

**Towards beating the curse of dimensionality for gravitational waves using reduced basis**Scott E. Field,<sup>1</sup> Chad R. Galley,<sup>2,3</sup> and Evan Ochsner<sup>4</sup><sup>1</sup>*Department of Physics, Maryland Center for Fundamental Physics, Joint Space Sciences Institute, Center for Scientific Computation and Mathematical Modeling, University of Maryland, College Park, Maryland 20742, USA*<sup>2</sup>*Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California 91109, USA*<sup>3</sup>*Theoretical Astrophysics, California Institute of Technology, Pasadena, California 91125, USA*<sup>4</sup>*Center for Gravitation and Cosmology, University of Wisconsin-Milwaukee, Milwaukee, Wisconsin 53201, USA*

(Received 29 May 2012; published 23 October 2012)

Using the reduced basis approach, we efficiently compress and accurately represent the space of waveforms for nonprecessing binary black hole inspirals, which constitutes a four-dimensional parameter space (two masses, two spin magnitudes). Compared to the nonspinning case, we find that only a *marginal* increase in the (already relatively small) number of reduced basis elements is required to represent any nonprecessing waveform to nearly numerical round-off precision. Most parameters selected by the algorithm are near the boundary of the parameter space, leaving the bulk of its volume sparse. Our results suggest that the full eight-dimensional space (two masses, two spin magnitudes, four spin orientation angles on the unit sphere) may be highly compressible and represented with very high accuracy by a remarkably small number of waveforms, thus providing some hope that the number of numerical relativity simulations of binary black hole coalescences needed to represent the entire space of configurations is not intractable. Finally, we find that the *distribution* of selected parameters is robust to different choices of seed values starting the algorithm, a property which should be useful for indicating parameters for numerical relativity simulations of binary black holes. In particular, we find that the mass ratios  $m_1/m_2$  of nonspinning binaries selected by the algorithm are mostly in the interval  $[1, 3]$  and that the median of the distribution follows a power-law behavior  $\sim(m_1/m_2)^{-5.25}$ .

DOI: [10.1103/PhysRevD.86.084046](https://doi.org/10.1103/PhysRevD.86.084046)

PACS numbers: 04.80.Nn, 04.25.dg

**I. INTRODUCTION**

The upcoming generation of advanced-sensitivity ground-based gravitational wave interferometer detectors (i.e., advanced LIGO, advanced Virgo, Indigo, and KAGRA) [1–4] brings an increasing demand to accurately and efficiently represent gravitational waveforms from generic precessing compact binary sources [5–12]. Such waveforms for quasicircular inspirals are parametrized by a set of eight intrinsic physical quantities: two masses, two spin magnitudes, and four spin orientation angles on the unit sphere (8D) [13].

The large dimensionality of this parameter space makes gravitational wave searches, parameter estimation, and modeling prohibitively expensive and computationally unfeasible with most methods. This problem is called the “curse of dimensionality” [16] and, in particular, is a major hurdle for modeling astrophysical gravitational wave sources. In this setting, the cost of numerical relativity simulations of the full Einstein equations describing the inspiral, merger, and ringdown of binary black hole coalescences is so expensive that an optimal or nearly optimal criterion for selecting which points in parameter space to simulate is thus critical. This is also the case when numerical simulations are used to calibrate or build effective-one-body [17–20] or phenomenological [21–23] models.

To address these issues, the reduced basis (RB) approach was introduced to gravitational wave physics in Refs. [24,25] and shown to efficiently compress the space

of waveforms with a very small loss of accuracy, typically of the order of numerical round-off, for any given bounded parameter domain. The compression is accomplished by determining a set of nearly optimal physical parameter values from which a basis is constructed to represent any given waveform within the same physical model through its projection onto this basis. Details of the algorithm, which is known as the *greedy algorithm*, can be found in Refs. [24,25].

The reduced basis approach has several uniquely appealing features such as its hierarchical compression of the waveform space, its ability to handle large numbers of parameters, and its ability to identify the most relevant points in the parameter space. The latter may be particularly important for guiding which numerical relativity simulations to perform in order to calibrate effective-one-body or to fit phenomenological models.

