is *defined* by the relations

$$\hat{t} := B \mathbf{t} B^{-1}, \quad (\text{A2a})$$

$$\hat{x} := B \mathbf{x} B^{-1}, \quad (\text{A2b})$$

$$\hat{y} := B \mathbf{y} B^{-1}, \quad (\text{A2c})$$

$$\hat{z} := B \mathbf{z} B^{-1}, \quad (\text{A2d})$$

where B is a Lorentz rotor:

$$B := e^{-\boldsymbol{\varphi} t / 2} = \cosh \frac{\varphi}{2} - \frac{\boldsymbol{\varphi} t}{\varphi} \sinh \frac{\varphi}{2}. \quad (\text{A3})$$

Here, we use the formalism of geometric algebra [25–27] to describe the boost. In particular, the term $\boldsymbol{\varphi} t$ represents the geometric (or Clifford) product between these two vectors. Because $\boldsymbol{\varphi}$ is a spatial vector, this product $\boldsymbol{\varphi} t$ is a pure bivector $\boldsymbol{\varphi} \wedge t$ representing the hyperplane spanned by $\boldsymbol{\varphi}$ and t . The quantity B is a mild generalization of a unit quaternion (also called a rotor), except that now “rotations” need not be confined to spatial planes; the vectors spanning the plane of rotation can now include time components—as in this case. For simplicity, we have also specialized to the case where there is no additional rotation. If there is some additional rotation, it can easily be absorbed by redefining the frame of \mathcal{O} , or simply replacing every occurrence of B with $B R_f$, where R_f represents the required (purely spatial) rotation of the \mathcal{O} frame.

Using the form of B given above, we have $\hat{t} = \gamma(\mathbf{t} + \mathbf{v})$, which agrees with our earlier statement, because any point at the spatial origin of $\hat{\mathcal{O}}$ will be of the form $\eta \hat{t} = \eta \gamma(\mathbf{t} + \mathbf{v}) = \eta(\mathbf{t} + \mathbf{v})$ for some $\eta = \eta' \gamma$. We should also note that $\hat{\mathcal{O}}$ observes \mathcal{O} moving with velocity $-\hat{\mathbf{v}} = -B \mathbf{v} B^{-1} = -\gamma(\mathbf{v} + \beta^2 \mathbf{t})$, which is a purely spatial vector in $\hat{\mathcal{O}}$, with magnitude β .

APPENDIX B: SPHERICAL COORDINATES

Spherical coordinates are defined as usual, so that a point on the sphere at position \mathbf{r} has coordinates (θ, ϕ) when the angle between \mathbf{r} and \mathbf{z} is θ , and the angle between \mathbf{x} and the projection of \mathbf{r} onto the x - y plane is ϕ . Then, any point \mathbf{r} may be represented by a rotor $R_{\theta, \phi}$ as

$$\mathbf{r} = R_{\theta,\phi} \mathbf{z} R_{\theta,\phi}^{-1}, \quad (\text{B1})$$

where

$$R_{\theta,\phi} := e^{\phi \mathbf{y} \mathbf{x} / 2} e^{\theta \mathbf{x} \mathbf{z} / 2}. \quad (\text{B2})$$

Similarly, the observer $\hat{\mathcal{O}}$ can represent a point as

$$\hat{\mathbf{r}} = \hat{R}_{\hat{\theta},\hat{\phi}} \hat{\mathbf{z}} \hat{R}_{\hat{\theta},\hat{\phi}}^{-1}, \quad (\text{B3})$$

where

$$\hat{R}_{\hat{\theta},\hat{\phi}} := e^{\hat{\phi} \hat{\mathbf{y}} \hat{\mathbf{x}} / 2} e^{\hat{\theta} \hat{\mathbf{x}} \hat{\mathbf{z}} / 2} = B R_{\theta,\phi} B^{-1}. \quad (\text{B4})$$

Note that the final form above is written using basis vectors from the frame of \mathcal{O} , but uses the coordinates measured by $\hat{\mathcal{O}}$.

This presentation of spherical coordinates in terms of the corresponding rotor is useful, not only because the point itself may be expressed as in Eqs. (B1) and (B3), but also because the corresponding tangent vectors are easily expressed. For example, if $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$ are the standard tangent vectors, we have

$$\mathbf{t} = R_{\theta,\phi} \mathbf{t} R_{\theta,\phi}^{-1}, \quad (\text{B5a})$$

$$\boldsymbol{\theta} = R_{\theta,\phi} \mathbf{x} R_{\theta,\phi}^{-1}, \quad (\text{B5b})$$

$$\boldsymbol{\phi} = R_{\theta,\phi} \mathbf{y} R_{\theta,\phi}^{-1}, \quad (\text{B5c})$$

$$\mathbf{r} = R_{\theta,\phi} \mathbf{z} R_{\theta,\phi}^{-1}. \quad (\text{B5d})$$

