

Nyquist-resolving gravitational waves via orbital frequency-based refinement

William K. Black^{✉,*}, David Neilsen[✉], and Eric W. Hirschmann[✉]

*Department of Physics and Astronomy, Brigham Young University,
N283 Carl F. Eyring Science Center, Provo, Utah 84602, USA*

David F. Van Komen[✉]

*Kahlert School of Computing, University of Utah,
50 Central Campus Dr., Room 3190, Salt Lake City, Utah 84112, USA*

Milinda Fernando[✉]

Oden Institute, The University of Texas at Austin, 201 East 24th Street, Austin, Texas 78712, USA



(Received 18 March 2025; accepted 21 April 2025; published 3 June 2025)

Adaptive mesh refinement efficiently facilitates the computation of gravitational waveforms in numerical relativity. However, determining precisely when, where, and to what extent to refine when solving the Einstein equations poses challenges; several *ad hoc* refinement criteria have been explored in the literature. This work introduces an optimized resolution baseline derived *in situ* from the inspiral trajectory (ORBIT). This method uses the binary’s orbital frequency as a proxy for anticipated gravitational waves to dynamically refine the grid, satisfying the Nyquist frequency requirements on grid resolution up to a specified spin-weighted spherical harmonic order. ORBIT sustains propagation of gravitational waves while avoiding the more costly alternative of maintaining high resolution across an entire simulation—both spatially and temporally. We find that enabling ORBIT decreases waveform noise by an order of magnitude and better resolves high-order wave amplitudes through merger. Combined with WAMR and other improvements, updates to DENDRO-GR decrease waveform noise, decrease constraint violations, and boost refinement efficiency each by factors of $\mathcal{O}(100)$, while reducing computational cost by a factor of 4. ORBIT and other recent improvements to DENDRO-GR begin to prepare us for gravitational wave science with next-generation detectors.

DOI: [10.1103/PhysRevD.111.124001](https://doi.org/10.1103/PhysRevD.111.124001)

I. INTRODUCTION

Mergers of black holes (BHs) and neutron stars in binary inspirals generate gravitational waves (GWs) detectable by the LIGO and Virgo GW observatories. By matching observed signals to GW templates—precomputed for likely GW sources—we can estimate the likely progenitors of those GW events. The reliability of binary merger parameter estimation hinges on the accuracy, precision, and coverage of those templates.

The next generation of GW detectors, including LISA [1], Cosmic Explorer [2], and the Einstein Telescope [3], are being planned for the next decade. These new detectors will be over an order of magnitude more accurate than current detectors (with higher signal-to-noise ratio), making them not only more sensitive to a wider range of GW sources, but also able to detect much longer portions of the inspiral signal. Each of these advances in GW science places

new demands on the extent and quality of the waveform libraries. We need waveform templates that are at least an order of magnitude more accurate, for a larger class of possible binaries, and for longer portions of the inspiral [4,5].

The nonlinear equations of motion for gravity make calculating waveform templates for astrophysical sources challenging. While some methods approximate solutions for the early inspiral or late ringdown of binary mergers, the only consistent method to calculate the full waveform is to compute solutions of the full nonlinear equations with numerical relativity (NR) codes. Computing NR waveforms is challenging and expensive; even for the lower accuracy needed by today’s detectors, waveforms have only been computed for a relatively small part of the binary merger parameter space [see, e.g., [6]]. The task of computing waveforms of sufficient accuracy, duration, and variety (in mass ratio, spin, orbital eccentricities) to satisfy the requirements of next generation GW detectors is beyond the scope of current NR codes [4,5]. Solving this problem will require novel computational approaches.

*Contact author: wkblack@byu.edu

Numerical solutions of binary inspirals require adequate resolution on multiple scales, ranging from the size of the compact objects to their orbital scale to the wavezone to the outer boundary of the computational domain. In dynamical problems, adaptive mesh refinement (AMR) efficiently handles these multiple scales by adjusting grid resolution to concentrate computational resources close to the compact objects and in regions of strong gravitational wave emission. This approach maintains high accuracy in critical regions and lower resolution in less dynamically important regions, optimizing computational efficiency. In practice, however, determining precisely where, when, and to what degree to refine can challenge and defy naive intuition.

Refinement strategies come with various levels of complexity and dependence on the solution of the equations solved. Broadly speaking, some refinement criteria use specific features known to occur in a solution, such as the location of a shock in hydrodynamics or a black hole's location in NR, while others rely on measures of the solution's convergence (which criteria can be solution independent). As an example of the first type of refinement criteria, NR codes can directly track black hole (BH) apparent horizons and puncture locations. Refinement can then be focused around these locations to ensure a certain resolution on the horizon and in its vicinity, structuring a nested grid of concentric spheres or boxes defined by empirically determined parameters [see, e.g. [7–11]]. Other proxies (e.g., functions of the conformal factor, the lapse, and the Hamiltonian and momentum constraints, and/or their derivatives) have also been widely used to determine refinement in binary inspiral simulations [see, e.g., [10]].

Refinement criteria based on solution quality have also been widely used in numerical relativity. Choptuik [12–15] was the first to use Berger-Oliger AMR [16] in NR, which estimates convergence using Richardson extrapolation to guide refinement. Choptuik also advocated for a more efficient measure of self-convergence using a shadow hierarchy [17,18], which has been widely implemented (see [19] for a partial bibliography). Recently, Rashti *et al.* [11] monitored solution convergence using Lagrange interpolation. Radia *et al.* [10] describes another method, monitoring truncation error between refinement levels as a criteria for additional refinement. Our approach to monitoring solution convergence uses a wavelet expansion of the solution variables. The wavelet basis is generated by shifting and scaling a mother wavelet, giving a compact basis set that naturally generates spatial refinement. Wavelet adaptive multiresolution (WAMR) expands in interpolating wavelets [20–22]; we implement this method in DENDRO-GR [23,24].

A practical AMR scheme almost always needs a combination of the two broad approaches listed above. Refinement schemes based on a few specific features can miss other important aspects of the solution, while general

schemes can waste resources by refining on physically unimportant regions or aspects of the solution. Thus, there always remains an element of intuition and experience in determining precisely how to parametrize each of these methods. This paper highlights our work using such a hybrid approach to refinement.

The current work introduces a novel refinement strategy which provides a resolution floor for GWs produced by compact binaries. Because the orbital frequency equals or supersedes the frequency of the emitted GWs (once rescaled appropriately), we can use it to create an *in situ* criteria for resolving the GW Nyquist frequency. We dub this method ORBIT (see Sec. II B 3) and implement it in the code DENDRO-GR. We will show that this baseline resolution sustains GWs as they propagate through the grid, capturing full waveform amplitudes to a higher precision.

In this paper, we show how recent improvements to DENDRO-GR (Sec. II), including a refinement floor based on tracking inspiral frequency (ORBIT, Sec. II B 3), permit us to simulate with improved accuracy, precision, and efficiency. Section III details these improvements, focusing on lower mass ratios $q = \{1, 4\}$; a future paper will catalog waveforms at larger mass ratios, $q \in [8, 32]$.

II. METHODS

The relativistic astrophysics code DENDRO-GR [22–27] uses wavelet-based adaptive multiresolution (WAMR) to capture sharp features on an unstructured octree grid. DENDRO-GR calculates wavelet coefficients for all 24 BSSN fields [28] and triggers refinement if the wavelet expansion coefficients exceed a threshold magnitude $\epsilon(t; \vec{x})$. Decreasing this threshold both increases the number of basis functions used to represent the solution while also increasing grid resolution in those regions which need it most, driving the solution toward the convergent regime. An insufficiently strict wavelet tolerance could then preclude access to the convergent regime, failing to model crucial features of the evolution. WAMR responds to the resolution requirements of the wavelet basis, refining regions near the BHs and regions with strong GW signals while keeping regions far from the BHs at a lower grid resolution, truncating the wavelet expansion. DENDRO-GR scales well: In massively parallel jobs, work per processor remains roughly constant. This structure allows us to complete full inspirals in a matter of days to weeks for $q \lesssim 16$.

