

Efficiently enclosing the compact binary parameter space by singular-value decompositionKipp Cannon,^{1,2,*} Chad Hanna,^{1,3,†} and Drew Keppel^{1,4,5,6,‡}¹*LIGO Laboratory, California Institute of Technology, Pasadena, California 91125, USA*²*Canadian Institute for Theoretical Astrophysics, 60 St. George Street, University of Toronto, Toronto, Ontario M5S 3H8, Canada*³*Perimeter Institute for Theoretical Physics, Waterloo, Ontario N2L 2Y5, Canada*⁴*Theoretical Astrophysics, California Institute of Technology, Pasadena, California 91125, USA*⁵*Albert-Einstein-Institut, Max-Planck-Institut für Gravitationsphysik, D-30167 Hannover, Germany*⁶*Leibniz Universität Hannover, D-30167 Hannover, Germany*

(Received 27 January 2011; published 4 October 2011)

Gravitational-wave searches for the merger of compact binaries use matched filtering as the method of detecting signals and estimating parameters. Such searches construct a fine mesh of filters covering a signal parameter space at high density. Previously it has been shown that singular-value decomposition can reduce the effective number of filters required to search the data. Here we study how the basis provided by the singular-value decomposition changes dimension as a function of template-bank density. We will demonstrate that it is sufficient to use the basis provided by the singular-value decomposition of a low-density bank to accurately reconstruct arbitrary points within the boundaries of the template bank. Since this technique is purely numerical, it may have applications to interpolating the space of numerical relativity waveforms.

DOI: [10.1103/PhysRevD.84.084003](https://doi.org/10.1103/PhysRevD.84.084003)

PACS numbers: 04.30.-w, 07.05.Kf

I. INTRODUCTION

Several broadband laser interferometer gravitational wave (GW) detectors are operating at high sensitivities and will continue to improve over the next decade [1–5]. As detectors improve, it is increasingly likely that GW astronomers will observe gravitational radiation emitted from the coalescence of compact binary systems involving neutron stars and/or stellar mass black holes [6].

Because compact binary coalescence (CBC) waveforms are well modeled, GW searches for such signals are conducted by matched-filtering the detectors' data with banks of template waveforms, chosen to adequately cover a region of the signal parameter space [7]. For GW signals from the merger of compact objects with negligible spin, this parameter space is defined by functions of the masses of the two objects. To search for signals within this parameter space, a bank of templates is constructed to sample the parameter space sufficiently densely such that there is minimal loss of signal-to-noise ratio (SNR). Traditionally, template banks used to search this two-dimensional signal parameter space have been constructed using the $(A_2)^*$ lattice [8], referred to as “hexagonally placed” template banks. This problem becomes more difficult in higher dimensions, where other types of template placement algorithms have recently been investigated [9–12].

In [13], the singular-value decomposition (SVD) was applied to CBC waveforms to show how hexagonally placed template banks with M templates could be implemented with $N' \ll 2M$ filters ($2M$ being the nominal

number of filters required for the M 2-phase templates). This was achieved by truncating the SVD of the matrix consisting of the time series of the template waveforms. Here we demonstrate that the bases identified by the SVD are effective at spanning the space of all CBC waveforms within the region of parameter space sampled by the original bank. We find that the SVD of a low-density bank provides a basis suitable for constructing all the waveforms from a higher-density bank, even waveforms at arbitrary locations within that region of parameter space. Such a basis could be used to reduce the computational cost of (1) performing hierarchical searches of parameter space that minimize waveform-mismatch errors associated with identifying signals in the data, (2) running parameter-estimation algorithms that take the inner product between the data and millions of waveforms, and (3) generating computationally costly waveforms.

This paper is organized as follows. Section II describes how we apply the SVD to approximately embed the signal manifold in a vector space. Section III tests this embedding by reconstructing various points in the manifold. Finally, Sec. IV expands on possible applications of this technique.

