Rapidly evaluating the compact-binary likelihood function via interpolation

R. J. E. Smith,[1,2] C. Hanna,[3,4] I. Mandel,[2] and A. Vecchio[2]

[1]*LIGO, Division of Physics, Mathematics and Astronomy, California Institute of Technology, 1200 East Caslifornia Boulevard, Pasadena, California 91125, USA*

[2]*School of Physics and Astronomy, University of Birmingham, Edgbaston, Birmingham B15 2TT, UK*

[3]*Perimeter Institute for Theoretical Physics, Waterloo, Ontario N2L 2Y5, Canada*

[4]*Pennsylvania State University, University Park, Pennsylvania 16802, USA*

(Received 23 May 2013; revised manuscript received 29 January 2014; published 29 August 2014)

Bayesian parameter estimation on gravitational waves from compact-binary coalescences (CBCs) typically requires millions of template waveform computations at different values of the parameters describing the binary. Sampling techniques such as Markov chain Monte Carlo and nested sampling evaluate likelihoods and, hence, compute template waveforms, serially; thus, the total computational time of the analysis scales linearly with that of template generation. Here we address the issue of rapidly computing the likelihood function of CBC sources with nonspinning components. We show how to efficiently compute the continuous likelihood function on the three-dimensional subspace of parameters on which it has a nontrivial dependence—the chirp mass, symmetric mass ratio and coalescence time—via interpolation. Subsequently, sampling this interpolated likelihood function is a significantly cheaper computational process than directly evaluating the likelihood; we report improvements in computational time of two to three orders of magnitude while keeping likelihoods accurate to $\lesssim 0.025\%$. Generating the interpolant of the likelihood function over a significant portion of the CBC mass space is computationally expensive but highly parallelizable, so the wall time can be very small relative to the time of a full parameter-estimation analysis.

DOI: 10.1103/PhysRevD.90.044074 PACS numbers: 04.30.-w, 04.80.Nn

I. INTRODUCTION

The direct detection of gravitational waves will initiate an entirely new kind of astronomy, offering an unprecedented probe of relativistic astrophysics and strong-field gravity. Ground-based gravitational-wave interferometers LIGO [1] and Virgo [2] are undergoing upgrades to their second generation designs, Advanced LIGO/Virgo (aLIGO/AdV), and are expected to be operational around 2015 [3]. In addition two new detectors in India and Japan, LIGO India [4] and KAGRA [5], respectively, are expected to be operational around 2020. These advanced instruments will be an order of magnitude more sensitive than their predecessors [6] and are expected to usher in routine detections of gravitational waves [7]. The coalescence of compact binaries, consisting of neutron stars and/or black holes, are prime targets for gravitational-wave observatories, with realistic estimates of detection rates ~ 1–$100\ \mathrm{yr}^{-1}$ [7].

Estimating the parameters of CBC sources—e.g., their masses, spins, and sky location—is a crucial aspect of gravitational-wave astronomy. Accurate determination of the component masses can distinguish black holes from neutron stars, thereby helping to address the black hole "mass gap" issue [[8] and references therein]. Measurements of the location of binary mergers on the sky will potentially allow for searches for electromagnetic counterparts of mergers; these will, for example, make it possible to ascertain whether compact binaries are the progenitors of short, hard gamma-ray bursts [9]. Inference on the distribution of the source population from multiple gravitational-waves observations can be used to constrain compact binary formation scenarios, which are currently highly speculative [e.g., [10]]. Thus, parameter estimation to extract this information is central gravitational-wave astronomy, but remains challenging in practice.

The goal of Bayesian parameter estimation is to compute the posterior probability density function (PDF) of a set of parameters, $\vec{\theta}$, which underlie a model assumed to describe a data set d. The PDF is related to the likelihood function and prior probability via Bayes' theorem,

$$p(\vec{\theta}|d) = \frac{\mathcal{P}(\vec{\theta})\mathcal{L}(d|\vec{\theta})}{p(d)}, \qquad (1)$$

where $\mathcal{L}(d|\vec{\theta})$ is the likelihood and $\mathcal{P}(\vec{\theta})$ is the prior probability which encodes our *a priori* belief in the distribution of $\vec{\theta}$. The quantity in the denominator, $p(d)$, is known as the "evidence". Computing (1) requires evaluating the likelihood.