In this paper we apply the RB approach to *nonprecessing* binary black hole inspirals that are nonspinning (2D), have equal spin magnitudes (3D), or have unequal spin magnitudes (4D). To our knowledge, this is the first paper, along with Refs. [25,26], applying reduced order modeling techniques to gravitational waves with a larger number of parameters than previously considered.

The results presented in this paper point towards the possibility that the eight-parameter space of waveforms for precessing binary inspirals might admit a remarkably compact representation. For the nonprecessing inspiral

waveforms, we find a very moderate increase in the number of RB elements as the dimensionality of the parameter space is increased from 2D to 4D. While this number could and might change significantly for precessing binaries, former results in Ref. [27] show dimensionality reduction in the precessing case (with respect to the spin dynamics).

## II. SPIN-ALIGNED POST-NEWTONIAN WAVEFORMS

Throughout this paper we use the so-called restricted TaylorF2 post-Newtonian (PN) waveforms (see Ref. [28] and references therein for a more detailed discussion of TaylorF2 and other PN approximants). These waveforms use the stationary-phase approximation to construct analytic frequency-domain waveforms of the form

$$\tilde{h}(f) = \mathcal{A} f^{-7/6} e^{i\Psi(f)}, \quad (1)$$

where  $\Psi(f)$  is expressed as a polynomial in the PN expansion parameter  $v = (\pi M f)^{1/3}$  and  $\mathcal{A}$  is a constant such that the  $\tilde{h}$  is normalized to unity. Spin-independent corrections to  $\Psi(f)$  are currently known to  $v^7$ , or 3.5PN order [29]. The dimensionless spin vectors of the component masses,  $\tilde{\chi}_i = \tilde{S}_i/m_i^2$  with  $0 \leq |\tilde{\chi}_i| \leq 1$ , can also enter the corrections to  $\Psi(f)$  and these spin-dependent corrections are known to 2.5PN order. These include spin-orbit corrections of the form  $\tilde{\chi}_i \times \hat{L}_N$  (where  $\hat{L}_N$  is the ‘‘Newtonian’’ orbital angular momentum, which is transverse to the orbital plane) at 1.5PN and 2.5PN order, and spin-spin and self-spin corrections of the forms  $\tilde{\chi}_1 \cdot \tilde{\chi}_2$  and  $\tilde{\chi}_i^2$  at 2PN order [30–33].

For a generic binary system the spins and orbital plane precess, so the relative orientations among  $\tilde{\chi}_1$ ,  $\tilde{\chi}_2$  and  $\hat{L}_N$ , and thus the spin-dependent PN corrections of  $\Psi(f)$ , vary on a precessional time scale. Currently, there are no frequency-domain template families for precessing spins, although there is at least one such effort under way [34]. In the ‘‘spin-aligned’’ case, when each spin is aligned or antialigned with  $\hat{L}_N$ , the binary does not precess and the relative orientations of  $\tilde{\chi}_1$ ,  $\tilde{\chi}_2$  and  $\hat{L}_N$  [and the spin-dependent corrections to  $\Psi(f)$ ] remain constant throughout inspiral. Thus the simple TaylorF2 model is valid and does not need to be augmented with precession equations. The spin-dependent corrections are expressed in terms of the (constant) projection of each spin along the orbital angular momentum:  $-1 \leq \chi_i \equiv \tilde{\chi}_i \cdot \hat{L}_N \leq 1$ .

In the *restricted* TaylorF2 approximation, the amplitude is expanded only to leading order while the phase is expanded to a higher order. In all cases, the spin-independent contributions to the phase are here included up to 3.5PN order, as given by Eq. (3.18) of Ref. [28]. In cases where spin is included, the 3.5PN nonspinning phase is augmented with spin-orbit, spin-spin, and self-spin corrections through 2.5PN order, as can be read from Eqs. (6.22)–(6.25) of Ref. [33].

