Accelerated gravitational-wave parameter estimation with reduced order modeling

Priscilla Canizares,[1] Scott E. Field,[2] Jonathan Gair,[1] Vivien Raymond,[3] Rory Smith,[3] and Manuel Tiglio[2,4,5]

[1] *Institute of Astronomy, Madingley Road, Cambridge, CB30HA, United Kingdom*
[2] *Department of Physics, Joint Space Sciences Institute, Maryland Center for Fundamental Physics. University of Maryland, College Park, MD 20742, USA*
[3] *LIGO, California Institute of Technology, Pasadena, CA 91125, USA*
[4] *Center for Scientific Computation and Mathematical Modeling, University of Maryland, College Park, MD 20742, USA*
[5] *Theoretical Astrophysics, California Institute of Technology, Pasadena, CA, 91125, USA*

(Dated: April 24, 2014)

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 parameter estimation approaches for such scenarios can lead to computationally intractable problems in practice. 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 studies. By implementing a reduced order quadrature scheme within the LIGO Algorithm Library, we show that Bayesian inference on the 9-dimensional parameter space of non-spinning binary neutron star inspirals can be sped up by a factor of 30 for the early advanced detectors' configurations. This speed-up will increase to about 150 as the detectors improve their low-frequency limit to 10Hz, reducing to hours analyses which would otherwise take months to complete. Although these results focus on gravitational detectors, the techniques are broadly applicable to any experiment where fast Bayesian analysis is desirable.

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, often unrealistic 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. Given the expected detection rates, there is a clear need for new approaches which can estimate the astrophysical GW source parameters in feasible timescales.

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 the Bayes' theorem: $p(\vec{\theta}|d) \propto \mathcal{P}(\vec{\theta})\ \mathcal{L}(d|\vec{\theta})$.

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

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

where $d = h_t(\vec{\theta}_t) + \mathrm{n}$ and $(a|b)$ is a weighted inner product for discretely sampled noisy data

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

In this equation $\tilde{d}(f_k)$ and $\tilde{h}(\vec{\theta};f_k)$ are the discrete Fourier transforms at frequencies $\{f_k\}_{k=1}^{L}$, * 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 $(f_{\tt high} - f_{\tt low})$ there are

$$L = \texttt{int}\left(\left[f_{\tt high} - f_{\tt low}\right] T\right) \quad (3)$$

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

In general, smoothly parameterized models are amenable to dimensional reduction which, in turn, provides computationally efficient representations. The specific application of dimensional reduction we consider in this paper 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 [9, 10].

In this paper we demonstrate an 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 [11], that ROQs provide a factor of ~ 30 speedup for the early advanced detector's configuration [12]. This speedup will raise to ~ 150 when the sensitivity band is lowered to a target of 10Hz, allowing for param in realistic timescales.

Compressed likelihood evaluations– Compared to previous work on which this paper is based [13, 14], parameter estimation of gravitational waves from binary neutron stars carries a number of challenges unique to large datasets 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 third and last one is startup – i.e., performed at the beginning of the parameter estimation analysis – and data-dependent. (1) Construct a reduced basis, i.e. a set of N basis whose span approximates 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 $\{F_k\}_{k=1}^{N}$ [15–17]. (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 bases only needs to be built over the space of intrinsic parameters for the waveform family. Therefore, the representation of the waveform family can be used with any PSD in the noise-weighted inner product (2).

Basis generation in this paper proceeds in two stages. A greedy algorithm first identifies a preliminary basis suitable for any value of Δf [19]. Next, this preliminary basis is evaluated at L equally spaced frequency samples appropriate for the detector. These "resampled basis" are neither orthogonal nor linearly independent, and so a second similar dimensional reduction is necessary. Through these steps appropriately conditioned numerical algorithms [20] are used to avoid poor conditioning that for large values of L 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. [13] 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. [9])

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

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),

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

with the ROQ weights given by

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

Our approach for the dependence of (6) on t_c is through domain decomposition: an estimate for the time window W centered around the coalescence time $t_{\mathtt{trigger}}$ is given by the GW search pipeline. This suggests a prior interval $[t_{\mathtt{trigger}} - W, t_{\mathtt{trigger}} + W]$ 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). Offline steps (1) and (2) are independently generated. Our approach guarantees, though, that those steps need be to carried out only once for each waveform model.

