Accelerated Gravitational Wave Parameter Estimation with Reduced Order Modeling

Priscilla Canizares,[1,6] Scott E. Field,[2,5] Jonathan Gair,[1] Vivien Raymond,[3] Rory Smith,[3] and Manuel Tiglio[4]
[1]*Institute of Astronomy, Madingley Road, Cambridge CB3 0HA, United Kingdom*
[2]*Department of Physics, Joint Space Sciences Institute, Maryland Center for Fundamental Physics, University of Maryland, College Park, Maryland 20742, USA*
[3]*LIGO, California Institute of Technology, Pasadena, California 91125, USA*
[4]*Center for Astrophysics and Space Sciences, Center for Computational Mathematics, San Diego Supercomputer Center, University of California San Diego, 9500 Gilman Drive, La Jolla, California 92093-0424, USA*
[5]*Center for Radiophysics and Space Research, Cornell University, Ithaca, New York 14853, USA*
[6]*Department of Applied Mathematics and Theoretical Physics, Wilberforce Road, Cambridge CB3 0WA, United Kingdom*
(Received 27 April 2014; published 20 February 2015)

Inferring the astrophysical parameters of coalescing compact binaries is a key science goal of the upcoming advanced LIGO-Virgo gravitational-wave detector network and, more generally, gravitational-wave astronomy. However, current approaches to parameter estimation for these detectors require computationally expensive algorithms. Therefore, there is a pressing need for new, fast, and accurate Bayesian inference techniques. In this Letter, we demonstrate that a reduced order modeling approach enables rapid parameter estimation to be performed. By implementing a reduced order quadrature scheme within the LIGO Algorithm Library, we show that Bayesian inference on the 9-dimensional parameter space of nonspinning binary neutron star inspirals can be sped up by a factor of ~30 for the early advanced detectors' configurations (with sensitivities down to around 40 Hz) and ~70 for sensitivities down to around 20 Hz. This speedup will increase to about 150 as the detectors improve their low-frequency limit to 10 Hz, reducing to hours analyses which could otherwise take months to complete. Although these results focus on interferometric gravitational wave detectors, the techniques are broadly applicable to any experiment where fast Bayesian analysis is desirable.

DOI: 10.1103/PhysRevLett.114.071104 PACS numbers: 04.80.Nn, 04.30.-w, 95.55.Ym

Introduction.—Advanced LIGO (aLIGO) [1] and advanced Virgo (AdV) [2] are expected to yield the first direct detections of gravitational waves (GWs) from astrophysical sources in the next few years. Compact binary coalescences (CBCs) are the most promising GW sources, with expected detection rates between a few and tens per year [3]. Effective parameter estimation for CBCs has been demonstrated [4–6], but approaches to date carry high computational costs for the cases of interest, even when using efficient algorithms such as Markov chain Monte Carlo (MCMC) or nested sampling [7]. For the advanced detectors, which will start taking data within a year or two, current approaches will lead to months or years of computational wall (clock) time for the analysis of each detected signal. Therefore, given the expected detection rates, new techniques which can estimate the astrophysical source parameters in short time scales are highly desirable. Such techniques are also important for large-scale mock data challenges.

In parameter estimation studies, the posterior probability density function (PDF) of a set of parameters, $\vec{\theta}$, is computed from a GW model, $h(\vec{\theta})$, assumed to describe the detector's signal d. The PDF is related to the likelihood function, $\mathcal{L}(d|\vec{\theta})$, and the prior probability on the model parameters, $\mathcal{P}(\vec{\theta})$, via Bayes' theorem, $p(\vec{\theta}|d) \propto \mathcal{P}(\vec{\theta})\mathcal{L}(d|\vec{\theta})$.