In this paper, we use DENDRO-GR with much of the same setup as in Fernando *et al.* [24] (this code version is hereafter referred to as “v2022”): Utilizing the octree code Dendro for refinement, we generate initial data with TwoPunctures [29] and solve the BSSN equations.

Unlike Fernando *et al.* [24], we make several alterations to the equations (Sec. II A) and grid structure (Sec. II B). Our current formalism now incorporates slow-start lapse [[30], Sec. II A 1] and customized Hamiltonian damping (HD; Sec. II A 2) into the BSSN equations. Three key

strategies improve our grid structure: structural adjustments within the orbital radius (Sec. II B 1), causal wavelet adaptive mesh refinement (Sec. II B 2), and an evolving Nyquist refinement floor dictated by BH trajectories (ORBIT; see Sec. II B 3). The following subsections detail these enhancements to Dendro’s core functionality.

A. Formalism

We solve the BSSN equations as in Fernando *et al.* [24] but with two additions: “slow-start lapse” (SSL) and custom Hamiltonian damping (HD). Both work to reduce noise and error in the solution: SSL at early times by decreasing noise generated by the initial lapse wave, and HD at later times, damping Hamiltonian constraint violation over time.

1. Slow-start lapse (SSL)

A common challenge with the standard $1 + \log$ slicing and Γ -driver shift conditions is the emergence of an initial gauge wave packet propagating outward from the BHs with superluminal speed $\sqrt{2}c$ [31]. If WAMR is equally applied across the space that the sharp gauge pulse propagates through, it will trigger extensive refinement along the two-dimensional surface associated with the outward-traveling wave. Dendro most efficiently resolves zero-dimensional (i.e., pointlike) features, so resolving this 2D surface incurs substantial computational cost. As this refinement occurs on an unphysical feature, it would be preferable to avoid such inefficient expenditures.

Recent results from Etienne [30] show that adding a “slow-start lapse” (SSL; see Sec. II A 1) condition to BSSN weakens the initial lapse wave. Particularly at high q , SSL spreads out the wave packet, reducing lapse wave sharpness and amplitude. This smoothing of the lapse wave decreases refinement necessary to resolve the peak and reduces constraint violations by 2–6 orders of magnitude. Together with other improvements, numerical noise in the waveform decreased by a factor of ~ 4.3 , revealing previously obscured higher-order modes.

Etienne [30] defines SSL as adding the following term to the lapse evolution equation:

$$\partial_t \alpha = [\dots] - W[h e^{-\frac{1}{2}t^2/\sigma^2}](\alpha - W), \quad (1)$$

where $W = \chi^{1/2}$ is the square root of the conformal factor (DENDRO-GR evolves χ), and the dimensionful constants h and σ are found via numerical experiment; in this work, we use the default values of $h = 0.6/M$ and $\sigma = 20M$. This additional term drives the initial lapse wave toward $\alpha = W$. SSL decreases the lapse wave’s amplitude and frequency such that its Gaussian curvature (functional; not spacetime) is reduced by a factor of $\sim 8q$; this both reduces noise on the grid and reduces refinement required by feature-dependent strategies like WAMR.

We find that enabling SSL for a $q = 1$ run reduces constraint violations by over an order of magnitude at early times and near merger decreases the necessary mesh size by a factor of 3 (due to reduced noise on the grid). Overall, this results in waveforms with roughly four times less noise. As discussed in Etienne [30], SSL may shift merger time: We found for a low-resolution $q = 1$ run a $\sim -4M$ shift (which decreased with higher resolution).

2. Hamiltonian damping (HD)

Our Hamiltonian damping scheme (HD) builds off previous work [30,32,33], which added the diffusive term $C\mathcal{H}$ (where C is some dimensionful constant and \mathcal{H} is the Hamiltonian constraint violation) to the evolution of the conformal exponent ϕ . They found a Courant condition where $2C\Delta t/(\Delta x)^2 \leq 1$. We replace C with $\frac{1}{4}c_H(\Delta x)^2/\Delta t$ to account for dependence on local grid spacing and time step. As we evolve the conformal factor $\chi = e^{-4\phi}$ instead of ϕ , we need to include an extra factor of χ in our diffusive term:

$$\partial_t \chi = [\dots] + c_H \chi \frac{(\Delta x)^2}{\Delta t} \mathcal{H} \quad (2)$$

(we use $c_H = 0.077$). This modification approaches zero in the continuum limit.

In our runs, HD rapidly lowers \mathcal{H} by an order of magnitude, but these gains level off. The constant remeshing caused by the BHs traveling across the grid generates new \mathcal{H} with each new interpolation of data (whether from refining or coarsening). Errors reach an equilibrium during inspiral, as the same amount of \mathcal{H} generated at each remesh is removed by the next remesh step, sustaining a somewhat steady \mathcal{H} floor. As with SSL, HD causes a slight shift in merger time ($\sim +1.5M$ for a low-resolution $q = 1$ run, decreasing with higher resolution).

We have found that the two modifications above (SSL + HD) allow DENDRO-GR to run at about half the computational cost while reducing constraint violations by over a factor of 20.

B. Grid structure

DENDRO-GR is built atop the computational framework Dendro, which uses an octree to represent the computational domain as an adaptive unstructured grid. Dendro subdivides the initial computational domain into octants dependent on features present in the initial data through a wavelet-based approach, iterating until refinement is no longer required for accurate computation or until a maximum refinement level is reached. Dendro subjects all octants to a 2:1 grid balancing condition—this ensures that any given octant differs at most by a single level from its neighbors. Fernando *et al.* [24] provides an in-depth explanation of the initial grid generation and refinement procedures.

Our modifications require the grid structure to include hard-coded elements depending on current BH locations (Onion; see Sec. II B 1), the current time of the simulation (causal WAMR; see Sec. II B 2), and current BH velocities (ORBIT; see Sec. II B 3).

1. Onion: Recursive inner grid structure

Different parts of a grid call for different refinement strategies. We must ensure sufficient resolution while striving for an efficient implementation. The BHs in the central region of the grid each require sufficient resolution throughout the simulation to accurately generate waveforms. We therefore enforce a base, core refinement around each BH. However, we cannot sustain such a high level of resolution across the entire simulation—we must relax that refinement as we recede from the BHs; the rate at which that derefinement occurs can affect the waveform. Furthermore, at the orbital scale, we find that we cannot have reverted to too little refinement. We discuss our refinement approaches at each of these scales.

Core refinement. In the vicinity of each BH, we manually set the refinement level to ensure a minimum of 50 grid points across each (nonspinning) horizon (see Appendix A). This core, maximum refinement level $\ell_{\text{BH},i}$ encapsulates each BH to a radius of $R_{\text{AMR},i} \sim 2m_i$.

Neighboring refinement. As mentioned above, Dendro’s octree structure has a 2:1 derefinement criteria, which requires that cell neighbors differ at most by a single refinement level [24,34]. While this prevents *immediate* derefinement by more than one level, physically proximal regions could still drastically differ in resolution. Because coarser grids have a higher effective impedance for signals (acting like stiff barriers), high-frequency signals can suffer significant back-reflection on encountering refinement boundaries of lower resolution regions [10]. Without gradual derefinement about the BHs, the outward propagating initial lapse wave will tend to immediately back-reflect onto the BHs, generating large constraint violations and potentially causing an unphysical influence on the BH trajectories. To help alleviate these rapid shifts in grid spacing, we construct buffer regions, further limiting the rate at which the grid coarsens beyond the 2:1 condition. In Radia *et al.* [10], the authors implement buffer regions by demanding a minimum number of cells between consecutive refinement level boundaries, thus spatially slowing derefinement. We elect instead to use a physical criteria for grid derefinement, dependent on the distance from each BH.