For binaries with nonspinning components $\vec{\theta}$ is nine dimensional. Exploring such a high dimensional space requires sophisticated stochastic Bayesian inference techniques [11–13] which preferentially sample the parameter space in regions of high posterior probability. The bulk of the computational cost of evaluating the likelihood function comes from computing template waveforms. Analyses on first-generation interferometer data require computing $\mathcal{O}(10^6)$ such waveforms [12,14]. Sampling techniques such as Markov chain Monte Carlo (MCMC) [12,13] and nested sampling [11,15] evaluate likelihoods, and hence compute template waveforms, serially. Thus the total computational time to fully sample the parameter space scales linearly with the total time spent generating template waveforms. It can take hours to weeks to analyze a single stretch of data of a few seconds in duration, depending on the choice of the template waveform family. This problem will be exacerbated when analysing second-generation interferometer data as the waveforms will be forty times longer in duration if the starting frequency $f_{\min}$ is changed from 40 Hz to 10 Hz.

For binaries with nonspinning components, the waveform has a convenient representation in the frequency domain:

$$\tilde{h}(\vec{\theta}; f) = \sum_{\mu=+,\times} A_\mu(\vec{\theta}_L)\tilde{h}_0(\mathcal{M}, \eta; f) e^{2\pi i f t_c}, \tag{2}$$

where $A_{+,\times}$ denotes the (scalar) amplitudes of the "plus-" and "cross-" polarization states of the waveform. In general $\tilde{h}_0$ depends on the waveform family being used and can be computed by Fourier transforming the associated time-domain representation of the waveform family. The parameters which describe the binary are the chirp mass and symmetric mass ratio, $\mathcal{M}$ and η, the time at coalescence t_c and a set of parameters which describe the location, orientation of, and distance to the binary, all of which enter $\vec{\theta}_L$.

Evaluating the likelihood function on the three-dimensional subspace of parameters $(\mathcal{M}, \eta, t_c)$ represents the largest computational burden to parameter estimation on gravitational waves from CBC sources with nonspinning components because a new waveform calculation is required to compute the likelihood for a new set of these parameters. In [14], we considered interpolation between waveforms over the mass parameter space as a way to reduce computational cost. Here, we demonstrate that the evaluation of an interpolated likelihood function over the $(\mathcal{M}, \eta, t_c)$ subspace is a much faster computational procedure than the standard calculation of the likelihood (3) by using either full or interpolated waveforms. For the purposes of parameter estimation, one is not interested in template waveforms per se, but rather in the posterior probability distributions of the underlying parameters of the template waveforms that are assumed to describe the data. By directly using interpolated likelihood functions, one effectively bypasses template waveform generation during the sequential steps of an MCMC. This likelihood-interpolation technique is robust and could, in principle, be generalized to arbitrary template waveform families, in particular those that describe CBCs with spinning components.

II. DIRECTLY INTERPOLATING THE LIKELIHOOD FUNCTION

We wish to generate a representation of the likelihood function over the continuous $\mathcal{M}$, η, and t_c subspace. To achieve this we will interpolate the likelihood function over $\mathcal{M}$, η and t_c. The likelihood function that describes the probability of observing a data stream $d = h + n$ containing a given gravitational-wave signal $h(\vec{\theta}; t)$ and Gaussian and stationary noise $n(t)$ is [11]

$$\log \mathcal{L}(d|\vec{\theta}) = (d|h(\vec{\theta})) - \frac{1}{2}[(h(\vec{\theta})|h(\vec{\theta})) + (d|d)], \tag{3}$$

where $(a|b)$ is the usual noise-weighted inner product [16]. We define the complex-valued time series corresponding to the inner product between two time series $a(t)$ and $b(t)$ as one is shifted by an amount t_c with respect to the other,

$$z[a, b](t_c) \coloneqq 4 \int_{f_{\min}}^{f_{\max}} df \frac{\tilde{a}(f)\tilde{b}^*(f)}{S_n(f)} e^{-2\pi i f t_c}. \tag{4}$$

In the above, $\tilde{a}(f)$ is the Fourier transform of $a(t)$ and $S_n(f)$ is the detector noise power spectral density (PSD). The limits of integration are in general specified by the bandwidth over which an analysis is being conducted. In terms of $z(t_c)$ the inner products in (3) are succinctly expressed as

$$(h(\vec{\theta})|h(\vec{\theta})) = \Re\mathcal{A}(\vec{\theta}_L) z[h_0(\mathcal{M}, \eta), h_0(\mathcal{M}, \eta)](0), \tag{5}$$

$$(d|h(\vec{\theta})) = \Re\mathcal{B}(\vec{\theta}_L) z[d, h_0(\mathcal{M}, \eta)](t_c), \tag{6}$$

and $\mathcal{A}(\vec{\theta}_L)$ and $\mathcal{B}(\vec{\theta}_L)$ are known projections which contain the $\vec{\theta}_L$ dependence in the likelihood function [12].

