

Black hole hunting with LISA

Neil J. Cornish¹ and Kevin Shuman¹

*eXtreme Gravity Institute, Department of Physics, Montana State University,
Bozeman, Montana 59717, USA*

 (Received 7 May 2020; accepted 26 May 2020; published 5 June 2020)

The Laser Interferometer Space Antenna (LISA) will be able to detect massive black hole mergers throughout the visible Universe. These observations will provide unique information about black hole formation and growth, and the role black holes play in galaxy evolution. Here we develop several key building blocks for detecting and characterizing black hole binary mergers with LISA, including fast heterodyned likelihood evaluations, and efficient stochastic search techniques.

DOI: [10.1103/PhysRevD.101.124008](https://doi.org/10.1103/PhysRevD.101.124008)

I. INTRODUCTION

The first detection by LIGO of gravitational waves from a binary black hole merger [1] has been followed by dozens of additional detections [2] that are revealing insights into stellar evolution and black hole formation [3]. In the next decade, the launch of the Laser Interferometer Space Antenna (LISA) [4,5] will allow for similar studies of much more massive black holes, potentially allowing us to unravel the interplay between massive black hole growth and galaxy evolution.

The detection and characterization of binary black hole mergers was investigated two decades ago [6–14] as part of the planning for an earlier incarnation of the LISA mission. In the intervening years there have been several improvements in both the modeling of the signals and the techniques used to detect and characterize the signals. The most significant changes are that the signal models now include inspiral, merger and ringdown, as opposed to just the inspiral, and a range of techniques have been developed that greatly speed up the calculation of the likelihood function, which plays a central role in the analysis. These techniques have recently been used in a study of Bayesian parameter estimation for LISA observation of massive black hole binaries, with an emphasis on the impact of including higher harmonics in the signal model [15,16].

The goal of this work is to develop several key building blocks for detecting massive black hole mergers with LISA and inferring their physical properties. The effort is part of the LISA Data Challenge [17], a successor to the original series of Mock LISA Data Challenges [18–20], where simulated LISA data is used as a playground for developing analysis algorithms that can be used once the mission is operational. The LISA Data Challenges are following a staged development, starting with relatively simple data sets and progressively building in additional realism. In the first round of the new challenges, dubbed *Radler*, the data

sets are broken out by source type. For massive black holes the simulated data set contains a single binary merger in uninterrupted stationary, Gaussian noise. The next challenge, *Sangria*, will include simulated data with multiple black hole binaries. Future challenges will add gaps, nonstationary and non-Gaussian noise, and will include multiple signals of different types, as well as increasing the complexity of the simulated signals. The techniques described here are sufficient to handle the *Radler* and *Sangria* Challenges, and will serve as a foundation for the development of the more advanced techniques needed to handle more realistic data sets that will ultimately form part of the global solution that simultaneously models thousands of overlapping signals of different types. Our approach is similar to that in Ref. [15], but with a greater emphasis on the initial search. Another key difference is that our analysis accommodates instrument noise, while the analysis in Ref. [15] is limited to noise-free data. The GPU accelerated likelihood approached used in Ref. [16] is able to account for instrument noise.

To avoid getting bogged down in details, most of the technical aspects of the analysis, such as the instrument response function, noise spectra etc., are relegated to appendixes. Geometric units with $G = c = 1$ are used throughout. The code used to produce the results shown here can be downloaded from GitHub [21].

II. THE QUARRY

The black hole mergers we are considering have total masses between $10^5 M_\odot$ and $10^8 M_\odot$. Lower mass systems, including the stellar origin black holes detected by LIGO and Virgo, will require slightly different search techniques due to their longer duration. For the signal model we use a phenomenological model (PhenomD), which describes the dominant harmonic of a quasicircular binary with spins aligned with the orbital angular




FIG. 1. The time interval before merger to accumulate 99% of the signal-to-noise ratio squared, T_{99} , as a function of the detector-frame total mass for two different mass ratios. For binaries with total mass above $10^5 M_{\odot}$ the bulk of the signal-to-noise is accumulated in less than one month.

momentum [22,23]. The model attaches an augmented post-Newtonian inspiral to a parameterized merger and ringdown, and the model is calibrated against a suite of numerical relativity simulations. Including additional harmonic content in the signal model will have little impact on the search strategy, however allowing for misaligned spins and orbital precession would require modifications to the search.

The search strategy is guided by the fact that systems in the mass range being considered only spend a short amount of time in the LISA band relative to the mission duration (years) and orbital modulation time scale (months). It is conventional to define “time in band” as the time to merger from some fiducial frequency, but a more meaningful measure is the time until merger during which a large fraction of the signal-to-noise ratio, or SNR, is accumulated. Since the Bayes factor between signal and noise scales as the signal-to-noise ratio squared, we define the time in-band, T_{99} , to be the time before merger taken to accumulate 99% of the SNR². Figure 1 shows the time in-band as a function of the detector-frame total mass M for two different mass ratios $q = m_2/m_1$. The source was simulated with ecliptic co-latitude $\theta = \pi/3$ and longitude $\phi = 0$, with polarization angle $\psi = \pi/3$ and inclination angle $\iota = \pi/3$. The choice of source location and orientation is largely irrelevant here since the time in-band is so short—the LISA antenna is roughly constant on durations shorter than one month. Here we are using the LISA mission configuration and noise model described in the LDC manual [24]. Details of the noise model are given in Appendix A.

III. THE HUNT

Because the time in-band is short for systems with total mass greater than $10^5 M_{\odot}$, an efficient search strategy is to

analyze shorter stretches of data. During the mission this could be done on a rolling basis, with the data segment being advanced day-by-day as the data arrives in an effort to provide low-latency alerts to aid searches for electromagnetic counterparts. Note though that there is little hope of providing advanced warning of a merger for systems with moderate mass ratios and detector frame total masses above $3 \times 10^5 M_{\odot}$.

The short duration of the signals allows us to ignore the antenna response in the first stage of the search, cutting the search space from eleven dimensions (two masses, two spins, merger time and phase, distance, sky location and orbital orientation) to just four—the two masses m_1, m_2 and the two dimensionless spins χ_1, χ_2 . The waveform amplitude, phase and merger time are maximized over analytically using the methods described in Section 8 of the LIGO Data Analysis guide [25]. In the LISA setting we analyze two channels of data, the signal-orthogonal A and E time-delay-interferometry (TDI) channels. The merger time maximization is performed simultaneously for both channels, while the amplitude and phase maximization is performed individually in each channel. The search over masses and spins could be performed using a LIGO-style template bank, but we prefer to use a stochastic search that is a variant of the Markov Chain Monte Carlo (MCMC) method we use for parameter estimation.

The PhenomD [22,23] waveform code provides the frequency domain amplitude $\mathcal{A}(f)$ and phase $\Phi(f)$ for the gravitational wave signal $h(f) = \mathcal{A}(f)e^{i\Phi(f)}$. To convert this to the fractional-frequency TDI response used in the LDC data sets we have to multiply the amplitude by a factor of $8(f/f_*) \sin(f/f_*)$, where $f_* = 1/2\pi L \simeq 19.1$ mHz is the transfer frequency and L is the arm-length. The factor of $2 \sin(f/f_*)$ accounts for the time-delay interferometry, while the factor of $4f/f_*$ accounts for the fractional frequency response. The time delay interferometry also introduces a phase shift of $\pi/2$ and a time shift of L , but these are taken care of by the analytic maximization.