Compressed norm evaluations. 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 [21].

Overall likelihood-compression. By design, weight generation is computed in the startup stage for each detection-triggered data set, requiring N full inner product (2) evaluations. 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

Input: $d, S_n, \{B_j\}_{j=1}^{N}, \Delta f, t_{\mathtt{trigger}}, W, \Delta t_c$.

Set $n_c = \mathtt{int}\left((2W)/\Delta t_c\right) + 1$
for $j = 1 \rightarrow n_c$ **do**
  $T_j = t_{\mathtt{trigger}} - W + (j-1)\,\Delta t_c$
  for $k = 1 \rightarrow N$ **do**
    Compute $\omega_k(T_j)$ via Eq. (6)
  end for
end for

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 [22], which can be described by the closed-form TaylorF2 approximation [23]. 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 (see [24]) and can therefore be neglected. Even for this simple to evaluate family of waveforms, inference on a single data set requires significant computational walltime 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 N needed to represent this model for any pair of BNS masses. In our studies we fix $f_{\mathtt{high}}$ to 1024Hz while $f_{\mathtt{low}}$ varies between 10Hz and 40Hz.

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

$$T_{\mathtt{BNS}} = \left[6.32 + 2.07 \times 10^6/(f_{\mathtt{low}}^3 + 5.86 f_{\mathtt{low}}^2)\right] \mathrm{s}\,, \qquad (7)$$

is empirically found by generating a $(1+1)\ M_\odot$ waveform (directly given in the frequency domain) and Fourier transforming to the time domain. Equation (7) and subsequent fits were found using a genetic algorithm-based symbolic regression software, `Eureqa` [25, 26]. 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 cut-offs [27]. 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

$$N_{\mathtt{BNS}} = 3.12 \times 10^5 f_{\mathtt{low}}^{-1.543}\,, \qquad (8)$$

and is depicted in the middle panel of Fig. 1.We have found that increasing the high-frequency cutoff to 4096Hz 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 seem to be needed for the merger and ringdown regimes [9].

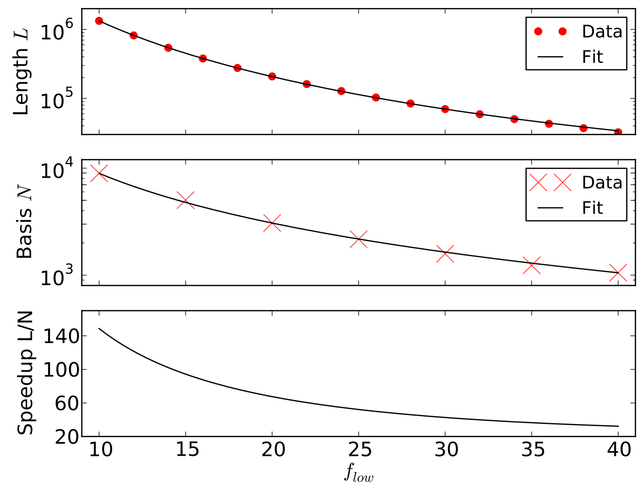

FIG. 1. **Top**: Length L (red dots) of a typical binary neutron star inspiral waveform, with the solid black curve connecting this data implied by the fit (7). **Middle**: Number of reduced basis waveforms (red crosses), with the solid black curve given by the fit (8). **Bottom**: Speedup implied by operation counts, as given by equation (9).