Assuming that the detector data d contains the source's signal $h(\vec{\theta}_{\rm true})$ and stationary Gaussian noise n, the likelihood function is given by

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

where $d = h(\vec{\theta}_{\rm true}) + n$ and $(a|b)$ is a weighted inner product for discretely sampled noisy data

$$(d|h(\vec{\theta})) = 4\Re\Delta f \sum_{k=1}^{L} \frac{\tilde{d}^*(f_k)\tilde{h}(\vec{\theta}; f_k)}{S_n(f_k)}, \quad (2)$$

where $\tilde{d}(f_k)$ and $\tilde{h}(\vec{\theta}; f_k)$ are the discrete Fourier transforms at frequencies $\{f_k\}_{k=1}^{L}$ (with units of Hz), * denotes complex conjugation, and the power spectral density (PSD) $S_n(f_k)$ characterizes the detector's noise.

For a given observation time $T = 1/\Delta f$ and detection frequency window $(f_{\rm high} - f_{\rm low})$, there are

$$L = {\rm int}([f_{\rm high} - f_{\rm low}]T) \quad (3)$$

sampling points in the sum (2). When L is large, as in the cases of interest for this Letter, there are two major bottlenecks: (i) evaluation of the model at each f_k and (ii) assembly of the likelihood (1).

In general, smoothly parametrized models are amenable to dimensional reduction which, in turn, provides computationally efficient representations. The specific application of dimensional reduction we consider in this Letter tackles the two aforementioned bottlenecks by permitting the inner product (2) to be computed with significantly fewer terms. In summary, if a reduced set of $N < L$ basis can be found which accurately spans the model space, it is possible to replace the inner product (2) with a reduced order quadrature (ROQ) rule (5) containing only N terms, reducing the overall parameter estimation analysis cost by a factor of L/N, provided the waveforms can be directly evaluated. For other models, in particular those described by partial or ordinary differential equations, direct evaluation may be accomplished using surrogates [8,9].

In this Letter, we demonstrate a ROQ accelerated GW parameter estimation study. While the approach is applicable to any GW model, here we focus on binary neutron star (BNS) inspirals, as these are expected to have the highest detection rates with the lowest uncertainty [3]. We show, both through operation counts and an implementation in the LIGO Algorithm Library (LAL) pipeline [10], that ROQs provide a factor of ∼30 speedup for the early advanced detectors' configuration [11] and ∼70 as the detectors' low-frequency sensitivity reaches 20 Hz. This speedup will rise to ∼150 as the sensitivity band is lowered to a target of 10 Hz, allowing for significant reduction of the computational cost of Bayesian parameter estimation analyses on BNS sources.

Compressed likelihood evaluations.—Compared to previous work on which this Letter is based [12,13], parameter estimation of gravitational waves from binary neutron stars carries a number of challenges unique to large data sets which contain long gravitational waveforms with many in-band wave cycles. We briefly summarize the construction of ROQs while focusing on technical but essential solutions to these challenges.

Reduced order quadratures can be used for fast parameter estimation whenever the waveform model is amenable to dimensional reduction, through three steps. The first two are carried out offline, while the final, data-dependent step is performed at the beginning of the parameter estimation analysis. (1) Construct a reduced basis, i.e., a set of N elements whose span reproduces the GW model within a specified precision. (2) Construct an empirical interpolant by requiring it to exactly match any template at N carefully chosen frequency subsamples $\{F_k\}_{k=1}^N$ [14–16]. (3) The empirical interpolant is used to replace, without loss of accuracy, inner product evaluations (2) by ROQ compressed ones (5).

Step 1: The reduced basis set only needs to be built over the space of intrinsic parameters for the waveform family. Furthermore, if the basis is generated using a PSD of unity, the representation of the waveform family can be used with any PSD whenever the weights are built as in Eq. (6).

Basis generation in this Letter proceeds in two stages. A greedy algorithm first identifies a preliminary basis suitable for any value of Δf [17]. Next, this preliminary basis is evaluated at L equally spaced frequency samples appropriate for the detector. The elements of this "resampled basis" are neither orthogonal nor linearly independent, and so a second similar dimensional reduction is necessary. During these steps, appropriately conditioned numerical algorithms [18] are used to avoid poor conditioning, which, for large values of L, would otherwise lead to bases with no accuracy whatsoever.