The full $2^{22} \times 10$ seconds of the LDC massive black hole data set was divided into 16 chunks, each roughly a month in duration ($T_{\text{seg}} = 2621440$ seconds). Each chunk was searched using a variant of the replica exchange MCMC [26] that we use for parameter estimation. The search used a total of twelve chains, geometrically spaced in “temperature” by a factor of 1.5. The parameters of the cold chain were cloned to the hottest chain every 100 iterations. A mixture of proposals were used, including uniform draws from the prior range for each parameter, and draws along eigendirections of the Fisher information matrix, scaled by the inverse square root of the eigenvalues. The Fisher matrix was computed using the masses and spins, in addition to the merger time and phase. The merger time and phase were included in the Fisher matrix calculation, even though they were maximized over in the likelihood, since there are covariances between them and the masses




FIG. 2. Trace plots of the cold chain from the stochastic search. The horizontal solid lines indicate the true parameter values. The search rapidly locks onto the signal.

and spins. Leaving out the merger time and phase in the Fisher matrix results in jump proposals that are inefficient. The maximization over merger time t_c in the calculation of the likelihood was restricted to be at most $\pm T_{\text{seg}}/8.0$ so as not to go too far outside the frequency range of the reference signal, which was set to be from $f(T_{\text{start}})$ to $f(T_{\text{start}} + T_{\text{seg}})$ where T_{start} is the start time of the segment [the mapping between time and frequency is given in Eq. (B1)]. The prior for the merger time was set to be $t_c \in [T_{\text{start}}, T_{\text{start}} + 2T_{\text{seg}}]$, so that the merger could occur in the current segment or the next segment over, allowing for the possibility of picking up the signal prior to merger. The stochastic search is not Markovian (reversible) since the likelihood is maximized and the proposal densities are not included in the Metropolis-Hastings jump acceptance probability.

Figure 2 shows a trace plot of the merger time, masses and spins from the cold chain in the stochastic search of data chunk 10 of 16. As expected, the search rapidly locked onto the merger time, followed by the masses and then the spins. Searches of the other 15 segments yielded no additional significant candidate signals with $\text{SNR} > 8$. The search of each segment takes less than four minutes using a quad-core 2.9 GHz Macbook Pro. The 12 chains were run in parallel using OpenMP.

The initial stage of the search provides a starting solution for the masses, spins and the detector frame merger time. The next stage of the search finds a good starting solution for the sky location (θ, ϕ) , luminosity distance D_L , Barycenter merger time t_c , merger phase ϕ_c , and orbital orientation (expressed in terms of the polarization angle ψ and inclination angle ι). The intrinsic parameters (masses and spins) are held fixed, while the other parameters are explored using a stochastic search algorithm. For the second stage of the search we need to apply the full instrument response to the PhenomD templates. The

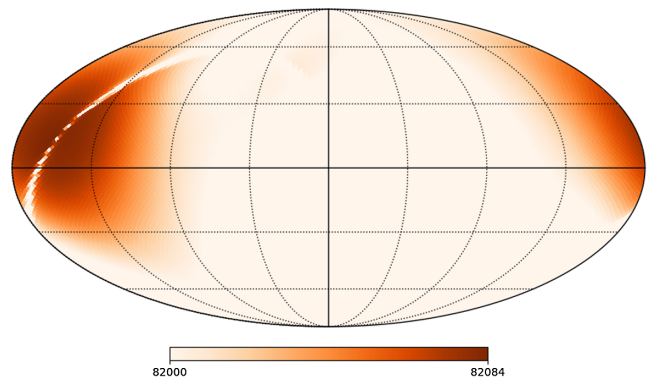


FIG. 3. A sky map showing the F-statistic likelihood computed using the values for the masses and spins found during the first phase of the search.

response is computed directly in the frequency domain using the method described in Appendix B. The likelihood is computed using the F-statistic [27] maximization described in Ref. [8], which uses a set of four filters found by setting $\iota = \pi/2$ and (i) $(\phi_c, \psi) = (0, 0)$; (ii) $(\phi_c, \psi) = (\pi/2, \pi/4)$; (iii) $(\phi_c, \psi) = (3\pi/4, 0)$; (iv) $(\phi_c, \psi) = (\pi/4, \pi/4)$; in the full response. The F-statistic maximizes the likelihood over the merger phase, luminosity distance, polarization angle and inclination angle, thus reducing the search to be over the sky location and barycenter merger time. The barycenter and detector frame merger times are related by $t_c^B = t_c^D + \hat{k} \cdot \mathbf{x}_0(t_c)$, where \hat{k} is the direction of propagation of the gravitational wave and $\mathbf{x}_0(t)$ is the center of the LISA constellation. For a given sky location for the source, this mapping can be used to estimate the barycenter merger time. We allow the barycenter merger time to vary a little from this value to account for the time delays introduced by the full instrument response. Figure 3 shows a sky map of the F-statistic likelihood computed using the masses and spins found during the first stage of the search. The F-statistic based search typically locks onto the true sky location in less than a hundred iterations.

The same search technique can be used on data sets containing multiple black hole mergers. The maximum likelihood solution from the previous pass is subtracted from the data and the search repeated. This process is repeated until no additional significant signals are found. The collection of maximum likelihood solutions from the search can be used as a starting point for more refined parameter estimation, or even better, can be turned into proposal distributions for performing the full multi-source global fit.

The simulated signals in the *Radler* data set only included the dominant harmonic in the waveform. If the data contains signals with full harmonic content, it is likely sufficient to perform the search with templates that only include the dominant harmonic, since for quasi-circular systems, the dominant harmonic contributes most of the SNR. It is possible that some of the chains might

temporarily lock on to the subharmonics before switching to the dominant mode.

IV. DRESSING OUT

With the rapid search phase complete and one or more sources identified, the next step is to refine the estimates for the source parameters. In reality this will be done while simultaneously inferring the parameters of many sources, including hundreds of massive black holes [28–30], tens of thousands of ultracompact galactic binaries [31,32], hundreds of extreme mass ratio inspirals [33], and a handful of stellar origin black hole binaries [34]. Additionally it will be necessary to model the instrument noise and residual galactic signal, which accounting for gaps in the data and other real-world complications.

