

Binary tree approach to template placement for searches for gravitational waves from compact binary mergers

Chad Hanna,^{1,2,3,4} James Kennington,^{1,2,*} Shio Sakon,^{1,2} Stephen Privitera,^{5,†} Miguel Fernandez,^{1,2} Jonathan Wang,⁶ Cody Messick,¹ Alex Pace,^{1,2} Kipp Cannon,⁷ Prathamesh Joshi,^{1,2} Rachael Huxford,^{1,2} Sarah Caudill,⁸ Chiwai Chan,⁷ Bryce Cousins,^{1,4} Jolien D. E. Creighton,⁹ Becca Ewing,^{1,2} Heather Fong,^{7,10} Patrick Godwin,^{1,2} Ryan Magee,^{1,2} Duncan Meacher,⁹ Soichiro Morisaki,¹¹ Debnandini Mukherjee,^{1,2} Hiroaki Ohta,⁷ Surabhi Sachdev,^{1,2,12} Divya Singh,^{1,2} Ron Tapia,^{1,4} Leo Tsukada,^{7,10} Daichi Tsuna,^{7,10} Takuya Tsutsui,⁷ Koh Ueno,⁷ Aaron Viets,⁹ Leslie Wade,¹³ and Madeline Wade¹³

¹*Department of Physics, The Pennsylvania State University, University Park, Pennsylvania 16802, USA*

²*Institute for Gravitation and the Cosmos, The Pennsylvania State University, University Park, Pennsylvania 16802, USA*

³*Department of Astronomy and Astrophysics, The Pennsylvania State University, University Park, Pennsylvania 16802, USA*

⁴*Institute for Computational and Data Sciences, The Pennsylvania State University, University Park, Pennsylvania 16802, USA*

⁵*Albert-Einstein-Institut, Max-Planck-Institut für Gravitationsphysik, D-14476 Potsdam-Golm, Germany*

⁶*Department of Physics, University of Michigan, Ann Arbor, Michigan 48109, USA*

⁷*RESCEU, The University of Tokyo, Tokyo, 113-0033, Japan*

⁸*Nikhef, Science Park, 1098 XG Amsterdam, Netherlands*

⁹*Leonard E. Parker Center for Gravitation, Cosmology, and Astrophysics, University of Wisconsin-Milwaukee, Milwaukee, Wisconsin 53201, USA*

¹⁰*Graduate School of Science, The University of Tokyo, Tokyo 113-0033, Japan*

¹¹*Institute for Cosmic Ray Research, The University of Tokyo, 5-1-5 Kashiwanoha, Kashiwa, Chiba 277-8582, Japan*

¹²*LIGO Laboratory, California Institute of Technology, MS 100-36, Pasadena, California 91125, USA*

¹³*Department of Physics, Hayes Hall, Kenyon College, Gambier, Ohio 43022, USA*



(Received 23 February 2023; accepted 26 July 2023; published 14 August 2023)

We demonstrate a new geometric method for fast template placement for searches for gravitational waves from the inspiral, merger and ringdown of compact binaries. The method is based on a binary tree decomposition of the template bank parameter space into nonoverlapping hypercubes. We use a numerical approximation of the signal overlap metric at the center of each hypercube to estimate the number of templates required to cover the hypercube and determine whether to further split the hypercube. As long as the expected number of templates in a given cube is greater than a given threshold, we split the cube along its longest edge according to the metric. When the expected number of templates in a given hypercube drops below this threshold, the splitting stops and a template is placed at the center of the hypercube. Using this method, we generate aligned-spin template banks covering the mass range suitable for a search of Advanced LIGO, Advanced Virgo and KAGRA data. The aligned-spin bank was generated in ~ 24 hours using a single CPU core and produced 2 million templates. Our primary motivation for developing this algorithm is to produce a bank with useful geometric properties in the physical parameter space coordinates. Such properties are useful for population modeling and parameter estimation.

DOI: [10.1103/PhysRevD.108.042003](https://doi.org/10.1103/PhysRevD.108.042003)

I. INTRODUCTION

Banks of template gravitational-wave signals are central tools in the matched-filter detection of gravitational-wave signals from compact binary coalescence [1–3]. The general

compact binary gravitational-wave signal depends on at least 15 parameters: two mass parameters, six spin parameters, distance, time, and five angles defining binary orientation with respect to the gravitational-wave antenna. The parameter space can be even larger if, for instance, matter or eccentricity effects are included. Since we do not know the source parameters *a priori*, we must search the data over all possible source parameters.

