

Metric assisted stochastic sampling search for gravitational waves from binary black hole mergers

Chad Hanna,^{1,2,3,4} Prathamesh Joshi^{1,2,*}, Rachael Huxford,^{1,2} Kipp Cannon,⁵ Sarah Caudill,⁶ Chiwai Chan,⁵ Bryce Cousins,^{1,4} Jolien D. E. Creighton,⁷ Becca Ewing,^{1,2} Miguel Fernandez,^{1,2} Heather Fong,^{5,8} Patrick Godwin,^{1,2} Ryan Magee,⁹ Duncan Meacher,⁷ Cody Messick,¹⁰ Soichiro Morisaki,¹¹ Debnandini Mukherjee,^{1,2} Hiroaki Ohta,⁵ Alexander Pace,^{1,2} Stephen Privitera,¹² Surabhi Sachdev,^{1,2,9} Shio Sakon,^{1,2} Divya Singh,^{1,2} Ron Tapia,^{1,4} Leo Tsukada,^{5,8} Daichi Tsuna,⁵ Takuya Tsutsui,⁵ Koh Ueno,⁵ Aaron Viets,⁷ Leslie Wade,¹³ Madeline Wade,¹³ and Jonathan Wang¹⁴

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

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

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

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

⁵*RESCEU, The University of Tokyo, Tokyo 113-0033, Japan*

⁶*Nikhef, Science Park, 1098 XG Amsterdam, Netherlands*

⁷*Leonard E. Parker Center for Gravitation, Cosmology, and Astrophysics, University of Wisconsin–Milwaukee, Milwaukee, Wisconsin 53201, USA*

⁸*Graduate School of Science, The University of Tokyo, Tokyo 113-0033, Japan*

⁹*LIGO Laboratory, California Institute of Technology, MS 100-36, Pasadena, California 91125, USA*

¹⁰*LIGO Laboratory, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*

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

¹²*Albert-Einstein-Institut, Max-Planck-Institut für Gravitationsphysik, D-14476 Potsdam-Golm, Germany*

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

¹⁴*Department of Physics, University of Michigan, Ann Arbor, Michigan 48109, USA*



(Received 1 November 2021; accepted 4 October 2022; published 20 October 2022)

We present a novel gravitational-wave detection algorithm that conducts a matched-filter search stochastically across the compact binary parameter space rather than relying on a fixed bank of template waveforms. This technique is competitive with standard template-bank-driven pipelines in both computational cost and sensitivity. However, the complexity of the analysis is simpler, allowing for easy configuration and horizontal scaling across heterogeneous grids of computers. To demonstrate the method we analyze approximately one month of public LIGO data from July 27 00:00 2017 UTC–Aug 25 22:00 2017 UTC and recover eight known confident gravitational-wave candidates. We also inject simulated binary black hole signals to demonstrate the sensitivity.

DOI: [10.1103/PhysRevD.106.084033](https://doi.org/10.1103/PhysRevD.106.084033)

I. INTRODUCTION

Advanced LIGO directly detected gravitational waves (GWs) for the first time in 2015 from the merger of two black holes each about 30 times the mass of our Sun [1]. The second confident binary black hole (BBH) observation came just three months later [2]. Since then, the LIGO and Virgo Collaborations have detected a total of 90 compact binary mergers [3–6], including two neutron star mergers [7,8] and two neutron star–black hole mergers [9]. LIGO

and Virgo have made their data public [10], resulting in several new BBH discoveries by the community [11–16].

Historically, gravitational-wave searches for compact binary coalescence have relied on matched filtering [17–19], with several groups building on matched filtering as the foundation for their algorithms [14,20–24]. These techniques rely on fixed banks of templates [18,25,26] and are known to scale poorly to high-dimensional spaces [27]. Stochastic sampling methods were first proposed to address gravitational-wave detection in future searches for gravitational waves with the Laser Interferometer Space Antenna [28] but have not been widely used for detection in LIGO and Virgo data. Stochastic sampling techniques

*ppj5075@psu.edu

are, however, state of the art for the estimation of compact binary parameters once detections have been made [29,30].

In this work we blend aspects of traditional matched-filter searches, bank placement techniques, and stochastic sampling to create a new bankless matched-filter search for gravitational waves. While it remains to be seen what the broad applications of these techniques could be, we demonstrate a useful case study here by analyzing LIGO data from the Hanford and Livingston detectors from August 2017 [31] to search for binary black hole mergers. We recover eight known gravitational-wave candidates.

II. MOTIVATION

Our goal is to develop an *offline* compact binary search pipeline which is designed to detect gravitational waves in archival, LIGO, Virgo, and KAGRA data based on the GstLAL framework [32–35]. We distinguish that an offline analysis has less strict time-to-solution requirements (hours or days) compared to low-latency analysis where the time to solution needs to be seconds. We will not strive to reach the time-to-solution needs of low-latency analysis with the algorithm we present here. Our motivation for revisiting offline matched-filter detection for gravitational waves is to more easily parallelize and deploy analysis across heterogeneous resources such as multiple concurrent sites on the LIGO and Virgo data grids, the Open Science Grid, campus resources, and commercial clouds. We aim to achieve this by having a simpler workflow than competing pipelines such as GstLAL. We also wish to simplify the setup required to conduct an analysis and to improve usability for new researchers wanting to learn about gravitational-wave detection at scale. The intersection of these desires led us to consider new algorithmic approaches to searching the compact binary parameter space.

The Open Science Grid defines criteria for opportunistic computing as an application that “does not require message passing... has a small run time between 1 and 24 hours... can handle being unexpectedly killed and restarted...” and “... requires running a very large number of small jobs rather than a few large jobs” [36]. Our proposed workflow consists of parallel jobs that each search a small amount of gravitational-wave data from LIGO, Virgo and KAGRA without any interdependency between jobs. To contrast, the current GstLAL analysis workflow consists of a directed acyclic graph with more than ten levels of interdependent jobs. In this new approach, we target an ~ 1 – 12 h run time for each job, the use of one CPU core per job, and ~ 2 GB of RAM required per job in order to maximize throughput on opportunistic compute resources. Each job implements a flexible checkpointing procedure allowing work to be periodically saved.