As we move away from the highly resolved BHs, we only permit the refinement level ℓ to decrement once the radial distance from the BHs has increased by a geometric factor $r_{\ell-1} \sim \gamma r_\ell$. Explicitly, the resolution floor follows

$$\ell \geq \ell_{\text{BH},i} - \log_\gamma \max(1, r_i/R_{\text{AMR},i}), \quad (3)$$

where again $\ell_{\text{BH},i}$ defines the core refinement level within radius $R_{\text{AMR},i}$ about each BH; the geometric ratio γ then relaxes between levels. We need a factor of $\gamma > 1$ in order for consecutive radii to expand and provide a buffer region to reduce back-reflections. If we were to use a factor of $\gamma > 2$, then the radii expand faster than the octree 2:1 grid structure, and BH positions could influence refinement in very distant, spacelike separated regions of the grid. We choose the golden ratio ($\gamma \approx 1.618$) as a number between these bounds; this marginally improves grid efficiency over $\gamma = 2$. We find that this gradual derefinement suppresses spikes in constraint violations (lowering total constraint violation by $\sim 4\times$ for a low-resolution $q = 1$ run), buffering the initial solution against early, high-amplitude echoes of the lapse wave on the BHs, reducing net constraint violation.

Orbital scale. At the orbital scale, we enforce a wide base level of refinement to minimize changes in grid structure induced by changes in BH locations. As mentioned in Sec. II A 2, each change to grid resolution generates Hamiltonian constraint violations \mathcal{H} ; we therefore desire to minimize changes to grid structure. To this end, we keep all points within a certain coordinate radius at a higher refinement level, ℓ_{orbit} . We set this coordinate radius at the orbital scale to be

$$R_{\text{orbit}} = \max(r_1, r_2) + B, \quad (4)$$

where r_i is the current radial distance of the i th BH from the grid center, and B is a buffer length. For our particular grids and grid sizes, we use $\ell_{\text{orbit}} = 9$ (which corresponds to a $\Delta x \approx 0.26M$), which gives roughly a tenth the resolution of the BHs themselves. The buffer length permits the high refinement immediately about each BH to derefine to ℓ_{orbit} ; for the scales in our simulations, we use $B = 8M$. In addition to minimizing grid level changes within R_{orbit} , this base level of refinement also keeps the outer regions of the grid from changing structure with BH rotation. (As an example of this, in the absence of this floor, if the BHs were at angles 45° and 225° in the xy plane, then Dendro would instruct the octants at 135° and 315° to derefine—then re-refine after the BHs moved another 45° .) This base level of refinement thus limits the changes to refinement level in both the inner regions and outer regions of the grid as the BHs orbit the origin. We find that enabling this orbital-scale refinement floor reduces Hamiltonian constraint violation \mathcal{H} by about a factor of 2 compared to runs without it.

2. Causal WAMR refinement

The initial lapse wave (see Sec. II A 1) posed a significant challenge in our previous paper [24]. The 2022 version of the code used wavelets across the entire grid, keeping the orbital region highly refined with wavelets from the start of the simulation. This invoked a large computational cost at high mass ratios: The initial sharp lapse pulses propagated

outward from the BHs in spherical shells, demanding refinement on their two-dimensional surfaces (while Dendro most efficiently resolves pointlike, zero-dimensional objects). Enabling WAMR globally at simulation start caused a steep mountain of initial refinement as the lapse wave dispersed.

While SSL ameliorates the issue by smoothing the lapse wave, it does not eliminate the issue entirely. We use a new refinement function, which only activates refinement at times causally connected to the BHs from the initial time onward. We additionally avoid the noisy wake of the lapse wave. These two constraints activate WAMR only at times

$$t > \max\left(r, \frac{r + L_\alpha}{\sqrt{2}}\right), \quad (5)$$

where t is the current time of the simulation, r is the radial coordinate, and $L_\alpha = 120M$ is the width of the lapse wave. We linearly fade in over a time $t_{\text{fade}} = 100M$ from a wavelet tolerance of $\log_{10} \epsilon = -3$ (effectively disabling wavelet refinement) to a run-time value of $\log_{10} \epsilon = -5$.

This refinement strategy enables WAMR in only clean, causally connected regions of the grid, disabling it in regions of the grid spacelike separated to the BH evolution and regions influenced by the lapse wave's propagation (see Fig. 1). By avoiding causally disconnected and noisy regions, we can increase the sensitivity of our wavelet tolerance ϵ across the grid, thereby lowering constraint violations and generating more accurate waveforms.

In addition to the gradual activation of WAMR, we spatially decrease wavelet tolerance ϵ with respect to radius. We use a constant value of 10^{-5} within the orbital radius $r \leq 8M$ and the value of 10^{-3} in outer regions

$r \geq 400M$. We linearly interpolate $\log \epsilon$ with respect to $\log r$ between those boundaries.

3. ORBIT resolution floor

The Nyquist frequency [35] determines the maximum frequency (minimum wavelength) one can measure on a grid. If there are fewer than two points per wavelength of a signal, it will not transmit properly, distorted by aliasing. In audio processing, 1.1 to 5 times the Nyquist frequency is often used as a sampling rate to provide “lossless” signals [36]. Beyond the Nyquist frequency, error goes down with sampling frequency squared—sampling at twice the Nyquist frequency decreases errors by about a half decade compared to sampling at the Nyquist frequency (see Appendix B 1).

We measure GWs with spin-weighted spherical harmonics (l, m) of the Weyl scalar Ψ_4 [39]. Properly transmitting GWs demands grid resolution $\Delta x \leq \lambda_{\text{Nyquist}} = \lambda_{\text{GW}}/2$. Because we will not always know the precise GW structure of a simulation *a priori*, we cannot easily use the GWs themselves to make the grid Nyquist compliant. During the early inspiral, the orbital frequency of the black holes matches the gravitational wave frequency. In particular, the orbital frequency f_{orbit} will match the GW frequencies rescaled by their order m : $f_{\text{GW}}^{(l,m)}/m$. Approaching merger, the black holes' orbital frequency outpaces that of the GWs they emit (see, e.g., Fig. 7), and the BHs pass into a common horizon; beyond this point, their coordinate positions causally disconnect from the rest of the solution. Because $f_{\text{orbit}} \geq f_{\text{GW}}^{(l,m)}/m$, rescaling the BH orbital frequency to some desired mode m provides a sufficient standard for estimating GW frequencies. We may thus satisfy the Nyquist criterion by tracking BH orbital




FIG. 1. Cartoon illustrating two novel refinement strategies recently implemented in DENDRO-GR. Left: causal WAMR, a new refinement function dictating wavelet tolerances ϵ across the grid. We only activate WAMR in the causally connected and clean regions of the grid, ignoring the lapse wave (red) and waiting for the GW signal (green) to fill the grid with physical data before refining. After these two events propagate outward, we gradually lower ϵ to its target value in a given region, requiring strict refinement in the blue-tinted region. Right: ORBIT, a new refinement floor based on orbital frequency; lighter colors indicate higher refinement levels. We only activate ORBIT in regions causally connected to the event of the gravitational wave ringdown striking the largest extraction radius [indicated by dotted white and black lines; see Eq. (7)].

frequency, calculating the corresponding refinement requirement and propagating that requirement outward through the grid at the speed of light (matching the outward propagating GWs).

In summary, to ensure Nyquist resolution of a gravitational wave mode m_{\max} , it is sufficient to mandate at radius r and time t that the grid spacing $\Delta x(t, r)$ relates to retarded orbital frequency $f_{\text{orbit}}(t - r)$ as follows:

$$\Delta x(t, r) \leq \frac{c}{2m_{\max}f_{\text{orbit}}(t - r)} \leq \lambda_{\text{Nyquist}}^{(l, m_{\max})}(t, r). \quad (6)$$

While this prevents *total* back-reflection of low frequencies $f < m_{\max}f_{\text{orbit}}$, partial back-reflection will still occur to some degree [with error following Eq. (B1)]. To mitigate back-reflections, one should pick m_{\max} somewhat larger than the target spin-weighted spherical harmonic order. In this paper, we use twice the frequency, Nyquist-resolving $m_{\max} = 12$ to clean $m = 6$. We dub this refinement scheme *ORBIT*: an optimized resolution based on the inspiral trajectory.