Step 2: Given an N-size basis, it is possible to uniquely and accurately reconstruct any waveform from only N evaluations $\{\tilde{h}(\vec{\theta}; F_k)\}_{k=1}^N$. The special frequencies $\{F_k\}_{k=1}^N$, selected from the full set $\{f_i\}_{i=1}^L$, can be found from Algorithm 5 of Ref. [12] without modification. This step provides a near-optimal compression strategy in frequency which is complimentary to the parameter one of step 1. The model's empirical interpolant, valid for all parameters, can be written as (cf. Eq. (19) of Ref. [8])

$$\tilde{h}(\vec{\theta}; f_i) \approx e^{-2\pi i t_c f_i} \sum_{j=1}^{N} B_j(f_i) \tilde{h}(\vec{\theta}, t_c = 0; F_j), \qquad (4)$$

a sum over the basis set $\{B_j\}_{j=1}^N$ and where, for the sake of the discussion below, we have temporarily isolated the coalescence time t_c from the other parameters.

Step 3: All extrinsic parameters, except the coalescence time t_c, do not affect the frequency evolution of the binary and simply scale the inner product (2), thereby sharing the same ROQs. The coalescence time, however, requires special treatment. Substituting Eq. (4) into Eq. (2),

$$(d|h(\vec{\theta}, t_c)) = \sum_{k=1}^{N} \omega_k(t_c) \tilde{h}(\vec{\theta}, t_c = 0; F_k), \qquad (5)$$

with the ROQ weights given by

$$\omega_k(t_c) = 4\Re \Delta f \sum_{i=1}^{L} \frac{\tilde{d}^*(f_i) B_k(f_i)}{S_n(f_i)} e^{-2\pi i t_c f_i}. \qquad (6)$$

Our approach for the dependence of Eq. (6) on t_c is through domain decomposition: an estimate for the time window W centered around the coalescence time $t_{\rm trigger}$ is given by the GW search pipeline. This suggests a prior interval $[t_{\rm trigger} - W, t_{\rm trigger} + W]$ be used for t_c. The prior interval is then split into n_c equal subintervals of size Δt_c. The number of subintervals is chosen so that the discretization error is below the measurement uncertainty on the coalescence time. Finally, on each subinterval a unique set of ROQ weights is constructed.

Since step (3) is currently implemented in the LAL pipeline, we summarize it in Algorithm (1). The offline steps (1) and (2) are carried out independently. Our approach guarantees, though, that those steps need to be carried out only once for each waveform model.

To quickly compute the likelihood, we also need an inexpensive rule for $(h(\vec{\theta})|h(\vec{\theta}))$, whose evaluation no longer depends on the data stream or coalescence time. Consequently, such expressions are typically simple. For example, the norm of the restricted TaylorF2 gravitational waveform model considered below is exactly computable [19]. In the general case, building a basis for $\{h^2\}$ and an associated ROQ provides fast norm evaluations. The TaylorF2 gravitational waveform model's norm is computable as a 1-term ROQ rule, for example.

By design, weight generation is computed in the startup stage for each detection-triggered data set, requiring N full inner product (2) evaluations for each t_c interval. This cost is negligible, while each likelihood is subsequently calculated millions of times, leading to significant speedups in parameter estimation studies. The latter scales as the fractional reduction L/N of the number of terms in the quadrature rules (2) and (5).