III. METHODS

In this work, we will conduct a matched-filter search for binary black holes with the goal of identifying the maximum likelihood parameters for candidate events over 4 s coalescence-time windows using an analysis that foregoes the use of a precomputed template bank and instead employs stochastic sampling of the binary parameter space. Our workflow consists of two stages. The first stage executes N parallel jobs that conduct the bulk of the CPU-intensive work—in this study, this first stage consisted of 2974 such jobs. The results of these parallel jobs are returned to a single location at which point a second stage is run to combine results, assess candidate significance, estimate the search sensitivity and visualize the results. This second stage requires significantly lower computing power than the first stage, but is I/O intensive and is designed to be run potentially on local resources after grid jobs have completed.

In stage one, we begin by reading in gravitational-wave data from each observatory. Next, we measure the data noise power spectrum and whiten the data using the inferred spectrum. We then stochastically sample the data by proposing jumps governed by a parameter space metric as described in Sec. III D. For each jump, we generate the appropriate template waveform and then compute the matched-filter signal-to-noise ratio (SNR) over a 6 s stretch of time using 122 s of data per calculation.

Within a 4 s time window, we identify peaks in the matched-filter output, known as triggers, for each detector that is being analyzed. For each collection of triggers, we perform signal consistency checks [32] and calculate a likelihood ratio ranking statistic [37]. If the new sample has a larger SNR than the previous sample, it is stored—otherwise, a new jump is proposed.

A local estimate of the noise background is obtained by forming synthetic events from disjoint windows. This causes the time and phase difference between detectors of a single background event to be uniformly distributed, which is what we expect from noise events. This is a somewhat hybrid approach between the time-slide method [22] and sampling methods [38] already employed in GW searches. The second stage gathers the candidate events, results of the simulated GW search, and the background samples to produce a final summary view of the analysis results. In order to estimate the sensitivity of our methods to detecting gravitational waves, we conduct a parallel analysis over the same data with simulated signals added and repeat the same process as described above.

The remainder of this section describes key elements of our methods in more detail.

A. Data

We assume a linear model for the gravitational-wave strain data [17], s , which is a vector of discretely sampled time points for a gravitational-wave detector, j :

$$\mathbf{s}_j = \mathbf{h}(\vec{\theta}_j) + \mathbf{n}_j, \quad (1)$$

where $\mathbf{h}(\vec{\theta}_j)$ is an unknown gravitational waveform accurately modeled as a function of $\vec{\theta}_j = \{m_1, m_2, a_{1z}, a_{2z}, t_j, \phi_j, A_j\}$ ¹ with m_1 and m_2 being the component masses, a_{1z} and a_{2z} being the orbital-angular-momentum-aligned component spins and t_j , ϕ_j , and A_j being the time of coalescence, phase of coalescence, and amplitude, respectively, all of which depend on exactly where the binary is with respect to the j th gravitational-wave detector. $\mathbf{n}_j$ is a realization of detector noise. As a concrete example, in this work each job analyzes 800 s stretches of data, divided into 4 s windows sampled at 2048 Hz. Thus, after including the Fourier transform block length (124 s), the dimension of each vector in the work described in this manuscript is 262144 sample points. In addition, each job also contains start padding (128 s) and stop padding (32 s). The templates have at least 6 s of zero padding, which makes their length no more than 122 s.

We assume that the noise samples are entirely uncorrelated between the gravitational-wave observatories but that the signals are correlated between observatories. In fact, we make the simplifying assumption that the gravitational waveform is identical between detectors except for an overall amplitude A_j , time shift Δt_j , and phase shift $\Delta \phi_j$ [20]:

$$\mathbf{h}(\vec{\theta}_j) = \text{Re}(\mathcal{F}^{-1}[A_j e^{2\pi i f \Delta t_j + \Delta \phi_j} \mathcal{F}[\mathbf{h}(\vec{\lambda}, t, \phi)]]), \quad (2)$$

where $\mathcal{F}[\cdot]$ denotes the unitary Fourier transform and $\vec{\lambda} = \{m_1, m_2, a_{1z}, a_{2z}\}$.

The exact realization of noise $\mathbf{n}$ is not possible to predict, but we will assume it is well characterized as a multivariate normal distribution with a diagonal covariance matrix in the frequency domain, i.e., that it is stationary. However later on, particularly in Sec. III D 6, we account for the fact that the data are often not stationary.

B. Spectrum estimation and whitening

We rely on the same spectrum estimation methods as described in [32]. Namely, we use a median-mean, stream-based spectrum estimation technique that adjusts to changes in the noise spectrum on $\mathcal{O}(\min)$ timescales. The data are divided into 8 s blocks with 6 s overlap, and the spectrum $\mathbf{S}_n$ is estimated by windowing the input blocks with 2 s of zero padding on each side of the window. Since we analyze only 800 s of data per job, we use a fixed spectrum over the job duration.

¹These parameters are adequate to describe the measurable gravitational-wave parameters for a nonprecessing, circular binary black hole system with only 2–2 mode emission in a single gravitational-wave detector.

From here forward, we will work in a whitened basis for the data, namely, that

$$\mathbf{s}_j \rightarrow \mathcal{F}[\mathbf{s}] \circ (\mathbf{S}_n)^{-1/2}, \quad (3)$$

which implies that all components of $\mathbf{s}$ are transformed by the inverse noise amplitude spectrum. Therefore, if the amplitude of $\mathbf{h}$ is zero, $\mathbf{s}$ has components that satisfy $p(s_i) = (2\pi)^{-1/2} e^{-s_i^2/2}$ with $\langle s_i, s_j \rangle = \delta_{ij}$. In this whitened data basis, an inner product between two vectors is the dot product $\mathbf{u} \cdot \mathbf{v}$, and unit vectors are denoted as $\hat{\mathbf{u}}$. We adopt a normalization such that $\mathbf{h} \cdot \mathbf{h} = 1$ and $\langle \mathbf{n} \cdot \mathbf{n} \rangle = \dim \mathbf{n}$. With these choices the SNR is given by $\rho(\vec{\theta}_j) = \mathbf{h}(\vec{\theta}_j) \cdot \mathbf{s}_j$. We can evaluate the SNR for the unknown phase and time of coalescence by defining a complex SNR:

$$\rho(\vec{\lambda}, t_j, \phi_j) = \mathcal{F}^{-1}[\mathbf{h}(\vec{\lambda}) \cdot \mathbf{s}_j] + i \mathcal{F}^{-1}[\mathbf{h}(\vec{\lambda}, \pi/2) \cdot \mathbf{s}_j], \quad (4)$$

which is a *valid* matched-filter output for a duration of time equal to the length of the data minus the length of the template. With at least 6 s of zero padding, the template length is 122 s, and with each window using 128 s of data, the matched-filter output is valid for a duration of 6 s.