Early in the simulation, low orbital frequencies require minimal refinement, but at plunge, orbital frequencies rapidly increase (and then flatline; see Fig. 7), resulting in a stricter refinement criterion postmerger. This high refinement postmerger need only be used late in the simulation and only in past timelike regions relative to the event where the last clean signal reaches the outermost radius at which GWs are measured $R_{\text{GW}, \max}$. In particular, we disable ORBIT $\sim 122M$ after merger at the largest gravitational wave extraction radius $R_{\text{GW}, \max} = 100M$ and in all points not causally influencing that final point. That is, for all regions

$$|r - R_{\text{GW}, \max}| > (t - t_{\text{end}}), \quad (7)$$

using $t_{\text{end}} = t_{\text{merge}} + R_{\text{GW}, \max} + t_{\text{ring}}$, with ringdown duration $t_{\text{ring}} = \{22, 72\}M$ past merger for mass ratios $q = \{1, 4\}$ [40]. ORBIT therefore gradually increases the refinement of the grid from the inside out as we approach merger and dials down refinement in the outer and inner regions postmerger. As compared to retaining postmerger resolution globally, ORBIT’s causal activation and deactivation reduces mesh size by an order of magnitude, improving our ability to place refinement only where (and when) it is needed.

Figure 1 illustrates our two novel refinement strategies: causal WAMR and ORBIT. The left plot shows how causal WAMR acts on a clean and data-filled grid, in the future lightcone of the simulation. The right plot shows how ORBIT requires a resolution level dictated by the orbital frequency. The two refinement strategies combine to clean waveform noise by an order of magnitude compared to our previous work.

III. RESULTS

In this section, we present improvements achieved by enabling ORBIT and other adjustments as we upgrade DENDRO-GR from its 2022 version (“v2022”) to its 2025 version (“v2025”).

In this paper, we focus on mass ratios $q = 1$ and $q = 4$. Mass ratio $q = 1$ is of particular interest for ORBIT, as $q = 1$ has the fastest orbital frequency (of nonspinning, noneccentric mergers) and therefore stricter refinement constraints than higher mass ratios. We include a mass ratio $q = 4$ run, leaning toward larger mass ratios. We have not reduced eccentricity for either of the mass ratios; this allows for one-to-one comparison to v2022. Future work will reduce eccentricities and delve into $q \geq 8$ in more detail.

For mass ratios $q = \{1, 4\}$, we compare enabling ORBIT (“ORBIT on”; $m_{\max} = 12$, aiming to resolve well $m = 6$) to disabling ORBIT (“ORBIT off”; $m_{\max} = 0$). For v2022, only $q = 4$ is available for comparisons (without ORBIT) and then only spherical harmonic modes $|m| \leq l \leq 4$.

We first show that ORBIT enhances waveform precision: reducing noise and better propagating higher-order modes (Sec. III A). We also show how v2025 reduces constraint violations and improves efficiency of grid structure (Sec. III B). Finally, we compare run times between v2022 and v2025, illustrating reductions in both wall time and total cpu hours (Sec. III C).

A. ORBIT enhances waveform fidelity

In this section, we showcase the improved fidelity of our measured gravitational waveforms, focusing in particular on the effects of enabling ORBIT, our Nyquist-compliant orbital-frequency-based resolution floor. Figure 2 displays Weyl scalar Ψ_4 waveform amplitudes for $q = 1$ and $q = 4$ mass ratios for ORBIT disabled (blue; “ORBIT off”), for ORBIT enabled (green; “ORBIT on”), and for the previous code version (orange; “v2022”; ORBIT unavailable). Both left and right panels show spin-weighted spherical harmonic modes $(l, m) = (x, -x)$ with $x \in \{2, 4, 6, 8\}$ (though v2022 only has $l \leq 4$ available). Shaded regions approximately indicate portions of the signal made noisy by the initial conditions (early times) and regions where perturbation theory can take over from numerical relativity (late times).

We see a clear reduction in waveform noise on enabling ORBIT, not only in the ringdown (where ORBIT demands its highest refinement), but also in the early inspiral (where ORBIT makes minimal demands). Noise reduction occurs in all modes but is most apparent at higher orders. Focusing on the $(l, m) = (6, -6)$ mode (the target of ORBIT’s m_{\max} criterion); Fig. 3 shows roughly an order of magnitude reduction in waveform noise in the early inspiral for $q = 1$ (we see about half an order of magnitude reduction for $q = 4$).




FIG. 2. Waveform amplitudes (measured at $R_{\text{GW}} = 50M$) with and without ORBIT enabled ($m_{\text{max}} = 12$ to enable; 0 to disable); the end time of the right plot shows at a glance that recent improvements lower the noise floor by two decades. Left and right plots give Weyl scalar $\Psi_4^{(l,m)}$ spherical harmonic amplitudes for mass ratios $q = 1$ and $q = 4$, respectively. We here show modes $(l, m) = (x, -x)$ for $x \in \{2, 4, 6, 8\}$ in order from top to bottom. The right plot additionally includes results from the 2022 version of DENDRO-GR (orange; “v2022”) for comparison to the new 2025 version of DENDRO-GR (blue, green); this only has modes $x \in \{2, 4\}$ available for comparison. Shaded regions on the left of each plot indicate the causally disconnected regime ($t < R_{\text{GW}}$) and the region corrupted by the initial lapse wave ($t/M < (L_\alpha + R_{\text{GW}})/\sqrt{2} \doteq 120.2M$). Shaded regions on the right of each plot indicate where we terminate ORBIT support. Note that none of the simulations were eccentricity reduced; the $q = 1$ run has large eccentricity modulating the amplitude.




FIG. 3. Zoom in on several notable regions of the left plot of Fig. 2 ($q = 1$; identical coloring), focusing on the $(l, m) = (6, -6)$ mode. ORBIT improves waveform quality both in the inspiral (left) as well as at merger and ringdown (right).

After plunge, ORBIT ensures that the proper peak amplitude of high-order modes propagates outward. Disabling ORBIT results in a 24% reduction in merger amplitude for the $(6, -6)$ mode for $q = 1$. For $q = 4$, the same mode has a higher amplitude, and disabling ORBIT results in only a 6% reduction. Without ORBIT’s added resolution floor, higher order modes (which have lower amplitudes) fail to propagate through the grid to the GW extraction radii. After transmitting peak amplitude, ORBIT continues, carrying the ringdown outward. WAMR (our primary refinement scheme) ignores low-amplitude signals, so the exponential decay of ringdown amplitude is lost without ORBIT’s sustaining influence.

Figure 4 highlights waveform improvements gained on upgrading the code from v2022 to v2025. For the $(l, m) = (4, -4)$ mode of $q = 4$, we see roughly a factor of 4




FIG. 4. Zoom in on several notable regions of the right plot of Fig. 2 ($q = 4$; identical coloring), focusing on the $(l, m) = (4, -4)$ mode. Compared to the v2022 (the old code version), v2025 reduces noise in inspiral waveform (left), better captures peak waveform at merger (middle), and lowers the noise floor by nearly 2 orders of magnitude (right).

reduction in waveform noise in the early inspiral (left panel). At merger (center panel), v2025 better captures peak waveform amplitude, matching trends discussed in Appendix C. Post-ringdown, the noise floor (right panel) drops by roughly a factor of 10 on moving from v2022 to ORBIT off v2025, then by an additional factor of 10 on enabling ORBIT. As next-generation detectors lower the noise floor on GW observations, the systematic errors of NR simulations must drop a commensurate amount. This shows progress in that direction.