Here we start with the simpler problem of inferring the parameters of a single massive black hole binary merger in gapfree data, with stationary, Gaussian noise with a known power spectrum. We apply Bayesian inference to compute the posterior distribution for the source parameters. Our method of choice is the replica exchange (parallel tempered) Markov Chain Monte Carlo (PTMCMC) algorithm. Over the years we have developed a standard recipe [35] for implementing PTMCMCs that uses a combination of local and global proposal distributions, and we adopt that approach here, while also adding a new ingredient—maximized jumps—that significantly improves the sampling. The individual chains are advanced using the Metropolis-Hastings algorithm, whereby a chain is advanced from parameters x to parameters y with acceptance probability

$$H(y|x) = \min \left(1, \frac{p(d|y)p(y)q(x|y)}{p(d|x)p(x)q(y|x)} \right). \quad (1)$$

Here $p(d|x)$ is the likelihood of observing data d given model parameters x , $p(x)$ is the prior distribution for parameters x , and $q(y|x)$ is the proposal density for drawing a new set of parameters y given the current set of parameters x .

One drawback of stochastic algorithms such as PTMCMCs is that they require large numbers of likelihood evaluations. Using the fast frequency domain technique described in Appendix B, applied to the full *Radler* data set, each likelihood evaluation takes roughly one second on a single 2.6 GHz CPU core. A variety of techniques can be used to speed up the likelihood evaluation, including reduced order models [36], reduced-order quadratures [37] and computational approaches such as GPU acceleration [16]. Here we use a different technique [38] that allows us to compute the likelihood without having to generate any waveforms aside from the one reference waveform that is used to heterodyne the data. The heterodyne likelihood has been rediscovered and used in LIGO/Virgo data analysis, though there the technique has been called “relative binning” [39]. In the current application, the heterodyned likelihood takes ~ 1 ms to compute—a factor

of one thousand times faster than direct evaluation with the (already fast) frequency domain waveforms. The method used to compute the heterodyned likelihood is described in Appendix C.

A. Priors

We assumed uniform priors on all the parameters. The detector frame individual masses were taken to be uniform in the range $m_1, m_2 \in [5 \times 10^4 M_\odot, 10^8 M_\odot]$. The dimensionless spins were taken to be uniform in the range $\chi_1, \chi_2 \in [-1, 1]$. The luminosity distance was taken to be uniform in the range $D_L = [0.1 \text{ Gpc}, 400 \text{ Gpc}]$. The merger time t_c was taken to be uniform in the range $[0, 2T_{\text{obs}}]$, where T_{obs} is the observation time. The cosine of the ecliptic colatitude, $\cos \theta$, and the cosine of the inclination, $\cos i$, were taken to be uniform in the range $[-1, 1]$. The orbital phase at merger ϕ_c , and the polarization angle ψ were taken to be uniform in the range $[0, \pi]$. The ecliptic longitude, ϕ , was taken to be uniform in the range $[0, 2\pi]$.

B. Proposals

We used a PTMCMC with a geometrically spaced temperature ladder, with inverse temperatures scaling as $\beta_i = \alpha^{-i}$. Samples are recorded from the chain with $\beta_0 = 1$. A total of $N_c = 16$ sixteen chains were used, with each chain running on a separate computational core. The geometric temperature spacing $\alpha > 1$ was set so as to give an effective signal to noise ratio of $\text{SNR}_{\text{eff}} = \text{SNR}/\beta_{N_c-1} = 5$. More efficient sampling could probably be achieved using an adaptive temperature spacing. We did check that all the chains remained “connected,” that is, that the exchange rate between neighboring chains remained above zero throughout the simulation and for all temperatures.

A mixture of proposal distributions were used to advance the chains. As per our standard recipe [35], the mix included local and global proposals, with the choice of proposal at each iteration drawn randomly. To the standard mix we also added a new technique that incorporates maximization over parameters in a way that maintains detailed balance (reversibility) in the chains.

The ultimate proposal density would be posterior distribution itself, but lacking that, we instead use two local approximations to the posterior: proposals derived from previous chain samples and proposals that use the quadratic Fisher information matrix approximation to the likelihood. A variant of the differential evolution approach [40] is used to propose jumps based on previously collected samples. A running history of past samples is collected at each temperature level. Jumps are proposed from the current position to a new location found by adding the vector connecting two points drawn randomly from the history. The proposal can be shown to be asymptotically Markovian [40]. Heuristically, the asymptotic reversibility can be understood from the observation that as the samples accumulate, the chain history approaches the stationary

target distribution. Differential evolution is very effective at exploring strong parameter degeneracies, as the vectors connecting past samples tend to line up along the degenerate directions. The Fisher information matrix provides a quadratic approximation to the likelihood, which in turn is a good approximation to the posterior distribution so long as curvature of the prior is less than the curvature of the likelihood. The Fisher matrix can be computed using a fast spline integration method [38]. Writing the signal as $h(f) = \mathcal{A}(f)e^{i\Phi(f)}$, the Fisher matrix $\Gamma_{ij} = (h_i|h_j)$ can be expressed as

$$\Gamma_{ij} = 4 \sum_{I=A,E} \int \frac{\mathcal{A}_i^I \mathcal{A}_j^I + \mathcal{A}^{I2} \Phi_i^I \Phi_j^I}{S_n^I(f)} df. \quad (2)$$

Here \mathcal{A}^I and Φ^I are the amplitude and phase in the I th data channel and $S_n^I(f)$ is the noise spectral density in the I th data channel. The sum is over the $I = A, E$ TDI channels. All of the terms appearing in (2) vary slowly in frequency and can be evaluated on a coarse grid using the spline integration method described in Appendix C. The eigenvalues and eigenvectors of the Fisher matrix are computed and used to propose jumps by first randomly selecting an eigendirection, then drawing the jump size from a normal distribution with variance equal to the inverse of the corresponding eigenvalue. The effectiveness of these proposals is predicated on the Fisher matrix providing a reasonable approximation to the posterior distribution. Figures 4 and 5 compare the Fisher matrix approximation to the posterior to the MCMC derived posterior distributions for the noiseless *Radler* data set. We see that the Fisher matrix provides a good approximation for the parameters shown in Fig. 4 that enter directly into the gravitational wave phase, but the approximation is poor for the extrinsic parameters, such as sky location, luminosity distance and inclination angle shown in Fig. 5. In part, the poor showing for the extrinsic parameters is because the distributions are multi-modal whereas the Fisher matrix approximation is monomodal. The agreement is better in practice since the Fisher information matrices are updated as the simulation progresses, so all the modes get covered. On the other hand, the agreement seen in Figure 4 is deceiving, as the likelihoods computed using the Fisher matrix approximation are often very different from those computed using the full likelihood, mostly due to inaccuracies in the merger time and merger phase. These inaccuracies were found to severely limit the acceptance rate for jumps along certain eigendirections. Two strategies were used to improve the acceptance of the Fisher matrix based proposals. The first was to break the Fisher matrix into blocks and sometimes just propose jumps in the subset of parameters that are well approximated, the second was to use the likelihood maximization procedure described below.