C. Simulation capabilities

We use the GstLAL data source module [35], which provides an interface into the LAL Simulation package [39]. By providing a LIGO-LW XML format document containing simulation parameters, we can inject simulated strain into each of the currently operating ground-based gravitational-wave detectors, LIGO, Virgo, KAGRA and GEO-600. When operating the pipeline in a simulation mode, gravitational-wave events are reconstructed around a ± 2 s interval around the GPS second of the geocentric arrival time of the gravitational-wave peak strain.

D. Parameter space sampling

The gravitational-wave parameter space is explored stochastically, with Gaussian jump proposals and refinement steps that gradually reduce the jump size as the peak in SNR is identified. We will refer to this procedure as “sampling.” Our proposal distribution has a covariance matrix that depends on the location in the parameter space and the refinement level. It relies on computing the parameter space metric $\mathbf{g}$ [18], which is described more in the next section. We define a sequence of two parameters that control how the sampling is done, namely, σ_k , which controls the jump size, and N_k , which controls the number of samples to reject at each level, k , before moving on to the next. How exactly to define these parameters is certainly a topic for future research. Our choices here were determined empirically for the particular search we have done. We define

$$\sigma_k = 10^{1-k}, \quad (5)$$

$$N_k = \begin{cases} 500 & (k = 0), \\ 100 & (k > 0), \end{cases} \quad (6)$$

for $k = 0 \dots n$, where k is terminated based on the mismatch as in step 7 below. We define the characteristic jump proposal distance as

$$\delta_k(\vec{\lambda}) := \sigma_k |\mathbf{g}(\vec{\lambda})|^{1/8}, \quad (7)$$

where $\vec{\lambda}$ is the set of intrinsic parameters as defined before and δ_k^2 is the template mismatch.

Gravitational waves are searched over 4 s *windows* of coalescence time using the following procedure.

- (1) Establish a bounding box in the physical parameter space to search over.
- (2) Pick a starting parameter point somewhere in the middle of the parameter space. We use the approximate expression for template count in [18] to estimate a good central point.
- (3) Evaluate the SNR at this point and set a counter to zero.
- (4) Sample from a sampling function $\Theta(\delta_k, \vec{\lambda})$, which is described in detail below in Sec. III D 4.
- (5) Check that the new point lies within boundaries established in step 1 and apply any constraint functions. If the point fails to fall within the constraints, go to step 4.
- (6) Evaluate the SNR at the new point. If the point has a higher SNR than the previous sample, update the sample and reset the counter to zero. If the point has a lower SNR, increment the counter.
- (7) If the counter is less than N_k , go back to step 4. If the counter is greater than or equal to N_k , check δ^2 , where δ^2 is the template mismatch between the current and previous sample point. If $\delta^2 < 0.1$, terminate the sampling. Otherwise, increment k , reset the counter, and proceed to step 4.

1. Computation of the binary parameter space metric

We define the match between adjacent compact binary waveforms in the space of intrinsic parameters as

$$m(\vec{\lambda}, \vec{\lambda} + \vec{\Delta\lambda}) = \max_{\phi_c, t_c, A} [\hat{\mathbf{h}}(\vec{\lambda}) \cdot \hat{\mathbf{h}}(\vec{\lambda} + \vec{\Delta\lambda})], \quad (8)$$

where the maximum is over extrinsic parameters $\{t_c, \phi_c, A\}$. Note that $m(\vec{\lambda}, \vec{\lambda}) = 1$. We also introduce a shorthand for computing the match along a deviation in only one coordinate as

$$m(\vec{\lambda}, \vec{\lambda} + \vec{\Delta\lambda}_i) = \max_{\phi_c, t_c, A} [\hat{\mathbf{h}}(\vec{\lambda}) \cdot \hat{\mathbf{h}}(\vec{\lambda} + \vec{\Delta\lambda}_i)], \quad (9)$$

where it is assumed that $\Delta\vec{\lambda}_i$ is nonzero only along a given coordinate direction.

It has previously been shown [18] that is possible to derive a metric on the space of intrinsic parameters describing compact binary waveforms by expanding our definition of the match locally, e.g., about $\Delta\lambda = 0$, as follows:

$$m(\vec{\lambda}, \vec{\lambda} + \vec{\Delta\lambda}) \approx 1 + \frac{1}{2} \frac{\partial^2}{\partial \Delta\lambda_i \partial \Delta\lambda_j} m(\vec{\lambda}, \vec{\lambda} + \vec{\Delta\lambda})|_{\vec{\Delta\lambda}=0} \Delta\lambda_i \Delta\lambda_j, \quad (10)$$

which suggests the metric

$$g_{ij}(\vec{\lambda}) = -\frac{1}{2} \frac{\partial^2}{\partial \Delta\lambda_i \partial \Delta\lambda_j} m(\vec{\lambda}, \vec{\lambda} + \vec{\Delta\lambda})|_{\vec{\Delta\lambda}=0}. \quad (11)$$

The mismatch between templates, $\delta^2 = 1 - m$, becomes

$$\delta(\vec{\lambda}, \vec{\Delta\lambda})^2 \approx \vec{\Delta\lambda}^T \mathbf{g}(\vec{\lambda}) \vec{\Delta\lambda}. \quad (12)$$

In this work, the components of the metric are evaluated with second-order finite differencing:

$$g_{ii}(\vec{\lambda}) = -\frac{1}{2} \left[\frac{m(\vec{\lambda}, \vec{\lambda} + \vec{\Delta\lambda}_i) + m(\vec{\lambda}, \vec{\lambda} - \vec{\Delta\lambda}_i) - 2}{|\vec{\Delta\lambda}_i|^2} \right] \quad (13)$$

and

$$g_{ij}(\vec{\lambda}) = -\frac{1}{2} \frac{1}{4|\vec{\Delta\lambda}_i||\vec{\Delta\lambda}_j|} [m(\vec{\lambda}, \vec{\lambda} + \vec{\Delta\lambda}_i + \vec{\Delta\lambda}_j) - m(\vec{\lambda}, \vec{\lambda} + \vec{\Delta\lambda}_i - \vec{\Delta\lambda}_j) - m(\vec{\lambda}, \vec{\lambda} - \vec{\Delta\lambda}_i + \vec{\Delta\lambda}_j) + m(\vec{\lambda}, \vec{\lambda} - \vec{\Delta\lambda}_i - \vec{\Delta\lambda}_j)], \quad (14)$$