*james.kennington@ligo.org

†stephen.privitera@ligo.org

We are often able to quickly maximize the signal-to-noise ratio (SNR) over a subset of the parameters either analytically or by efficient numerical techniques. For instance, some parameters enter only into the overall amplitude of the signal which is normalized away by the matched-filter definition of SNR. Which parameters these are depends on the assumptions made about the signal. For instance, nonprecessing binaries have a constant inclination angle, which enters into the gravitational-wave signal only in the overall scale of the waveform, whereas precessing binaries have a time-dependent inclination, leading to modulation in the waveform phase and amplitude. The coalescence time enters into the waveform as a frequency-dependent phase shift which can efficiently be searched over using widely available fast Fourier transform routines. Considering only dominant $(\ell, |m|) = (2, 2)$ modes of gravitational-wave signals, the coalescence phase can also be maximized over analytically.

Given the approximations, assumptions and techniques described above, a subset of parameters, $\vec{\lambda}$, the *template bank* parameters, are generally relevant for template placement. We search over these parameters by laying down a discrete set of points in the parameter space and repeating the matched-filter calculation for each template. The set of points must be chosen as a compromise between optimal SNR recovery and available computational resources. Placing templates finely in the template parameter space leads to high SNR recovery, but can quickly make the search prohibitively expensive. In particular, the number of templates required to cover a D -dimensional parameter space such that no more than a fraction M of the SNR is lost to any potential signal scales as $M^{-D/2}$ [2].

In the case of nonspinning binaries, lattice placement strategies based on an approximate analytic expression for the signal space “distance” between two nearby templates have been shown to be effective for covering the template parameter space [4,5]. To guarantee efficiency of the placement, these methods require that the metric $\mathbf{g}(\vec{\lambda})$, which defines the distance between nearby templates, is very nearly constant throughout the parameter space. For waveforms involving spin, in which a metric is either unavailable or varies rapidly throughout the parameter space, stochastic template placement has proven to be effective in covering the parameter space [6–10]. The stochastic placement technique works by randomly selecting a large number of points in parameter space and keeping only those points which fall sufficiently far away from points which have already been accepted into the bank. This technique, while robust, is computationally inefficient, although recent implementations have made significant strides towards optimization [9,11,12].

Geometric techniques have also been applied to generate aligned-spin template banks [13–15]. In Ref. [13], the authors demonstrate a geometric template bank for neutron-star–black-hole binaries. The authors find satisfactory

coverage for this parameter space by stacking two two-dimensional lattices, taking advantage of the fact that the parameter space is “thin” in the third dimension. This placement strategy was used in conjunction with ordinary stochastic placement [11] to cover the full compact binary parameter space searched in the recent LIGO-Virgo searches [16,17]. In Ref. [14], the authors consider an interesting extension of this technique which starts with a true three-dimensional lattice, and falls back to the stochastic approach when the lattice approach breaks down. In Ref. [12], the authors also consider a hybrid stochastic-geometric technique, similar to the algorithm we propose here; however, the notion of lattice adjacency the authors used is Cartesian whereas we incorporate the intrinsic geometry of the parameter manifold.

These solutions continue to rely at least partially on stochastic placement methods, which scales poorly with the number of templates. The required number of template parameters to cover a parameter space at a given minimal match threshold increases dramatically with the bandwidth of the interferometer and the dimension of the target signal space, both of which are ever increasing in ground-based gravitational wave searches [11,18,19]. Currently used aligned-spin template banks have four template parameters (two masses and two spins) and over 1 million templates at maximal mismatches between 1%–3% [20]. Precessional effects add five more parameters (four spin components and the binary inclination at some reference frequency) and an additional order of magnitude in templates [19]. At high mass ratios, subdominant modes may also be important for detection, which can only further increase the template bank size. Presently template bank generation with stochastic methods may be computationally slow. Future larger banks will require more computing resources to generate as gravitational wave detector sensitivity improves. This can be problematic if banks are generated often.

