

Towards beating the curse of dimensionality for gravitational waves using Reduced Basis

Scott E. Field,¹ Chad R. Galley,^{2,3} Frank Herrmann,¹ Evan Ochsner,⁴ and Manuel Tiglio¹

¹*Department of Physics, Maryland Center for Fundamental Physics, Joint Space Sciences Institute, Center for Scientific Computation and Mathematical Modeling, University of Maryland, College Park, MD 20742, USA*

²*Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109 USA*

³*TAPIR, California Institute of Technology, Pasadena, CA 91125, USA*

⁴*Center for Gravitation and Cosmology, University of Wisconsin-Milwaukee, Milwaukee, WI 53201, USA*

Using the Reduced Basis approach, we efficiently compress and accurately represent the space of waveforms for non-precessing binary black hole inspirals, which constitutes a four dimensional parameter space (two masses, two spin magnitudes). Compared to the non-spinning case, we find that only a *marginal* increase in the (already relatively small) number of reduced basis elements is required to represent any non-precessing 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 non-spinning binaries selected by the algorithm are mostly in the interval $[1, 3]$ and that the median of the distribution follows a power-law behavior $\sim (m_1/m_2)^{-5.25}$.

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 [5–8]. Such waveforms for quasi-circular inspirals are parameterized by a set of eight intrinsic physical quantities – two masses, two spin magnitudes, and four spin orientation angles on the unit sphere (8D) [35].

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” [9] 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 [10–12] or phenomenological [13–15] models.

To address these issues, the Reduced Basis (RB) approach was introduced to gravitational wave physics in Refs. [16, 17] 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 [16, 17].

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 *non-precessing* binary black hole inspirals that are non-spinning (2D), have equal spin magnitudes (3D), or have unequal spin magnitudes (4D). To our knowledge, this is the first paper, along with Ref. [17], 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 non-precessing 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 [18] 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 TaylorF2 post-Newtonian (PN) waveforms (see [19] 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 [20]. The dimensionless spin vectors of the component masses, $\vec{\chi}_i = \vec{S}_i/m_i^2$ with $0 \leq |\vec{\chi}_i| \leq 1$, 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 $\vec{\chi}_i \cdot \hat{L}_N$ (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 $\vec{\chi}_1 \cdot \vec{\chi}_2$ and $\vec{\chi}_i^2$ at 2PN order [21–24].

For a generic binary system the spins and orbital plane precess, so the relative orientations among $\vec{\chi}_1$, $\vec{\chi}_2$ and \hat{L}_N , and thus the spin-dependent PN corrections of $\Psi(f)$, vary on a precessional timescale. A naive implementation of the TaylorF2 waveforms that does not take into account this precessional motion is not valid. There is at least one effort underway to implement such a frequency-domain PN model for precessing binaries [25], but at present only time-domain PN models for precessing binaries have been implemented. In the ‘‘spin-aligned’’ case, when each spin is aligned or anti-aligned with \hat{L}_N , the binary does not precess and the relative orientations of $\vec{\chi}_1$, $\vec{\chi}_2$ 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: $-1 \leq \chi_i \equiv \vec{\chi}_i \cdot \hat{L}_N \leq 1$.

In this paper we use *restricted* TaylorF2 waveforms, meaning that 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 included up to 3.5PN order, as given by Eq. (3.18) of [19]. In cases where spin is included, the 3.5PN non-spinning 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 [24].

III. RESULTS

We consider binary black holes [36] 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 (PSD) for advanced LIGO [37] as given by the fitting formula in Ref. [26]. 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 non-spinning case, the extra dimensionality of the spin parameter space (for non-precessing 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 Section 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 sets, corresponding to 2D, 3D, and 4D parameter spaces of non-precessing inspiral waveforms, as a function of the number of RB elements. For each waveform \tilde{h} in the training set we compute the representation error [38] $\|\tilde{h} - P_N[\tilde{h}]\|^2$ and then plot its maximum over the entire training set, as the number of RB elements is increased. In all cases Fig. 1 shows a similar behavior, namely, an initially slow fall-off 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 2×10^6 (4×10^6) 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 Section 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 $\approx 1,725$, the 3D case $\approx 1,824$, and the 4D case $\approx 1,839$ 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.