This suggests the use of rotors more generally as a better way to keep track of a basis frame than retaining all four vectors separately. We also take this opportunity to define another frame:

$$\mathbf{t}' := R_{\hat{\theta},\hat{\phi}} \mathbf{t} R_{\hat{\theta},\hat{\phi}}^{-1} \equiv \mathbf{t}, \quad (\text{B6a})$$

$$\boldsymbol{\theta}' := R_{\hat{\theta},\hat{\phi}} \mathbf{x} R_{\hat{\theta},\hat{\phi}}^{-1}, \quad (\text{B6b})$$

$$\boldsymbol{\phi}' := R_{\hat{\theta},\hat{\phi}} \mathbf{y} R_{\hat{\theta},\hat{\phi}}^{-1}, \quad (\text{B6c})$$

$$\mathbf{r}' := R_{\hat{\theta},\hat{\phi}} \mathbf{z} R_{\hat{\theta},\hat{\phi}}^{-1}. \quad (\text{B6d})$$

In the frame of \mathcal{O} , the last three are pure spatial vectors. Once they are boosted they will be pure spatial vectors in the frame of $\hat{\mathcal{O}}$ and, along with $\hat{\mathbf{t}}$, will comprise the correct frame for a point on the sphere at coordinates $(\hat{\theta}, \hat{\phi})$, as measured by $\hat{\mathcal{O}}$:

$$\hat{\mathbf{t}} = B \mathbf{t}' B^{-1}, \quad (\text{B7a})$$

$$\hat{\boldsymbol{\theta}} = B \boldsymbol{\theta}' B^{-1}, \quad (\text{B7b})$$

$$\hat{\boldsymbol{\phi}} = B \boldsymbol{\phi}' B^{-1}, \quad (\text{B7c})$$

$$\hat{\mathbf{r}} = B \mathbf{r}' B^{-1}. \quad (\text{B7d})$$

This shows that we can extend Eq. (B5) to use Lorentz rotors (generalizing from pure spatial rotors) to keep track of all possible basis frames related by a Lorentz transformation, rather than retaining all four vectors separately.

APPENDIX C: ROTOR OF A BOOST

Any future-directed null vector may be represented by \mathcal{O} up to a positive scaling as

$$\mathbf{l} := \mathbf{r} + \mathbf{t} \equiv R_{\theta,\phi} (\mathbf{z} + \mathbf{t}) R_{\theta,\phi}^{-1}. \quad (\text{C1})$$

Note that the rotor has no effect on \mathbf{t} , as it is an entirely spatial rotor. Similarly, observer $\hat{\mathcal{O}}$ may express any future-directed null vector via

$$\hat{\mathbf{l}} := \hat{\mathbf{r}} + \hat{\mathbf{t}} \equiv \hat{R}_{\hat{\theta},\hat{\phi}} (\hat{\mathbf{z}} + \hat{\mathbf{t}}) \hat{R}_{\hat{\theta},\hat{\phi}}^{-1} \quad (\text{C2a})$$

$$= B R_{\theta,\phi} (\mathbf{z} + \mathbf{t}) R_{\theta,\phi}^{-1} B^{-1}. \quad (\text{C2b})$$

The final expression represents the conjugation by B of a vector expressed entirely in the basis of \mathcal{O} , though using coordinates as measured by $\hat{\mathcal{O}}$.

We need to know the coordinates (θ, ϕ) given $(\hat{\theta}, \hat{\phi})$ such that \mathbf{l} is a positive scalar multiple of $\hat{\mathbf{l}}$, which is possible if and only if $\hat{\mathbf{l}} \cdot \mathbf{l} = 0$. (Again, juxtaposition of the vectors \mathbf{l} and $\hat{\mathbf{l}}$ denotes the geometric product.) For now, let us assume that B is a boost along \mathbf{z} . Then clearly $\phi = \hat{\phi}$ since \mathbf{y} and \mathbf{x} are unaffected. We can calculate

$$\hat{\mathbf{u}} = R_{\theta,\phi} (\mathbf{z} + \mathbf{t}) R_{\theta,\phi}^{-1} B R_{\hat{\theta},\hat{\phi}} (\hat{\mathbf{z}} + \hat{\mathbf{t}}) R_{\hat{\theta},\hat{\phi}}^{-1} B^{-1} \quad (\text{C3a})$$

$$= e^{\theta \mathbf{x} \mathbf{z} / 2} (\mathbf{z} + \mathbf{t}) e^{-\theta \mathbf{x} \mathbf{z} / 2} B e^{\hat{\theta} \mathbf{x} \mathbf{z} / 2} (\hat{\mathbf{z}} + \hat{\mathbf{t}}) e^{-\hat{\theta} \mathbf{x} \mathbf{z} / 2} B^{-1}. \quad (\text{C3b})$$