Discussed more in Appendix C, mode coupling strengthens toward more extreme mass ratios, increasing the relative importance of higher-order modes. Figure 2 exemplifies this in the decreased distance between each of the

$q = 4$ modes (less than an order of magnitude between modes) compared to $q = 1$ (over an order of magnitude between modes during inspiral). The trend with mass ratio continues, such that higher-frequency modes are of greater importance (to reconstruct observed waveforms) for higher mass ratios than they are for lower mass ratios. We must therefore resolve higher-frequency modes well for asymmetric mass ratios if we are to faithfully reconstruct the outgoing waveform.

B. Code updates efficiently reduce constraints

Updating our code from v2022 to v2025 both improves waveform quality and reduces net constraint violations. While globally refining the v2022 run would certainly reduce constraint violations, we quantify how well our improvements outpace the expected improvements of adding refinement to a sixth order code like ours, showing that our improved strategies refine more judiciously than before.

The upper panel of Fig. 5 shows reductions in constraint violations between code versions for $q = 4$ in moving




FIG. 5. Comparison of total constraint violation (upper panel; L_2 norm of \mathcal{H} & \mathcal{M}_i across grid) and grid efficiency [lower; see Eq. (8)] between code versions for $q = 4$. We see 10–100× lower constraint violations as we upgrade from v2022 to v2025. These improvements are executed efficiently, outperforming by 100× the expected gains from globally refining a sixth order code. Coloring as in Fig. 2; shading indicates ORBIT disabling at gravitational extraction radii of $R_{\text{GW}} = \{50, 100\}M$.

from v2022 to v2025. Net constraint violation decreases by 2 orders of magnitude while Hamiltonian constraint decreases by up to 3 orders of magnitude. One may then ask whether this reduction of constraint violation is superficial—merely an expected outcome of increasing resolution (not yielding a true improvement in accuracy per cpu hour spent).

Recall that Dendro uses an octree grid, subdividing each cell in half with each division. Thus, the grid resolution dx of a given cell relates to the mesh size μ (the number of mesh elements or cells in the grid) as $dx \propto \mu^{-1/3}$. We currently use sixth order spatial derivatives, so error decreases as $\epsilon \propto dx^6$. Globally refining the grid would thus decrease errors as $\epsilon \propto \mu^{-2}$; we would thus expect $\epsilon\mu^2$ to remain constant. As WAMR refines the grid locally rather than globally, refining an unimportant region of the grid will do less to reduce global errors than refining an important region of the grid. Therefore, lower values of $\epsilon\mu^2$ between runs imply improved refinement schemes: not merely that we have increased refinement, but that we have done so efficiently, allocating resources where they are most needed. To this end, we define a notion of “grid efficiency” as

$$\eta \equiv C_{\text{tot}}^{-1} \mu^{-2}, \quad (8)$$

where C_{tot} is the L_2 norm of both Hamiltonian and momentum constraint violations, \mathcal{H} and \mathcal{M}_i , across the entire grid, and μ is again the mesh size. Comparing η thus measures the relative efficiency of grid subdivision schemes: Do our choices for refinement result in appropriate error reduction, or do they show some relative (in) efficiency?

The lower panel of Fig. 5 compares grid efficiency η between v2022 and v2025. Not only do we reduce constraint violations by 1 to 2 orders of magnitude in the upgrade, we see similar scale improvements in grid refinement efficiency, accounting for the expected accuracy differences due to refinement. This shows that these gains are not due to superficial changes in refinement.

While premerger refinement primarily serves to prevent errors (as WAMR does explicitly), postmerger refinement called for by ORBIT serves to propagate GWs out to the detection radii. For this reason, we expect the loss of efficiency in the ORBIT-enabled run, particularly as the high-frequency ringdown (here at $t/M \sim 625$) propagates out to the most distant GW detection radius (100M). The slight inefficiency incurred by enabling ORBIT thus points to the difference in *purpose* of ORBIT refinement compared to WAMR refinement.

C. Code updates accelerate runs

Recent changes to our algorithm achieve dramatic improvements in run cost—both in human time (wall hours) and in computer time (cpu hours). Figure 6 shows

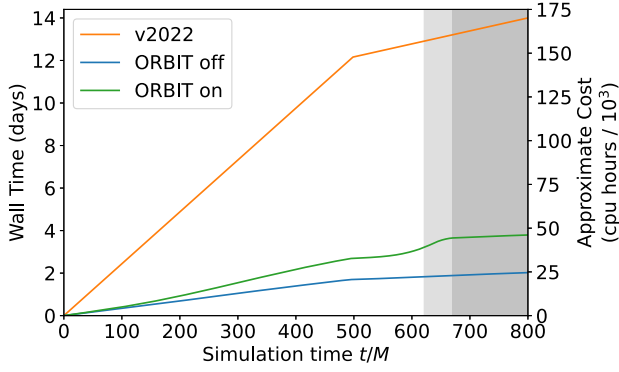


FIG. 6. Comparison of $q = 4$ run cost between code versions, both in wall time as well as in approximate cpu hours. Both humans and computers see four times lower run cost in v2025 than in v2022. Coloring and shading as in Fig. 5.

that while v2022 took over two weeks to run, a current v2025 run takes less than four *days*, cutting down wall time by a factor of 3.7. Furthermore, computational cost is cut down by a factor of almost 4. (This is a nontrivial distinction, as hypothetically, the wall time could have been decreased by submitting identical code to four times the number of processors.)

The “ORBIT on” run shows a slight bump in wall time around $625M$, as ORBIT propagates out to sustain the high-frequency ringdown. As illustrated in Figs. 1 and 7, postmerger refinement is the most costly, as it fills a 3D sphere with high refinement. While enabling ORBIT makes the run almost twice as expensive as disabling it, we find that the resulting gains in waveform quality (see, e.g., Fig. 3) justify the cost.

Updates to the code thus not only improve waveform quality and reduce constraint violations, but they do so efficiently and with decreased net computational cost.

IV. CONCLUSIONS

We introduced ORBIT, an optimized resolution baseline derived from BH inspiral trajectories, and showcase recent advances in DENDRO-GR over the past few years. These improvements work toward fulfilling the need of next-generation GW detectors for higher accuracy in NR simulations.

ORBIT causally refines Dendro’s grid to sustain propagation of high-order spherical harmonic modes. The Nyquist criterion—combined with the relationship between BH orbital frequency $\omega_{\text{orbit}} = 2\pi f_{\text{orbit}}$ and GW frequency—gives a resolution floor of

$$\Delta x \leq c/(2mf_{\text{orbit}}) \leq \lambda_{\text{Nyquist}} \quad (9)$$

as a sufficient criteria for Nyquist-resolving GW spin-weighted spherical harmonics of order m . ORBIT captures low-amplitude features otherwise missed by error-driven refinement methods like WAMR. As compared to

maintaining high resolution across the entire duration of the simulation, this dramatically reduces computational cost while still maintaining necessary resolution.

This 2025 version of DENDRO-GR, with ORBIT and other improvements, demonstrates significant progress:

- (i) Waveform noise has decreased by 1 to 2 orders of magnitude.
- (ii) Constraint violations are reduced by 2 to 3 orders of magnitude.
- (iii) Refinement efficiency has improved by 2 orders of magnitude.
- (iv) Wall time and total run cost is reduced by a factor of 4.

In addition to ORBIT, other improvements include implementing slow-start lapse (SSL), our custom Hamiltonian Damping scheme (HD), causal WAMR, and central grid structure updates (see Sec. II for details on each of these changes). By reducing the sharpness of the initial lapse wave, SSL relaxes grid refinement requirements and reduces noise on the grid. Though the constant remeshing required as BHs travel across the grid generates Hamiltonian constraint violation \mathcal{H} (as grid patches interpolate to different resolutions), our HD scheme damps total \mathcal{H} by an order of magnitude, preventing its accumulation during inspiral. Our new causal WAMR structure puts refinement where we need it by ignoring spacelike regions of the grid (causally disconnected to the BH evolution) and by avoiding regions contaminated by the (dampened) lapse wave. In total, these improvements allow DENDRO-GR to run at about one-fourth the computational cost while reducing constraint violations by orders of magnitude. These gains position DENDRO-GR as an increasingly capable tool for simulating binary BH mergers, particularly at higher mass ratios.