The posterior distributions encountered in gravitational wave astronomy are often multimodal. To fully explore all

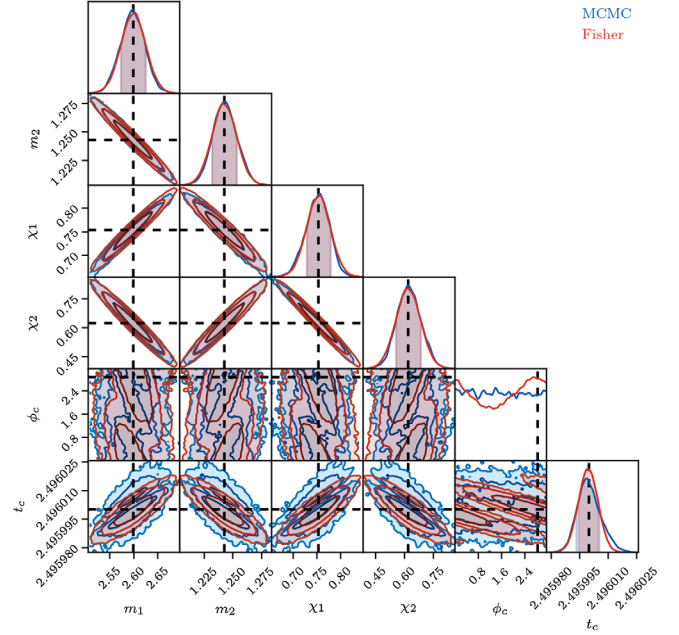


FIG. 4. Corner plot comparing the MCMC (blue) and Fisher matrix (red) estimates for the posterior distribution for the masses, spins, merger phase and merger time. The true values are indicated by dashed lines. The two estimates agree well for the masses and spins, but less so for the merger time and phase.




FIG. 5. Corner plot comparing the MCMC (blue) and Fisher matrix (red) estimates for the posterior distribution for the sky location, luminosity distance, inclination and polarization. The true values are indicated by dashed lines. The Fisher matrix provides a poor approximation to the MCMC derived posterior distribution for these extrinsic parameters.

the modes, our approach [35] is to use global proposals combined with replica exchange (parallel tempering). Here we use a global proposal based on the F-statistic likelihood. A similar approach has been used for ultracompact galactic binaries [41], but here the implementation is different. Rather than precomputing a collection of F-statistic maps as was done in Refs. [41,42], here we compute the F-statistic likelihood on-demand, and use it as part of a maximized likelihood proposal.

The maximized proposal technique is a general method that can be used with any form of likelihood maximization. In the current context we applied the technique to the F-statistic likelihood and also to likelihoods that are maximized with respect to time offset, overall phase and amplitude. The method is illustrated in Fig. 6. The first step is to draw a new set of parameters from some distribution. For example, when using the F-statistic, a new sky location is drawn while holding the masses and spins fixed. The remaining parameters, $\{t_c, \phi_c, \psi, i, D_L\}$ are analytically maximized using the F-statistic. The barycenter merger time is adjusted to keep the detector frame merger time fixed. The new sky location is either drawn from the prior, or from a wide normal distribution centered on the current sky location. When coupled with the Fisher matrix proposal, the maximization is performed on $\{t_c, \phi_c, D_L\}$. If uncorrected for, the maximization over parameters would violate detailed balance and bias the posterior distribution. To restore detailed balance a second step is added to the proposal: the Fisher information matrix is computed at y_{\max} for the subset of parameters that are maximized over. A second jump Δy is drawn from a normal distribution with covariance matrix equal to the inverse of the Fisher matrix yielding the proposed point $y = y_{\max} + \Delta y$. The full proposal density $q(y|x)$ is then given by $q(y|x) = q(\Delta y|y_{\max})q(y_{\max}|x)$. The proposal density for the reverse move, $q(x|y) = q(\Delta x|x_{\max})q(x_{\max}|y)$ is computed after finding the point x_{\max} by maximizing the likelihood using the fixed parameters at x .




FIG. 6. An illustration of the maximized jump proposal technique. A jump from the current location x to a new location y proceeds in two steps. First, a new location is proposed and parameters such as merger time and phase are analytically maximized over, yielding the new point y_{\max} . Next, a jump is drawn from a normal distribution from y_{\max} to y . The proposal densities for the forward and reverse jumps are included in the Metropolis-Hastings ratio to ensure detailed balance.

The maximized proposal technique dramatically improves the mixing of the chains and the discovery of secondary posterior modes. For example, even without using dedicated proposals that exploit the symmetries of the LISA instrument response, maximized jumps using uniform draws on the sky location were able to quickly find all the secondary modes.

The acceptance rates of the proposals and the autocorrelation length of the Markov chains depend somewhat on the source being analyzed, sometimes by as much as a factor of two or three. For the *Radler* data set, the average acceptance rate for each proposal were: Fisher matrix 0.6; differential evolution 0.04; and F-statistic sky move 0.1. The autocorrelation length of the chains was 150 for the masses and spins, 300 for the sky location and 500 for the inclination and distance. A typical run went for two million iterations, producing thousands of independent samples.

C. Results

The black hole binary system used in the *Radler* simulation has detector frame masses $m_1 = 2.599137 \times 10^6 M_\odot$, $m_2 = 1.242860 \times 10^6 M_\odot$ and dimensionless spins $\chi_1 = 0.75348$ and $\chi_2 = 0.62159$. The system was placed at a luminosity distance of $D_L = 56.006$ Gpc, corresponding to a redshift of $z = 5.7309$ for the assumed cosmological model. The source frame masses are a much more modest $m_1^s = 3.8615 \times 10^5 M_\odot$, $m_2^s = 1.8465 \times 10^5 M_\odot$. The relatively high detector frame total mass for this system results in it becoming detectable less than a day before merger: the signal reaches $\text{SNR} = 10$ just 11 hours prior to merger. The modulation of the amplitude and phase of the signal due to the LISA orbit is essentially irrelevant since very little SNR is accumulated prior to the last few hours before merger. Figure 7 shows the whitened signal amplitude for this system in the TDI A channel, with an inset showing the amplitude modulation that occurs at low frequencies where the signal is undetectable. Consequently, the signal is effectively a short duration burst, and most of the directional information comes from differences in the time of arrival of the signal at each spacecraft, much like the situation for bursts from cosmic string cusps and kinks [43,44] or generic short duration bursts [45].

Unless data is downloaded from the LISA constellation every few hours or so, it is highly unlikely that systems such as this one will be detected prior to merger. If data were available in advance, the sky localization would be poor. To investigate this possibility, we smoothly truncated the time domain data using a cosine window of the form

$$W(t) = \begin{cases} 1 & t \leq T_{\text{cut}} - \Delta T \\ 0 & t > T_{\text{cut}} \\ \frac{1}{2} \left(1 - \cos\left(\frac{\pi(t - T_{\text{cut}})}{\Delta T}\right) \right) & \text{otherwise} \end{cases} \quad (3)$$




FIG. 7. The whitened signal amplitude in the TDI A channel for the *Radler* black hole binary system. The inset highlights the amplitude modulation due to the LISA orbital motion.

with $T_{\text{cut}} = t_c - 6.7 \times 10^3 \text{ s}$ and $\Delta T = 5 \times 10^4 \text{ s}$. The same window was applied to the frequency domain waveforms using the time-frequency mapping $t(f)$. This choice of parameters removes the late inspiral, merger and ring-down, and reduces the signal-to-noise ratio to $\text{SNR} = 11.1$. Note that more rapid truncations with smaller ΔT result in unacceptable spectral leakage and large Gibbs oscillations in the frequency domain signals.