Algorithm 1 Computing the ROQ integration weights

1: **Input:** d, S_n, $\{B_j\}_{j=1}^{N}$, Δf, $t_{\rm trigger}$, W, Δt_c
2: Set $n_c = \mathrm{int}((2W)/\Delta t_c) + 1$
3: **for** $j = 1 \to n_c$ **do**
4: $\quad T_j = t_{\rm trigger} - W + (j-1)\Delta t_c$
5: $\quad$ **for** $k = 1 \to N$ **do**
6: $\quad\quad$ Compute $\omega_k(T_j)$ via Eq. (6)
7: $\quad$ **end for**
8: **end for**
9: **Output:** $\{T_j\}_{j=1}^{n_c}$, $\{\{\omega_k(T_j)\}_{k=1}^{N}\}_{j=1}^{n_c}$

Parameter estimation acceleration for binary neutron star signals.—The majority of a binary neutron star's GW signal will be in the inspiral regime [20], which can be described by the closed-form TaylorF2 approximation [21]. While TaylorF2 does not incorporate spins or the merger-ringdown phases of the binary's evolution, these should not be important for BNS parameter estimation and can therefore be neglected [22]. Even for this simple-to-evaluate waveform family, inference on a single data set requires significant computational wall time with standard parameter estimation methods [6]. We now report on the anticipated speedup L/N achieved by ROQ compressed likelihood evaluations. First, we compute the observation time T required to contain a typical BNS signal. Next, we find the number of reduced basis elements N needed to represent this model for any pair of BNS masses. In our studies, we fix $f_{\rm high}$ to 1024 Hz while $f_{\rm low}$ varies between 10 Hz and 40 Hz.

The time taken for a BNS system with an initial GW frequency of $f_{\rm low}$ to inspiral to 1024 Hz,

$$T_{\rm BNS} = \left[6.32 + 2.07 \times \frac{10^6}{(f_{\rm low}/{\rm Hz})^3 + 5.86(f_{\rm low}/{\rm Hz})^2}\right] \text{ s}, \tag{7}$$

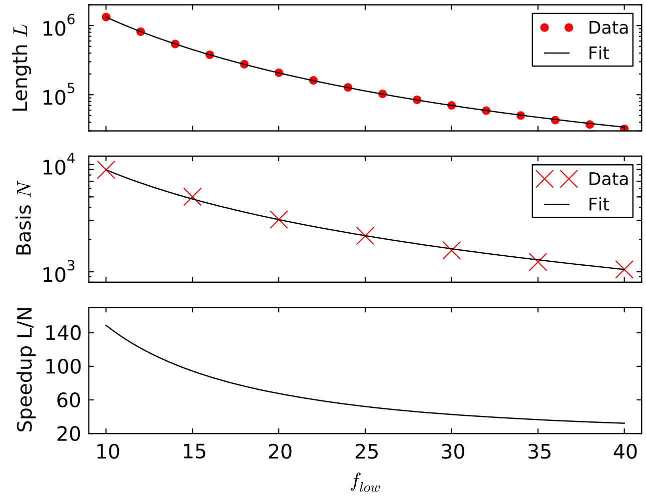

FIG. 1 (color online). Top: length L (red dots) of a typical binary neutron star inspiral waveform, with the solid black curve connecting this data implied by the fit (7). Middle: number of reduced basis waveforms (red crosses), with the solid black curve given by the fit (8). Bottom: speedup implied by operation counts, as given by Eq. (9).

is empirically found by generating a $(1+1)M_\odot$ waveform (directly given in the frequency domain) and Fourier transforming to the time domain where we measure the duration up to when the waveform's evolution terminates. Equation (7) and subsequent fits were found using a genetic algorithm-based symbolic regression software, EUREQA [23,24]. The length L, as implied by Eq. (3), is plotted in the top panel of Fig. 1.

As discussed, each basis only needs to be constructed over the space of intrinsic parameters—in this case the two-dimensional space of component masses in the range $[1,4]M_\odot$. This range is wider than expected for neutron stars, but ensures that the resulting PDFs do not have sharp cutoffs [25]. The number of reduced basis required to represent the TaylorF2 model within this range with a representation error around double precision ($\sim 10^{-14}$) can be fit by

$$N_{\rm BNS} = 3.12 \times 10^5 (f_{\rm low}/{\rm Hz})^{-1.543}, \tag{8}$$

and is depicted in the middle panel of Fig. 1. We have found that increasing the high-frequency cutoff to 4096 Hz only adds a handful of basis elements, while L changes by a factor of 4, thus indicating that the speedup for an inspiral-merger-ringdown model might be higher, especially given that not many empirical interpolation nodes are needed for the merger and ringdown regimes [8].