Here, we demonstrate a new method for template placement based on a binary tree decomposition of the parameter space which is purely geometric originally explored here [21]. The algorithm relies on a numerical estimation of the parameter space metric and uses this metric to determine how to grow the binary tree. This algorithm requires $\mathcal{O}(2^n D^2)$ overlap calculations, where n is the bifurcation number of the parameter space, i.e., how many times a characteristic cell is split, and dim is the dimension of the resulting template bank. We demonstrate this method by constructing a bank suitable for Advanced LIGO [22], Advanced Virgo [23] and KAGRA [24] data analysis.

II. MOTIVATION

Beyond general interest in pursuing novel template placement algorithms, our motivation for pursuing this work is threefold based on experiences analyzing LIGO and Virgo data during the third observing run [25]. First,

in order to apply a population model to gravitational wave detection, it is important to account for template placement [26] in a way that may account for the coordinate volume that a template occupies [25,27–29]. The binary tree approach that we have taken guarantees that each template ends up in a hyper-rectangle in the physical coordinates making coordinate volume calculations easy. Second, in order to ensure a high availability of service for online compact binary searches we run searches at two different data centers. The goal is to split the parameter space in a way that if one site goes down the other is still efficient at detecting a broad class of binary signals. The binary tree approach allows us to use a bank derived from the “right” and “left” splits separately. Finally, having a bank that is gridlike in physical coordinates is generally useful for template interpolation [30] and rapid parameter estimation [31] problems and we are interested in exploring this as future work.

III. METHODS

Our method, whose implementation we refer to as *treebank*, relies on having an accurate approximation of the template space metric $\mathbf{g}(\vec{\lambda})$, which gives a measure of the distance between nearby templates. For our work $\vec{\lambda} \equiv \{t_c, \log m_1, \log m_2, \chi_{\text{eff}}\}$, where t_c is the coalescence time, m_1 is the primary component’s mass, m_2 is the secondary component’s mass, and $\chi_{\text{eff}} \equiv (m_1 a_{1z} + m_2 a_{2z}) / (m_1 + m_2)$ with a_{1z}, a_{2z} as the dimensionless spins of each component along the orbital angular momentum [32]. We define the mismatch δ^2 between two nearby gravitational-wave templates, $h(\vec{\lambda})$ and $h(\vec{\lambda} + \vec{\Delta\lambda})$, according to

$$\delta(\vec{\lambda}, \vec{\Delta\lambda})^2 = 1 - \langle \hat{h}(\vec{\lambda}) | \hat{h}(\vec{\lambda} + \vec{\Delta\lambda}) \rangle, \quad (1)$$

$$\langle a | b \rangle \equiv \left| \int_{-f_N}^{f_N} \frac{\tilde{a}(f) \tilde{b}^*(f)}{S_n(f)} df \right|, \quad (2)$$

where the template a or b is taken to be complex valued containing both the sine and cosine phases, thereby maximizing over phase, and f_N is the Nyquist frequency. δ^2 can be expressed in terms of a metric tensor $\mathbf{g}$ on the template signal manifold as

$$\delta(\vec{\lambda}, \vec{\Delta\lambda})^2 = \vec{\Delta\lambda}^T \mathbf{g}(\vec{\lambda}) \vec{\Delta\lambda}. \quad (3)$$

From the metric, we can also compute a local volume element and thereby estimate the number of templates required to fill a given hypercube cell in the binary tree decomposition [2]:

$$\mathcal{N}_C(\vec{\lambda}) = \frac{\int \sqrt{|\det g(\vec{\lambda})|} dV}{V_T}, \quad (4)$$

where V_T is the volume of a template in mismatch space. We use the definition by Owen for the metric components in terms of the mismatch [2]

$$g_{ij} = -\frac{1}{2} \left[\frac{\partial^2 \delta^2(\vec{\lambda}, \vec{\Delta\lambda})}{\partial \Delta\lambda^i \partial \Delta\lambda^j} \right]_{\Delta\lambda^k=0}. \quad (5)$$