## III. RESULTS

We consider binary black holes [35] with individual component masses  $m_i \in [3, 30]M_\odot$  and dimensionless spin magnitudes  $\chi_i \in [-1, 1]$ . Each mass direction (or dimension) in the training space (any discretization of the continuum space of parameters) is sampled with  $n_m$  points and each dimensionless spin magnitude direction with  $n_s$  points. After some numerical experimentation, we found that the majority of the selected mass components have values at the lower end of the considered range. Thus, we use a training space with points logarithmically spaced in the  $m_1$ – $m_2$  plane to provide a sufficiently dense cluster of points in the training space at those lower mass values. The inner products between any two waveforms are weighted by the reciprocal of the power spectral density for advanced LIGO [36] as given by the fitting formula in Ref. [37], with a lower cutoff frequency of 10 Hz and a maximum one corresponding to the highest innermost stable circular orbit frequency. For clarity, when we refer to the number of RB elements we mean the number of basis vectors needed to represent the training space to within the specified tolerance, typically round-off ( $\sim 10^{-14}$ – $10^{-12}$ ). The number of RB elements also equals the number of selected parameter tuples of mass and spin magnitudes.

Compared to the nonspinning case, the extra dimensionality of the spin parameter space (for nonprecessing inspiral waveforms, as considered here) has a remarkably small effect on the final number of RB elements needed to represent the *entire* training space with numerical round-off precision. In fact, as will be discussed in Sec. III C, the same RB represents *any* waveform, and not necessarily a member of the training space, in the same range of parameters to within essentially round-off precision. These are the most important results of this paper.

### A. Effect of increasing dimensionality

Figure 1 shows the maximum representation error for training spaces, corresponding to 2D, 3D, and 4D parameter spaces of nonprecessing inspiral waveforms, as a function of the number of RB elements. For each waveform  $\tilde{h}$  in the training space we compute the representation error [38]  $\|\tilde{h} - P_N[\tilde{h}]\|^2$  and then plot its maximum over the entire training space, as the number of RB elements is increased. In all cases Fig. 1 shows a similar behavior; namely, an initially slow falloff in the representation error and a later rapid convergence to round-off.

The 2D runs were done with up to  $n_m = 400$  points in each mass dimension, for a maximum size of the training space of  $n_m^2 = 1.6 \times 10^5$  total samples. For the 3D (4D) runs we used up to  $n_m = 200$  (100) and  $n_s = 50$  (20) samples in the spin direction for a training space with a maximum size of  $2 \times 10^6$  ( $4 \times 10^6$ ) samples. In order to directly compare the 3D and 4D cases in Fig. 1, we use  $n_m = 100$  samples for both since, as discussed in

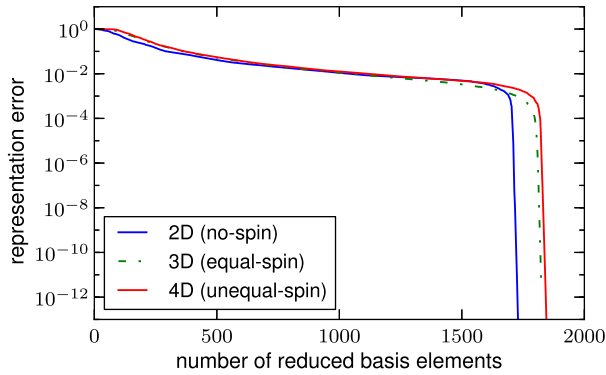


FIG. 1 (color online). The reduced basis representation error as a function of the number of reduced basis for different dimensionalities. The greedy error shows the rapid exponential convergence to round-off. The number of reduced basis elements increases very mildly from the 2D ( $\approx 1725$  RB) to the 3D ( $\approx 1824$  RB) case and to the 4D ( $\approx 1839$  RB) case. This shows that the increasing dimensionality may not be a major obstacle for the reduced basis approach.

Sec. III C, more samples are not needed for high accuracy reconstruction of arbitrary 3D or 4D waveforms.