Recalling Eq. (3), the speedup from standard to ROQ-compressed likelihood evaluations is given by

$$\frac{L}{N} \approx (1024\ {\rm Hz} - f_{\rm low}) \frac{T_{\rm BNS}}{N_{\rm BNS}}, \tag{9}$$

with $T_{\rm BNS}$ and $N_{\rm BNS}$ given by Eqs. (7) and (8). This speedup is shown in Fig. 1 (bottom), with a reduction in computational cost and time of $\sim$30 for the initial detectors (with a cutoff of $f_{\rm low} = 40$ Hz) and $\sim$150 once the advanced detectors reach $f_{\rm low} \sim 10$ Hz.

Implementation and numerical studies.—We have implemented compressed likelihood evaluations and Algorithm (1) in the LAL parameter estimation pipeline, known as LALInference [10,26], naming the resulting variation LALInference_ROQ.

Next, we compare MCMC parameter estimation results using the standard version of LALInference to ROQ accelerated studies using LALInference_ROQ. We consider gravitational waveforms emitted from binary neutron star systems with TaylorF2 as the waveform model. We inject synthetic signals embedded in simulated Gaussian noise into the LAL pipeline, for settings anticipating the initial configuration of aLIGO, which should be online within the next two years, using the zero detuned high power PSD [27] and the initial configuration of AdV, using Eq. (6) of [28], with in both cases $f_{\rm low} = 40$ Hz.

We take $W = 0.1$ s as the typical time window for the coalescence time t_c of a binary neutron star signal centered around the trigger time [5,6]. Following the procedure discussed above, LALInference_ROQ discretizes this prior into $n_c = 2000$ subintervals, each of size $\Delta t_c = 10^{-5}$ s, for which it constructs a unique set of ROQ weights on each subinterval. A width of 10^{-5} s ensures that this discretization error is below the measurement uncertainty on the coalescence time, which is typically $\sim 10^{-3}$ s [26].

We found that, as expected, the ROQ and standard likelihood approaches produce statistically indistinguishable results for posterior probability density functions over the full nine-dimensional parameter space. Figure 2 and Table I describe results for the intrinsic mass parameters obtained in one particular MCMC simulation; other simulations were qualitatively similar. It is also useful to quantify the fractional difference in the 9D likelihood function computed using ROQs and the standard approach. We have observed this fractional error to be

$$\Delta \log \mathcal{L} = 1 - \left(\frac{\log \mathcal{L}}{\log \mathcal{L}_{\rm ROQ}}\right) \lesssim 10^{-6}$$

in all cases. That is, both approaches are indistinguishable for all practical purposes.

In addition to providing indistinguishable results, ROQ accelerated inference is significantly faster: the ROQ-based MCMC study with the discussed settings takes $\sim$1 h, compared to $\sim$30 h using the standard likelihood approach, in remarkable agreement with the expected savings based on operation counts. The wall time of the analysis is proportional to the total number of posterior samples of the MCMC simulation, which in this case was $\sim 10^7$. The startup stage required to build the ROQ weights has negligible cost and is completed in near real time, $\sim$30 s, which is equivalent to $\sim$0.028% of the total cost of a standard likelihood parameter estimation study.