Posterior distributions for the masses and spins using the truncated premerger signal are shown in Figure 8. The mass




FIG. 8. Posterior distributions for the masses and spins using the truncated *Radler* data. The truncated signal has $\text{SNR} = 11.1$. The two-dimensional distribution for the masses exhibits the characteristic “banana” shape along the line of constant chirp mass. The instrument noise pushes the masses away from their true values.

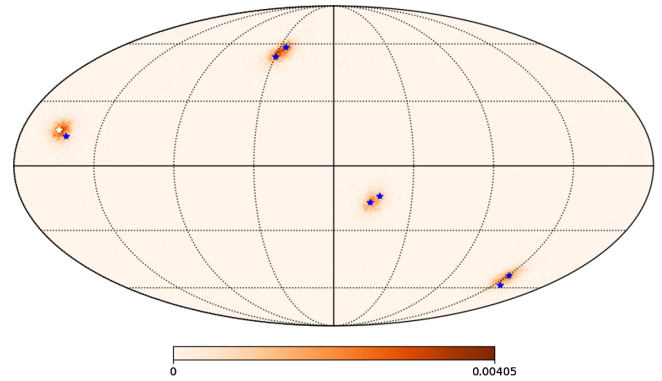


FIG. 9. Sky map found using the truncated premerger *Radler* data. The truncated signal has $\text{SNR} = 11.1$. The sky map exhibits the typical multimodality of a short duration signal. The true source location is indicated by a white star, while the blue stars mark the locations of the seven secondary maxima expected for a low frequency, short duration burst signal.

distribution follows a line of constant chirp mass, as expected for the inspiral-only portion of the signal. The spin of the more massive system is already quite well constrained. The instrument noise pushes the masses and spins away from their true values. The posterior distribution for the ecliptic latitude and longitude of the truncated signal are shown in Fig. 9. The short duration, burst-like nature of




FIG. 10. Slices through the posterior distribution for the masses, spins, merger phase and merger time. The true values are indicated by dashed lines. Here the offsets from the true values is mostly due to projection effects and not noise—the peak of the posterior in the full $D = 11$ dimensions can align with the true values but appear offset in the one and two dimensional projections shown here.




FIG. 11. Slices through the posterior distribution for the masses, spins, merger phase and merger time. The true values are indicated by dashed lines. The peak of the posterior distribution is offset from the true values by the noise. The offsets are relatively small for the particular noise realization used in the *Radler* data set.

the signal results in four distinct modes. The 90% credible interval covers 630 square degrees. The four modes follow the eight-fold symmetry of the static, low frequency antenna pattern [15], (the seven secondary modes are marked by blue stars in Fig. 9).

Including the merger and ringdown boosts the signal-to-noise ratio to $\text{SNR} = 405$ and reduces the 90% credible interval for the sky location to 47 square degrees. Of the original four modes for the sky location only one survives, with the surviving mode resolving into two closely spaced modes as seen in Fig. 10. The instrument noise moves the extrinsic parameters away from their true values, but the effect on the projected posterior distribution is small compared to the apparent displacements caused by the projections. To see this, compare the noise-free distributions shown in Fig. 5 to the analysis with noise shown in Fig. 10. The luminosity distance is significantly impacted by these projection effects.

The parameters that directly enter the phase, shown in Fig. 11 are shifted slightly by including the instrument noise. The smallness of the shifts is just luck of the draw—repeating the analysis with different noise realizations yielded larger shifts on average.

V. DISCUSSION

We have presented an early prototype for detecting and characterizing massive black holes with LISA. A key

feature of the methods we have developed is speed. All the analyses were conducted on a 2016 vintage laptop, with the search stage taking tens of minutes and the characterization stage taking a few hours.

The search phase is accelerated by using smaller segments of data and a fast approximate likelihood that ignores the LISA instrument response and maximizes over amplitude, merger time and merger phase. The characterization phase is accelerated by using a heterodyning procedure to speed up the likelihood evaluation by three orders of magnitude. A novel jump proposal was introduced to improve the mixing of the MCMC algorithm. The new proposal employs a two stage scheme where some parameters are first maximized over, followed by a stochastic jump in those parameters to preserve detailed balance.

The problem tackled here is considerably simpler than will be encountered with real LISA data. In reality we will have to contend with nonstationary and non-Gaussian noise, data gaps, more complex waveforms and multiple overlapping signals. Some of the techniques we have used will need to be generalized to handle these complications. For example, spin precession causes the orientation of the orbital plane, and hence the inclination and polarization angles, to vary with time, thus rendering the F-statistic approximate. Data gaps and nonstationary noise will require a shift from frequency domain to time-frequency (wavelet) domain methods, necessitating a generalization to the heterodyned likelihood approach. We aim to tackle each of these complications in future work that builds on the foundation we have laid here.

ACKNOWLEDGMENTS

We appreciate the support of the NASA LISA foundation Science Grant No. 80NSSC19K0320. This work was initiated while N.J.C. was on sabbatical at the Observatoire de la Côte d’Azur, kindly hosted by Nelson Christensen, and supported in part by the Centre National d’Études Spatiales. We have benefited from many exchanges with Stanislav Babak and Tyson Littenberg.

APPENDIX A: NOISE MODEL

The noise spectral density in the fractional-frequency A , E TDI channels is modeled as

$$S_n(f) = \frac{64}{3L^2} \left(\frac{f}{f_*} \sin\left(\frac{f}{f_*}\right) \right)^2 \left[\left(2 + \cos\left(\frac{f}{f_*}\right) \right) S_{\text{ps}} + \left(6 + 4 \cos\left(\frac{f}{f_*}\right) + 2 \cos\left(\frac{2f}{f_*}\right) \right) \times \frac{S_{\text{acc}}}{(2\pi f)^4} \left(1 + 16 \left(\frac{10^{-4}}{f} \right)^2 \right) \right]. \quad (\text{A1})$$

with position noise level $S_{\text{ps}} = 2.25 \times 10^{-22} \text{ m}^2 \text{ Hz}^{-1}$ and acceleration noise level $S_{\text{acc}} = 9 \times 10^{-30} \text{ m}^2 \text{ s}^{-4} \text{ Hz}^{-1}$.

APPENDIX B: FREQUENCY DOMAIN INSTRUMENT RESPONSE

The LISA constellation cartwheels around the Sun resulting in a time dependent instrument response function. Since the likelihood is computed in the frequency domain, and the PhenomD waveform is already expressed in the frequency domain, it is most efficient to map the frequency to time and compute the response directly in the frequency domain, as was first proposed by Cutler [46]. The mapping is given by

$$t(f) = \frac{1}{2\pi} \frac{d\Phi(f)}{df} + t_c. \quad (\text{B1})$$

The time-frequency mapping is computed by taking numerical derivatives of the PhenomD phase function. In the expressions below it is understood that the time t is mapped to the frequency f by $t(f)$.