We have previously interpolated template waveforms over the mass parameters [14,17], and here we show that the same technique can be applied to interpolating the time series $z[d, h_0](t_c)$.[1] The interpolation of $z[d, h_0](t_c)$ is based on the singular value decomposition (SVD) of a set of (discretely sampled) time series distributed on a two-dimensional grid. In this case the two-dimensional grid spans $\mathcal{M}$ and η and the time parameter is t_c. We use the notation $\vec{z}[d, h_0]$ to describe the discretely sampled $z[d, h_0](t_c)$. Recall that the SVD of the discretely sampled time series $\vec{z}[d, h_0]$ allows it to be written as a linear superposition of orthonormal basis vectors $\vec{u}_\mu$ and projection coefficients M_μ [19]:

[1] Mitra *et al.* [18] considered interpolating the signal-to-noise ratio as a scalar quantity; interpolating z as a function of time allows us to further reduce computational costs by exploiting correlations along the t_c direction.

$$\vec{z}[d, h_0(\mathcal{M}, \eta)] = \sum_\mu M_\mu(\mathcal{M}, \eta)\vec{u}_\mu. \tag{7}$$

The coefficients M_μ can be interpolated over $\mathcal{M}$ and η and we follow the method in [17] which uses Chebyshev polynomials of the first kind.

Interpolation of $z[h_0, h_0](0)$ over $\mathcal{M}$ and η is simple as it is scalar valued and we again use Chebyshev polynomials of the first kind. Below we provide an example of the interpolation technique outlined here.

III. LIKELIHOOD INTERPOLATION: EXAMPLES

We compare interpolated likelihood functions to those generated by direct evaluation of waveforms and inner products. We consider two test cases: (i) the coalescence of binary black holes and (ii) the coalescence of binary neutron stars.

We generate a discretely sampled, simulated data set $\vec{d}$ for a single interferometer consisting of Gaussian and stationary noise $\vec{n}$ and a gravitational-wave signal $\vec{h}$. The data set is 32 s in duration and has a sampling rate in the time domain of 4096 Hz. For binary black holes we model the gravitational-wave signal $\vec{h}$ using the effective one-body approach calibrated to numerical relativity simulations (EOBNR) [20]. Such a gravitational-wave signal describes the full inspiral, merger, and ringdown phases of coalescence. For binary neutron stars we model the gravitational waveform using a post-Newtonian (PN) model computed to 3.5 PN order in phase [21], which describes the inspiral phase of the coalescence only. We use an implementation of EOBNR and post-Newtonian waveforms from the LSC Algorithms Library (LAL) [22] corresponding to the approximants EOBNRv2 and TaylorT4, respectively.

Generating the interpolant of the likelihood function requires the following stages: (i) patch the mass space into smaller domains, (ii) generate a set of waveforms over a dense grid in each patch, (iii) filter the data with the template waveforms to compute the likelihoods, (iv) pack the likelihoods into a matrix and perform the SVD, and (v) build the interpolant in each patch. Only after these stages have been completed can the interpolated likelihood function be sampled.

We first construct a discrete, uniform grid of template waveforms in $\mathcal{M} - \eta$ parameter space. We will use a small region around the parameters of the signal, as $\mathcal{M}$ and η are typically constrained to $\lesssim 1\%$ and $\lesssim 10\%$, respectively, depending on the signal parameters and signal-to-noise ratio (SNR) [14,23]. The region in $\mathcal{M} - \eta$ where the posterior has significant support can be found quickly during the burn-in phase of the MCMC, which requires a small fraction of the total number of samples necessary to evaluate the posterior probability distribution function.