The key expression here is the Lorentz rotor

$$L = e^{-\theta \mathbf{x} \mathbf{z} / 2} e^{-\varphi \mathbf{z} \mathbf{t} / 2} e^{\hat{\theta} \mathbf{x} \mathbf{z} / 2} \quad (\text{C4a})$$

$$= \cosh \frac{\varphi}{2} \cos \frac{\theta - \hat{\theta}}{2} - \mathbf{z} \mathbf{t} \sinh \frac{\varphi}{2} \cos \frac{\theta - \hat{\theta}}{2} \\ + \mathbf{x} \mathbf{t} \sinh \frac{\varphi}{2} \sin \frac{\theta + \hat{\theta}}{2} - \mathbf{x} \mathbf{z} \cosh \frac{\varphi}{2} \sin \frac{\theta - \hat{\theta}}{2}. \quad (\text{C4b})$$

In particular, if $\hat{\mathbf{l}}$ is a scalar, we have $\hat{\mathbf{l}} = (z + \mathbf{t})L(z + \mathbf{t})L^{-1}$, and if the latter expression is to be a scalar, $L(z + \mathbf{t})L^{-1}$ must have no \mathbf{x} component. A simple argument from geometric algebra shows that L can only have terms involving \mathbf{x} of the form $\mathbf{x}(z + \mathbf{t})$; terms of the form $\mathbf{x}(z - \mathbf{t})$ must vanish. Using the coefficients of $\mathbf{x}\mathbf{t}$ and $\mathbf{x}\mathbf{z}$ above, some simple algebra shows us that this implies that

$$\tan \frac{\theta}{2} = e^{-\varphi} \tan \frac{\hat{\theta}}{2}. \quad (\text{C5a})$$

We can repeat this analysis for a past-directed null vector, and find the condition that $L(z - \mathbf{t})L^{-1}$ must have no \mathbf{x} component, which implies that

$$\tan \frac{\theta}{2} = e^{\varphi} \tan \frac{\hat{\theta}}{2}. \quad (\text{C5b})$$

This is equivalent to flipping the sign or direction of the boost in Eq. (C5a). Note that Eq. (C5b) is the standard formula for stellar aberration due to a boost,¹² because an observer detects photons moving into the future along past-directed null vectors. Put another way, an observer receiving null rays assigns a direction to a ray according to where it *came from*, rather than where it *is going*; an emitter assigns directions according to where the ray is going, rather than where it would have come from—this is the reason for the sign difference.

By looking at this more geometrically, we can eliminate the requirement that the boost be in the z direction. We first dispense with the trivial case for which \mathbf{v} and \mathbf{r} are parallel or antiparallel, in which case $\mathbf{r} = \mathbf{r}'$. Assuming henceforth the situation is not so trivial, we note that \mathbf{v} , \mathbf{r} , and \mathbf{r}' all lie in the same plane, and angles between them are governed by Eq. (C5). To be specific, define $\hat{\Theta}$ to be the angle measured by $\hat{\Theta}$ between $\hat{\mathbf{v}}$ and $\hat{\mathbf{r}}$, and similarly for Θ . We can calculate $\hat{\Theta}$ in the \mathcal{O} frame as

$$\hat{\Theta} := \arccos[\hat{\mathbf{v}} \cdot \hat{\mathbf{r}}] \equiv \arccos \left[\mathbf{v} \cdot \left(R_{\hat{\theta}, \hat{\phi}} z R_{\hat{\theta}, \hat{\phi}}^{-1} \right) \right]. \quad (\text{C6})$$

The corresponding value of Θ for future-directed (respectively past-directed) null rays is simply

$$\tan \frac{\Theta}{2} = e^{\mp \varphi} \tan \frac{\hat{\Theta}}{2}. \quad (\text{C7})$$

Using this equation, we can find another useful relation between \mathbf{l} and $\hat{\mathbf{l}}$: the latter can be rotated into the former with a rotation that is purely spatial in \mathcal{O} . Essentially, we

simply rotate by $\Theta - \hat{\Theta}$ in the $\mathbf{r}\text{-}\mathbf{v}$ plane. The rotor that does this is

$$B' := \exp \left[\frac{\Theta - \hat{\Theta}}{2} \frac{\mathbf{r} \wedge \mathbf{v}}{|\mathbf{r} \wedge \mathbf{v}|} \right]. \quad (\text{C8})$$

With this rotor, we have

$$R_{\theta, \phi}(z + \mathbf{t})R_{\theta, \phi}^{-1} = B'R_{\hat{\theta}, \hat{\phi}}(z + \mathbf{t})R_{\hat{\theta}, \hat{\phi}}^{-1}B'^{-1}. \quad (\text{C9})$$