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

$$\frac{L}{N} = (1024 - f_{\mathtt{low}})\,\frac{T_{\mathtt{BNS}}}{N_{\mathtt{BNS}}}\,. \qquad (9)$$

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

Implementation and numerical studies– We have implemented compressed likelihood evaluations and Algorithm (1) in the LAL parameter estimation pipeline, known as LALInference [11, 18], naming the resulting variation LALInference_ROQ.

Next we describe comparisons between MCMC parameter estimation studies using the standard version of LALInference and LALInference_ROQ, for binary neutron stars with TaylorF2 as 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 with sensitivities of $f_{\mathtt{low}} = 40\mathrm{Hz}$, using the zero detuned high power PSD [28].

A typical time window for the coalescence time t_c of a binary neutron star signal, centered around the trigger time, is $W = 0.1s$ [6, 29]. Following the procedure dis-

	$\mathcal{M}_c\,(M_\odot)$	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

TABLE I. Intrinsic parameters (chirp mass $\mathcal{M}_c$, symmetric mass ratio η, masses m_1 and m_2) and Signal-to-Noise Ratio (SNR) of the analysis from Figure 2. The first line give the injected values. The last two lines give median value and 90% credible intervals, for the same parameters with the standard likelihood (second line) and the compressed likelihood using ROQs (third line). The SNR was then computed with $\text{Likelihood}_{\max} \approx \text{SNR}^2/2$. 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.

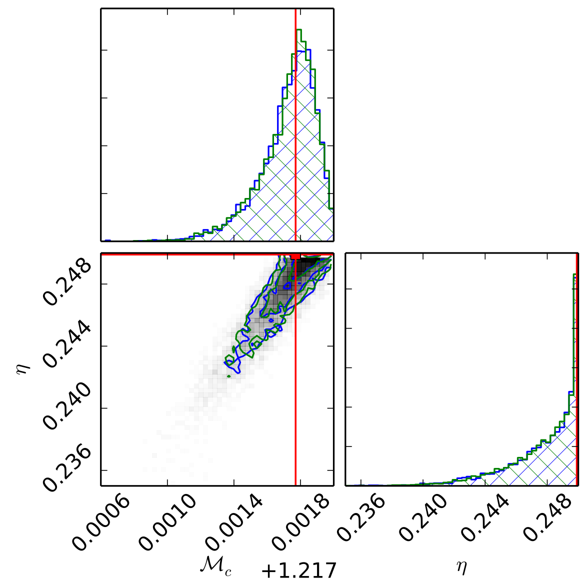

FIG. 2. Probability density function for the chirp mass $\mathcal{M}_c$ and symmetric mass ratio η of a simulated event in LIGO/Virgo data. In green as obtained in ~ 30 hours by the standard likelihood, and in blue as obtained in 1 hour with the ROQ. The injection values are in red, and are listed in Table I. The overlap region of the sets of PDFs is the hatched region.

cussed above, LALInference_ROQ discretizes this prior into $n_c = 2,000$ sub-intervals, each of size $\Delta t_c = 10^{-5}$s, for which it constructs a unique set of ROQ weights. This 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 [18].

We found that, as expected, the ROQ and standard likelihood approaches, through their LAL implementations, produce statistically indistinguishable results for posterior probability density functions over the full 9-dimensional parameter space. As examples, results for the two intrinsic mass parameters for the injection parameters in Table I are shown in Figure 2.

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

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

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

However, the compressed likelihood using ROQs is significantly faster: the ROQ-based MCMC study with the discussed settings takes $\sim$ 1 hour, compared to $\sim$ 30 hours 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, which in this case was $\sim 10^7$. The startup stage of building the ROQ has negligible cost and is done in near real-time: $\sim$ 30s, which is equivalent to $\sim$ 0.028% of the total cost of a standard likelihood parameter estimation study.