The key result from Fig. 1 is the very moderate increase in the maximum number of RB elements needed as the dimensionality increases from 2D to 3D to 4D. At an error of  $10^{-11}$  the 2D case requires  $\approx 1725$ , the 3D case  $\approx 1824$ , and the 4D case  $\approx 1839$  basis vectors. The number of basis vectors that result by increasing the dimensionality of the parameter space from 2D to 4D is only a 6.6% increase and from 3D to 4D requires only 15 more basis vectors. Despite doubling the dimensionality of the parameter space, the number of basis vectors needed to represent the four-parameter waveform space is nearly the same as for non-spinning binaries.

Naïve scaling arguments would suggest an increase in the number of RB elements by orders of magnitude. If the 2D case requires  $\sim 2 \times 10^3$  basis vectors then one might be lead to think that doubling the number of parameters could require more by a factor the size of the training space in the spin directions, which would here be about  $20^2$  for a total of  $\sim 8 \times 10^5$ . The fact that we observe an exceedingly small relative increase in the number of RB elements lends hope that the number of basis vectors needed to represent the full 8D parameter space for inspiral waveforms is also much smaller than what might be anticipated by estimating the volume of the parameter space or by using naive (such as equally spaced or random) sampling techniques. This is only suggestive since precession has significant effects on the structure of the inspiral waveforms [6,30,33]. But if it holds, the curse of dimensionality might be beaten. In fact, indications that the problem, in the presence of precession, is amenable to dimensional reduction, have already been found through a principal component analysis of the precessing dynamics of compact binary inspirals [27]. There it was found that for the case of a random selection of

precessing binaries with the same total mass there are three combinations of spin orientations that are semiconserved (in a statistical sense) throughout the inspiral. The presence of such combinations implies that the dynamical configurations can be parametrized by a reduced number of independent parameters than the fiducial seven (one mass, six spin components).

## B. Parameter values selected

The reduced basis-greedy algorithm approach possesses several unique features. One of them is being able to identify, in a precise mathematical sense, a nearly optimal set of points in the physical parameter space [39]. In this section we analyze the structure of the selected parameter values for the models and scenarios here studied.

### 1. Mass parameters (2D, 3D, 4D)

Figure 2 shows the distribution of selected points for the mass ratio  $m_1/m_2$  in the nonprecessing (4D) case. There we plot  $\log_{10}(m_1/m_2)$  to obtain a symmetric representation in  $m_1$  and  $m_2$ . Since in this paper we focus on the range  $m_1, m_2 \in [3, 30]M_\odot$ , the mass ratio has a range  $m_1/m_2 \in [1/10, 10]$ . The selected values by the greedy algorithm have a strong peak at the equal-mass case. This follows from the selected individual mass components, which are clustered at low values (see also Fig. 3).

We have found similar results for the distribution of  $m_1/m_2$  in the 2D and 3D cases, also resulting from the low-mass selection. This is also visible from Fig. 3, which shows the selected parameter values in the  $m_1$ - $m_2$  plane. Notice that there is little qualitative difference in going from 2D to 3D to 4D in Fig. 3. This indicates that the problem might admit some kind of dimensional separability. For example, it might be possible to initially select  $(m_1, m_2)$  values through a 2D greedy strategy and then use only those to populate the  $m_1$ - $m_2$  plane in the construction of very efficient and compact training spaces for building




FIG. 2 (color online). Histogram of the selected parameter choices for the mass ratio  $m_1/m_2$  in the 4D case. Note the large spike near equal masses and the rapid falloff to small counts for increasing disparity in the component masses.




FIG. 3. The reduced basis parameter choices in the  $m_1$ - $m_2$  plane. The overall structure is similar for all three (2D, 3D, 4D) nonprecessing cases. The gray scale indicates the number of selected points in a bin on a logarithmic scaling.

reduced basis in higher dimensional parameter cases. This is beyond the scope of this paper, but it might provide a dramatic reduction in training space sizes and, if so, a possible avenue for expansions to higher dimensions and the full 8D problem.