For a lower cutoff frequency of $f_{\rm low} \sim 20$ Hz, the speedup reduction is from a couple of weeks to hours. Once the advanced detectors have achieved their target sensitivity, with $f_{\rm low} \sim 10$ Hz, the longest BNS signals will last around 2048 s in duration. Assuming a fiducial high frequency cutoff of 1024 Hz, which is approaching the upper limit of where aLIGO-AdV will be sensitive, we estimate data sets as large as $L \sim 1024\ {\rm Hz}^{-1} \times 2048\ {\rm s} \sim 10^6$. Assuming that the advanced detectors will require at least $\sim 10^7$ posterior samples, this implies runtimes upwards of

TABLE I. Chirp mass $\mathcal{M}_c$, symmetric mass ratio η, component masses m_1 and m_2, and signal-to-noise ratio (SNR) of the analysis from Fig. 2. Median value and 90% credible intervals are provided for both the standard likelihood (second line) and the ROQ compressed likelihood (third line). The SNR is empirically measured from likelihood$_{\rm max} \approx {\rm SNR}^2/2$. The differences between the two methods are dominated by statistics from computing intervals with a finite number of samples. In our analysis, the masses are subject to the constraint $m_1 < m_2$, leading to the true values (where $m_1 = m_2$) being at the edge of the confidence interval.

	$\mathcal{M}_c(M_\odot)$	η	$m_1(M_\odot)$	$m_2(M_\odot)$	SNR
Injection	1.2188	0.25	1.4	1.4	11.4
Standard	$1.2188^{1.2189}_{1.2184}$	$0.249^{0.250}_{0.243}$	$1.52^{1.66}_{1.41}$	$1.30^{1.39}_{1.18}$	12.9
ROQ	$1.2188^{1.2189}_{1.2184}$	$0.249^{0.250}_{0.243}$	$1.52^{1.66}_{1.41}$	$1.30^{1.39}_{1.19}$	12.9

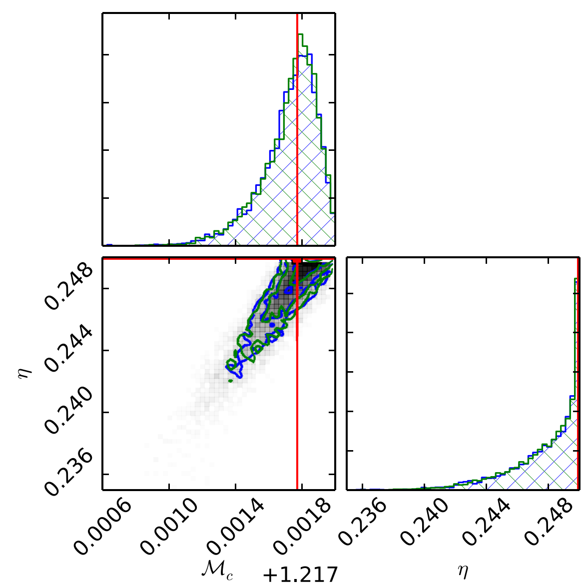

FIG. 2 (color online). Probability density function for the chirp mass $\mathcal{M}_c$ and symmetric mass ratio η of a simulated event in LIGO-Virgo data. In green as obtained in $\sim$30 h by the standard likelihood, and in blue as obtained in 1 h with the ROQ. The injection values are in red, and are listed in Table I. The overlap region of the sets of PDFs is the hatched region. Plotting based on [29].

∼100 days and 1 petabyte worth of model evaluations using the standard approach. The results of this Letter indicate that a ROQ approach will reduce this to less than a day. Remarkably, this approach when applied to the advanced detectors operating at design sensitivity will be faster than even the standard likelihood one used for the initial detectors. Additionally, with parallelization of the sum in each likelihood evaluation essentially real-time full Bayesian analysis could be achieved. More details can be found in Ref. [30].

Outlook.—For detectors operating at $f_{\rm low} = 40$ Hz, around three weeks of real (wall) time are needed to perform a precessing-spin parameter estimation study on a single data stream with standard likelihood evaluations [26]. With a cutoff of $f_{\rm low} = 10$ Hz, these analyses could take months to years, so techniques for accelerated inference on these models, such as the one presented in this Letter, are expected to play a central role in gravitational-wave astronomy for extracting the full science potential of the upcoming advanced gravitational-wave detectors.