We have implemented two numerical schemes for estimating the metric component values that we call the iterative and deterministic methods. The iterative method is a standard convergence scheme for numerical differentiation leveraging the Python package NUMDIFFTOOLS [33]. The deterministic method uses definitions of the metric components as partial derivatives of the mismatch to compute the preliminary metric $\gamma_{\mu\nu}$ in a single step:

$$\begin{aligned} \gamma_{\mu\mu} &= \frac{\delta^2(\vec{\lambda}, \vec{\Delta\lambda})}{\Delta\lambda^{\mu 2}} \\ \gamma_{\mu\nu} &= \frac{\delta^2(\vec{\lambda}, \vec{\Delta\lambda}) - \gamma_{\mu\mu} \Delta\lambda^{\mu 2} - \gamma_{\nu\nu} \Delta\lambda^{\nu 2}}{2 \Delta\lambda^{\mu} \Delta\lambda^{\nu}}. \end{aligned} \quad (6)$$

The preliminary metric is calculated using the faster deterministic method. We then postprocess the metric in two steps. First, we minimize $\gamma_{\mu\nu} \Delta\lambda^{\mu} \Delta\lambda^{\nu}$ with respect to the time lag between signals $\Delta\lambda^0$ by projecting out the time component of the metric estimate. This results in the adjusted, spatial metric components

$$g_{ij} = \gamma_{ij} - \frac{\gamma_{0i} \gamma_{0j}}{\gamma_{00}}, \quad (7)$$

where we use the term *spatial* above to mean nontemporal, as in the familiar 3 + 1 decomposition. Second, we use an eigenvalue decomposition to check for numerical stability and validity of the estimated metric. If a negative eigenvalue is found, which would incorrectly imply a negative spatial signature, we attempt a reevaluation of the metric with a perturbed set of intrinsic parameters $\vec{\lambda} \rightarrow \vec{\lambda}'$. At the moment our perturbation simply removes precision from the definition of intrinsic parameters, e.g., changing 1.1234 to 1.123, etc., if after removing precision the metric evaluation still fails, we attempt to use the NUMDIFFTOOLS method. If that fails then the code would fatally crash. So far, we find that these numerical error mitigation methods are adequate but highlight it as an area for future improvement.

The template-bank algorithm then works as follows:

- (1) Initialize a hyper-rectangle bounding the parameter space one wishes to cover, e.g., a bounding box in component masses.
- (2) Compute the metric $\mathbf{g}(\vec{\lambda})$ numerically at the center of the hyper-rectangle. Alternatively, skip this step if the metric is sufficiently constant, the condition for which is described below.

- (3) From the metric, estimate the number of templates $\mathcal{N}_C$ needed to cover this hyper-rectangle via Eq. (4).
- (4) If $\mathcal{N}_C$ is greater than the user-supplied threshold $\mathcal{N}_C^*$, compute the side lengths of the hyper-rectangle according to the metric and split the cube along its largest side in two children cells A and B . Call the algorithm recursively on A and B . Choices of $\mathcal{N}_C$ are described in more detail below.
- (5) If $\mathcal{N}_C < \mathcal{N}_C^*$, place a template at the center of the cell and stop splitting. Choices $\mathcal{N}_C^*$ are described below.

The splitting stops when all rectangles have $\mathcal{N}_C < \mathcal{N}_C^*$ or alternatively if the user specifies a maximum coordinate volume. In Fig. 1, we illustrate the decomposition.

In step 2, we determine if a metric is sufficiently constant by placing a threshold on $\epsilon \equiv |1 - \frac{1}{2} \mathcal{V}_{i-2}/\mathcal{V}_{i-1}|$, where $\mathcal{V}$ is the proper volume of a cell. If the volume element of the previous two iterations ($i-2$, $i-1$) is sufficiently unchanged so that the cell and its parent's volume differ by approximately a factor of 2, the user may decide to skip this step. Setting epsilon to 0 forces the metric to be recomputed. The user also can add an additional maximum mismatch which must be achieved in order to skip recomputing the metric that is applied in addition to the constraint on ϵ . We find that for a bank targeting a 3% mismatch it is reasonable to skip metric reevaluation when a cell has achieved a 10% mismatch and the volumes are not changing by more than $\epsilon = \pm 10\%$. We note that skipping the metric evaluation is only used for computational efficiency. It can be disabled for accuracy and for typical template banks




FIG. 1. Example hyper-rectangle bifurcation in two arbitrary dimensions, x_1 , x_2 . Each boxed number represents a place where the metric, $\mathbf{g}$, was computed at the n th stage of the bifurcation. This example results in nine hyper-rectangles, which is less than the maximum value of 2^4 after four bifurcations.

recomputing every metric changes the run-time by at most a factor of a few.