Once the advanced detectors have achieved their target sensitivity, with $f_{\texttt{low}} \sim 10$Hz, the longest BNS signals will last around 2048s in duration. Assuming a fiducial high frequency cut-off of 1024 Hz, which is approaching the upper limit of where aLIGO/AdV will be sensitive, we estimate datasets as large as $L \sim 1024\text{Hz}^{-1} \times 2048\text{s} \sim 10^6$. Assuming that the advanced detectors will require at least $\sim 10^7$, this implies runtimes upwards of $\sim$ 100 days and one Petabyte worth of model evaluations using the standard approach. The results of this paper indicate that an ROQ approach will reduce this to less than a day. With parallelization of the sum in each likelihood evaluation run-times could be significantly reduced further to essentially real time. Remarkably, even without parallelization, this approach when applied to the advanced detectors having reached design sensitivity will be faster than even the standard likelihood one used for the initial detectors.

Outlook– Around three weeks of real (wall) time, for $f_{\texttt{low}} = 40$Hz are needed to perform a precessing-spin parameter estimation study on a single data stream with standard likelihood evaluations [18]. With a cutoff of $f_{\texttt{low}} = 10$Hz, these analyses could take months to years, so techniques for accelerated inference on these models, such as the one presented in this paper, are essential for realistic gravitational-wave astronomy and extracting the full science potential of the upcoming advanced gravitational-wave detectors.

In this paper we have addressed the issue of fast likelihood evaluation for non-spinning 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. [19] it was found that the number of reduced basis waveforms barely changes as spins are included – at least in the non-precessing 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 [30] have shown that ultra-compact bases can also be constructed for fully precessing systems. This might provide a means for constructing fast to evaluate models of precessing binary inspirals using surrogates models [9, 10].

Acknowledgments– We thank Collin Capano, William Farr, Chad Galley, Larry Price, Leo Singer, Alan Weinstein and the LVSC CBC and parameter estimation groups for many useful discussions and encouragement through this project. PC's work was supported by a Marie Curie Intra-European Fellowship within the 7th European Community Framework Programme (PIEF-GA-2011-299190), SEF thanks hospitality at and financial support from the Institute of Astronomy at Cambridge, UK, where part of this work was done, JG's work was supported by the Royal Society, VR by a Richard Chase Tolman fellowship at the California Institute of Technology (Caltech) and overall this work was supported in part by NSF grants PHY-1208861 and PHY-1316424 to the University of Maryland (UMD). 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 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. MT thanks Tryst DC Coffeehouse Bar and Lounge, where part of this work was done, for its hospitality.

[1] G. M. Harry and the LIGO Scientific Collaboration, Classical and Quantum Gravity **27**, 084006 (2010), URL http://stacks.iop.org/0264-9381/27/i=8/a=084006.

[2] The Virgo Collaboration, Tech. Rep. (2009), URL https://tds.ego-gw.it/itf/tds/file.php?callFile=VIR-0027A-09.pdf.

[3] J. Abadie et al. (LIGO Scientific Collaboration, Virgo Collaboration), Class.Quant.Grav. **27**, 173001 (2010), 1003.2480.

[4] V. Raymond, M. V. van der Sluys, I. Mandel, V. Kalogera, C. Rover, and N. Christensen, Class. Quant. Grav. **26**, 114007 (2009).

[5] R. J. E. Smith, K. Cannon, C. Hanna, D. Keppel, and I. Mandel, Phys. Rev. D **87**, 122002 (2013), URL http://link.aps.org/doi/10.1103/PhysRevD.87.122002.

[6] J. Aasi, J. Abadie, B. P. Abbott, R. Abbott, T. D. Abbott, M. Abernathy, T. Accadia, F. Acernese, C. Adams, T. Adams, et al. (LIGO-Virgo Scientific Collaboration), Phys. Rev. D **88**, 062001 (2013), URL http://link.aps.org/doi/10.1103/PhysRevD.88.062001.

[7] J. Skilling, AIP Conf. Ser. Vol. 735 p. 395 (2004).