In this Letter, we have addressed the issue of fast likelihood evaluation for nonspinning binary neutron star inspirals. Results of previous work indicate that significant computational savings are to be expected for other waveform models. For example, in Ref. [17] it was found that the number of reduced basis waveforms barely changes as spins are included—at least in the nonprecessing case—and, by construction, neither does the number of ROQ evaluations. While waveform evaluation is the dominant cost for models which incorporate precession, recent results [31] have shown that ultracompact bases can also be constructed for fully precessing systems. This might provide a means for constructing fast-to-evaluate surrogate models of precessing binary inspirals [8,9].

We thank Collin Capano, Peter Diener, Will Farr, Chad Galley, Larry Price, Jorge Pullin, Leo Singer, Alan Weinstein and the LVC CBC and parameter estimation groups for many useful discussions and encouragement throughout this project. P. C.'s work was supported by a Marie Curie Intra-European Fellowship within the 7th European Community Framework Programme PIEF-GA-2011-299190, FP7-PEOPLE-2011-CIG Grant No. 293412 and STFC Consolidator Grant No. ST/L000636/1. S. E. F. thanks the Institute of Astronomy at Cambridge, UK, where part of this work was done, for hospitality and financial support. J. G.'s work was supported by the Royal Society and V. R. by the LIGO Laboratory and the California Institute of Technology (Caltech). This work was supported in part by NSF Grant No. PHY-1208861 and No. PHY-1316424 to the University of Maryland (UMD), NSF Grant No. PHY-1500818 to the University of California at San Diego, NSF Grants No. PHY-1306125 and No. AST-1333129 to Cornell University, and the Sherman Fairchild Foundation. The authors also gratefully acknowledge the support of the U.S. National Science Foundation for the construction and operation of the LIGO Laboratory under cooperative agreement NSF-PHY-0757058. This paper carries LIGO Document Number LIGO-P1400038. Some of the computations were carried out at the Center for Scientific Computation and Mathematical Modeling cluster at UMD and the LIGO Laboratory computer cluster at Caltech. Portions of this research were conducted with high performance computing resources provided by Louisiana State University [32].

[1] Gregory M. Harry and LIGO Scientific Collaboration, Advanced LIGO: The next generation of gravitational wave detectors, Classical Quantum Gravity **27**, 084006 (2010).

[2] Virgo Collaboration, Advanced Virgo Baseline Design, Technical Report No. VIR027A0 (2009).

[3] J. Abadie *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Predictions for the rates of compact binary coalescences observable by ground-based gravitational-wave detectors, Classical Quantum Gravity **27**, 173001 (2010).

[4] V. Raymond, M. V. van der Sluys, I. Mandel, V. Kalogera, C. Rover, and N. Christensen, Degeneracies in sky localisation determination from a spinning coalescing binary through gravitational wave observations: A Markov-chain Monte-Carlo analysis for two detectors, Classical Quantum Gravity **26**, 114007 (2009).

[5] R. J. E. Smith, K. Cannon, C. Hanna, D. Keppel, and I. Mandel, Towards rapid parameter estimation on gravitational waves from compact binaries using interpolated waveforms, Phys. Rev. D **87**, 122002 (2013).

[6] J. Aasi *et al.* (LIGO-Virgo Scientific Collaboration), Parameter estimation for compact binary coalescence signals with the first generation gravitational-wave detector network, Phys. Rev. D **88**, 062001 (2013).

[7] J. Skilling, Nested sampling, AIP Conf. Proc. **735**, 395 (2004).

[8] S. E. Field, C. R. Galley, J. S. Hesthaven, J. Kaye, and M. Tiglio, Fast Prediction and Evaluation of Gravitational Waveforms using Surrogate Models, Phys. Rev. X **4**, 031006 (2014).

[9] Michael Pürrer, Frequency domain reduced order models for gravitational waves from aligned-spin black-hole binaries, Classical Quantum Gravity **31**, 195010 (2014).