II. ENCLOSING THE SIGNAL SPACE WITH SINGULAR-VALUE DECOMPOSITION

In this section we explore how the number of basis vectors required to reconstruct a template bank scales with the initial density of the template bank. We define a template bank of signal waveforms covering a patch \mathcal{P} of the signal manifold, which is used to test for the presence and strength of signals from \mathcal{P} in the detectors' data. We construct a signal matrix in the same manner as [13]. Specifically, we create a real-valued matrix \mathbf{H} by

*kipp.cannon@ligo.org

†chad.hanna@ligo.org

‡drew.keppel@ligo.org

alternately filling its rows with the real and imaginary parts (cosine and sine) of the template waveform time series from a CBC template bank covering \mathcal{P} , $\mathbf{H} = \{H_{\alpha j}\} = \{\Re \vec{h}_1, \Im \vec{h}_1, \Re \vec{h}_2, \Im \vec{h}_2, \dots, \Re \vec{h}_M, \Im \vec{h}_M\}^T$.

As in [13], we constructed the template matrix with chirp masses $M_c = M\eta^{3/5}$, where $M = m_1 + m_2$ is the total mass and $\eta = m_1 m_2 / M^2$ is the symmetric mass ratio, of $1.125M_\odot \leq M_c < 1.240M_\odot$ and component masses of $1M_\odot \leq m_1, m_2 < 3M_\odot$. Template banks covering this region are created using the template-placement algorithms of the LSC Algorithm Library [14]. Template placement is done in the (τ_0, τ_3) plane, where τ_0 and τ_3 are defined as

$$\tau_0 = \frac{5}{256} (\pi f_0)^{-8/3} M_c^{-5/3}, \quad (1)$$

$$\tau_3 = \frac{\pi}{8} (\pi f_0)^{-5/3} M_c^{-2/3} \eta^{-3/5}, \quad (2)$$

and where f_0 is some fiducial frequency, which we choose to be $f_0 = 60$ Hz.

The nonspinning waveforms for each template are produced to 3.5 post-Newtonian order, sampled at 2048 Hz, up to the Nyquist frequency of 1024 Hz. The last 10 seconds of each waveform, whitened with the initial LIGO amplitude spectral density, are used to construct \mathbf{H} . The SVD is then applied to \mathbf{H} , decomposing the matrix into two unitary matrices, \mathbf{V} and \mathbf{U} , and a diagonal matrix $\mathbf{\Sigma}$

$$\mathbf{H} = \mathbf{V}\mathbf{\Sigma}\mathbf{U}^T, \quad (3)$$

where \mathbf{U} is a matrix composed of basis vectors (i.e., unit-norm time-series vectors), \mathbf{V} is a matrix composed of reconstruction coefficients, and $\mathbf{\Sigma}$ is a matrix containing the singular values of \mathbf{H} .

In [13], it was demonstrated that truncating the reconstruction of \mathbf{H} to use only the N' -basis vectors with the largest singular values results in an average fractional SNR loss $\langle \delta\rho/\rho \rangle$ proportional to the sum of the discarded singular values squared. In this investigation, we truncate these reconstruction matrices at $\langle \delta\rho/\rho \rangle = 10^{-7}$. This corresponds roughly to the truncation error of Institute of Electrical and Electronics Engineers 754 32-bit floating-point numbers.

We explore how the number of basis vectors changes as the number of rows in \mathbf{H} is increased by generating template banks for \mathcal{P} with increasing density (i.e., increasing minimal match). We confirm that the number of basis vectors required to reconstruct \mathbf{H} saturates at a particular value of minimal match. Figure 1 shows that, as the minimal match of the template bank is increased, resulting in denser samplings of \mathcal{P} , the number of basis vectors needed to reconstruct \mathbf{H} to the required accuracy saturates around a minimal match of $\sim 89.9\%$. This indicates that \mathcal{P} is able to be embedded—to an accuracy of 1 part in 10^7 —in a vector space consisting of ~ 150 dimensions.