for the off diagonal terms. However, we use a more efficient formula for the off diagonal terms, in which the number of template evaluations is the same, but the number of match calculations is reduced:

$$g_{ij}(\vec{\lambda}) = -\frac{1}{2} \frac{1}{2|\vec{\Delta\lambda}_i||\vec{\Delta\lambda}_j|} [m(\vec{\lambda}, \vec{\lambda} + \vec{\Delta\lambda}_i + \vec{\Delta\lambda}_j) - m(\vec{\lambda}, \vec{\lambda} + \vec{\Delta\lambda}_i) - m(\vec{\lambda}, \vec{\lambda} + \vec{\Delta\lambda}_j) + 2 - m(\vec{\lambda}, \vec{\lambda} - \vec{\Delta\lambda}_i) - m(\vec{\lambda}, \vec{\lambda} - \vec{\Delta\lambda}_j) + m(\vec{\lambda}, \vec{\lambda} - \vec{\Delta\lambda}_i - \vec{\Delta\lambda}_j)]. \quad (15)$$

The sampling method described in Sec. IV below will not make jumps in coalescence time; therefore, the time component is projected out [18]:

$$g_{ij}(\vec{\lambda}) \rightarrow g_{ij}(\vec{\lambda}) - \frac{g_{ti}(\vec{\lambda})g_{tj}(\vec{\lambda})}{g_{tt}(\vec{\lambda})}. \quad (16)$$

2. Choice of coordinates

We sought out a coordinate system that maps the masses and spins to be in the interval $[-\infty, \infty]$. We also want to choose well-measured physical parameters for mass and spin in at least one dimension each. Therefore, we use the following coordinates to evaluate the metric:

$$\lambda_1 = \log_{10} \left[\frac{(m_1 m_2)^{3/5}}{(m_1 + m_2)^{1/5}} \right], \quad (17)$$

$$\lambda_2 = \log_{10}(m_2), \quad (18)$$

$$\lambda_3 = \tan \left[\left(\frac{\pi}{2} \right) \left(\frac{a_{1z} m_1 + a_{2z} m_2}{m_1 + m_2} \right) \right], \quad (19)$$

$$\lambda_4 = \tan \left[\left(\frac{\pi}{2} \right) \left(\frac{a_{1z} m_1 - a_{2z} m_2}{m_1 + m_2} \right) \right]. \quad (20)$$

3. Pathologies of the numerical metric

For certain regions of the parameter space the metric is nearly singular, which leads to numerical errors causing a nonpositive-definite matrix. To fix this, we conduct an eigenvalue decomposition of g_{ij} :

$$g_{ij} = q_{ik} \beta^k q_{kj}^{-1}. \quad (21)$$

We then define a new set of eigenvalues:

$$\beta_{\min} \equiv \max_k [\beta^k] \times \epsilon, \quad (22)$$

$$(\beta^k)' = \begin{cases} \beta_{\min}, & \beta^k < \beta_{\min}, \\ \beta^k, & \text{otherwise,} \end{cases} \quad (23)$$

where ϵ is a parameter which we will call the aspect ratio. We define the new metric as

$$g_{ij} \rightarrow q_{ik} (\beta^k)' q_{kj}^{-1}. \quad (24)$$

In practice we find that sampling is better when we artificially distort the metric by setting $\epsilon = 0.1$ for the broadest refinement level and $\epsilon = 0.0001$ for all other levels, and we have done so in this work, though this should be a direction of future work.

4. Drawing random samples from $\Theta(\delta_k, \vec{\lambda})$

When sampling, we desire to have a jump proposal distribution that effectively probes the space by not making jumps that are either too near or too far. The calculation of the parameter space metric $\mathbf{g}$ enables that. We wish to propose a jump from $\vec{\theta} \rightarrow \vec{\theta} + \Delta\vec{\theta}$ such that the expected mismatch is δ^2 . The metric described in previous sections only applies to the intrinsic parameters. For the extrinsic

parameters, our jump proposal will always choose those values of t and ϕ which maximize the SNR. At every accepted jump point, the metric is calculated locally, which requires 21 template evaluations, including the diagonal and off diagonal terms, as specified in Sec. III D 1. However, we can afford to calculate coarse versions of the template waveform, since the match we need to calculate is between two adjacent templates. This means the waveform calculation cost is not high. The distance between adjacent templates to calculate the match at $\Delta\vec{\lambda}_i$ as defined in Sec. III D 1 is hard coded and is the same for all iterations of the sampling procedure.

To facilitate jumping in the intrinsic parameters, we make a coordinate basis transformation in which the new basis has a Euclidian metric. The transformation matrix $\mathbf{M}$ will then be used to transform the coordinates

$$\vec{\lambda}' = \mathbf{M} \vec{\lambda}. \quad (25)$$

To solve for $\mathbf{M}$ we rely on the fact that distance is invariant, giving

$$\delta^2 = \vec{\Delta\lambda}^T \mathbf{g} \vec{\Delta\lambda} \quad (26)$$

$$= (\vec{\Delta\lambda}')^T \mathbf{g}' \vec{\Delta\lambda}' \quad (27)$$

$$= (\vec{\Delta\lambda}')^T \mathbf{M}^T \mathbf{g} \mathbf{M} \vec{\Delta\lambda}'. \quad (28)$$

Setting $\mathbf{g}' = \mathbf{I}$ gives

$$\mathbf{I} = \mathbf{M}^T \mathbf{g} \mathbf{M}, \quad (29)$$

$$\mathbf{M}^{-1} (\mathbf{M}^T)^{-1} = \mathbf{g}, \quad (30)$$

$$\mathbf{g}^{-1} = \mathbf{M}^T \mathbf{M}. \quad (31)$$

The last line implies that we can solve for $\mathbf{M}$ by taking the Cholesky decomposition of the inverse metric tensor. Once obtaining $\mathbf{M}$ we can produce random samples with an expected mismatch by defining

$$\Theta(\delta_k, \vec{\lambda}) := \delta_k, \quad \vec{\lambda} \rightarrow \vec{\lambda} + \delta_k \mathbf{M}^T \vec{\mathcal{N}}, \quad (32)$$

where $\vec{\mathcal{N}}$ is a four-dimensional vector with random components satisfying $p(\mathcal{N}_i) = \sqrt{1/2\pi} \exp[-\mathcal{N}_i^2/2]$.