[10] LIGO Algorithm Library, https://www.lsc-group.phys.uwm.edu/daswg/projects/lalsuite.html.

[11] J. Aasi *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Prospects for localization of gravitational wave transients by the advanced LIGO and advanced Virgo observatories, arXiv:1304.0670.

[12] H. Antil, S. E. Field, F. Herrmann, R. H. Nochetto, and M. Tiglio, Two-step greedy algorithm for reduced order quadratures, J. Sci. Comput. **57**, 604 (2013).

[13] P. Canizares, S. E. Field, J. R. Gair, and M. Tiglio, Gravitational wave parameter estimation with compressed likelihood evaluations, Phys. Rev. D **87**, 124005 (2013).

[14] M. Barrault, Y. Maday, N. C. Nguyen, and A. T. Patera, An empirical interpolation method: Application to efficient

reduced-basis discretization of partial differential equations, C. R. Acad. Sci. Paris Ser. IV **339**, 667 (2004).

[15] Y. Maday, N. C. Nguyen, A. T. Patera, and S. H. Pau, A general multipurpose interpolation procedure: The magic points, Commun. Pure Appl. Anal. **8**, 383 (2009).

[16] S. Chaturantabut and D. C. Sorensen, Nonlinear model reduction via discrete empirical interpolation, SIAM J. Sci. Comput. **32**, 2737 (2010).

[17] S. E. Field, C. R. Galley, and E. Ochsner, Towards beating the curse of dimensionality for gravitational waves using reduced basis, Phys. Rev. D **86**, 084046 (2012).

[18] W. Hoffmann, Iterative algorithms for Gram-Schmidt orthogonalization, Computing **41**, 335 (1989).

[19] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. E. Creighton, Findchirp: An algorithm for detection of gravitational waves from inspiraling compact binaries, Phys. Rev. D **85**, 122006 (2012).

[20] B. S. Sathyaprakash and B. F. Schutz, Physics, astrophysics and cosmology with gravitational waves, Living Rev. Relativity **12**, 2 (2009).

[21] A. Buonanno, B. Iyer, E. Ochsner, Y. Pan, and B. S. Sathyaprakash, Comparison of post-Newtonian templates for compact binary inspiral signals in gravitational-wave detectors, Phys. Rev. D **80**, 084043 (2009).

[22] Leo P. Singer, The first two years of electromagnetic follow-up with advanced LIGO and Virgo, Astrophys. J. **795**, 105 (2014).

[23] Eureqa, http://www.nutonian.com.

[24] M. Schmidt and H. Lipson, Distilling free-form natural laws from experimental data, Science **324**, 81 (2009).

[25] I. Mandel, C. P. L. Berry, F. Ohme, S. Fairhurst, and W. M. Farr, Parameter estimation on compact binary coalescences with abruptly terminating gravitational waveforms, Classical Quantum Gravity **31**, 155005 (2014).

[26] J. Aasi *et al.*, Parameter estimation for compact binary coalescence signals with the first generation gravitational-wavedetector network, Phys. Rev. D **88**, 062001 (2013).

[27] LIGO Scientific Collaboration, Advanced LIGO Anticipated Sensitivity Curves, Technical Report No. LIGO-T0900288-v3 (2010).

[28] A. Manzotti and A. Dietz, Prospects for early localization of gravitational-wave signals from compact binary coalescences with advanced detectors, arXiv:1202.4031.

[29] D. Foreman-Mackey, A. Price-Whelan, G. Ryan, E. Rice, M. Smith, K. Barbary, D. W. Hogg, and B. J. Brewer, triangle.py v0.1.1 (2014).

[30] LAL-ROQ, http://www.ligo.caltech.edu/~vraymond/rom/20Hz/.

[31] J. Blackman, B. Szilagyi, C. R. Galley, and M. Tiglio, Sparse Representation of Gravitational Waves from Precessing Compact Binaries, Phys. Rev. Lett. **113**, 021101 (2014).

[32] http://www.hpc.lsu.edu