The fractional frequency shift imparted by a gravitational wave to the laser light propagating along one spacecraft to another is given in Eq. (B11) of Ref. [47], which can be written in full generality as

$$\frac{\delta\nu_{ij}(t)}{\nu_0} = \frac{\hat{r}_{ij} \otimes \hat{r}_{ij}}{2(1 - \hat{k} \cdot \hat{r}_{ij})} : (\mathbf{h}(t, x_j) - \mathbf{h}(t - L, x_i)). \quad (\text{B2})$$

This expression describes the Doppler shift of the laser light in going from the spacecraft at \mathbf{x}_i to the spacecraft at \mathbf{x}_j , arriving at barycenter time t . Here $\hat{r}_{ij} = (\mathbf{x}_j - \mathbf{x}_i)/L$, where L is the armlength, which we are assuming is constant. The gravitational wave is propagating in the \hat{k} direction with surfaces of constant phase given by $\xi = t - \hat{k} \cdot \mathbf{x}$. The full TDI response is formed out of a linear combination of one-arm Doppler shifts along various arms at various times. In order to be able to add together these Doppler shifts in a consistent way, it is helpful to reference all of expression to the gravitational wave signal at the center of the constellation, given by $\mathbf{x}_0 = (\mathbf{x}_1 + \mathbf{x}_2 + \mathbf{x}_3)/3$. Working in the rigid-adiabatic approximation [47] for a GW signal with instantaneous frequency f we have

$$\begin{aligned} \frac{\delta\nu_{ij}(t)}{\nu_0} &= i(\hat{r}_{ij} \otimes \hat{r}_{ij}) : \mathbf{h}(t, \mathbf{x}_0) \left[\text{sinc}\left(\frac{f}{2f_*}(1 - \hat{k} \cdot \hat{r}_{ij})\right) \right. \\ &\quad \left. \times e^{-i\frac{f}{2f_*}(1 + \hat{k} \cdot \hat{r}_{ij} - \frac{2}{\sqrt{3}}\hat{k} \cdot \hat{r}_{i0})} \right] \end{aligned} \quad (\text{B3})$$

where $f_* = 1/(2\pi L)$ is the transfer frequency. The gravitational wave signal is given by

$$\mathbf{h}(\xi) = h_+(\xi)\mathbf{e}^+ + h_\times(\xi)\mathbf{e}^\times. \quad (\text{B4})$$

For the leading order 22-mode of a nonprecessing circular binary we have

$$\mathbf{h}(\xi) = h(\xi)(A_+\mathbf{e}^+ + iA_\times\mathbf{e}^\times), \quad (\text{B5})$$

where

$$A_+ = \frac{1 + \cos^2\iota}{2}, \quad A_\times = -\cos\iota, \quad (\text{B6})$$

and ι is the inclination of the binary orbit.

In the rigid adiabatic approximation the time delay interferometry introduces an overall transfer function given by

$$\mathcal{T} = (1 - e^{-4\pi ifL}) = 2ie^{-i\frac{f}{f_*}} \sin(f/f_*). \quad (\text{B7})$$

Putting all the pieces together, the X channel TDI variable extracted from vertex 1 is given by

$$\begin{aligned} X(t) &= -\frac{f}{f_*} e^{-i\frac{f}{f_*}} \sin\left(\frac{f}{f_*}\right) [(\hat{r}_{12} \otimes \hat{r}_{12})\mathcal{T}_{12}(t) \\ &\quad - (\hat{r}_{13} \otimes \hat{r}_{13})\mathcal{T}_{13}(t)] : \mathbf{h}(t, \mathbf{x}_0) \end{aligned} \quad (\text{B8})$$

where

$$\begin{aligned} \mathcal{T}_{ij}(t) &= \text{sinc}\left(\frac{f}{2f_*}(1 - \hat{k} \cdot \hat{r}_{ij})\right) e^{-i\frac{f}{2f_*}(3 + \hat{k} \cdot \hat{r}_{ij} - \frac{2}{\sqrt{3}}\hat{k} \cdot \hat{r}_{i0})} \\ &\quad + \text{sinc}\left(\frac{f}{2f_*}(1 + \hat{k} \cdot \hat{r}_{ij})\right) e^{-i\frac{f}{2f_*}(1 + \hat{k} \cdot \hat{r}_{ij} - \frac{2}{\sqrt{3}}\hat{k} \cdot \hat{r}_{i0})}. \end{aligned} \quad (\text{B9})$$

If we define $d_{ij}^{+,\times} = (\hat{r}_{ij} \otimes \hat{r}_{ij}) : \mathbf{e}^{+,\times}$, the 22-mode response can be written as

$$X(t) = [F_X^+(t)A_+ + iF_X^\times(t)A_\times]h(t, \mathbf{x}_0), \quad (\text{B10})$$

where the complex antenna patterns are given by

$$F_X^{+,\times} = -\frac{f}{f_*} e^{-i\frac{f}{f_*}} \sin\left(\frac{f}{f_*}\right) (d_{12}^{+,\times}\mathcal{T}_{12} - d_{13}^{+,\times}\mathcal{T}_{13}). \quad (\text{B11})$$

Expressions for the Y and Z channels follow by cyclic permutation of the labels (1,2,3) in the expression for X. The signal orthogonal A, E, T can be formed out of linear combinations of X, Y, Z:

$$\begin{aligned} A &= \frac{1}{3}(2X - Y - Z) \\ E &= \frac{1}{\sqrt{3}}(Z - Y) \\ T &= \frac{1}{3}(X + Y + Z). \end{aligned} \quad (\text{B12})$$

For the A, E channels used in the analysis, the miss-match between the noiseless *Radler* data and the frequency

domain rigid adiabatic waveforms was $MM_A = 3.6 \times 10^{-7}$ for the A channel and $MM_E = 7.7 \times 10^{-7}$ for the E channel. These systematic mis-matches are well below the expected statistical mis-matches due to noise, $E[MM] = (D-1)/(2SNR^2)$, for a signal with dimension $D = 11$ and $SNR_A = 360.9$, $SNR_E = 184.9$.

APPENDIX C: HETERODYNED LIKELIHOOD

The log-likelihood in Gaussian noise is given by

$$\log L = \frac{1}{2}(d-h|d-h) - \frac{1}{2} \int 2\pi S_n(f) df, \quad (C1)$$

where d is the data, h is the waveform model, $S_n(f)$ is the noise spectral density and the notation $(a|b)$ indicates that usual noise weighted inner product. In an effort to reduce clutter in the notation we will suppress the sum over channels in what follows.

The idea behind the heterodyned likelihood [38] is that given a good reference model \bar{h} , such as the maximum likelihood waveform found in the search phase, the likelihood for waveforms “close” to h can be computed by heterodyning the residual, $\bar{r} = d - \bar{h}$, against the signal \bar{h} . This results in a likelihood that can be computed very cheaply using a coarse spline interpolation of the amplitude and phase. In an MCMC, any parameters that are in the central 99.9+% of the posterior will generate waveforms that are close enough to the reference waveform for the heterodyned likelihood to be used. In fact, the heterodyned likelihood itself is *exact*. It is the approximations used to make the computation fast that introduce error, and the size of the error can be controlled by how many terms are kept in the splines and FFTs used to speed up the evaluation.