We anticipate additional improvements to DENDRO-GR. For these first pass, proof-of-concept results, our implementation of ORBIT was hardcoded (see Appendix B); the method is now softcoded for greater flexibility. Additionally, we expect a need to extend ORBIT further as we test it more thoroughly with systems having high spins or large eccentricities. These will certainly present new challenges.

In conclusion, ORBIT and other recent improvements to DENDRO-GR reduce waveform noise while lowering computational costs. These and other planned advances bring us closer to meeting the demands of next-generation GW detectors and to expanding our understanding of BBH systems across a wider range of systems.

ACKNOWLEDGMENTS

This research was supported by NSF Grants No. PHY-2207615 (BYU) and No. PHY-2207616 (Utah), as well as NASA Grant No. 80NSSC20K0528. W. K. B. is supported by the College of Computational, Mathematical, and Physical Sciences at BYU. Supercomputer time was provided by ACCESS [42] on Stampede3 [43]. We thank

our collaborators Yosef Zlochower and Chloe Malinowski for many discussions during this work. W. K. B. thanks J. K. Black for assistance in formulating Eq. (A1) and thanks Z. B. Etienne for conversations on constraint damping.

DATA AVAILABILITY

The data that support the findings of this article are openly available [44].

APPENDIX A: CENTRAL REFINEMENT ABOUT THE BHS

Here we determine core refinement about each black hole. To ensure we have N points across the horizon of a BH of radius r_{BH} on an octree grid of size L with element order n (determining the number of points across a cell), we must keep the level of refinement ℓ (where larger ℓ indicate higher resolution) about the BH to be at least

$$\ell \geq 2 + \log_2 \left(\frac{L}{n \cdot r_{\text{BH}}} \left\lfloor \frac{N+1}{2} \right\rfloor \right), \quad (\text{A1})$$

(where $\lfloor * \rfloor$ is the floor operator). In our evolution coordinates, $R_i = m_i$ [45]; the larger BH has a mass of $m_1 = M/(1+q)$ while the smaller BH has a mass of $m_2 = M/(1+1/q)$. Additionally accounting for our box size $L = 800M$, element order $n = 6$, and our goal of $N \geq 50$ gives us level requirements about each BH of

$$\ell_1 \geq 13.7 + \log_2(1+q), \quad (\text{A2})$$

$$\ell_2 \geq 13.7 + \log_2(1+1/q), \quad (\text{A3})$$

(where $m_{\text{BH},1} \leq m_{\text{BH},2}$). These then dictate $\ell_{\text{BH},i} \geq 15$ for $q = 1$, while for $q = 4$, it dictates 14 and 16 for the larger and smaller BHs, respectively. More generally, as we approach extreme mass ratios, we find we need cell level 14 and $14 + \log_2 q$ as $q \rightarrow \infty$ given our grid size and element order. We note there exists a two-level offset in the cell level parameters ℓ_i from the resulting level on the grid, such that the above equations enforce cell level 13 for $q = 1$.

APPENDIX B: MORE ON ORBIT

Here we present more details regarding our Nyquist-compliant refinement floor ORBIT.

1. Choice of target order

The Nyquist criterion dictates the minimum refinement necessary to capture and reconstruct a frequency; lower refinements will preclude propagation of the signal. Fractional error of a signal frequency f_{signal} as a function of sampling rate f_{sample} is given by

$$\varepsilon = 1 - \frac{1}{\sqrt{1 + (f_{\text{signal}}/f_{\text{sample}})^2}}, \quad (\text{B1})$$

where $f_{\text{Nyquist}} = 2f_{\text{signal}}$. At high sampling rate (in the regime where $f_{\text{sample}} \gg f_{\text{signal}}$), this error approaches $\varepsilon = \frac{1}{2}(f_{\text{signal}}/f_{\text{sample}})^2$. Beyond f_{Nyquist} , error falls off as sampling frequency squared. Thus, if we Nyquist-resolve mode m (i.e., $f_{\text{sample}} = 2f_{\text{GW}}^{(l,m)} = f_{\text{Nyquist}}$), we have $\varepsilon \sim 10\%$. Nyquist-resolving to $2m$ ($f_{\text{sample}} = 2f_{\text{Nyquist}}$) would then have only error $\varepsilon \sim 3\%$ on mode m while Nyquist-resolving $7m$ would yield $\varepsilon \sim 1\%$ (lower-frequency waveforms would improve even more so). In this paper, we Nyquist-resolve $m = 12$, decreasing aliasing-induced errors in the target $(l, m) = (6, -6)$ waveform to 3%.

2. Calculation from orbital frequency

We next discuss the relationship between BH orbital frequency and GW wavelength. The upper panel of Fig. 7 shows angular orbital frequencies ω for the BH orbit as well as for the emitted GWs; the lower panel shows the corresponding rescaled wavelengths $m\lambda$ (with orbital wavelength $\lambda_{\text{orbit}} \equiv 2\pi c/\omega_{\text{orbit}}$) each overlapping. Angular frequencies for the BHs are calculated from the progression of the relative angle between the BHs in the orbital plane and for the GWs from the phase angle in complex space. The lower plot then rescales by m , showing universal agreement in waveform across m in the early inspiral. As we approach merger, $\lambda_{\text{orbit}} < m\lambda_{\text{GW}}^{(l,m)}$, implying that BHs inspiral faster than their emitted wavelengths. At merger, the BHs form a common horizon and their coordinate locations decorrelate from the emitted GWs. This works to our favor, as refining to grid spacing $\Delta x = \lambda_{\text{orbit}}/2|m|$ gives a stricter refinement criterion than refining to $\Delta x = \lambda_{\text{GW}}^{(l,m)}/2$; using rescaled orbital wavelength therefore *anticipates* refinement requirements, ensuring the grid already has sufficient resolution for the GWs before they reach a given portion of the grid. As wavelengths cannot always be known *a priori*, this provides us with a sufficient resolution floor to satisfy the Nyquist criterion *in situ*. ORBIT thus adapts to support the emitted GWs mid-simulation.

3. Our formulation for cell level

Given an input orbital wavelength, one can then determine the particular refinement level necessary to support GWs up to order $|m|$. Determining the target resolution level ℓ depends on several simulation-specific factors: We use a sixth-order derivative scheme ($n = 6$) on a box size of $L = 800M$; Dendro resolution level ℓ gives the subdivision of the octree: $\ell = 0$ would be a single cube (low resolution), while higher levels repeatedly subdivide that base cube by factors of 2. Dividing the target wavelength λ_{orbit} by two for the Nyquist

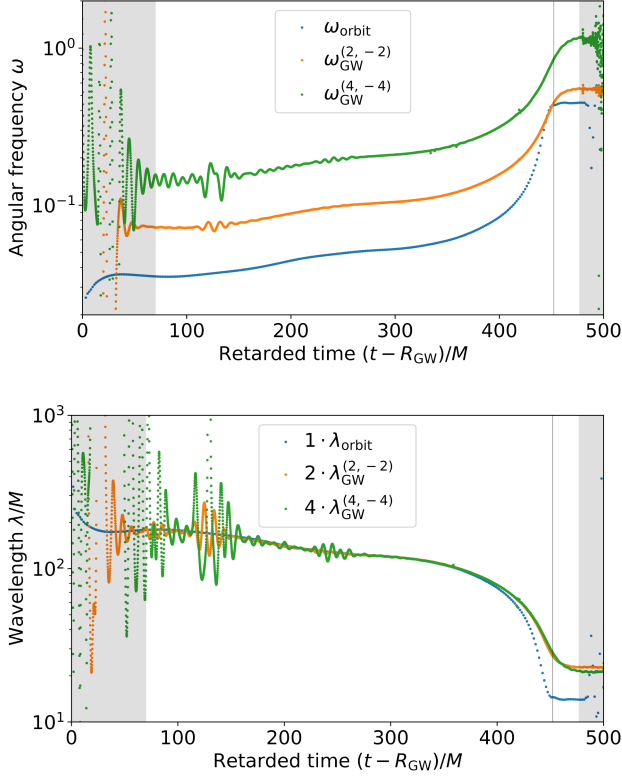


FIG. 7. Angular frequencies (upper plot) and corresponding rescaled wavelengths (lower plot) from a $q = 1$ binary inspiral. In blue are curves from BH coordinate positions; the vertical line indicates the point where BH separation distance is $ds = 1M$ (near this time, the BHs form a common horizon and their coordinate positions decorrelate with GW emission). In orange and green are data from the $(l, m) = (2, -2)$ and $(4, -4)$ GW modes. Because $\lambda_{\text{orbit}} \lesssim m\lambda_{\text{GW}}^{(l,m)}$, refining to $\Delta x = \lambda_{\text{orbit}}/2|m|$ is sufficient to Nyquist-resolve GWs of order $\leq |m|$. Shaded regions as in Fig. 2.