We use the Chebyshev interpolation described in [17] to interpolate $z[h_0, h_0](0)$ for waveforms across the grid. To interpolate $\vec{z}[d, h_0]$ we first find the basis vectors $\vec{u}_\mu$ by constructing a matrix from the set of $\{\vec{z}[d, h_0]\}$, the columns of which correspond to a unique $\vec{z}[d, h_0]$ on the grid of waveforms, which we factor using the SVD. After performing the SVD, we truncate the number of basis vectors such that on average the norm of each $\vec{z}$ is conserved to one part in 10^5 [17]. This can significantly reduce the number of basis vectors. We then apply the Chebyshev interpolation [17] to interpolate projection coefficients across the $\mathcal{M} - \eta$ grid.

A. Example 1: High-mass binary black holes

The signal is parametrized by $\vec{\theta}^s = (\mathcal{M} = 15.01 M_\odot, \eta = 0.205, D = 100\,\mathrm{Mpc}, \iota = 0, \psi = 0, \alpha = 0, \delta = 0, t_c = 0.1\,\mathrm{s}, \phi_c = 0)$. We use a noise PSD typical of initial LIGO [1]. The signal has an SNR of ≈ 15. In order to interpolate the likelihood function across $\mathcal{M}$, η and t_c, we work within a small region of $\mathcal{M} - \eta$ space whose boundaries are given by $14.6 M_\odot \leq \mathcal{M} \leq 15.6 M_\odot$ and $0.14 \leq \eta \leq 0.25$. We have restricted the prior ranges for ease of implementation, which will suffice to demonstrate the efficacy of our technique. In practice, we can quickly narrow in on the approximate parameters of interest. Our prior range corresponds to a $\sim 3\sigma$ range about the signal value (assuming a statistical measurement uncertainty on $\mathcal{M}$ and η of 1% and 10%, respectively), so we are analyzing the full region where the posterior is significant. Note that we cannot go above $\eta = 0.25$ in the η interval. We further restrict our range in t_c to be in a ± 0.2 s window about the trigger time, which is a common time prior in Bayesian parameter estimation [11].

In Fig. 1 we compare a likelihood function generated via direct evaluation of inner products, to one which we have generated via SVD interpolation. We find that we are able to reconstruct the log likelihood function by interpolation to within a fractional percentage error of at most 0.025%. While we have only plotted an interpolated likelihood function at the signal values of $\mathcal{M}$ and η, the errors quoted here are typical across the mass range we have considered. Meanwhile, for this waveform model and parameters, computing the likelihood via the interpolation procedure is around two orders of magnitude faster than generating a template waveform and directly evaluating the inner products in (3).

B. Example 2: Binary neutron stars

The signal is parametrized by $\vec{\theta}^s = (\mathcal{M} = 1.217 M_\odot, \eta = 0.2497, D = 20\,\mathrm{Mpc}, \iota = 0, \psi = 0, \alpha = 0, \delta = 0, t_c = 0.1\,\mathrm{s}, \phi_c = 0)$. We again use a noise PSD typical of initial LIGO [1], and the signal has SNR ≈ 15. We interpolate the likelihood function over a small region of $\mathcal{M} - \eta$ space whose boundaries are given by $1.199 M_\odot \leq \mathcal{M} \leq 1.235 M_\odot$ and $0.21 \leq \eta \leq 0.25$. Assuming a statistical measurement uncertainty of 0.5% on $\mathcal{M}$ and 5% on η, these parameter ranges correspond to a $\sim 3\sigma$ range about the