Naive 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 [21, 24, 27]. 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 [18]. 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 semi-conserved (in a statistical sense) throughout the inspiral. The presence of such combinations implies that the dynamical configurations can be parameterized by a reduced number of independent parameters than the fiducial seven (1 mass, 6 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 [28]. 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 non-precessing (4D) case.




FIG. 1: 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 (no-spin) to the 3D (equal-spin) case and to the 4D (generic non-precessing spins) case. This shows that the increasing dimensionality may not be a major obstacle for the Reduced Basis approach.

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).




FIG. 2: 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 fall-off to small counts for increasing disparity in the component masses.

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 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.




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) non-precessing cases. The grey-scale bar utilizes a logarithmic scaling.

2. Mass and spin parameters (3D)

The 3D equal-spin parameter space ($m_1, m_2, \chi_1 = \chi_2 = \chi$) 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 anti-aligned 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. 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).




FIG. 4: The reduced basis parameter choices for the 3D case (m_1, m_2, χ). 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.

3. Spin parameters (3D, 4D)

Fig. 5(b) shows the 1,839 spin magnitudes, χ_1 and χ_2 , selected by the greedy algorithm in the generic non-precessing 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 anti-aligned with the orbital angular momentum are more relevant than those that are aligned.




FIG. 5: Fig. (a) shows a histogram of spin values in the 3D case. Spins anti-aligned with the orbital angular momentum are predominantly chosen. Fig. (b) shows a two-dimensional histogram of spin values in the 4D case. Similarly to Fig. (a), anti-aligned systems $\chi_i = -1$ are predominantly chosen. The grey-scale bar utilizes a logarithmic scaling.

The figure is asymmetric in χ_1, χ_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 se-

lected 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 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.




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

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 (non-spinning). The seeds are chosen to correspond to the $n_m = 200$ 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. [17] that the representation error is robust to different choices of seed values by running the greedy algorithm multiple times, once for each possible value in the training space for the seed [39]. Here, we discuss how the distribution of points is affected by different seed choices.




FIG. 7: Parameters chosen by the greedy algorithm (m_1 , left and m_1/m_2 , right) when using 200 seed values corresponding to equal mass, non-spinning 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.

Fig. 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 non-precessing, spinning binaries. We have shown (see Figs. 1 and 6) that the number of reduced basis vectors needed to represent the full space of non-precessing 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 [18] for precessing inspiral dynamics and [17] for quasi-normal mode ringing, that a full representation of the eight-parameter space of inspiral-merger-ringdown waveforms might actually be achievable with a relatively compact reduced basis. In addition, the results of this paper may help in guiding the parameter values to choose for numerical simulations of binary black hole coalescences. These two things taken together suggest that the number of such simulations needed 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.

As a reconstruction and compression technique, RB has a number of similarities with Singular Value Decomposition (SVD) (e.g., output includes basis vectors in both schemes). Recently, SVD of gravitational waves [29] has been implemented in a realistic search pipeline and achieved very low-latency [30] for prompt follow-up searches of potential electromagnetic counterparts to candidate gravitational wave signals. 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., [31]) and rapid evaluation of high accuracy quadratures for parametrized problems [32]. For the problem at hand, a splitting of dimensions as discussed in Sec. III B 1 might be useful.

V. ACKNOWLEDGMENTS

This work has been supported by NSF Grants PHY1208861 and PHY1005632 to the University of Maryland, and NSF Grant 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 Dianne O'Leary for very helpful suggestions.