criterion and dividing by the order m of a target mode $\lambda_{\text{GW}}^{(l,m)}$ yields a minimum resolution level requirement of

$$\ell_{\text{Nyquist}}(m) = \left\lceil \log_2 \frac{2L/n}{\lambda_{\text{GW}}^{(l,m)}} \right\rceil \leq \left\lceil \log_2 \frac{800M/3}{\lambda_{\text{orbit}}/m} \right\rceil, \quad (\text{B2})$$

(where $\lambda_{\text{GW}}^{(l,m)}$ is the target wavelength, λ_{orbit}/m Nyquist-resolves mode m , and $\lceil * \rceil$ is the ceiling operator).

While DENDRO-GR now softcodes ORBIT (actively *in situ* tracking BH orbital frequency), these first-pass proof-of-concept runs were prototyped with an *a priori* hardcoded formulation (modeling BH orbital frequency from past high-quality runs). Following quasicircular orbit results, we used

$$\lambda_{\text{orbit}} \approx \begin{cases} \tau > 0: \max(A\tau^{3/8}, \lambda_{\text{min}}) \\ \tau < 0: \lambda_{\text{min}}, \end{cases} \quad (\text{B3})$$

TABLE I. Fit parameters for Eq. (B3), modeling BH orbital wavelength λ_{orbit} . Amplitude of inspiral given by A , merger time estimated as τ_0 , minimum orbital wavelength given by λ_{min} , and ORBIT end time given by t_{end} [see Eq. (7)].

Run	A	τ_0	λ_{min}	t_{end}
$q = 1$	18.8	445.0	14.49	580.14
$q = 4$	17.0	490.0	18.10	670.40

(where $\tau \equiv \tau_0 - t$ is time to merger); Table I gives best fit parameters. Adding the first post-Newtonian term changes the formula to $\lambda = A\tau^{3/8}(1 + B\tau^{-2/8})$, but B is consistent with zero for $q = 1$, and for $q = 4$, the fit works well enough with $B = 0$ for most of the inspiral; near plunge the simpler fit outpaces the more complex fit, giving a slightly more conservative refinement criterion. Rather than adding further post-Newtonian expansion terms to improve the fit, we will soon have a version of DENDRO-GR, which tracks past BH orbital frequency midsimulation, giving live updates to refinement for future, heretofore unexplored mass ratios.

APPENDIX C: MODE COUPLING OF WAVEFORMS AT MERGER

Mode coupling strengthens with mass ratio q : spherical harmonic modes lie closer together at higher mass ratios. The more extreme the mass ratio, the more nonlinear effects plague the simulation. We can quantify this mode coupling by observing how the spin-weighted spherical harmonic amplitude at merger $|\Psi_4^{(l,m)}|_{\text{max}}$ decreases with degree l (it *increases* with order $|m|$). The ratio of modes follows steady patterns, where $|\Psi_4^{(l,l)}|_{\text{max}}/|\Psi_4^{(l+1,l+1)}|_{\text{max}}$ is roughly constant (as well as $|\Psi_4^{(l,m)}|_{\text{max}}/|\Psi_4^{(l,m+1)}|_{\text{max}}$). We can quantify this trend across many mass ratios, measuring the log slope of $|\Psi_4^{(l,m)}|_{\text{max}}$ as a function of l , m , and q . This allows extrapolation of high-frequency modes, avoiding the expense of direct calculation.

Figure 8 shows relative GW amplitudes at merger (normalized at the maximum amplitude) for both $q = 1$ and $q = 4$. The left and right plots have different vertical scales: Most of the modes for $q = 4$ are of higher relative amplitude than those of $q = 1$; they therefore have far smaller errors (see also Fig. 2).

There are a few ways to estimate error on these values. The symmetries of our $q = 1$ run imply that all odd m modes (colored red) should have null amplitude; their nonzero values thus roughly estimate epistemic error. Mirroring $\pm m$ mode amplitudes should have equal amplitude; differences between values thus roughly estimates aleatoric error. Accounting for these uncertainties, clear trends emerge across l , m , and q .

Fitting $l = |m|$ (excluding odd m for $q = 1$) amplitudes for nonspinning, nonelliptical mergers, we find a steady




FIG. 8. Mode amplitudes at merger for $q = 1$ (left) and $q = 4$ (right). Within each vertical cell, spherical harmonic order $|m|$ increases toward the right (with odd m colored red for visibility). The first cell thus displays mode amplitudes $(2,0)$, $(2, \pm 1)$, and $(2, \pm 2)$; the final cell displays $(8,0)$, $(8, \pm 1)$, ... $(8, \pm 8)$. Errors are small enough that the $+m$ and $-m$ modes lie on top of each other for sufficiently high amplitude modes.

flattening of log slope as a function of symmetric mass ratio $\nu \equiv q/(1+q)^2$. This shows that higher-degree modes are increasingly important for unequal mass ratios. Fitting the results of $q \in \{1, 2, 4, 8\}$, we find a clean power law relation of slope versus symmetric mass ratio:

$$\frac{d}{dl} \log_{10} |\Psi_4^{(l,l)}|_{\max} = -e^{(7.23 \pm 0.07)\nu - (2.91 \pm 0.01)}. \quad (\text{C1})$$

Thus, we have relative mode strength (quantified as the magnitude of the power law slope of the spherical harmonic decomposition main modes) increasing in amplitude toward more equal mass ratios ($\nu \rightarrow 1/4$) and decreasing in amplitude toward more extreme mass ratios ($\nu \rightarrow 0$). While at $q = 1$, an increment in l results in roughly a factor of 2 loss of relative amplitude, as $q \rightarrow \infty$, incrementing l results in only a $\sim 10\%$ loss of amplitude. While for $q = 1$, the $(l, m) = (2, \pm 2)$ mode makes up $\lesssim 1/2$ of the total signal, the same mode for $q \rightarrow \infty$ only makes up $\lesssim 10\%$ of the total signal. This then quantifies the relative importance of higher-amplitude modes at more extreme mass ratios at merger.

Not only do we see clear trends in $|\Psi_4^{(l,l)}|_{\max}$, we also see rich structure within lower- m modes. The right panel of Fig. 8 (with $q = 4$) has larger amplitude modes and is thus less contaminated by noise; this better resolves the non-dominant modes, revealing an inner structure to the mode amplitudes. The patterns in (l, m) mode amplitudes hint at the ability to extrapolate to infinite degree (traveling out on the top bar of $l = |m|$), then filling in the inner structure with decreasing $|m|$, given only the trends from the first few (l, m) modes. This would incur relatively small extrapolation error, as modes grow progressively weaker with increased extrapolation. As estimation of mode amplitudes grows increasingly expensive (and error fraught) toward high l , extrapolating higher modes would decrease computational expense. This may well produce more accurate waveforms than using the raw measured waveform amplitudes (which are infiltrated by noise around 10^{-3} relative amplitude).