Note that this equation does *not* imply $R_{\theta, \phi} = B'R_{\hat{\theta}, \hat{\phi}}$; instead we have¹³

$$R_{\theta, \phi} e^{\lambda \mathbf{x}\mathbf{y}/2} = B'R_{\hat{\theta}, \hat{\phi}}, \quad (\text{C10})$$

for some angle λ . It turns out that this angle is the spin phase described in Sec. II A. Though it will never be necessary to compute this directly (except for the purposes of visualizations like Fig. 2), we can rearrange Eq. (C11) and express λ as

$$\lambda = 2 \log [R_{\theta, \phi}^{-1} B' R_{\hat{\theta}, \hat{\phi}}] \mathbf{y}\mathbf{x}. \quad (\text{C11})$$

Of course, rather than computing this angle to evaluate spin-weighted functions, we can just use the right-hand side of Eq. (C10) directly, and evaluate the spin-weighted function on that rotor.

It may be helpful to see why this spin phase is a meaningful quantity under a boost. The fact that \mathbf{v} , \mathbf{r} , and \mathbf{r}' lie in the same plane and the fact that angles between them are governed by Eq. (C7) are purely geometric statements; they are independent of our basis frame. We can use these facts to express the value of a spin-weighted function in the boosted frame in terms of the spin-weighted function in the original frame. Assuming \mathbf{v} and \mathbf{r} are not proportional to each other, we know that the products $\mathbf{t} \wedge \mathbf{v} \wedge \mathbf{r}$ and $\mathbf{t} \wedge \mathbf{v} \wedge \mathbf{r}'$ are the same up to some nonzero scalar multiple; they represent the same hyperplane. Under the boost $e^{-\varphi \mathbf{t}/2}$, the three vectors \mathbf{t} , \mathbf{v} , and \mathbf{r}' transform among themselves, which means that $\hat{\mathbf{t}} \wedge \hat{\mathbf{v}} \wedge \hat{\mathbf{r}}$ also represents the same hyperplane.¹⁴ There is a unique axis orthogonal to this hyperplane. In fact, we can construct a unique unit vector $\hat{\Phi}$ along this axis by defining

¹³We know that the extra factor in this equation is the most general possible such factor because: (1) it must be an even-grade element of unit norm, since all other elements of this equation have even grade and unit norm; (2) it must be purely spatial in \mathcal{O} , since all other elements are purely spatial; (3) it must commute with $z + \mathbf{t}$. Thus, it can only be a rotation in the $x\text{-}y$ plane.

¹⁴That is, the latter product is the same as the former two products up to some other nonzero scalar multiple. In fact, a straightforward calculation shows that $\csc \hat{\Theta} \hat{\mathbf{t}} \wedge \hat{\mathbf{v}} \wedge \hat{\mathbf{r}} = \csc \Theta \mathbf{t} \wedge \mathbf{v} \wedge \mathbf{r}$.

¹²See, e.g., Eq. (1.3.5) of Ref. [35].

$$\Phi := \frac{\mathbf{t} \wedge \mathbf{v} \wedge \mathbf{r}}{\beta \sin \Theta} \mathbf{t} \wedge \boldsymbol{\theta} \wedge \boldsymbol{\phi} \wedge \mathbf{r}, \quad (\text{C12a})$$

$$= \frac{\mathbf{t} \wedge \mathbf{v} \wedge \mathbf{r}'}{\beta \sin \hat{\Theta}} \mathbf{t} \wedge \boldsymbol{\theta} \wedge \boldsymbol{\phi} \wedge \mathbf{r}, \quad (\text{C12b})$$

$$= \frac{\hat{\mathbf{t}} \wedge \hat{\mathbf{v}} \wedge \hat{\mathbf{r}}}{\beta \sin \hat{\Theta}} \hat{\mathbf{t}} \wedge \hat{\boldsymbol{\theta}} \wedge \hat{\boldsymbol{\phi}} \wedge \hat{\mathbf{r}}. \quad (\text{C12c})$$

When $\mathbf{v} = \beta\mathbf{z}$, we have $\Phi = \boldsymbol{\phi}$ —the usual basis vector. But the definition given in Eq. (C12) is geometrically invariant. By construction Φ is orthogonal to $\mathbf{v} \wedge \mathbf{t}$, and so is invariant under boosts. More specifically, Φ is a purely spatial vector for both observers, orthogonal to the velocity, and lies in the tangent space of the sphere at $\hat{\mathbf{r}}$ for $\hat{\mathcal{O}}$ and at \mathbf{r} for \mathcal{O} . We can therefore use it to compare directions in the tangent spaces for our spin-weighted functions.