5. Parameter space constraints

The previously defined sampling function can produce samples that, while physical, may be outside of the desired search range. We implement a series of user-defined constraints that will reject samples drawn from $\Theta(\delta_k, \vec{\lambda})$. These are as follows.

- m_1, m_2, a_1, a_2 .**—The user can specify a bounding box in component masses and z -component spins. Samples outside this bounding box are rejected.
- η .**—The user can specify a minimum symmetric mass ratio $\eta \equiv (m_1 m_2)/(m_1 + m_2)^2$, below which samples will be rejected.
- $\mathcal{M}$.**—The user can specify a chirp mass range, outside of which samples will be rejected.

6. Glitch rejection

Glitches [40,41] are nonstationarity and non-Gaussian transient noise artifacts of instrumental or environmental origin found in the data. We employ a novel data-driven technique to reject short-duration glitches, using two parameters, the bandwidth and the effective spin parameter χ . The bandwidth is the standard deviation of frequency weighted by template amplitude. It is defined as [42]

$$\text{bandwidth}^2 = \frac{\int |\hat{\mathbf{h}}(\vec{\lambda})|^2 f^2 df / \mathbf{S}_n}{\int |\hat{\mathbf{h}}(\vec{\lambda})|^2 df / \mathbf{S}_n} - \left(\frac{\int |\hat{\mathbf{h}}(\vec{\lambda})|^2 f df / \mathbf{S}_n}{\int |\hat{\mathbf{h}}(\vec{\lambda})|^2 df / \mathbf{S}_n} \right)^2, \quad (33)$$

whereas χ is defined as

$$\chi = \frac{m_1 a_{1z} + m_2 a_{2z}}{m_1 + m_2}. \quad (34)$$

It has been found that short-duration glitches ring up templates which exclusively occupy the low bandwidth–low χ region in bandwidth– χ space and that this region is not occupied by gravitational-wave signals. This is illustrated in Fig. 1. As part of the simulation campaign we performed (refer to Sec. IV C for details), we found that

only 28 injections out of 112526 fell into the glitch region. Minimizing this number by fine-tuning the boundary of the glitch region would be a direction for future work. We define the glitch region as

$$\text{bandwidth} \times (1 + \chi) \leq 20. \quad (35)$$

Any trigger which falls in this region is not considered as a gravitational-wave candidate. Similarly, any time-slid background samples falling in the glitch region are eliminated and not used for background estimation. Triggers are explained in more detail in the next subsection, whereas background estimation is explained in Sec. III E.

7. Computing the log-likelihood ratio $\mathcal{L}$

We generally follow the same procedure for ranking candidates as described in [37,38,43] with a couple of notable exceptions. First, we only implement a subset of the terms used in the GstLAL-inspiral pipeline—it will be the subject of future work to include more. Second, we approximate some of the data-driven noise terms with analytic functions. Third, we adopt a normalization so that for signals the log-likelihood ratio $\mathcal{L}$ is approximately $\rho^2/2$, where ρ is the network matched-filter SNR defined as the square root of the sum of the squares of the SNRs found in each observatory. We use the following terms in the log-likelihood ratio.

$\mathcal{L}(\vec{\rho}, \vec{\xi}^2)$.—We approximate this term of the log-likelihood ratio as

$$\mathcal{L}(\vec{\rho}, \vec{\xi}^2) = \sum_i \mathcal{L}_i(\rho_i, \xi_i^2), \quad (36)$$

$$\text{with } \mathcal{L}_i(\rho_i, \xi_i^2) = \rho_i^2 e^{-4\xi_i^2/2} - 4\xi_i^2, \quad (37)$$




FIG. 1. Triggers found in one month's data for Hanford (left) and Livingston (right), colored by their log-likelihood ratio. All the bright points to the right of the boundary are known gravitational-wave candidates, and all those to the left of the boundary are glitches and so not considered gravitational-wave candidates and not used for background estimation.

where $x_i \equiv \max\{0, \xi_i^2 - 1 - 0.0005\rho_i^2\}$ for each of the i detectors which are assumed to be independent. The $-4x_i^2$ term acts as a penalty for high ξ_i^2 values and helps eliminate glitches.

$\mathcal{L}(\Delta\vec{t}, \Delta\vec{\phi}, \Delta\vec{\rho})$.—For this term we follow the procedure in [43] with two changes. We do not include the ρ^{-4} term. We do this because we are not constructing a data-driven noise term like GstLAL-inspiral, so it is not necessary to have the corresponding signal term. We also normalize the result to be 0 when only one detector is operating. This is useful for achieving the normalization discussed above. These changes have the effect of making this term ~ 0 for things that are consistent with signals.

$\mathcal{L}(\vec{T})$.—This term quantifies the probability of having “triggered” the combination of the gravitational-wave detectors in which the event was found and is a function of the detectors’ sensitivity. We will describe triggering in more detail below. For example, it is unlikely that only the least-sensitive detector would be triggered for a real gravitational-wave event, so this term would be negative in that case. This term is complementary to the previous term but accounts for events lacking triggers.

$\mathcal{L}(\vec{D}_H)$.—This term quantifies the relative likelihood of detecting an event based on the detector horizon BNS distances $(D_H)_i$. We normalize to the horizon distance of LIGO Livingston during O3–315 Mpc:

$$\mathcal{L}(\vec{D}_H) = \ln \left(\frac{\max_i \{(D_H)_i\}}{315 \text{ Mpc}} \right)^3. \quad (38)$$

The log-likelihood ratio $\mathcal{L}$ is then given by

$$\mathcal{L} = \mathcal{L}(\vec{\rho}, \vec{\xi}^2) + \mathcal{L}(\Delta\vec{t}, \Delta\vec{\phi}, \Delta\vec{\rho}) + \mathcal{L}(\vec{T}) + \mathcal{L}(\vec{D}_H). \quad (39)$$

For each sample drawn in step 4, we construct a template waveform $\mathbf{h}(\vec{\lambda})$ and filter that waveform against the data in each detector stream producing an SNR time series over a 6 s period, including 1 s padding on either side. We then find the peak SNR in the middle 4 s window in each detector and record the time, phase, SNR, and ξ^2 of each peak, which we call a “trigger.” For the collection of triggers, we cycle through every detector combination—for example, if analyzing $\{H, L, V\}$, we cycle through $\{HLV, HL, HV, LV, H, L, V\}$ and evaluate the likelihood ratio for each combination. We then keep the maximum $\mathcal{L}$ found over these detector combinations. This is done to mitigate the effect of bad data (noisy data and possibly also glitchy data) in one detector. Hence, triggers are obtained by maximizing SNR over 4 s windows, whereas the detectors to be considered for the trigger are obtained by

maximizing the likelihood ratio over all possible detector combinations. Note that the SNR maximization for updating the sample discussed in step 6 is a separate procedure from either of these.