## 2. Mass and spin parameters (3D)

The 3D equal-spin parameter space ( $m_1, m_2, \chi_1 = \chi = \chi_2$ ) can be visualized directly. In Fig. 4 we plot the parameter values selected for the reduced basis. Comparing this to Fig. 3 one sees again the usual selection of primarily low-mass systems. For the lowest mass configurations, a large range of spin values is selected. The figure also shows that very little of the bulk of the parameter space volume is used. The greedy algorithm mostly picks values at some edges and corners in the  $m_1$ - $m_2$  plane and preferentially picks binaries with both spins antialigned to the orbital angular momentum ( $\chi = -1$ ); see also Fig. 5(a) showing a histogram of the selected equal-spin magnitudes.

The global aspect of the RB-greedy algorithm allows an identification of the underlying sparsity clearly visible in





FIG. 4. The reduced basis parameter choices for the 3D case ( $m_1, m_2, \chi$ ). Comparing to Fig. 3 one can also see the selection of primarily low-mass systems. For the lowest mass systems a large number of spin values are selected. Few systems from the bulk volume of the parameter space are chosen.



(a) Spin magnitude distribution (3D case) (b) Spin magnitudes distribution (4D case)

FIG. 5 (color online). Figure 5(a) shows a histogram of spin values in the 3D case. Spins antialigned with the orbital angular momentum are predominantly chosen. Figure 5(b) shows a two-dimensional histogram of spin values in the 4D case. Similarly to Fig. 5(a), antialigned systems  $\chi_i = -1$  are predominantly chosen. The gray scale indicates the number of selected points in a bin on a logarithmic scaling.

Fig. 4. This provides another hint that: 1) the binary coalescence problem is amenable to dramatic dimensional reduction with respect to the physical parameters describing it, and 2) specific techniques for training space construction exploiting sparsity might provide dramatic computational cost advantages when applied to higher dimensional problems (we elaborate more on this in Sec. IV).

## 3. Spin parameters (3D, 4D)

Figure 5(b) shows the 1839 spin magnitudes  $\chi_1$  and  $\chi_2$  selected by the greedy algorithm in the generic nonprecessing 4D case. The training space corresponding to this figure has  $n_s = 20$  points in each spin direction. As found in the 3D case, binaries with both spins mostly antialigned with the orbital angular momentum are more relevant than those that are aligned.

The figure is asymmetric in  $\chi_1, \chi_2$  for the following reason. Although the representation error of any two waveforms is symmetric in  $(m_1, \chi_1) \leftrightarrow (m_2, \chi_2)$ , the greedy algorithm needs to select one set of parameters from the two possible choices introducing an asymmetry. This effect is visually enhanced by the logarithmic scale used in the figure.

As discussed in Sec. III B 1, the distribution of masses selected does not change strongly as one goes from 2D to 4D. Similarly here, the distribution of spin values selected seems qualitatively similar in the 3D and 4D cases, suggesting that it might be possible to construct training spaces in a more efficient manner.

## C. Monte Carlo representation error studies

To test the accuracy of the RB to represent waveforms that are not necessarily members of the training space, we randomly sample the corresponding parameter space used to build each basis, and for each sampled waveform we compute its representation error by projecting the waveform onto the RB. We performed this test for all models (2D, 3D, 4D), but for definiteness we show the more




FIG. 6 (color online). Monte Carlo reconstruction error study for the 4D case. The histogram shows the distribution of the representation error  $\sigma = \|\tilde{h} - P_N[\tilde{h}]\|^2$  for  $10^7$  randomly selected parameter values. Notice the logarithmic scale on the vertical axis.

interesting 4D case in Fig. 6. Shown there is the computed waveform representation error for  $10^7$  randomly chosen values in the 4D parameter space. The histogram shows that the overwhelming majority of sampled waveforms have a representation error near double-precision round-off values, with a few isolated cases with an error slightly larger than  $10^{-11}$ .