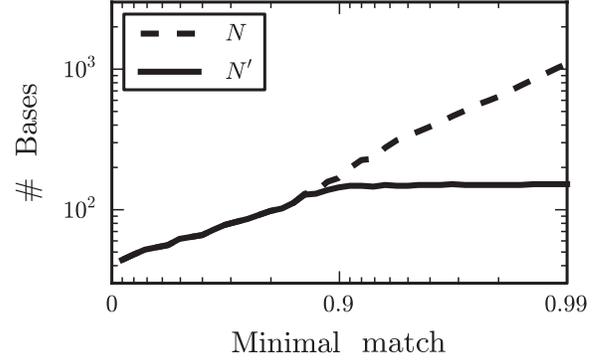


FIG. 1. The number of filters as a function of minimal match, which increases with the density of the template bank. The total number of filters in the template bank, N , is shown by the dashed line. The number of filters needed to reconstruct the template matrix such that $\langle \delta\rho/\rho \rangle = 10^{-7}$, N' , is shown by the solid line. We find that the number of filters needed to reconstruct \mathbf{H} saturates when the minimal match reaches $\sim 89.9\%$.

In the next section, we demonstrate how the basis waveforms identified by the coarsely sampled bank can be used to reconstruct templates at arbitrary points on the signal manifold.

III. EFFICIENT RECONSTRUCTION OF WAVEFORMS IN THE MANIFOLD

In order to determine how well these waveforms can be reconstructed, we compute a quantity called the “average fractional SNR loss,” $\delta\rho_\alpha/\rho_\alpha$. This quantity can be thought of as the mismatch between the original waveform \vec{h}_α and the reconstructed waveform \vec{h}'_α , averaged over the phase angle. It tells us how far the reconstructed waveform is from the original waveform. This quantity can be split into two pieces,

$$\frac{\delta\rho_\alpha}{\rho_\alpha} = \left(\frac{\delta\rho_\alpha}{\rho_\alpha} \right)_\perp + \left(\frac{\delta\rho_\alpha}{\rho_\alpha} \right)_\parallel, \quad (4)$$

where $(\delta\rho_\alpha/\rho_\alpha)_\perp$ is due to the (in)completeness of the basis vectors, and $(\delta\rho_\alpha/\rho_\alpha)_\parallel$ is due to the truncation of the SVD reconstruction. The first piece is given in terms of the SVD quantities by

$$\left(\frac{\delta\rho_\alpha}{\rho_\alpha} \right)_\perp = 1 - \left[\frac{1}{2} \sum_{\mu=1}^N (v_{(2\alpha-1)\mu}^2 + v_{(2\alpha)\mu}^2) \sigma_\mu^2 \right]^{1/2}, \quad (5)$$

where $v_{(2\alpha-1)\mu}$ and $v_{(2\alpha)\mu}$ are the reconstruction coefficients for the real and imaginary parts, respectively, of the α th waveform associated with the μ th basis vector and are elements of \mathbf{V} ; σ_μ is the μ th element of $\mathbf{\Sigma}$; and the sum is over all of the terms of \mathbf{V} and $\mathbf{\Sigma}$. The second piece is given by Eq. (25) of [13] in terms of SVD quantities,

$$\left(\frac{\delta\rho_\alpha}{\rho_\alpha} \right)_\parallel = \frac{1}{4} \sum_{\mu=N'+1}^N (v_{(2\alpha-1)\mu}^2 + v_{(2\alpha)\mu}^2) \sigma_\mu^2, \quad (6)$$

where the sum is over the truncated terms of \mathbf{V} and $\mathbf{\Sigma}$.

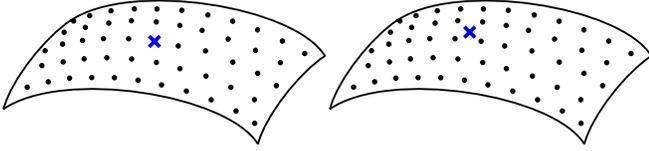


FIG. 2 (color online). A visual representation showing the two types of points of \mathcal{P} that we can choose to reconstruct. The left panel shows an example point that is in \mathbf{H} and thus in \mathcal{P} . The right panel shows an example point that is not part of \mathbf{H} but is in \mathcal{P} .