E. Background estimation

We treat windows recovered as single triggers and windows recovered as coincidences differently while estimating the background. For single trigger windows, the foreground sample itself is used as the background sample representing that window. To estimate the coincident background, we form false coincidences from a given job which analyzes 800 s of data in 200 coalescence time windows. To form false coincidences, we shift the windows in time with respect to each other. We then draw samples randomly from all single detector triggers. For each recovered false coincidence, we compute a $\mathcal{L}$ and histogram the result. This process is then repeated 100 times with different time offsets to increase the amount of background we have. This background is given an appropriate weight so that the ratio of singles to coincidences in the background and foreground is the same, as well as to ensure that the background is normalized. Using the $\mathcal{L}$ histogram for the background, false alarm rates (FARs) are assigned to all the triggers. One point to note is that the windows in which we detect events are not used to form combinations so as to not contaminate the background with signals.

IV. RESULTS

A. Dataset

We analyze public gravitational-wave data from LIGO taken from July 27 00:00 2017 UTC–Aug 25 22:00 2017 UTC during advanced LIGO’s second observing run. We choose segments of data with a minimum length of 1200 s for each of the LIGO detectors. From those segments we form coincident segments. Jobs require 128 s of start padding, 32 s of end padding and 124 s for the Fourier transform block to produce triggers. Thus, each job can analyze a minimum of 288 s (which produces triggers for a single 4 s window) and we choose a maximum duration of 1084 s to produce 800 s of triggers over 200 windows. Jobs are overlapped so that triggers are produced contiguously.

After accounting for the segment selection effects, we analyzed approximately 20.17 days of coincident data.

B. Search parameter space

We search for gravitational-wave candidates with component masses between 0.9 and 400 $M_\odot$ with z -component spins between -1 and 1 . We conduct the matched-filter integration between 10 and 1024 Hz.




FIG. 2. Distribution of component masses as measured at Earth for the BBH simulation set.

C. Simulation set

In order to ascertain our sensitivity to gravitational-wave signals of the type discovered in this data, we conducted a simulation campaign with 112526 simulated signals having a $32 M_{\odot}$ mean component mass and standard deviation of $4.0 M_{\odot}$ with aligned dimensionless spins up to 0.25 and a maximum redshift of 1 isotropically distributed in location. The injections were distributed uniformly in comoving volume. The redshifted component mass distribution is visualized in Fig. 2. The injection set was specifically created for the BBH parameter space. We do not make any claims about the sensitivity of our pipeline in other lower mass parameter spaces.

D. Candidate list

Our search results are summarized in Fig. 3 and Table I. Results from the entire search are shown in Fig. 3. In this plot, we show the observed distribution of all events as a function of $\sqrt{2\mathcal{L}}$, an expression proportional to the SNR, as well as the background distribution expected from noise during the same time. The detected events clearly stand out from the expected noise curve at $\sqrt{2\mathcal{L}}$ around 8, which suggests that the extra events at high $\mathcal{L}$ must be signal-like.

In Table I, we report the ten triggers with the smallest FARs. The first five of these events as well as the seventh were previously reported by the LIGO Collaboration and others [3,12,44] and labeled GW170817, GW170814, GW170809, GW170823, GW170729, and GW170818. These events are detected confidently with FARs of $5 \times 10^{-3} \text{ yr}^{-1}$ for the first five and $4 \times 10^{-2} \text{ yr}^{-1}$ for the seventh. GW170817 is recovered as a single detector




FIG. 3. Cumulative histograms of our search results as a function of likelihood ratio. The orange line represents the corresponding histogram expected from noise during the same time frame.

candidate in Hanford, since there is a simultaneous glitch in Livingston, and the resulting high ξ^2 in Livingston causes its log-likelihood ratio to be strongly penalized. We report many of the components masses of these events outside of confidence ranges reported by the LIGO Collaboration [3]. It is important to note that this is not a contradiction: we are not optimizing the posterior probability distribution, as is done during parameter estimation for the results reported by the LIGO Collaboration. Despite the differences in masses, we are able to recover each trigger to within tens of milliseconds of the reported values by the LIGO Collaboration and are confident they correspond to the respective gravitational-wave candidates.

We also recover one binary black hole event, GW170727, previously reported by other groups [12,44] as well as one, GW170817a, reported by Zackay *et al.* [11] which do not appear in the LIGO GWTC-2. We recover the GW170727 event with a FAR of $3 \times 10^{-1} \text{ yr}^{-1}$. We recover GW170817a in Livingston with a FAR of $5 \times 10^{-3} \text{ yr}^{-1}$ while Zackay reports it with a FAR of $8.7 \times 10^{-2} \text{ yr}^{-1}$. Zackay also reports the probability of it being of astrophysical origin at 86% [11], but we do not make that estimation here. As in the previous case, we recover both these events to within tens of milliseconds of the previously reported values and are confident that they correspond to the respective gravitational-wave candidates.

We make no claims regarding the possibility of the remaining two events we report being gravitational-wave candidates. These appear eighth and ninth in Table I. They are not recovered significantly, and it is likely they are noise.

TABLE I. Candidate gravitational-wave events with the ten smallest false alarm rates and largest SNRs. The first seven triggers as well as the last one correspond to known gravitational-wave candidates: GW170817, GW170814, GW170809, GW170823, GW170729, GW170817a, GW170818, and GW170727. The other two triggers have not been previously reported as gravitational-wave candidates.