In step 4, in most situations a user will set $\mathcal{N}_C^* = 1$ which is the condition for normal template bank generation, i.e., that one template should fit in the requested mismatch volume cell. However, it is useful to sometimes set $\mathcal{N}_C^*$ to be larger, e.g., 1000, in order to identify regions of the parameter space that would hold ~ 1000 templates at a requested match. This is useful to parallelize bank generation. By setting $\mathcal{N}_C^* = 1000$, one could quickly (in a matter of minutes) produce a “seed bank” that can define regions for further splitting. Through the use of high throughput computing resources, we have used this procedure to parallelize bank construction across a compute cluster with thousands of cores and reduce the time to create a template bank from hours down to minutes.

Other than waveform generation, the most computationally costly step of this process is the evaluation of the mismatch between two templates (2), which is needed to evaluate the metric coefficients (5). In the case where the template parameter space is bifurcated n times, there will be at most 2^n hyper-rectangles. If $\epsilon = 0$, then the metric will be evaluated for every cell and

$$\text{number of metric evaluations} = \sum_{i=0}^n 2^i = 2^{n+1} - 1. \quad (8)$$

Each metric evaluation requires $\mathcal{O}(D(D+1)/2)$ match calculations, where D is the dimension of the template parameter space, and the exact scaling depends on the finite differencing scheme chosen. This means that the total number of match calculations for a given bank assuming $\epsilon = 0$ is

$$\begin{aligned} \text{number of match evaluations} &= \frac{(2^{n+1} - 1)D(D+1)}{2} \\ &= \mathcal{O}(2^n D^2). \end{aligned} \quad (9)$$

Each hyper-rectangle will contain one template, which means that a well balanced tree will contain a bank of at most $\mathcal{N}_B = 2^n$ templates. Thus, the number of match calculations *per waveform* in the template bank is

$$\frac{\text{number of match evaluations}}{\text{number of templates } (\mathcal{N}_B)} = \mathcal{O}(D^2). \quad (10)$$

The above gives a worst case scenario. Under normal circumstances $\epsilon > 0$ and the metric is found to be sufficiently constant that it does not need to be evaluated at the final tree depth. This leads to typical scaling where there are *fewer match calculations than there are templates in the bank* $\mathcal{N}_B$.

By definition, the matches between waveforms used in the metric calculation are extremely high—approaching 1 minus floating point epsilon. Therefore, the function of




FIG. 2. Example template bank. This is a projection of the three-dimensional bank in coordinates $\{\log m_1, \log m_2, \chi_{\text{eff}}\}$ into the $\{\log m_1, \log m_2\}$ plane. The templates that appear to be outside of the region of interest have hyper-rectangles that overlap with the region. Note that the naive template density is directly related to local volume element magnitude, and varies accordingly.

frequency is extremely smooth and we evaluate waveforms and matches with extremely coarse spacing, typically 1 Hz.

IV. RESULTS

Software implementing the previously described algorithm has been developed in a package called MANIFOLD [34]. We used the MANIFOLD software to generate an advanced LIGO, Virgo and KAGRA template bank using projected O4 sensitivity estimates [35]. We used a chirp mass range from $0.87 - 174M_\odot$, a minimum secondary

mass of $0.98M_\odot$, a maximum mass ratio of 20 and a maximum total mass of $400M_\odot$. We specified an effective spin range, χ , from -0.99 to 0.99 but limited the spin of objects below $3M_\odot$ to be less than 0.05. We allowed the template low frequency to go down to 10 Hz, but specified a maximum duration of 128 s. We requested a maximum mismatch of 3%, but also set the maximum coordinate volume ($\Delta \log m_1 \times \Delta \log m_2 \times \Delta \chi$) to be less than 0.0001. This resulted in 2,083,547 templates as shown in Fig. 2.

We validated the template bank by injecting 16,000 simulated signals in the parameter space using the IMRPhenomD [32] approximant—the same approximant that was used for computing the template metric. This validation was done by choosing random points that satisfy the bounding-box constraints of the requested template bank and evaluating the overlap between those signals and the templates in the bank. The signals were chosen uniformly in the logarithm of component masses and uniformly in χ_{eff} . For each simulated signal, the best match was chosen among the nearest 30,000 templates in the bank defined by their absolute chirp mass. We find that the bank achieves the requested 97% match better than 99% of the time as shown in Fig. 3. The software required to perform this validation is included in the release described below.