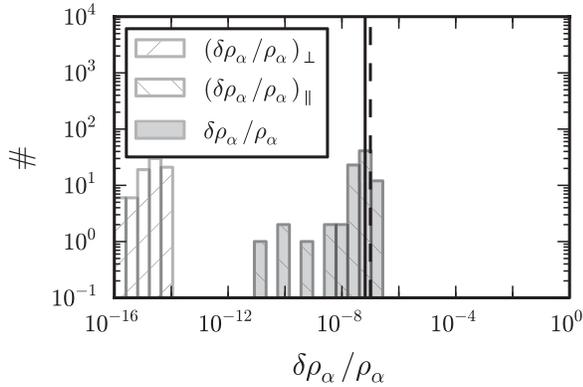
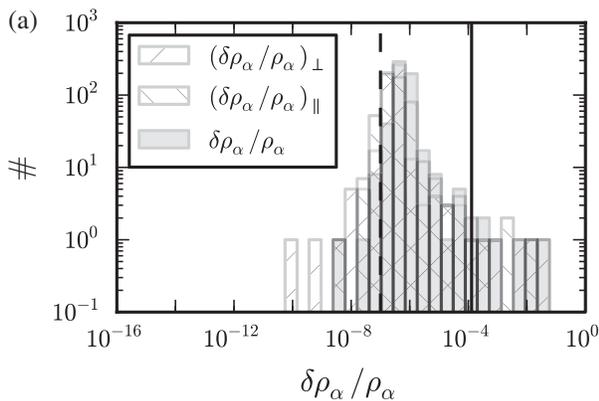


FIG. 3. Histograms of $(\delta\rho_\alpha/\rho_\alpha)_\perp$, $(\delta\rho_\alpha/\rho_\alpha)_\parallel$, and $\delta\rho_\alpha/\rho_\alpha$ for the waveforms that went into the construction of \mathbf{H} . As expected, $(\delta\rho_\alpha/\rho_\alpha)_\perp$ is around single-precision floating-point roundoff error, and the average $\delta\rho_\alpha/\rho_\alpha$ (solid vertical line) matches the expected fractional SNR loss (dashed vertical line).

We test this embedding of the signal manifold to see how well various points in the manifold can be reconstructed. The test points that we reconstruct are of two types: (1) those from the original signal matrix \mathbf{H} , and (2) those



absent from \mathbf{H} but within \mathcal{P} . These two types of tests are illustrated in Fig. 2.

A test of the first type is shown in Fig. 3. For these points, since the SVD provides an exact decomposition of \mathbf{H} , $(\delta\rho_\alpha/\rho_\alpha)_\perp$ should be identically zero, which is what is observed to numerical precision. We also see that the average reconstruction accuracy for points from \mathbf{H} agrees with our chosen value of 10^{-7} . This result is expected as it is an extension of the investigation from Fig. 4 of [13] applied to a more stringent reconstruction accuracy.

A test of the second type is shown in Fig. 4. To choose points uniformly from \mathcal{P} but absent from \mathbf{H} , we generate a denser template bank within the same region of parameter space described in Sec. II. Specifically, we generate this template bank with a minimal match of 99%. In order to test the reconstruction accuracy of these waveforms, we project the real and imaginary parts of the waveforms onto the basis vectors from the SVD of \mathbf{H} ,

$$v'_{\alpha\mu} = \frac{1}{\sigma_\mu} \sum_j h_{\alpha j} u_{\mu j}, \quad (7)$$

where $v'_{\alpha\mu}$ represents a reconstruction coefficient associated with the μ th basis vector for the real or imaginary part of the α th waveform from the denser template bank; σ_μ is the μ th element of $\mathbf{\Sigma}$; $h_{\alpha j}$ is the j th time sample of the real or imaginary part of the α th waveform from the denser template bank; and $u_{\mu j}$ is the j th time sample from the μ th basis vector in \mathbf{U} . The real and imaginary parts of the waveforms from the denser template bank are then reconstructed using

$$h'_{\alpha j} = \sum_\mu v'_{\alpha\mu} \sigma_\mu u_{\mu j}. \quad (8)$$

The distribution of $\delta\rho_\alpha/\rho_\alpha$ for these waveforms, Fig. 4(a), shows a tail extending to large mismatches.