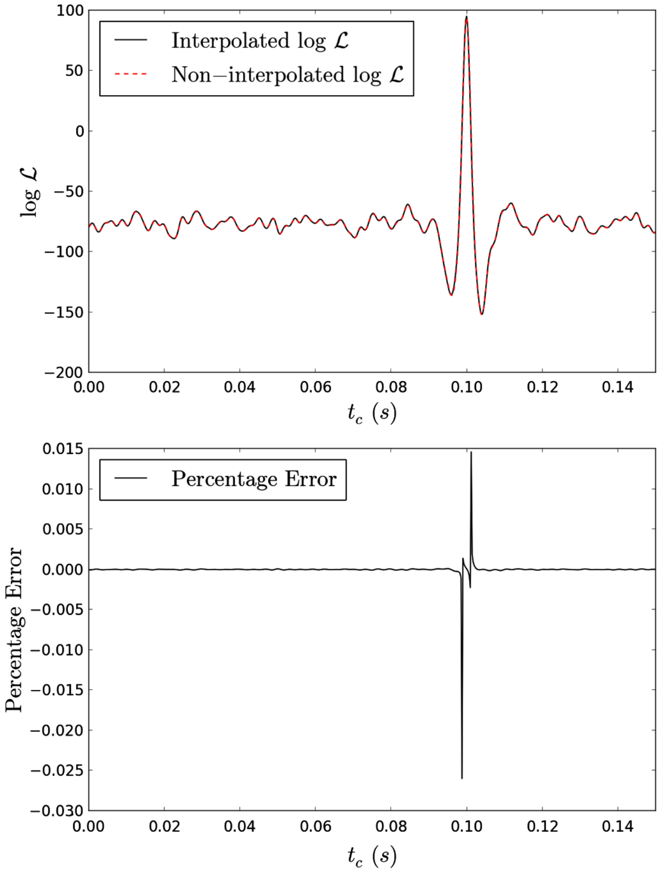

FIG. 1 (color online). Interpolated and noninterpolated log likelihoods (top), and percentage error (bottom) for a data set containing a gravitational-wave signal from the coalescence of binary black holes.

signal value. We restrict our range in t_c to be in a ± 0.2 s window about the trigger time.

Again, we find that we are able to reconstruct the log likelihood function to within a fractional percentage error of at most 0.025%. For the binary neutron star case, we find that computing the likelihood via interpolation is around three orders of magnitude faster than direct evaluation. This difference is larger than for the higher-mass binary black hole case because the waveform duration is significantly longer for binary neutron stars, whereas the cost of computing interpolated likelihoods remains fixed.

Below we discuss practical issues pertaining to incorporating interpolated likelihoods into real gravitational-wave parameter-estimation pipelines.

IV. PRACTICAL CONSIDERATIONS

For our interpolation technique to be viable for real data analyses, the total computational time of first constructing the interpolated likelihood function, and then sequentially sampling the interpolated likelihood function, must be less than the time for sequentially sampling the likelihood function directly.

To sample the interpolated likelihood function there is an additional upfront cost of constructing the interpolant of the likelihood function. This cost will depend on the region of the parameter space over which the likelihood function needs to be interpolated and template waveforms must be computed. However, building the interpolant is highly parallelizable and computing it over an extended region of parameter space could be split into multiple independent subsets. This could greatly reduce the wall time of computing the interpolant. We have noted that one can restrict the range in parameter space over which the interpolant is built by using an MCMC to sparsely explore the parameter space in regions of high posterior probability. In practice, the number of samples for this "burn-in" is often $\sim\mathcal{O}(10^4)$ [14], and the likelihood has significant support in a relatively small patch in parameter space. The likelihoods computed during the burn-in evaluation could thus be stored for future interpolation.

V. DISCUSSION AND CONCLUSION

We have demonstrated a method to sample the CBC likelihood function via interpolation, with improvements of two to three orders of magnitude in efficiency. Our method utilizes a SVD of the likelihood function on a three-dimensional subspace of parameters: the chirp mass $\mathcal{M}$, symmetric mass-ratio η, and time at coalescence t_c. The SVD factors the likelihood function into a set of scalar coefficients that describe a surface in $\mathcal{M}$ and η and a set of orthonormal basis vectors that describe how the surface is translated along t_c. The projection coefficients can be interpolated on the $\mathcal{M}-\eta$ plane and then trivially scaled by elements of the basis vectors to generate the likelihood at $(\mathcal{M}, \eta, t_c)$. This provides an efficient means to interpolate in three dimensions.

We note that while we have chosen an interpolation technique based on the SVD, it is by no means unique and other interpolation techniques have been applied to gravitational-wave data analysis [e.g., [24]]. The approach in [24] uses the so called "reduced basis method" to construct a basis which has been applied to a four-dimensional parameter space in the context of gravitational-wave templates [25]. It also utilized the "empirical interpolation method" [26], a robust interpolation technique in high-dimensional spaces [27]. A combination of these techniques may be useful for extending our technique to higher-dimensional spaces, which will be necessary to describe component spins. Higher-dimensional models could also be used to add parameters describing nonstationary or non-Gaussian noise [e.g., [28]].

Likelihood interpolation appears to be more robust than waveform interpolation [14], and so utilizing interpolated likelihood functions may also be a stepping stone to tackling the more difficult issue of rapidly estimating the parameters of binaries with spinning components.