These invariance properties of the $\Phi = \hat{\Phi}$ vector field allow us to identify the alignment of the tangent space. We choose the point on the sphere designated by $\hat{\mathbf{r}}$, with coordinates $(\hat{\theta}, \hat{\phi})$. This has a standard [39] alignment of the tangent space given by

$$\hat{\mathbf{m}} := \frac{1}{\sqrt{2}}(\hat{\boldsymbol{\theta}} + i\hat{\boldsymbol{\phi}}) \equiv BR_{\hat{\theta}, \hat{\phi}} \frac{\mathbf{x} + i\mathbf{y}}{\sqrt{2}} R_{\hat{\theta}, \hat{\phi}}^{-1} B^{-1}. \quad (\text{C13a})$$

Because this is a purely spatial vector in the frame of $\hat{\mathcal{O}}$, but has a time component in the frame of \mathcal{O} , direct comparison would be complicated. However, we can define the similar vectors

$$\mathbf{m}' := \frac{1}{\sqrt{2}}(\boldsymbol{\theta}' + i\boldsymbol{\phi}') \equiv B'R_{\theta, \phi} \frac{\mathbf{x} + i\mathbf{y}}{\sqrt{2}} R_{\theta, \phi}^{-1} B'^{-1}, \quad (\text{C13b})$$

$$\mathbf{m} := \frac{1}{\sqrt{2}}(\boldsymbol{\theta} + i\boldsymbol{\phi}) \equiv R_{\theta, \phi} \frac{\mathbf{x} + i\mathbf{y}}{\sqrt{2}} R_{\theta, \phi}^{-1}, \quad (\text{C13c})$$

and the products

$$\hat{m}_{\hat{\Phi}} := \hat{\mathbf{m}} \cdot \hat{\Phi}, \quad (\text{C14a})$$

$$m'_{\Phi} := \mathbf{m}' \cdot \Phi, \quad (\text{C14b})$$

$$m_{\Phi} := \mathbf{m} \cdot \Phi. \quad (\text{C14c})$$

Almost by definition, we have $\hat{m}_{\hat{\Phi}} \equiv m'_{\Phi}$. Thus, the non-trivial comparison is between m'_{Φ} and m_{Φ} .

To make this comparison, suppose that \mathbf{l} and \mathbf{l}' are as given in Eqs. (C1) and (C2). We know that the rotors involved in those expressions are related by Eq. (C10), so we can calculate the relative alignment of the tangent spaces as follows:

$$m'_{\Phi} := \mathbf{m}' \cdot \Phi, \quad (\text{C15a})$$

$$= \left(B'R_{\theta, \hat{\phi}} \frac{\mathbf{x} + i\mathbf{y}}{\sqrt{2}} \bar{R}_{\theta, \hat{\phi}} \bar{B}' \right) \cdot \Phi, \quad (\text{C15b})$$

$$= \left(R_{\theta, \phi} e^{\lambda xy/2} \frac{\mathbf{x} + i\mathbf{y}}{\sqrt{2}} e^{-\lambda xy/2} \bar{R}_{\theta, \phi} \right) \cdot \Phi, \quad (\text{C15c})$$

$$= e^{i\lambda} \left(R_{\theta, \phi} \frac{\mathbf{x} + i\mathbf{y}}{\sqrt{2}} \bar{R}_{\theta, \phi} \right) \cdot \Phi, \quad (\text{C15d})$$

$$= e^{i\lambda} m_{\Phi}. \quad (\text{C15e})$$

This relation is exactly the one implied by Newman and Penrose's original definition of spin at the beginning of Sec. III in Ref. [16]: they defined spin with respect to the transformation $\mathbf{m}' = e^{i\lambda} \mathbf{m}$. This describes the relative alignment of the $\hat{\mathbf{m}}$ and \mathbf{m} fields—except for any time component orthogonal to Φ produced by the relative boost. Those additional components cannot be accounted for simply by a rotation; they must be accounted for by mixing between different components of the tensor in question. This is why we find various Newman-Penrose scalars on the right-hand sides of Eqs. (17a) through (17d), for example.

This property of rotating the tangent space is very important, and is the primary motivation for this more geometric approach. That is because we are dealing with spin-weighted functions, which means that we need to know not only how points move around on the sphere, but also how the tangent space to the sphere changes at each point under a boost. Using rotors allows us to automatically track both the change of position and the change of the tangent basis.