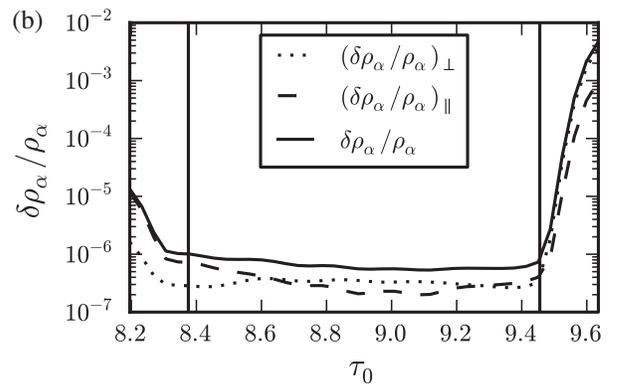


FIG. 4. (a) Histograms of $(\delta\rho_\alpha/\rho_\alpha)_\perp$, $(\delta\rho_\alpha/\rho_\alpha)_\parallel$, and $\delta\rho_\alpha/\rho_\alpha$ for waveforms from \mathcal{P} that were not in \mathbf{H} . $(\delta\rho_\alpha/\rho_\alpha)_\perp$ is seen to be of the same order of magnitude as $(\delta\rho_\alpha/\rho_\alpha)_\parallel$. The peak of $\delta\rho_\alpha/\rho_\alpha$ is above the expected fractional SNR loss for waveforms from \mathbf{H} (vertical dashed line); however, the average $\delta\rho_\alpha/\rho_\alpha$ is skewed to large values by a tail in the distribution. (b) How these mismatches vary across \mathcal{P} , averaged over the τ_3 direction. The largest mismatches come from near the borders of the template bank in the τ_0 direction. Figure 5 restricts our attention to the central 75% of the domain of \mathcal{P} , whose boundaries are shown by the vertical lines here in Fig. 4(b).

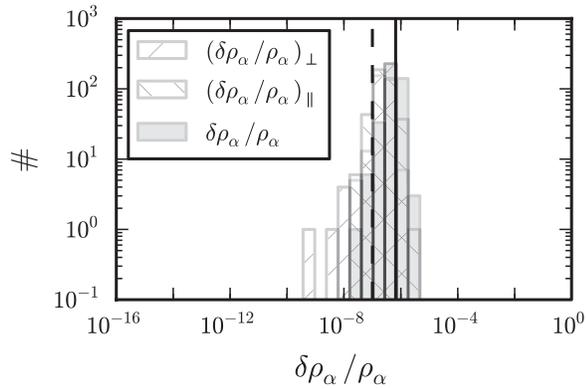


FIG. 5. Histograms of $(\delta\rho_\alpha/\rho_\alpha)_\perp$, $(\delta\rho_\alpha/\rho_\alpha)_\parallel$, and $\delta\rho_\alpha/\rho_\alpha$ for waveforms from \mathcal{P} that were not in \mathbf{H} , similar to Fig. 4(b), eliminating, however, test points near the boundaries in the τ_0 direction. The average $\delta\rho_\alpha/\rho_\alpha$ (solid vertical line) accuracy is slightly worse than the expected fractional SNR loss (dashed vertical line).

Examining where these large mismatches are located in parameter space, we find they originate from near the boundaries of \mathcal{P} . Removing the test points near the boundaries in the τ_0 direction, shown in Fig. 5, we find that the tail of large mismatches disappears.