For the particular reconstruction test shown in Fig. 6 we used a RB built out of a training space with  $n_m = 100$  points for each mass component and  $n_s = 10$  for each spin magnitude, as described in Sec. III A. For the 2D case, this number of mass samples is not enough to achieve such a low representation error. This indicates that there is a trade-off between sampling density in the mass and spin dimensions that can be used to reduce the representation error. The effectiveness of the reduced basis approach is such that, for example, even though a relatively small number of spin values are used in the training space to build the RB, a waveform with *any* spin value is represented within essentially machine precision.

#### D. Seed dependence

Lastly, we investigate the dependence of the parameters selected by the greedy algorithm on the choice of initial seed for the 2D case (nonspinning). The seeds are chosen to correspond to the 200 equal-mass binaries in the training space. For a given seed and tolerance error, the greedy algorithm picks a set of nearly optimal parameters, or a *greedy chain*. Different seeds give rise to different greedy chains and the chosen parameters will not necessarily be part of other greedy chains.

We showed in Ref. [25] that the representation error is robust to different choices of seed values by running the greedy algorithm multiple times, once for each possible




FIG. 7 (color online). Parameters chosen by the greedy algorithm ( $m_1$  left and  $m_1/m_2$  right) when using 200 seed values corresponding to equal mass, nonspinning binaries with different total masses. The lines are quantile curves representing the fraction of selected parameters with values greater than the value of the quantile. Black corresponds to the median and the bounding black curves correspond to the 0% (top) and 100% (bottom) quantiles. The red line is a fit to the median. Despite the different choice of greedy parameters that are selected, the *distribution* of points is rather robust to the choice of seed.

value in the training space for the seed [40]. Here, we discuss how the distribution of points is affected by different seed choices.

Figure 7 shows the probability density over all 200 seeds for a given value of  $m_1$  (left) and for  $m_1/m_2$  (right), where the mass ratio is taken to be such that if  $m_1 > m_2$  then the 1 and 2 labels are interchanged so that  $m_1/m_2 \geq 1$ . The gray-level shaded lines correspond to quantile curves, with black indicating the median of the distribution. Both plots indicate that the *distribution* of points is rather robust to different seed choices. This reflects the global nature of the greedy algorithm to select parameters that are distributed similarly even though the selected parameter values themselves exhibit variations from one chain to another. These distributions also show that the majority of points selected, and overwhelmingly so, are those with low masses independently of the seed value. A similar distribution of points exists for the selected values of  $m_2$ . The red line in the right panel of Fig. 7 is a power-law fit to the median and is given by  $\approx 1.83(m_1/m_2)^{-5.25} + 0.05$ . Most mass ratios selected lie within the range  $m_1/m_2 \in [1, 3]$ .

#### IV. CONCLUDING REMARKS

The main focus of this work has been the study of reduced basis in higher dimensional parameter spaces; more specifically, the study of nonprecessing, spinning binaries. We have shown (see Figs. 1 and 6) that the number of reduced basis vectors needed to represent the full space of nonprecessing inspiral waveforms with

essentially round-off precision increases very mildly with the number of parameters from 2D to 4D.

While the work presented here is limited to the non-precessing case, our results provide the first hints, together with Ref. [27] for precessing inspiral dynamics (see also Ref. [12]), Ref. [25] for quasinormal mode ringing, and studies using a singular value decomposition (SVD) technique [26,41,42] that a full representation of the eight-parameter space of inspiral-merger-ringdown waveforms (either numerical, effective-one-body, or phenomenological ones) might actually be achievable with a relatively compact reduced basis. If this is the case, then there is hope that the number of simulations needed to represent the space of precessing binary inspirals might be relatively small (perhaps on the order of several thousand for advanced LIGO but not orders of magnitude larger than this), thus allowing for a tractable number of numerical simulations if the parameters are chosen with malice of forethought from our reduced basis studies.

Recently, SVD of gravitational waves [41] has been implemented in a realistic, low-latency search pipeline [43]. Such a pipeline could be used to improve upon the work in Ref. [44] to allow prompt follow-up searches for electromagnetic counterparts to candidate gravitational wave signals in the advanced detector era. As a reconstruction and compression technique, RB has a number of similarities with SVD (e.g., output includes basis vectors and reconstruction coefficients in both schemes). Thus the RB results of this

paper could be implemented in such a pipeline, in principle. Such practical data analysis implementations of RB will be left for future work.