APPENDIX D: CONFORMAL FACTOR OF A BOOST

We can use spatial directions to label all the null directions from a point, which has the topology of a sphere. We define the metric on this null sphere as the metric induced on the sphere of spatial directions. In that case, a boost induces a conformal transformation of the null sphere, which means that we can find the conformal factor of the boost. In particular, for future-directed null rays, by simply applying the transformation of θ given by Eq. (C5a) we can calculate

$$d\hat{\theta}^2 + \sin^2 \hat{\theta} d\hat{\phi}^2 = \left[\frac{1}{\gamma(1 - \mathbf{v} \cdot \mathbf{r})} \right]^2 (d\theta^2 + \sin^2 \theta d\phi^2), \quad (\text{D1a})$$

and similarly for past-directed null rays with Eq. (C5b),

$$d\hat{\theta}^2 + \sin^2\hat{\theta}d\hat{\phi}^2 = \left[\frac{1}{\gamma(1 + \mathbf{v} \cdot \mathbf{r})} \right]^2 (d\theta^2 + \sin^2\theta d\phi^2). \quad (\text{D1b})$$

So we define the conformal factors for future-directed (k_+) and past-directed (k_-) null spheres as

$$k_{\pm} := \frac{1}{\gamma(1 \mp \mathbf{v} \cdot \mathbf{r})}. \quad (\text{D2})$$

We know that the form of the metric is invariant under rotations, and the form of this conformal factor is clearly invariant under rotations, so this is the correct conformal factor for boosts in any direction. In this paper, we are always dealing with future-directed null rays, so we drop the subscript and just use $k := k_+$.

APPENDIX E: ESTIMATING TRANSLATION AND BOOST IN SIMULATIONS

The coordinate center of mass of a simulated compact binary presents an imperfect representation of its motion. Obviously, this can be tainted by gauge effects, especially because the data are drawn from the most dynamical and nonlinear portion of the simulated spacetime. While this may be a topic ripe for improvement, it is nonetheless useful to have some way to illustrate the methods of this paper for the waveform data of the SXS catalog. In that spirit, this Appendix presents a simple method for estimating the translation and boost, given the coordinate tracks and Christodoulou masses of the black holes. As noted in Sec. VI, the data can be obtained from the HORIZONS.H5 file accompanying each waveform in the SXS catalog.

Denoting by $\mathbf{x}_{\text{CoM}}(t)$ the coordinate location of the center of mass, as a function of the coordinate time, and in units where the total mass of the system is 1, we can define the quantity

$$\Xi(\delta\mathbf{x}, \mathbf{v}) = \int_{t_i}^{t_f} |\mathbf{x}_{\text{CoM}} - (\delta\mathbf{x} + \mathbf{v}t)|^2 dt. \quad (\text{E1})$$

This measures the distance between the origin and the center of mass of a system transformed by $(\delta\mathbf{x}, \mathbf{v})$, integrated over some range of times. We can minimize this quantity over the transformation to find the optimal

transformation. This minimum can be found analytically by defining two moments of the center of mass:

$$\mathbf{x}_0 = \int_{t_i}^{t_f} \mathbf{x}_{\text{CoM}}(t) dt \quad \text{and} \quad \mathbf{x}_1 = \int_{t_i}^{t_f} t \mathbf{x}_{\text{CoM}}(t) dt. \quad (\text{E2})$$

Then, the minimum is given by

$$\delta\mathbf{x} = \frac{4(t_f^2 + t_f t_i + t_i^2)\mathbf{x}_0 - 6(t_f + t_i)\mathbf{x}_1}{(t_f - t_i)^3}, \quad (\text{E3a})$$

$$\mathbf{v} = \frac{12\mathbf{x}_1 - 6(t_f + t_i)\mathbf{x}_0}{(t_f - t_i)^3}. \quad (\text{E3b})$$

The moments can be computed by numerical integration of the data, and simply plugged into these formulas to find the desired transformation.

The only free parameters in this case are the limits of integration, t_i and t_f . In principle, these could span the entire time for which there are two separate apparent horizons in the data. In some cases, as when a simulation needs to be aligned with another simulation or an analytical waveform, for example, it would likely be better to restrict this time span to the same times over which the waveforms are being aligned. In this case, however, where we are simply interested in finding estimates for the motion of the systems, we can be somewhat more liberal. The initial time should be delayed slightly, to allow junk radiation to settle down, so that the black holes can be measured accurately. For simplicity and definiteness, we will set t_i to be 1% of the entire time for which data is available. On the other hand, bizarre features are sometimes present in the SXS catalog close to merger. To avoid these, and to lessen the impact of true physical recoils that develop close to merger, we similarly set t_f to be 10% before the end of the data.

With this simple recipe in hand, we can apply it to the entire SXS catalog very easily. The results are shown for SXS:BBH:0004 in Eq. (46), and are aggregated for the entire catalog in Fig. 11. Again, this is a very crude and gauge-sensitive measure of the motion of the system. It should no doubt be improved in future work. But for the purposes of illustration in this paper, it seems to be sufficient.