The heterodyne method can be applied to signals with multiple harmonics from higher order modes or orbital eccentricity. The heterodyning procedure is applied mode-by-mode and the results are summed to give the full likelihood. The result is exact when the overlap between modes vanishes, and is a very good approximation in general.

In contrast to the reduced order quadrature method for accelerating the likelihood evaluation, the heterodyned likelihood is able to accommodate updates to the noise model. While the noise model was held fixed in the current application, noise updates are included here for completeness. Given a reference waveform \bar{h} and noise model $\bar{S}_n(f)$. The second term in the likelihood can be computed directly using a spline integration. The first term in the likelihood requires more attention:

$$\begin{aligned} (d-h|d-h) &= 4 \int \frac{(\bar{r} + \Delta h)(\bar{r} + \Delta h)^*}{S_n(f)} df \\ &= 4 \int \frac{\Delta h \Delta h^* + (\bar{r} \Delta h^* + \bar{r}^* \Delta h) + \bar{r} \bar{r}^*}{S_n(f)} df, \end{aligned} \quad (C2)$$

where $\Delta h = \bar{h} - h$. Writing $h = \mathcal{A}(f)e^{i\Phi(f)}$, and similarly for \bar{h} , we have

$$\int \frac{\Delta h \Delta h^*}{S_n(f)} df = \int \frac{\bar{\mathcal{A}}^2 + \mathcal{A}^2 - 2\bar{\mathcal{A}}\mathcal{A} \cos \Delta\Phi}{S_n(f)} df. \quad (C3)$$

with $\Delta\Phi(f) = \bar{\Phi}(f) - \Phi(f)$. This expression is exact. So long as h is close to \bar{h} , all the terms in the integrand are slowly varying and can be evaluated using a coarse spline interpolation. Figure 12 shows the integrand $(f\Delta h \Delta h^*)/S_n(f)$ for the TDI A-channel evaluated between the parameters for the maximum likelihood solution from the search and the true source parameters. To better illustrate the structure of the integrand we have switched to a logarithmic scale in frequency and correspondingly scaled the integrand by f . The spline points used to compute the integrand are shown. The method used to select the spline points is described below.

The next term in the likelihood can be written as

$$\int \frac{\bar{r} \Delta h^* + \bar{r}^* \Delta h}{S_n(f)} df = \int (\bar{r}_w \Delta h_w^* + \bar{r}_w^* \Delta h_w) df, \quad (C4)$$

where

$$\bar{r}_w = \frac{\bar{r} e^{-i\bar{\Phi}(f)}}{\bar{S}_n^{1/2}(f)} \quad (C5)$$

is the whitened reference residual heterodyned by the reference phase and

$$\Delta h_w = \frac{(\bar{\mathcal{A}}(f) - \mathcal{A}(f)e^{-i\Delta\Phi(f)})\bar{S}_n^{1/2}(f)}{S_n(f)} \quad (C6)$$




FIG. 12. The integrand $(f\Delta h \Delta h^*)/S_n(f)$ for the TDI A-channel evaluated between the parameters for the maximum likelihood solution from the search and the true source parameters. The spline points used to compute the integral are indicated by crosses.




FIG. 13. The integrand $(\bar{r}_w(\tau)\Delta h_w^*(\tau) + \bar{r}_w^*(\tau)\Delta h_w(\tau))$ for the TDI A-channel evaluated between the parameters for the maximum likelihood solution from the search and the true source parameters.

is the heterodyned and whitened difference in the waveforms. Writing

$$\bar{r}_w(f) = \int e^{2\pi i\tau} \bar{r}_w(\tau) d\tau, \quad (\text{C7})$$

and similarly for $\Delta h_w(f)$ we have

$$\int \frac{\bar{r}\Delta h^* + \bar{r}^*\Delta h}{S_n(f)} df = \int (\bar{r}_w(\tau)\Delta h_w^*(\tau) + \bar{r}_w^*(\tau)\Delta h_w(\tau)) d\tau. \quad (\text{C8})$$

This expression is exact. It can be approximated by using a FFT to compute the Fourier transform and using a restricted range for the τ parameter. Note that the more expensive to compute $\bar{r}_w(\tau)$ can be evaluated once and stored. Figure 13 shows the integrand of equation (C8) for the TDI A-channel evaluated between the parameters for the maximum likelihood solution from the search and the true source parameters. The integrand is rapidly oscillating within an exponentially damped envelope. The fast decay of the envelope allows us to truncate the τ integral to one thousandth of its full extent, which accounts for the

increase in speed between the full and heterodyned likelihood.

The final term in the likelihood can be handled in a similar fashion to the previous term:

$$\int \frac{\bar{r}\bar{r}^*}{S_n(f)} df = \int R(f)S(f)df = \int R^*(\tau)S(\tau)d\tau, \quad (\text{C9})$$

where $R(f) = \bar{r}\bar{r}^*/\bar{S}_n(f)$, $S(f) = \bar{S}_n(f)/S_n(f)$ and $R(\tau)$ and $S(\tau)$ are their Fourier transforms. As with $\bar{r}_w(\tau)$, the expensive to compute $R(\tau)$ can be evaluated once and stored.

The reference integrals $\bar{r}_w(\tau)$ and $R(\tau)$ are calculated at the full sample cadence of the data, while the slowly varying terms such as (C3) are computed on a coarse spline in frequency. In the LISA context we want to ensure that the orbital motion of the constellation is adequately sampled, so we use the leading post-Newtonian expression for \dot{f} to set the frequency spacing df :

$$df = \dot{f}dT = (8\pi)^{8/3} \frac{3}{40} \mathcal{M}^{5/3} dT. \quad (\text{C10})$$

Setting $dT = 3 \times 10^5$ seconds yields rough 100 hundred samples per year during the early inspiral. To ensure that the dynamic frequency spacing is never too fine or too coarse we set $df_{\min} = 1/T_{\text{obs}}$ and $df_{\max} = f_{\max}/100$, where f_{\max} is the maximum frequency reached by the signal (set equal to $f_{\max} = 2f_{\text{ring}}$, where f_{ring} is the ringdown frequency, when the merger occurs within the observation time). With these choices, the frequency stencil typically has between 100 and 500 points for a one year data set.

For the discrete FFT used to compute $\Delta h_w(\tau)$ we settled on $N = 4096$ points in a trade-off between speed accuracy. This choice delivered at accuracy of order ± 0.3 for the $(n|\Delta h)$ term, with an evaluation time of 5 ms on a single 2.6 GHz core. The accuracy should be compared to the expected value and variance for this term, $E[(n|\Delta h)] \approx E[(n|h_{,i}(n|h_{,j})\Gamma^{ij})] = D$, and $\text{Var}[(n|\Delta h)] = D$. The standard deviation of the $(n|\Delta h)$ term, $\sigma_{(n|\Delta h)} = \sqrt{D} = 3.3$, is much larger than the numerical error.