ACKNOWLEDGMENTS

We thank Kipp Cannon and Drew Keppel for many helpful discussions. R. J. E. S. acknowledges support from a Perimeter Institute visiting graduate fellowship. Research at Perimeter Institute is supported through Industry Canada and by the Province of Ontario through the Ministry of Research and Innovation. The authors also gratefully acknowledge the support of the United States National Science Foundation for the construction and operation of the LIGO Laboratory under Cooperative Agreement No. NSF-PHY-0757058.

[1] D. Sigg and the LIGO Scientific Collaboration, Classical Quantum Gravity **25**, 114041 (2008).
[2] F. Acernese *et al.*, Classical Quantum Gravity **25**, 114045 (2008).
[3] G. M. Harry and the LIGO Scientific Collaboration, Classical Quantum Gravity **27**, 084006 (2010).
[4] B. Iyer, T. Souradeep, C. Unnikrishnan, S. Dhurandhar, S. Raja, A. Kumar, and A. S. Sengupta, Ligo-India Tech. Rep., https://dcc.ligo.org/cgi-bin/DocDB/ShowDocument?docid=75988.
[5] K. Somiya, Classical Quantum Gravity **29**, 124007 (2012).
[6] S. Hild, Classical Quantum Gravity **29**, 124006 (2012).
[7] J. Abadie and The LIGO Scientific Collaboration, Classical Quantum Gravity **27**, 173001 (2010).
[8] L. Kreidberg, C. D. Bailyn, W. M. Farr, and V. Kalogera, Astrophys. J. **757**, 36 (2012).
[9] E. Nakar, Phys. Rep. **442**, 166 (2007).
[10] I. Mandel and R. O'Shaughnessy, Classical Quantum Gravity **27**, 114007 (2010).
[11] J. Veitch and A. Vecchio, Phys. Rev. D **81**, 062003 (2010).
[12] V. Raymond, M. V. van der Sluys, I. Mandel, V. Kalogera, C. Rover, and N. Christensen, Classical Quantum Gravity **26**, 114007 (2009).
[13] M. van der Sluys, I. Mandel, V. Raymond, V. Kalogera, C. Roever, and N. Christensen, Classical Quantum Gravity **26**, 204010 (2009).
[14] R. Smith, K. Cannon, C. Hanna, D. Keppel, and I. Mandel, Phys. Rev. D**87**, 122002 (2013).
[15] J. Skilling, AIP Conf. Ser. Vol. **735**, 395 (2004).
[16] B. J. Owen, Phys. Rev. D **53**, 6749 (1996).
[17] K. Cannon, C. Hanna, and D. Keppel, Phys. Rev. D **85**, 081504 (2012).
[18] S. Mitra, S. V. Dhurandhar, and L. S. Finn, Phys. Rev. D **72**, 102001 (2005).
[19] K. Cannon, A. Chapman, C. Hanna, D. Keppel, A. C. Searle, and A. J. Weinstein, Phys. Rev. D **82**, 044025 (2010).
[20] Y. Pan, A. Buonanno, M. Boyle, L. T. Buchman, L. E. Kidder, H. P. Pfeiffer, and M. A. Scheel, Phys. Rev. D **84**, 124052 (2011).
[21] A. Buonanno, B. R. Iyer, E. Ochsner, Y. Pan, and B. S. Sathyaprakash, Phys. Rev. D **80**, 084043 (2009).
[22] The LIGO Scientific Collaboration, Lsc Algorithm Library (LAL), https://www.lsc-group.phys.uwm.edu/daswg/projects/lalsuite.html.
[23] The LIGO Scientific Collaboration and The Virgo Collaboration, Phys. Rev. D **88**, 062001 (2013).
[24] P. Canizares, S. E. Field, J. R. Gair, and M. Tiglio, Phys. Rev. D **87**, 124005 (2013).
[25] S. E. Field, C. R. Galley, F. Herrmann, J. S. Hesthaven, E. Ochsner, and M. Tiglio, Phys. Rev. Lett. **106**, 221102 (2011).
[26] M. Barrault, Y. Maday, N. C. Nguyen, and A. T. Patera, Comptes Rendus Mathematique **339**, 667 (2004).
[27] S. Chaturantabut and D. Sorensen, SIAM J. Sci. Comput. **32**, 2737 (2010).
[28] T. B. Littenberg, M. Coughlin, B. Farr, and W. M. Farr, Phys. Rev. D **88**, 084044 (2013).