The exploration of the 8D parameter space will require further development and/or implementation of technical but critical aspects, including the efficient and adaptive sampling techniques for large training spaces (see, e.g., Ref. [45]) and rapid evaluation of high accuracy quadratures for parametrized problems [46]. For the problem at hand, a splitting of dimensions as discussed in Sec. III B 1 might be useful.

## ACKNOWLEDGMENTS

This work has been supported by NSF Grants No. PHY1208861 and No. PHY1005632 to the University of Maryland, and NSF Grant No. PHY0970074 to the University of Wisconsin-Milwaukee. C. G. was supported by an appointment to the NASA Postdoctoral Program at the Jet Propulsion Laboratory administered by Oak Ridge Associated Universities through a contract with NASA. Copyright 2012. All rights reserved. We thank Alessandra Buonanno, Sarah Caudill, Tom Dent, Alexandre Le Tiec, Drew Keppel, and Dianne O’Leary for very helpful comments on the manuscript and/or suggestions. We especially thank Manuel Tiglio for helpful discussions and comments on the manuscript as well as generating much of the data used in this paper.

- 
- [1] LIGO, <http://www.ligo.caltech.edu/>.
  - [2] Virgo, <http://www.virgo.infn.it/>.
  - [3] GEO600, <http://www.geo600.uni-hannover.de/>.
  - [4] KAGRA, <http://gwcenter.icrr.u-tokyo.ac.jp/en/>.
  - [5] J. Abadie *et al.* (LIGO Scientific Collaboration), *Classical Quantum Gravity* **27**, 173001 (2010).
  - [6] T.A. Apostolatos, C. Cutler, G.J. Sussman, and K.S. Thorne, *Phys. Rev. D* **49**, 6274 (1994).
  - [7] D.A. Brown, A. Lundgren, and R. O’Shaughnessy, [arXiv:1203.6060](https://arxiv.org/abs/1203.6060).
  - [8] P. Ajith, *Phys. Rev. D* **84**, 084037 (2011).
  - [9] J.M. Centrella, J.G. Baker, B.J. Kelly, and J.R. van Meter, *Annu. Rev. Nucl. Part. Sci.* **60**, 75 (2010).
  - [10] A. Buonanno, Y. Chen, Y. Pan, and M. Vallisneri, *Phys. Rev. D* **70**, 104003 (2004).
  - [11] A. Buonanno, Y.-b. Chen, and M. Vallisneri, *Phys. Rev. D* **67**, 104025 (2003).
  - [12] Y. Pan, A. Buonanno, Y. Chen, and M. Vallisneri, *Phys. Rev. D* **69**, 104017 (2004).
  - [13] Alternative theories of gravity may have more parameters than general relativity; see, for example, Refs. [14,15] and references therein.
  - [14] K. Chatziioannou, N. Yunes, and N. Cornish, [arXiv:1204.2585](https://arxiv.org/abs/1204.2585).
  - [15] N. Yunes and F. Pretorius, *Phys. Rev. D* **80**, 122003 (2009).
  - [16] R.E. Bellman, *Dynamic Programming* (Dover Publications, New York, 2003).
  - [17] T. Damour and A. Nagar, *Phys. Rev. D* **79**, 081503 (2009).
  - [18] A. Buonanno and T. Damour, *Phys. Rev. D* **59**, 084006 (1999).
  - [19] Y. Pan, A. Buonanno, L.T. Buchman, T. Chu, L.E. Kidder, H.P. Pfeiffer, and M.A. Scheel, *Phys. Rev. D* **81**, 084041 (2010).
  - [20] A. Taracchini *et al.*, [arXiv:1202.0790](https://arxiv.org/abs/1202.0790).
  - [21] P. Ajith *et al.*, *Classical Quantum Gravity* **24**, S689 (2007).
  - [22] P. Ajith *et al.*, *Phys. Rev. Lett.* **106**, 241101 (2011).
  - [23] L. Santamaria *et al.*, *Phys. Rev. D* **82**, 064016 (2010).
  - [24] S.E. Field, C.R. Galley, F. Herrmann, J.S. Hesthaven, E. Ochsner, and M. Tiglio, *Phys. Rev. Lett.* **106**, 221102 (2011).
  - [25] S. Caudill, S.E. Field, C.R. Galley, F. Herrmann, and M. Tiglio, *Classical Quantum Gravity* **29**, 095016 (2012).
  - [26] D. Keppel, K. Cannon, M. Frei, and C. Hanna, <http://www.gravity.phys.uwm.edu/conferences/gwpaw/posters/keppel.pdf>.
  - [27] C.R. Galley, F. Herrmann, J. Silberholz, M. Tiglio, and G. Guerberoff, *Classical Quantum Gravity* **27**, 245007 (2010).
  - [28] A. Buonanno, B. Iyer, E. Ochsner, Y. Pan, and B. Sathyaprakash, *Phys. Rev. D* **80**, 084043 (2009).