[8] B. J. Owen, Phys. Rev. D **53**, 6749 (1996).

[9] S. E. Field, C. R. Galley, J. S. Hesthaven, J. Kaye, and M. Tiglio (2013), 1308.3565.

[10] M. Pürrer (2014), 1402.4146.

[11] LAL, lAL - https://www.lsc-group.phys.uwm.edu/daswg/projects/lalsuite.html.

[12] LIGO Scientific Collaboration, Virgo Collaboration, J. Aasi, J. Abadie, B. P. Abbott, R. Abbott, T. D. Abbott, M. Abernathy, T. Accadia, F. Acernese, et al., ArXiv e-prints (2013), 1304.0670.

[13] H. Antil, S. E. Field, F. Herrmann, R. H. Nochetto, and M. Tiglio, Journal of Scientific Computing pp. 1–34 (2013), ISSN 0885-7474, arXiv:1210.0577 [cs.NA], URL http://dx.doi.org/10.1007/s10915-013-9722-z.

[14] P. Canizares, S. E. Field, J. R. Gair, and M. Tiglio, Phys. Rev. **D87**, 124005 (2013), 1304.0462.

[15] M. Barrault, Y. Maday, N. C. Nguyen, and A. T. Patera, Comptes Rendus Mathematique **339**, 667 (2004), ISSN 1631-073X.

[16] Y. Maday, N. C. Nguyen, A. T. Patera, and S. H. Pau, Communications on Pure and Applied Analysis **8**, 383 (2009).

[17] S. Chaturantabut and D. C. Sorensen, SIAM Journal on Scientific Computing **32**, 2737 (2010), URL http://link.aip.org/link/?SCE/32/2737/1.

[18] J. Aasi, J. Abadie, B. P. Abbott, R. Abbott, T. D. Abbott, M. Abernathy, T. Accadia, F. Acernese, C. Adams, T. Adams, et al., Phys. Rev. D **88**, 062001 (2013), 1304.1775.

[19] S. E. Field, C. R. Galley, and E. Ochsner, Phys. Rev. D **86**, 084046 (2012), URL http://link.aps.org/doi/10.1103/PhysRevD.86.084046.

[20] W. Hoffmann, Computing **41**, 335 (1989), ISSN 0010-485X, URL http://dx.doi.org/10.1007/BF02241222.

[21] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. E. Creighton, Phys. Rev. D **85**, 122006 (2012), URL http://link.aps.org/doi/10.1103/PhysRevD.85.122006.

[22] B. Sathyaprakash and B. F. Schutz, Living Reviews in Relativity **12** (2009), URL http://www.livingreviews.org/lrr-2009-2.

[23] A. Buonanno, B. Iyer, E. Ochsner, Y. Pan, and B. Sathyaprakash, Phys.Rev. **D80**, 084043 (2009), 0907.0700.

[24] L. P. Singer, L. R. Price, B. Farr, A. L. Urban, C. Pankow, et al. (2014), 1404.5623.

[25] *Eureqa*, URL http://www.nutonian.com.

[26] M. Schmidt and H. Lipson, Science **324**, 81 (2009), http://www.sciencemag.org/content/324/5923/81.full.pdf, URL http://www.sciencemag.org/content/324/5923/81.abstract.

[27] I. Mandel, C. P. L. Berry, F. Ohme, S. Fairhurst, and W. M. Farr (2014), 1404.2382.

[28] LIGO Scientific Collaboration, Tech. Rep. (2010), LIGO-T0900288-v3.

[29] R. J. E. Smith, K. Cannon, C. Hanna, D. Keppel, and I. Mandel, Phys. Rev. D **87**, 122002 (2013), URL http://link.aps.org/doi/10.1103/PhysRevD.87.122002.

[30] J. Blackman, B. Szilagyi, C. R. Galley, and M. Tiglio

(2014), 1401.7038.