It is important to validate the binary tree template bank algorithm against other template placement algorithms in the context of a specific gravitational wave search and with independent tools. For that reason, some of the authors here and others have written a second manuscript describing in great detail bank construction and validation for one of the LIGO, Virgo, and KAGRA compact binary searches using the methods presented in our work [36] which includes an independent validation of a binary tree template bank using




FIG. 3. Template bank validation. The color bar indicates mismatch of simulated signal and nearest template. The injected signals were created using uniform distributions of the individual parameters $\{\log m_1, \log m_2, \chi_{\text{eff}}\}$. The bank achieves the requested 97% match 99% of the time and a better than 98% match 90% of the time. The evaluation method used here is likely to be conservative since it does not maximize the match of all templates in the bank but only the 30,000 templates that are closest according to the absolute value of their chirp mass.

the `sbank` [9] method. They conclude broadly similar results to the stochastic template placement methods in both template count and bank efficacy. They also conclude that the validation software provided by MANIFOLD produces consistent results to the bank validation from `sbank` within the considered parameter space.

V. CONCLUSION

We have described here a new method for fast template bank placement, and shown that the method works in three dimensions relevant to dominant-mode aligned-spin template searches. The `treebank` method is computationally efficient and we expect this method will scale to higher dimensional template placement, such as precessing or subdominant mode templates, but we leave this for future work. It should also have applications in producing high density banks for use in rapid parameter estimation [31].

ACKNOWLEDGMENTS

This research has made use of data, software and/or web tools obtained from the Gravitational Wave Open Science Center [37], a service of LIGO Laboratory, the LIGO Scientific Collaboration and the Virgo Collaboration. We especially made heavy use of the LVK Algorithm

Library [38,39]. LIGO Laboratory and Advanced LIGO are funded by the United States National Science Foundation (NSF) as well as the Science and Technology Facilities Council (STFC) of the United Kingdom, the Max-Planck-Society (MPS), and the State of Niedersachsen/Germany for support of the construction of Advanced LIGO and construction and operation of the GEO600 detector. Additional support for Advanced LIGO was provided by the Australian Research Council. Virgo is funded, through the European Gravitational Observatory (EGO), by the French Centre National de Recherche Scientifique (CNRS), the Italian Istituto Nazionale di Fisica Nucleare (INFN) and the Dutch Nikhef, with contributions by institutions from Belgium, Germany, Greece, Hungary, Ireland, Japan, Monaco, Poland, Portugal, and Spain. This work was supported by National Science Foundation Awards No. OAC-1841480, No. PHY-2011865, and No. OAC-2103662. Computations for this research were performed on the Pennsylvania State University's Institute for Computational and Data Sciences gravitational-wave cluster, partially funded by OAC-2201445. C. H. acknowledges generous support from the Eberly College of Science, the Department of Physics, the Institute for Gravitation and the Cosmos, the Institute for Computational and Data Sciences, and the Freed Early Career Professorship.