-
- [1] L. S. Finn, *Phys. Rev. D* **46**, 5236 (1992).
 [2] C. Cutler and É. E. Flanagan, *Phys. Rev. D* **49**, 2658 (1994).
 [3] A. Buonanno, G. B. Cook, and F. Pretorius, *Phys. Rev. D* **75**, 124018 (2007).

- [4] M. Boyle, D. A. Brown, L. E. Kidder, A. H. Mroue, H. P. Pfeiffer, M. A. Scheel, G. B. Cook, and S. A. Teukolsky, *Phys. Rev. D* **76**, 124038 (2007).
 [5] P. Schmidt, M. Hannam, S. Husa, and P. Ajith, *Phys. Rev. D* **84**, 024046 (2011).

- [6] R. O’Shaughnessy, B. Vaishnav, J. Healy, Z. Meeks, and D. Shoemaker, *Phys. Rev. D* **84**, 124002 (2011).
- [7] M. Boyle, R. Owen, and H. P. Pfeiffer, *Phys. Rev. D* **84**, 124011 (2011).
- [8] M. Boyle, L. E. Kidder, S. Ossokine, and H. P. Pfeiffer, [arXiv:1409.4431](https://arxiv.org/abs/1409.4431).
- [9] H. Bondi, J. G. M. van der Burg, and A. W. K. Metzner, *Proc. R. Soc. A* **269**, 21 (1962).
- [10] R. K. Sachs, *Proc. R. Soc. A* **270**, 103 (1962).
- [11] E. T. Newman and T. W. J. Unti, *J. Math. Phys. (N.Y.)* **3**, 891 (1962).
- [12] O. Moreschi, *Classical Quantum Gravity* **4**, 1063 (1987).
- [13] E. Deadman and J. M. Stewart, *Classical Quantum Gravity* **26**, 065008 (2009).
- [14] A. D. Helfer, *Phys. Rev. D* **81**, 084001 (2010).
- [15] R. Sachs, *Phys. Rev.* **128**, 2851 (1962).
- [16] E. T. Newman and R. Penrose, *J. Math. Phys.* **7**, 863 (1966).
- [17] J. Frauendiener, *Living Rev. Relativity* **7**, 1 (2004).
- [18] T. M. Adamo, C. Kozameh, and E. T. Newman, *Living Rev. Relativity* **12**, 6 (2009).
- [19] L. Gualtieri, E. Berti, V. Cardoso, and U. Sperhake, *Phys. Rev. D* **78**, 044024 (2008).
- [20] B. J. Kelly and J. G. Baker, *Phys. Rev. D* **87**, 084004 (2013).
- [21] A. H. Mroué, M. A. Scheel, B. Szilágyi, H. P. Pfeiffer, M. Boyle, D. A. Hemberger, L. E. Kidder, G. Lovelace, S. Ossokine, N. W. Taylor, A. Zenginoğlu, L. T. Buchman, T. Chu, E. Foley, M. Giesler, R. Owen, and S. A. Teukolsky, *Phys. Rev. Lett.* **111**, 241104 (2013).
- [22] SXS gravitational waveform database, 2015; <http://www.black-holes.org/waveforms/>.
- [23] S. Ossokine, F. Foucart, H. P. Pfeiffer, M. Boyle, and B. Szilágyi, *Classical Quantum Gravity* **32**, 245010 (2015).
- [24] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevD.93.084031> for python code to implement BMS transformations.
- [25] D. Hestenes, *New Foundations for Classical Mechanics*, 2nd ed. (Kluwer Academic Publishers, Boston, MA, 2002).
- [26] C. Doran and A. Lasenby, *Geometric Algebra for Physicists*, 4th ed. (Cambridge University Press, Cambridge, 2010).
- [27] D. Hestenes, *Space-Time Algebra*, 2nd ed. (Springer, New York, 2015).
- [28] K. M. Huffenberger and B. D. Wandelt, *Astrophys. J. Suppl. Ser.* **189**, 255 (2010).
- [29] K. M. Huffenberger, Fast and exact spin-s spherical harmonic transforms—spinsfast code (2013); <http://astrophysics.physics.fsu.edu/~huffenbe/research/spinsfast/index.html>.
- [30] R. Penrose, *Proc. R. Soc. A* **284**, 159 (1965).
- [31] S. W. Hawking and G. F. R. Ellis, *The Large Scale Structure of Space-Time* (Cambridge University Press, Cambridge, 1976).
- [32] S. Dain and O. M. Moreschi, *Classical Quantum Gravity* **17**, 3663 (2000).
- [33] C. N. Kozameh and G. D. Quiroga, [arXiv:1311.5854](https://arxiv.org/abs/1311.5854).