FAR (yr ⁻¹)	$\sqrt{2\mathcal{L}}$	ρ_{net}	$m_1 (M_\odot)$	$m_2 (M_\odot)$	a_1	a_2	ρ_{H1}	ξ_{H1}^2	ρ_{L1}	ξ_{L1}^2	Date (UTC)
5×10^{-03}	18.5	18.7	1.8	1.1	-0.3	0.7	18.7	0.8	...	...	2017-08-17 12:41:04
5×10^{-03}	16.3	17.1	38.7	24.5	0.7	-0.9	9.6	1.3	14.1	1.0	2017-08-14 10:30:43
5×10^{-03}	11.9	12.6	46.4	25.8	0.6	-1.0	6.5	1.3	10.8	0.7	2017-08-09 08:28:21
5×10^{-03}	11.7	11.8	51.5	38.6	0.4	-0.5	6.6	0.9	9.8	0.7	2017-08-23 13:13:58
5×10^{-03}	10.7	10.9	73.1	43.3	-0.1	1.0	7.9	1.1	7.5	1.0	2017-07-29 18:56:29
5×10^{-03}	10.6	10.7	122.7	45.5	0.9	-0.9	...	...	10.7	1.0	2017-08-17 03:02:46
4×10^{-02}	9.6	10.1	39.7	36.3	0.7	-0.8	...	...	10.1	1.2	2017-08-18 02:25:09
$6 \times 10^{+00}$	8.7	9.0	20.8	3.3	0.1	0.6	9.0	1.0	...	...	2017-08-03 05:59:03
$3 \times 10^{+01}$	8.6	8.8	53.1	1.2	-0.3	1.0	3.7	0.9	8.0	0.8	2017-08-14 07:35:04
$3 \times 10^{+01}$	8.5	8.7	51.5	43.6	0.4	-0.9	4.6	0.8	7.4	1.1	2017-07-27 01:04:30

The first seven events reported in Table I, as well as the last one, are excluded from the background, since all of them are previously reported gravitational-wave candidates.

E. Sensitivity estimate

The sensitivity of our new pipeline is demonstrated in Figs. 4 and 5. Figure 4 shows the distribution of all the injected events by SNR with a network SNR of 12 contours and detector SNR of seven contours added. This figure shows that the majority of loud injected events were recovered by our pipeline, with 70 missed in the region




FIG. 4. Distribution of injected SNRs for recovered injections above $\mathcal{L} = 35(\sqrt{2\mathcal{L}} = 8.37)$. Missed injections with network SNR above 12 and detector SNRs greater than 7 (indicated by the shaded contour) are discussed in Appendix B.

with network SNR above 12 and detector SNRs greater than 7. Only nine of these missed injections are because the pipeline could not adequately recover the injections. This shows that the pipeline only very rarely gets stuck at local peaks, instead of finding the global maxima, which will correspond to the injected signal. It is possible that as we move to a lower mass parameter space, the frequency of such occurrences will increase. All of the loud missed events are discussed in more detail in Appendix B.

Figure 5 shows the efficiency of the pipeline as a function of the injected network SNR of the synthetic gravitational-wave set described in Sec. IV C. This plot shows that without any data cleaning implementation, almost 90% of events at SNR 10 are recovered by the pipeline while that percentage only increases with the SNR and plateaus just short of 100% around SNR 13.




FIG. 5. Efficiency of recovering injections at different injected SNRs.

V. CONCLUSION

In this paper, we have described in detail a novel gravitational-wave detection algorithm. This algorithm searches stochastically over the chosen parameter space, saving the time and computing power required to generate large banks of template waveforms. The algorithm samples the parameter space by making jumps with a preestimated mismatch between templates informed by the parameter space metric and keeping those points which have a higher SNR. This method is shown to be of comparable accuracy in the recovery of gravitational-wave events at high masses as current template-based pipelines.

To demonstrate the validity of this method, we have presented an analysis of approximately one month of LIGO data from July 27 00:00 2017 UTC–Aug 25 22:00 2017 UTC exploring the binary black hole parameter space. We recovered six known gravitational-wave candidate events to within tens of milliseconds of previously reported coalescence times, as well as two gravitational-wave candidates previously reported.

Additionally, we conducted an injection campaign of compact binary mergers to prove the sensitivity of the pipeline to binary black hole merger events. We recovered almost 90% of events with SNR 10 and an increasing percentage at higher SNRs that plateaus just below 100% at SNR 13. The majority of the missing loud injections were due to the presence of glitches near the injected events.

In the future, we plan to extend our method to all regions of the parameter space. We expect that even though the algorithm will scale similarly to any search using template banks at lower mass, it will still retain its other advantages, such as simpler workflow and ease of setup. We plan to make our method competitive with other searches like GstLAL for LIGO’s fourth observing run. It remains an open project to get good convergence during the sampling process for all regions of the parameter space.

ACKNOWLEDGMENTS

This research has made use of data, software and/or web tools obtained from the Gravitational Wave Open Science Center, a service of LIGO Laboratory, the LIGO Scientific Collaboration and the Virgo Collaboration. 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 Grants No. OAC-1841480, No. PHY-2011865, and No. OAC-2103662. Computations for this research were performed on Pennsylvania State University’s Institute for Computational and Data Sciences gravitational-wave cluster. 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.

APPENDIX A: DATA RELEASE DETAILS AND CODE VERSIONS

A tarball containing the source code and data files necessary to reproduce the results and plots in this paper is available [45]. Instructions for installing the code and for using it to create the plots and results can be found in README.md inside the source_code directory in the tarball.

APPENDIX B: FOLLOW-UP OF MISSED INJECTIONS

In this appendix, we will discuss the particularly loud injections which were not recovered during the simulation mode of the pipeline. An injection is deemed to be recovered, if it was assigned a log-likelihood ratio $\mathcal{L}$ of 35 or greater. Out of the 112526 injections, 65361 were missed. Most of these (65291 out of 65361) were missed because the injected SNR was too low for them to be recovered significantly. Some, however, had a high injected SNR and were still missed. We will discuss the reasons for the same, for missed injections with network SNR above 12 and detector SNRs greater than 7. These injections are shown in Fig. 4, of which there are 70. Out of these, 33 were missed due to the data containing a glitch simultaneous to the injection, causing the glitch rejection mechanism to




FIG. 6. An example of a Q-transform plot showing a glitch in Livingston, causing a simultaneous injection to be missed.

reject that part of the data. The existence of a glitch in the data was verified by creating Q-transform plots of the data window. An example of such a glitch is shown in Fig. 6. Out of the remaining 37 loud missed injections, 28 fell into the glitch region as defined in Sec. III D 6 and, hence, were

rejected. The pipeline failed to recover only nine injections out of the original 112526. However, such problematic injections can be recovered by increasing N_k , the number of samples to reject at each level, k , before moving on to the next, at the cost of the run time of the pipeline.