- [29] L. Blanchet, G. Faye, B. R. Iyer, and B. Joguet, *Phys. Rev. D* **65**, 061501 (2002).
- [30] L. E. Kidder, *Phys. Rev. D* **52**, 821 (1995).
- [31] G. Faye, L. Blanchet, and A. Buonanno, *Phys. Rev. D* **74**, 104033 (2006).
- [32] L. Blanchet, A. Buonanno, and G. Faye, *Phys. Rev. D* **74**, 104034 (2006).
- [33] K. Arun, A. Buonanno, G. Faye, and E. Ochsner, *Phys. Rev. D* **79**, 104023 (2009).
- [34] P. Ajith and M. Favata (unpublished).
- [35] Because binary neutron stars can have lower component masses, possibly in the range  $[1, 3]M_{\odot}$ , there are many more cycles in band in the corresponding waveforms when compared to those of binary astrophysical black holes. This introduces a larger off-line computational cost to build reduced bases, but not a conceptual one. Hence, we use binary black holes for the foundation of our studies here and defer a detailed study of binary neutron stars once some additional computational developments to alleviate such off-line costs are in place (see the discussion at the end of Sec. IV).
- [36] As briefly discussed in Ref. [24], the form of the power spectral density is largely irrelevant for our purposes since our ability to represent waveforms with very high accuracy allows for one to represent waveforms of a template bank for other detectors, usually with very little loss of accuracy.
- [37] P. Ajith and S. Bose, *Phys. Rev. D* **79**, 084032 (2009).
- [38]  $P_N[\tilde{h}]$  is the best representation of  $\tilde{h}$  using  $N$  reduced basis elements, see Ref. [25] for more details.
- [39] P. Binev, A. Cohen, W. Dahmen, R. DeVore, G. Petrova, and P. Wojtaszczyk, *SIAM J. Math. Anal.* **43**, 1457 (2011).
- [40] Reference [25] was concerned with quasinormal mode ringdown waveforms. However, the result is independent of the specific waveform under consideration.
- [41] K. Cannon, A. Chapman, C. Hanna, D. Keppel, A. C. Searle, and A. J. Weinstein, *Phys. Rev. D* **82**, 044025 (2010).
- [42] K. Cannon, C. Hanna, and D. Keppel, *Phys. Rev. D* **84**, 084003, (2011).
- [43] K. Cannon *et al.*, *Astrophys. J.* **748**, 136 (2012).
- [44] LSC and Virgo Collaborations, [arXiv:1112.6005](https://arxiv.org/abs/1112.6005).
- [45] J. S. Hesthaven, B. Stamm, and S. Zhang, Brown University Report No. 2011-15, 2011 (unpublished).
- [46] H. Antil, S. Field, F. Herrmann, R. Nochetto, and M. Tiglio, [arXiv:1210.0577](https://arxiv.org/abs/1210.0577).