An additional test of the second type, which systematically explores the reconstruction accuracy near a point whose waveform went into \mathbf{H} , is shown in Fig. 6. Figure 6 (a) shows a set of three nearest-neighbor templates. We investigate how the reconstruction accuracy varies as one moves from point A to the central point, point B . Point B is assumed to have the largest mismatch between its waveform and the waveforms from any of the three

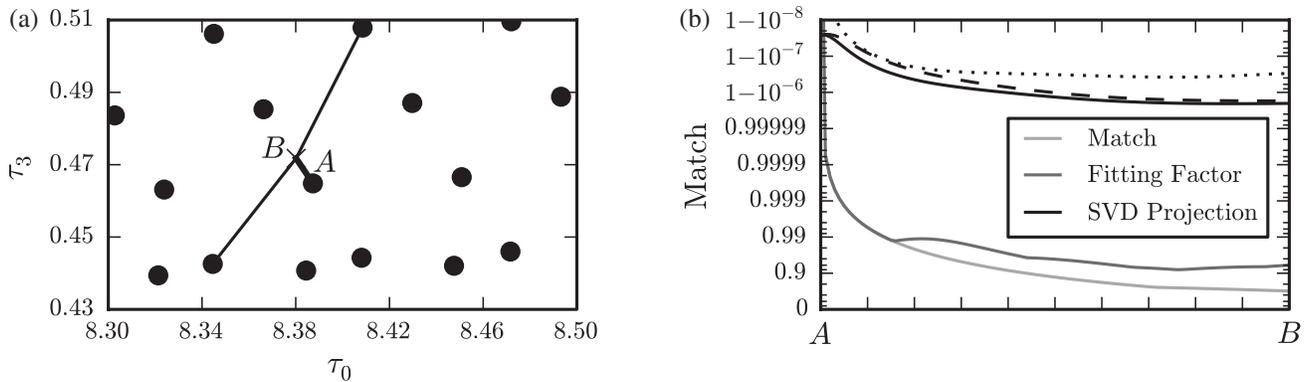


FIG. 6. (a) A plot showing a region of \mathcal{P} . The circles are points whose waveforms go into \mathbf{H} . The line segment \overline{AB} connects one of those points, A , with the point B , which is the central point of A and two of its nearest neighbors. Point B is situated such that it should have a 89.9% fitting factor with each of the surrounding points. (b) How well a waveform from a given point of \overline{AB} , \vec{h}_p , can be “matched.” The lowest, light-grey curve shows the *match*, which is the normalized inner product between \vec{h}_A and the waveform from the corresponding point along \overline{AB} . The next-lowest (dark-grey) curve shows the *fitting factor*, which is the match maximized over phase and time. Since we are using a 89.9% minimal match bank, it is expected that the fitting factor falls to around that value. The three upper curves are associated with the SVD projection. As before, the solid curve is $1 - \delta\rho_\alpha/\rho_\alpha$, the dashed curve is $1 - (\delta\rho_\alpha/\rho_\alpha)_\parallel$, and the dotted curve is $1 - (\delta\rho_\alpha/\rho_\alpha)_\perp$. The SVD-basis vectors are able to reconstruct to high accuracy all points along the line.

surrounding points. Figure 6(b) shows the accuracies of representing the waveforms along \overline{AB} with the waveform of point A with and without maximizing over phase and time, called the *fitting factor* and *match*, respectively, and with the SVD projection described by Eq. (8). The fitting factor falls to the minimal match of the template bank when comparing the waveforms from A and B , which is expected as the minimal match involves maximizing over phase and time. The reconstruction accuracy associated with SVD projection is consistently high and close to the chosen reconstruction accuracy of 1 part in 10^7 .

IV. DISCUSSION

These investigations show that the SVD can be used to find a set of basis vectors that not only span the signal matrix \mathbf{H} but also enclose the signal manifold \mathcal{P} sampled by \mathbf{H} .