-
- [1] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. Lett.* **116**, 061102 (2016).
 - [2] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. Lett.* **116**, 241103 (2016).
 - [3] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. X* **9**, 031040 (2019).
 - [4] R. Abbott, T. Abbott, S. Abraham, F. Acernese, K. Ackley, A. Adams, C. Adams, R. Adhikari, V. Adya, C. Affeldt *et al.*, *Phys. Rev. X* **11**, 021053 (2021).
 - [5] R. Abbott, T. Abbott, F. Acernese, K. Ackley, C. Adams, N. Adhikari, R. Adhikari, V. Adya, C. Affeldt, D. Agarwal *et al.*, [arXiv:2108.01045](https://arxiv.org/abs/2108.01045).
 - [6] R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, N. Adhikari, R. X. Adhikari *et al.* (T. L. S. Collaboration, the Virgo Collaboration, the KAGRA Collaboration), [arXiv:2111.03606](https://arxiv.org/abs/2111.03606).
 - [7] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. Lett.* **119**, 161101 (2017).
 - [8] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Astrophys. J. Lett.* **892**, L3 (2020).
 - [9] R. Abbott, T. Abbott, S. Abraham, F. Acernese, K. Ackley, A. Adams, C. Adams, R. Adhikari, V. Adya, C. Affeldt *et al.*, *Astrophys. J. Lett.* **915**, L5 (2021).
 - [10] A. Trovato (Ligo Scientific and The Virgo Collaborations), *Proc. Sci. Asterics2019* (2020) 082.
 - [11] B. Zackay, L. Dai, T. Venumadhav, J. Roulet, and M. Zaldarriaga, *Phys. Rev. D* **104**, 063030 (2021).
 - [12] T. Venumadhav, B. Zackay, J. Roulet, L. Dai, and M. Zaldarriaga, *Phys. Rev. D* **101**, 083030 (2020).
 - [13] B. Zackay, T. Venumadhav, L. Dai, J. Roulet, and M. Zaldarriaga, *Phys. Rev. D* **100**, 023007 (2019).
 - [14] T. Venumadhav, B. Zackay, J. Roulet, L. Dai, and M. Zaldarriaga, *Phys. Rev. D* **100**, 023011 (2019).
 - [15] A. H. Nitz, C. D. Capano, S. Kumar, Y.-F. Wang, S. Kasta, M. Schäfer, R. Dhurkunde, and M. Cabero, [arXiv:2105.09151](https://arxiv.org/abs/2105.09151).
 - [16] R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, A. Adams, C. Adams, R. X. Adhikari, V. B. Adya, C. Affeldt *et al.*, *Astrophys. J. Lett.* **915**, L5 (2021).
 - [17] L. S. Finn and D. F. Chernoff, *Phys. Rev. D* **47**, 2198 (1993).
 - [18] B. J. Owen, *Phys. Rev. D* **53**, 6749 (1996).
 - [19] B. J. Owen and B. S. Sathyaprakash, *Phys. Rev. D* **60**, 022002 (1999).
 - [20] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. E. Creighton, *Phys. Rev. D* **85**, 122006 (2012).
 - [21] K. Cannon *et al.*, *Astrophys. J.* **748**, 136 (2012).
 - [22] S. Babak, R. Biswas, P. Brady, D. A. Brown, K. Cannon, C. D. Capano, J. H. Clayton, T. Cokelaer, J. D. Creighton, T. Dent *et al.*, *Phys. Rev. D* **87**, 024033 (2013).
 - [23] A. H. Nitz, T. Dent, T. Dal Canton, S. Fairhurst, and D. A. Brown, *Astrophys. J.* **849**, 118 (2017).
 - [24] T. Adams, D. Buskulic, V. Germain, G. M. Guidi, F. Marion, M. Montani, B. Mours, F. Piergiovanni, and G. Wang, *Classical Quantum Gravity* **33**, 175012 (2016).
 - [25] I. W. Harry, B. Allen, and B. Sathyaprakash, *Phys. Rev. D* **80**, 104014 (2009).
 - [26] P. Ajith, N. Fotopoulos, S. Privitera, A. Neunzert, N. Mazumder, and A. Weinstein, *Phys. Rev. D* **89**, 084041 (2014).
 - [27] I. Harry, S. Privitera, A. Bohé, and A. Buonanno, *Phys. Rev. D* **94**, 024012 (2016).
 - [28] N. J. Cornish and J. Crowder, *Phys. Rev. D* **72**, 043005 (2005).
 - [29] J. Veitch *et al.*, *Phys. Rev. D* **91**, 042003 (2015).
 - [30] G. Ashton, M. Hübner, P. D. Lasky, C. Talbot, K. Ackley, S. Biscoveanu, Q. Chu, A. Divakarla, P. J. Easter, B. Goncharov *et al.*, *Astrophys. J. Suppl. Ser.* **241**, 27 (2019).
 - [31] R. Abbott *et al.* (LIGO Scientific Collaboration and V. Collaboration), *SoftwareX* **13**, 100658 (2021).
 - [32] C. Messick *et al.*, *Phys. Rev. D* **95**, 042001 (2017).
 - [33] S. Sachdev *et al.*, [arXiv:1901.08580](https://arxiv.org/abs/1901.08580).
 - [34] K. Cannon, S. Caudill, C. Chan, B. Cousins, J. D. Creighton, B. Ewing, H. Fong, P. Godwin, C. Hanna, S. Hooper *et al.*, *SoftwareX* **14**, 100680 (2021).
 - [35] GstLAL software: git.ligo.org/lscsoft/gstlal.
 - [36] Open Science Grid: Introduction, <https://opensciencegrid.org/about/introduction/> (2021).
 - [37] K. Cannon, C. Hanna, and J. Peoples, [arXiv:1504.04632](https://arxiv.org/abs/1504.04632).
 - [38] K. Cannon, C. Hanna, and D. Keppel, *Phys. Rev. D* **88**, 024025 (2013).
 - [39] Lvc algorithm library software: git.ligo.org/lscsoft/lalsuite.
 - [40] G. Ashton, S. Thiele, Y. Lecoecuche, J. McIver, and L. K. Nuttall, *Classical Quantum Gravity* **39**, 175004 (2022).
 - [41] D. Davis, L. V. White, and P. R. Saulson, *Classical Quantum Gravity* **37**, 145001 (2020).
 - [42] S. Fairhurst, *New J. Phys.* **11**, 123006 (2009).
 - [43] C. Hanna *et al.*, *Phys. Rev. D* **101**, 022003 (2020).
 - [44] A. H. Nitz, T. Dent, G. S. Davies, S. Kumar, C. D. Capano, I. Harry, S. Mozzon, L. Nuttall, A. Lundgren, and M. Tápai, *Astrophys. J.* **891**, 123 (2020).
 - [45] <https://dcc.ligo.org/T2100321>.