[1] B. P. Abbott *et al.* (Virgo, LIGO Scientific Collaborations), *Phys. Rev. Lett.* **116**, 061102 (2016).
 [2] B. Abbott *et al.* (LIGO Scientific, Virgo Collaborations), *Phys. Rev. X* **9**, 031040 (2019).
 [3] B. Abbott *et al.* (LIGO Scientific, Virgo Collaborations), *Astrophys. J.* **882**, L24 (2019).

[4] P. Amaro-Seoane *et al.* (LISA Collaboration), [arXiv:1702.00786](https://arxiv.org/abs/1702.00786).
 [5] J. Baker *et al.*, [arXiv:1907.06482](https://arxiv.org/abs/1907.06482).
 [6] N. J. Cornish and E. K. Porter, *Classical Quantum Gravity* **23**, S761 (2006).
 [7] N. J. Cornish and E. K. Porter, *Phys. Rev. D* **75**, 021301 (2007).

- [8] N. J. Cornish and E. K. Porter, *Classical Quantum Gravity* **24**, 5729 (2007).
- [9] N. J. Cornish and E. K. Porter, *Classical Quantum Gravity* **24**, S501 (2007).
- [10] C. Rover *et al.*, *Classical Quantum Gravity* **24**, S521 (2007).
- [11] S. Babak, M. Hannam, S. Husa, and B. F. Schutz, [arXiv:0806.1591](https://arxiv.org/abs/0806.1591).
- [12] E. K. Porter and N. J. Cornish, *Phys. Rev. D* **78**, 064005 (2008).
- [13] A. Petiteau, S. Yu, and S. Babak, *Classical Quantum Gravity* **26**, 204011 (2009).
- [14] J. R. Gair and E. K. Porter, *Classical Quantum Gravity* **26**, 225004 (2009).
- [15] S. Marsat, J. G. Baker, and T. Dal Canton, [arXiv:2003.00357](https://arxiv.org/abs/2003.00357) [*Phys. Rev. D* (to be published)].
- [16] M. L. Katz, S. Marsat, A. J. Chua, S. Babak, and S. L. Larson, [arXiv:2005.01827](https://arxiv.org/abs/2005.01827).
- [17] S. Babak and A. Petiteau, (the new) LISA Data Challenges, <https://lisa-ldc.lal.in2p3.fr>.
- [18] K. A. Arnaud *et al.*, *AIP Conf. Proc.* **873**, 619 (2006).
- [19] S. Babak *et al.*, *Classical Quantum Gravity* **25**, 184026 (2008).
- [20] S. Babak *et al.* (Mock LISA Data Challenge Task Force Collaborations), *Classical Quantum Gravity* **27**, 084009 (2010).
- [21] N. Cornish and K. Shuman, LISA Black Hole Search, <https://zenodo.org/badge/latestdoi/262135605>.
- [22] S. Husa, S. Khan, M. Hannam, M. Prer, F. Ohme, X. Jimnez Forteza, and A. Boh, *Phys. Rev. D* **93**, 044006 (2016).
- [23] S. Khan, S. Husa, M. Hannam, F. Ohme, M. Prer, X. Jimnez Forteza, and A. Boh, *Phys. Rev. D* **93**, 044007 (2016).
- [24] S. Babak and A. Petiteau, LISA Data Challenge Manual, <https://lisa-ldc.lal.in2p3.fr/static/data/pdf/LDC-manual-001.pdf>.
- [25] B. P. Abbott *et al.* (LIGO Scientific, Virgo Collaborations), *Classical Quantum Gravity* **37**, 055002 (2020).
- [26] R. H. Swendsen and J.-S. Wang, *Phys. Rev. Lett.* **57**, 2607 (1986).
- [27] P. Jaranowski, A. Krolak, and B. F. Schutz, *Phys. Rev. D* **58**, 063001 (1998).
- [28] A. Sesana, J. Gair, E. Berti, and M. Volonteri, *Phys. Rev. D* **83**, 044036 (2011).
- [29] A. Klein, E. Barausse, A. Sesana, A. Petiteau, E. Berti, S. Babak, J. Gair, S. Aoudia, I. Hinder, F. Ohme *et al.*, *Phys. Rev. D* **93**, 024003 (2016).
- [30] M. Bonetti, A. Sesana, F. Haardt, E. Barausse, and M. Colpi, *Mon. Not. R. Astron. Soc.* **486**, 4044 (2019).
- [31] N. Cornish and T. Robson, *J. Phys. Conf. Ser.* **840**, 012024 (2017).
- [32] A. Lamberts, S. Blunt, T. B. Littenberg, S. Garrison-Kimmel, T. Kupfer, and R. E. Sanderson, *Mon. Not. R. Astron. Soc.* **490**, 5888 (2019).
- [33] J. R. Gair, S. Babak, A. Sesana, P. Amaro-Seoane, E. Barausse, C. P. Berry, E. Berti, and C. Sopuerta, *J. Phys. Conf. Ser.* **840**, 012021 (2017).
- [34] C. J. Moore, D. Gerosa, and A. Klein, *Mon. Not. R. Astron. Soc.* **488**, L94 (2019).
- [35] N. J. Cornish, Black hole merging and gravitational waves, in *Black Hole Formation and Growth*, edited by R. Walter, P. Jetzer, L. Mayer, and N. Produit, Saas-Fee Advanced Course Vol. 48 (Springer, Berlin, Heidelberg, 2019), pp. 1–92, https://doi.org/10.1007/978-3-662-59799-6_1.
- [36] S. E. Field, C. R. Galley, F. Herrmann, J. S. Hesthaven, E. Ochsner, and M. Tiglio, *Phys. Rev. Lett.* **106**, 221102 (2011).
- [37] P. Canizares, S. E. Field, J. Gair, V. Raymond, R. Smith, and M. Tiglio, *Phys. Rev. Lett.* **114**, 071104 (2015).
- [38] N. J. Cornish, [arXiv:1007.4820](https://arxiv.org/abs/1007.4820).
- [39] B. Zackay, L. Dai, and T. Venumadhav, [arXiv:1806.08792](https://arxiv.org/abs/1806.08792).
- [40] C. J. Ter Braak, *Stat. Comput.* **16**, 239 (2006).
- [41] T. Littenberg, N. Cornish, K. Lackeos, and T. Robson, [arXiv:2004.08464](https://arxiv.org/abs/2004.08464) [*Phys. Rev. D* (to be published)].
- [42] B. Bcsy and N. J. Cornish, [arXiv:1912.08807](https://arxiv.org/abs/1912.08807).
- [43] J. Shapiro Key and N. J. Cornish, *Phys. Rev. D* **79**, 043014 (2009).
- [44] M. I. Cohen, C. Cutler, and M. Vallisneri, *Classical Quantum Gravity* **27**, 185012 (2010).
- [45] T. Robson and N. J. Cornish, *Phys. Rev. D* **99**, 024019 (2019).
- [46] C. Cutler, *Phys. Rev. D* **57**, 7089 (1998).
- [47] L. J. Rubbo, N. J. Cornish, and O. Poujade, *Phys. Rev. D* **69**, 082003 (2004).