-
- [1] B. S. Sathyaprakash and S. V. Dhurandhar, *Phys. Rev. D* **44**, 3819 (1991).
 - [2] B. J. Owen, *Phys. Rev. D* **53**, 6749 (1996).
 - [3] B. J. Owen and B. S. Sathyaprakash, *Phys. Rev. D* **60**, 022002 (1999).
 - [4] T. Cokelaer, *Phys. Rev. D* **76**, 102004 (2007).
 - [5] B. Abbott *et al.* (LIGO Scientific Collaboration), *Phys. Rev. D* **78**, 042002 (2008).
 - [6] I. W. Harry, B. Allen, and B. Sathyaprakash, *Phys. Rev. D* **80**, 104014 (2009).
 - [7] S. Babak, *Classical Quantum Gravity* **25**, 195011 (2008).
 - [8] G. M. Manca and M. Vallisneri, *Phys. Rev. D* **81**, 024004 (2010).
 - [9] P. Ajith, N. Fotopoulos, S. Privitera, A. Neunzert, and A. J. Weinstein, *Phys. Rev. D* **89**, 084041 (2014).
 - [10] S. Privitera, S. R. P. Mohapatra, P. Ajith, K. Cannon, N. Fotopoulos, M. A. Frei, C. Hanna, A. J. Weinstein, and J. T. Whelan, *Phys. Rev. D* **89**, 024003 (2014).
 - [11] C. Capano, I. Harry, S. Privitera, and A. Buonanno, *Phys. Rev. D* **93**, 124007 (2016).
 - [12] H. Fehrmann and H. J. Pletsch, *Phys. Rev. D* **90**, 124049 (2014).
 - [13] I. W. Harry, A. H. Nitz, D. A. Brown, A. P. Lundgren, E. Ochsner, and D. Keppel, *Phys. Rev. D* **89**, 024010 (2014).
 - [14] S. Roy, A. S. Sengupta, and N. Thakor, *Phys. Rev. D* **95**, 104045 (2017).
 - [15] S. Roy, A. S. Sengupta, and P. Ajith, *Phys. Rev. D* **99**, 024048 (2019).
 - [16] B. P. Abbott *et al.* (Virgo and LIGO Scientific Collaborations), *Phys. Rev. X* **6**, 041015 (2016).
 - [17] B. P. Abbott *et al.* (Virgo and LIGO Scientific Collaborations), *Astrophys. J.* **832**, L21 (2016).
 - [18] G. M. Harry (The LIGO Scientific Collaboration), *Classical Quantum Gravity* **27**, 084006 (2010).
 - [19] I. Harry, S. Privitera, A. Bohé, and A. Buonanno, *Phys. Rev. D* **94**, 024012 (2016).
 - [20] D. Mukherjee *et al.*, *Phys. Rev. D* **103**, 084047 (2021).
 - [21] J. Wang, TREEBANK: Differential geometric methods for fast template bank generation in searches for gravitational waves, Bachelor's thesis, The Pennsylvania State University, 2017.
 - [22] J. Aasi *et al.* (LIGO Scientific Collaboration), *Classical Quantum Gravity* **32**, 074001 (2015).
 - [23] F. Acernese *et al.* (Virgo Collaboration), *Classical Quantum Gravity* **32**, 024001 (2015).
 - [24] T. Akutsu *et al.*, *Prog. Theor. Exp. Phys.* **2021**, 05A101 (2020).
 - [25] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. X* **11**, 021053 (2021).

- [26] T. Dent and J. Veitch, *Phys. Rev. D* **89**, 062002 (2014).
- [27] H. K. Y. Fong, From simulations to signals: Analyzing gravitational waves from compact binary coalescences, Ph.D. thesis, University of Toronto (Canada), 2018.
- [28] R. Magee *et al.*, *Astrophys. J. Lett.* **878**, L17 (2019).
- [29] R. Abbott *et al.* (LIGO Scientific, Virgo, and KAGRA Collaborations), [arXiv:2111.03606](https://arxiv.org/abs/2111.03606) [*Phys. Rev. X* (to be published)].
- [30] K. Cannon, C. Hanna, and D. Keppel, *Phys. Rev. D* **85**, 081504 (2012).
- [31] C. Pankow, P. Brady, E. Ochsner, and R. O'Shaughnessy, *Phys. Rev. D* **92**, 023002 (2015).
- [32] P. Ajith, M. Hannam, S. Husa, Y. Chen, B. Brügmann, N. Dorband, D. Müller, F. Ohme, D. Pollney, C. Reisswig, L. Santamaría, and J. Seiler, *Phys. Rev. Lett.* **106**, 241101 (2011).
- [33] numdifftools, <https://pypi.org/project/numdifftools/> (2022).
- [34] <https://pypi.org/project/gwsci-manifold>.
- [35] M. Coughlin *et al.*, Noise curves for use in simulations pre-O4, Technical Report, LIGO Document T2200043-v3, 2022.
- [36] S. Sakon *et al.*, [arXiv:2211.16674](https://arxiv.org/abs/2211.16674).
- [37] <https://www.gw-openscience.org/>.
- [38] LSC Algorithm Library software packages lal, lalwrapper, and lalapps.
- [39] LIGO Scientific Collaboration, LSC Algorithm Library, <https://www.lsc-group.phys.uwm.edu/daswg/projects/lal.html>.