More generally, these trends could perhaps even be extrapolated out toward more extreme mass ratios, spins, and eccentricities, moving toward a universal fit to waveform amplitudes at merger.

-
- [1] P. Amaro-Seoane *et al.*, [arXiv:1702.00786](#).
 - [2] D. Reitze *et al.*, [arXiv:1907.04833](#).
 - [3] M. Maggiore, C. V. D. Broeck, N. Bartolo, E. Belgacem, D. Bertacca, M. A. Bizouard, M. Branchesi, S. Clesse, S. Foffa, J. García-Bellido, S. Grimm, J. Harms, T. Hinderer, S. Matarrese, C. Palomba, M. Peloso, A. Ricciardone, and M. Sakellariadou, *J. Cosmol. Astropart. Phys.* **03** (2020) 050.
 - [4] M. Pürrer and C.-J. Haster, *Phys. Rev. Res.* **2**, 023151 (2020).
 - [5] D. Ferguson, K. Jani, P. Laguna, and D. Shoemaker, *Phys. Rev. D* **104**, 044037 (2021).
 - [6] L. Pompili *et al.*, *Phys. Rev. D* **108**, 124035 (2023).
 - [7] E. Schnetter, S. H. Hawley, and I. Hawke, *Classical Quantum Gravity* **21**, 1465 (2004).
 - [8] Y. Zlochower, J. G. Baker, M. Campanelli, and C. O. Lousto, *Phys. Rev. D* **72**, 024021 (2005).
 - [9] M. Campanelli, C. O. Lousto, P. Marronetti, and Y. Zlochower, *Phys. Rev. Lett.* **96**, 111101 (2006).
 - [10] M. Radia, U. Sperhake, A. Drew, K. Clough, P. Figueras, E. A. Lim, J. L. Ripley, J. C. Aurekkoetxea, T. França, and T. Helfer, *Classical Quantum Gravity* **39**, 135006 (2022).

- [11] A. Rashti, M. Bhattacharyya, D. Radice, B. Daszuta, W. Cook, and S. Bernuzzi, *Classical Quantum Gravity* **41**, 095001 (2024).
- [12] M. W. Choptuik, A study of numerical techniques for radiative problems in general relativity, Ph.D. thesis, University of British Columbia, Vancouver, Canada, 1986.
- [13] M. W. Choptuik, in *Frontiers in Numerical Relativity*, edited by C. Evans, L. Finn, and D. Hobill (Cambridge University Press, Cambridge, England, 1989), pp. 206–221.
- [14] M. W. Choptuik, *Phys. Rev. D* **44**, 3124 (1991).
- [15] M. W. Choptuik, in *Approaches to Numerical Relativity*, edited by R. d’Inverno (Cambridge University Press, Cambridge, England, 1992), pp. 202–222.
- [16] M. J. Berger and J. Olinger, *J. Comput. Phys.* **53**, 484 (1984).
- [17] M. W. Choptuik, in *Computational Astrophysics: 12th Kingston Meeting on Theoretical Astrophysics*, Astronomical Society of the Pacific Conference Series Vol. 12, edited by D. A. Clarke and M. J. West (Astronomical Society of the Pacific, Halifax, 1997), p. 305.
- [18] F. Pretorius, Numerical simulations of gravitational collapse, Ph.D. thesis, University of British Columbia, 2002.
- [19] J. Thornburg, *Gen. Relativ. Gravit.* **43**, 1211 (2011).
- [20] S. Paolucci, Z. J. Zikoski, and D. Wirsaet, *J. Comput. Phys.* **272**, 814 (2014).
- [21] S. Paolucci, Z. J. Zikoski, and T. Grenga, *J. Comput. Phys.* **272**, 842 (2014).
- [22] J. DeBuhr, B. Zhang, M. Anderson, D. Neilsen, E. W. Hirschmann, T. Grenga, and S. Paolucci, *Astrophys. J.* **867**, 112 (2018).
- [23] M. Fernando, D. Neilsen, E. W. Hirschmann, and H. Sundar, in *Proceedings of the ACM International Conference on Supercomputing, ICS ’19* (Association for Computing Machinery, New York, USA, 2019), pp. 1–12.
- [24] M. Fernando, D. Neilsen, Y. Zlochower, E. W. Hirschmann, and H. Sundar, *Phys. Rev. D* **107**, 064035 (2023).
- [25] M. Fernando, D. Duplyakin, and H. Sundar, in *Proceedings of the 26th International Symposium on High-Performance Parallel and Distributed Computing, HPDC ’17* (Association for Computing Machinery, New York, USA, 2017), pp. 231–242.
- [26] M. Fernando, D. Neilsen, H. Lim, E. Hirschmann, and H. Sundar, *SIAM J. Sci. Comput.* **41**, C97 (2019).
- [27] M. Fernando, D. Neilsen, E. Hirschmann, Y. Zlochower, H. Sundar, O. Ghattas, and G. Biros, in *SC22: International Conference for High Performance Computing, Networking, Storage and Analysis* (2022), pp. 1–15.
- [28] We found that calculating wavelet coefficients for a subset of the fields was insufficient to capture proper dynamics of mergers. Scalar functions such as the conformal factor χ or lapse α were too smooth to provide adequate resolution for all other fields. Even after including \tilde{A} and other nonscalar fields, it was ultimately only when we included *all* fields that we achieved good behavior of our binary mergers [cf. [10]].
- [29] M. Ansorg, B. Brügmann, and W. Tichy, *Phys. Rev. D* **70**, 064011 (2004).
- [30] Z. B. Etienne, *Phys. Rev. D* **110**, 064045 (2024).
- [31] Z. B. Etienne, J. G. Baker, V. Paschalidis, B. J. Kelly, and S. L. Shapiro, *Phys. Rev. D* **90**, 064032 (2014).
- [32] H.-J. Yo, T. W. Baumgarte, and S. L. Shapiro, *Phys. Rev. D* **66**, 084026 (2002).
- [33] M. D. Duez, P. Marronetti, S. L. Shapiro, and T. W. Baumgarte, *Phys. Rev. D* **67**, 024004 (2003).
- [34] H. Sundar, R. S. Sampath, and G. Biros, *SIAM J. Sci. Comput.* **30**, 2675 (2008).
- [35] C. E. Shannon, *Proc. IRE* **37**, 10 (1949).
- [36] The upper range of human hearing is $\lesssim 20$ Hz; CDs sample at 44.1 or 48 kHz, about 1.1 to 1.2 \times the Nyquist frequency (40 kHz). National instruments suggests sampling at 1.5–2.5 \times the Nyquist frequency [37]. Robotae suggests $5f_{\text{Nyquist}}$ as a rule of thumb [38].
- [37] NI, Acquiring an analog signal: Bandwidth, nyquist sampling theorem, and aliasing, <https://www.ni.com/en/shop/data-acquisition/measurement-fundamentals/analog-fundamentals/acquiring-an-analog-signal--bandwidth--nyquist-sampling-theorem-.html> (2024).
- [38] K. Rathbone, Are you abusing nyquist?, <https://robotae.com/newsroom/are-you-abusing-nyquist/> (2021).
- [39] E. Newman and R. Penrose, *J. Math. Phys. (N.Y.)* **3**, 566 (1962).
- [40] The delay of $22M$ past merger comes from an analysis of ringdown modes in Giesler *et al.* [41]: Analyzing data displayed in their Fig. 8, we see at approximately $22M$ past merger the primary mode becomes 10 times larger than the secondary mode. Perturbation theory can then take over, modeling waveforms from there on out.
- [41] M. Giesler, M. Isi, M. A. Scheel, and S. A. Teukolsky, *Phys. Rev. X* **9**, 041060 (2019).
- [42] <https://access-ci.org/>.
- [43] <https://tacc.utexas.edu/systems/stampede3/>.
- [44] W. K. Black, 10.5281/zenodo.14649177 (2025).
- [45] The initial coordinates have $R_i = m_i/2$, but this quickly settles (after $t/M < 10$) into the evolution coordinates of $R_i = m_i$.