- [34] O. M. Moreschi, *Classical Quantum Gravity* **3**, 503 (1986).
- [35] R. Penrose and W. Rindler, *Spinors and Space-Time Vol. 1: Two-Spinor Calculus and Relativistic Fields* (Cambridge University Press, Cambridge, 1987).
- [36] E. Newman and R. Penrose, *J. Math. Phys.* **4**, 998 (1963).
- [37] R. Geroch, A. Held, and R. Penrose, *J. Math. Phys.* **14**, 874 (1973).
- [38] P. O’Donnell, *Introduction to Two Spinors in General Relativity* (World Scientific Publishing Co., River Edge, NJ, 2003).
- [39] P. Ajith, M. Boyle, D. A. Brown, S. Fairhurst, M. Hannam, I. Hinder, S. Husa, B. Krishnan, R. A. Mercer, F. Ohme, C. D. Ott, J. S. Read, L. Santamaria, and J. T. Whelan, [arXiv:0709.0093](https://arxiv.org/abs/0709.0093).
- [40] L. A. Tamburino and J. H. Winicour, *Phys. Rev.* **150**, 1039 (1966).
- [41] K. S. Thorne, in *Gravitational Radiation*, edited by N. Deruelle and T. Piran (North-Holland, Amsterdam, 1983), pp. 1–57.
- [42] A. D. Helfer, *Phys. Rev. D* **48**, 3625 (1993).
- [43] J. N. Goldberg, A. J. MacFarlane, E. T. Newman, F. Rohrlich, and E. C. G. Sudarshan, *J. Math. Phys.* **8**, 2155 (1967).
- [44] L. Blanchet, *Living Rev. Relativity* **9**, 4 (2006).
- [45] D. Christodoulou, *Phys. Rev. Lett.* **25**, 1596 (1970).
- [46] N. W. Taylor, M. Boyle, C. Reisswig, M. A. Scheel, T. Chu, L. E. Kidder, and B. Szilágyi, *Phys. Rev. D* **88**, 124010 (2013).
- [47] M. Boyle and A. H. Mroue, *Phys. Rev. D* **80**, 124045 (2009).
- [48] P. Kumar, I. MacDonald, D. A. Brown, H. P. Pfeiffer, K. Cannon, M. Boyle, L. E. Kidder, A. H. Mroué, M. A. Scheel, B. Szilágyi, and A. Zenginoğlu, *Phys. Rev. D* **89**, 042002 (2014).
- [49] A. Taracchini, A. Buonanno, Y. Pan, T. Hinderer, M. Boyle, D. A. Hemberger, L. E. Kidder, G. Lovelace, A. H. Mroué, H. P. Pfeiffer, M. A. Scheel, B. Szilágyi, N. W. Taylor, and A. Zenginoğlu, *Phys. Rev. D* **89**, 061502 (2014).
- [50] T. Damour and A. Nagar, *Phys. Rev. D* **90**, 044018 (2014).
- [51] A. Nagar, T. Damour, C. Reisswig, and D. Pollney, *Phys. Rev. D* **93**, 044046 (2014).
- [52] L. Santamaría, F. Ohme, P. Ajith, B. Brügmann, N. Dorband, M. Hannam, S. Husa, P. Mösta, D. Pollney, C. Reisswig, E. L. Robinson, J. Seiler, and B. Krishnan, *Phys. Rev. D* **82**, 064016 (2010).
- [53] T. Damour and A. Nagar, *Phys. Rev. D* **90**, 024054 (2014).
- [54] S. E. Field, C. R. Galley, J. S. Hesthaven, J. Kaye, and M. Tiglio, *Phys. Rev. X* **4**, 031006 (2014).
- [55] J. Blackman, S. E. Field, C. R. Galley, B. Szilágyi, M. A. Scheel, M. Tiglio, and D. A. Hemberger, *Phys. Rev. Lett.* **115**, 121102 (2015).
- [56] O. M. Moreschi, *Classical Quantum Gravity* **21**, 5409 (2004).
- [57] A. D. Helfer, *Gen. Relativ. Gravit.* **39**, 2125 (2007).
- [58] T. M. Adamo and E. T. Newman, *Phys. Rev. D* **83**, 044023 (2011).
- [59] D. A. Brown, Ph.D. thesis, The University of Wisconsin-Milwaukee, 2007.
- [60] M. Boyle, *Phys. Rev. D* **84**, 064013 (2011).
- [61] C. Capano, Y. Pan, and A. Buonanno, *Phys. Rev. D* **89**, 102003 (2014).
- [62] R. O’Shaughnessy, B. Farr, E. Ochsner, H.-S. Cho, V. Raymond, C. Kim, and C.-H. Lee, *Phys. Rev. D* **89**, 102005 (2014).
- [63] K. Kyutoku and N. Seto, *Mon. Not. R. Astron. Soc.* **441**, 1934 (2014).
- [64] P. Schmidt, D. A. Hemberger, M. Boyle, and M. A. Scheel, (unpublished).