-
- [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] LIGO Scientific, J. Abadie *et al.*, *Class. Quantum Grav.* **27**, 173001 (2010), arXiv:1003.2480.
 - [6] D. A. Brown, A. Lundgren, and R. O'Shaughnessy, (2012), arXiv:1203.6060.
 - [7] P. Ajith, *Phys. Rev.* **D84**, 084037 (2011), arXiv:1107.1267.
 - [8] J. M. Centrella, J. G. Baker, B. J. Kelly, and J. R. van Meter, *Ann.Rev.Nucl.Part.Sci.* **60**, 75 (2010), arXiv:1010.2165.
 - [9] R. E. Bellman, *Dynamic Programming* (Dover Publications, Incorporated, 2003).
 - [10] A. Buonanno and T. Damour, *Phys. Rev. D* **59**, 084006 (1999), arXiv:gr-qc/9811091.
 - [11] Y. Pan *et al.*, *Phys. Rev.* **D81**, 084041 (2010), arXiv:0912.3466.
 - [12] A. Taracchini *et al.*, (2012), arXiv:1202.0790.
 - [13] P. Ajith *et al.*, *Class. Quantum Grav.* **24**, S689 (2007), arXiv:0704.3764.
 - [14] P. Ajith *et al.*, (2009), arXiv:0909.2867.
 - [15] L. Santamaria *et al.*, *Phys. Rev.* **D82**, 064016 (2010), arXiv:1005.3306.
 - [16] S. E. Field *et al.*, *Phys. Rev.Lett.* **106**, 221102 (2011), arXiv:1101.3765.
 - [17] S. Caudill, S. E. Field, C. R. Galley, F. Herrmann, and M. Tiglio, *Class.Quant.Grav.* **29**, 095016 (2012), arXiv:1109.5642.
 - [18] C. R. Galley, F. Herrmann, J. Silberholz, M. Tiglio, and G. Guerberoff, *Class. Quantum Grav.* **27**, 245007 (2010), arXiv:1005.5560.
 - [19] A. Buonanno, B. Iyer, E. Ochsner, Y. Pan, and B. Sathyaprakash, *Phys.Rev.* **D80**, 084043 (2009), arXiv:0907.0700.
 - [20] L. Blanchet, G. Faye, B. R. Iyer, and B. Joguet, *Phys.Rev.* **D65**, 061501 (2002), arXiv:gr-qc/0105099.
 - [21] L. E. Kidder, *Phys. Rev.* **D52**, 821 (1995), gr-qc/9506022.
 - [22] G. Faye, L. Blanchet, and A. Buonanno, *Phys. Rev.* **D74**, 104033 (2006), gr-qc/0605139.
 - [23] L. Blanchet, A. Buonanno, and G. Faye, *Phys. Rev. D* **74**, 104034 (2006), arXiv:gr-qc/0605140.
 - [24] K. Arun, A. Buonanno, G. Faye, and E. Ochsner, *Phys. Rev.* **D79**, 104023 (2009), arXiv:0810.5336.

- [25] P. Ajith and M. Favata, in preparation.
- [26] P. Ajith and S. Bose, *Phys.Rev.* **D79**, 084032 (2009), arXiv:0901.4936.
- [27] T. A. Apostolatos, C. Cutler, G. J. Sussman, and K. S. Thorne, *Phys. Rev. D* **49**, 6274 (1994).
- [28] P. Binev *et al.*, *SIAM J. Math. Analysis* **43**, 1457 (2011).
- [29] K. Cannon *et al.*, *Phys. Rev.* **D82**, 044025 (2010), arXiv:1005.0012.
- [30] K. Cannon *et al.*, *Astrophys.J.* **748**, 136 (2012), arXiv:1107.2665.
- [31] J. S. Hesthaven, B. Stamm, and S. Zhang, Scientific Computing Group, Brown University Report No. 2011-15, 2011 (unpublished).
- [32] H. Antil, S. Field, F. Herrmann, R. Nochetto, and M. Tiglio, (2012), in preparation.
- [33] K. Chatziioannou, N. Yunes, and N. Cornish, (2012), arXiv:1204.2585.
- [34] N. Yunes and F. Pretorius, *Phys.Rev.* **D80**, 122003 (2009), arXiv:0909.3328.
- [35] Alternative theories of gravity may have more parameters than General Relativity; see, for example, [33, 34] and references therein.
- [36] 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 offline 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 offline costs are in place (see the discussion at the end of Section IV).
- [37] As briefly discussed in [16], the form of the PSD 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.
- [38] $P_N[\tilde{h}]$ is the best representation of \tilde{h} using N reduced basis elements, see Ref. [17] for more details.
- [39] Ref. [17] was concerned with quasi-normal mode ring-down waveforms. However, the result is independent of the specific waveform under consideration.