GW pipelines that search for known waveforms, such as GWs from CBCs, commonly compute waveform-consistency statistics that compare the observed response of a template-waveform filter to the data with what one would expect given the presence of that signal. These consistency statistics are found to perform better when the mismatch between the template waveform and the signal waveform is small [15]. Filtering with a fixed-density template bank can introduce mismatch between the nearest template and the signal. This mismatch can be greatly reduced if one is able to find the exact point in parameter space where the signal is located and filter the data using that point. Using the SVD-basis vectors, one could reconstruct a point closer to the point of the signal and improve the waveform consistency statistics.

Parameter-estimation techniques for GWs from CBCs often use Monte Carlo Markov chain algorithms to search the parameter space. This involves producing waveforms and filtering the data against many points of the parameter space, building up probability-density functions associated with each parameter. If one filtered the data using the basis vectors from the SVD, it would be simple to reconstruct to high accuracy the output one would have seen if one had filtered the data using any waveform from within the parameter space, which could then be used to map out the probability density functions implicitly.

Finally, some waveforms are computationally costly to produce. Such is the case to a limited extent for waveforms produced by solving post-Newtonian differential equations and, as an extreme case, solving the full Einstein equations using numerical relativity. The SVD could be used to generate numerically interpolated waveforms starting from a limited sampling of the parameter space.

In order to gain benefit from these applications, it would be necessary to determine the reconstruction coefficients in a computationally efficient manner. This paper has not tried to address this problem because (1) it has assumed

the target waveforms are known and (2) it computes the reconstruction coefficients using computationally expensive inner products. Generation of these reconstruction coefficients warrants future investigation, as the benefits derived from this technique would be substantial.

ACKNOWLEDGMENTS

The authors would like to acknowledge the support of the LIGO Lab, NSF Grants No. PHY-0653653 and No. PHY-0601459, and the David and Barbara Groce Fund at Caltech. LIGO was constructed by the California Institute of Technology and the Massachusetts Institute of Technology with funding from the National Science Foundation and operates under cooperative agreement No. PHY-0757058. Research at Perimeter Institute is supported through Industry Canada and by the Province of Ontario through the Ministry of Research and Innovation. K.C. was supported by the National Science and Engineering Research Council, Canada. D.K. was supported in part from the Max Planck Gesellschaft. This work has LIGO document No. LIGO-P1000039.

-
- [1] LIGO Scientific Collaboration, Report No. LIGO-T0900288-v2 (2009).
 - [2] Virgo Scientific Collaboration, Report No. VIR-NOT-DIR-1390-304 (2005).
 - [3] H. Grote (LIGO Scientific Collaboration), *Classical Quantum Gravity* **27**, 084003 (2010).
 - [4] F. Acernese *et al.*, *Classical Quantum Gravity* **25**, 114045 (2008).
 - [5] B.P. Abbott *et al.* (LIGO Scientific Collaboration), *Rep. Prog. Phys.* **72**, 076901 (2009).
 - [6] J. Abadie *et al.* (LIGO Scientific Collaboration), *Classical Quantum Gravity* **27**, 173001 (2010).
 - [7] B. J. Owen, *Phys. Rev. D* **53**, 6749 (1996).
 - [8] T. Cokelaer, *Phys. Rev. D* **76**, 102004 (2007).
 - [9] I. W. Harry, B. Allen, and B. S. Sathyaprakash, *Phys. Rev. D* **80**, 104014 (2009).
 - [10] G. M. Manca and M. Vallisneri, *Phys. Rev. D* **81**, 024004 (2010).
 - [11] C. Messenger, R. Prix, and M. A. Papa, *Phys. Rev. D* **79**, 104017 (2009).
 - [12] R. Prix, *Classical Quantum Gravity* **24**, S481 (2007).
 - [13] K. Cannon, A. Chapman, C. Hanna, D. Keppel, A. C. Searle, and A. J. Weinstein, *Phys. Rev. D* **82**, 044025 (2010).
 - [14] LIGO Scientific Collaboration, “LSC Algorithm Library” [<https://www.lsc-group.phys.uwm.edu/daswg/projects/lalsuite.html>].
 - [15] B. Allen, *Phys. Rev. D* **71**, 062001 (2005).