

A more effective coordinate system for parameter estimation of precessing compact binaries from gravitational waves

Benjamin Farr,^{1,2,*} Evan Ochsner,^{3,†} Will M. Farr,^{2,‡} and Richard O’Shaughnessy³

¹*Center for Interdisciplinary Exploration and Research in Astrophysics (CIERA) & Department of Physics and Astronomy, Northwestern University, 2145 Sheridan Road, Evanston, Illinois 60208, USA*

²*School of Physics and Astronomy, University of Birmingham, Edgbaston, Birmingham B15 2TT, United Kingdom*

³*Center for Gravitation and Cosmology, University of Wisconsin-Milwaukee, Milwaukee, Wisconsin 53201, USA*

(Received 29 April 2014; published 7 July 2014)

Ground-based gravitational wave detectors are sensitive to a narrow range of frequencies, effectively taking a snapshot of merging compact-object binary dynamics just before merger. We demonstrate that by adopting analysis parameters that naturally characterize this “picture,” the physical parameters of the system can be extracted more efficiently from the gravitational wave data and interpreted more easily. We assess the performance of Markov chain Monte Carlo parameter estimation in this physically intuitive coordinate system, defined by (a) a frame anchored on the binary’s spins and orbital angular momentum and (b) a time at which the detectors are most sensitive to the binary’s gravitational wave emission. Using anticipated noise curves for the advanced-generation LIGO and Virgo gravitational wave detectors, we find that this careful choice of reference frame and reference time significantly improves parameter estimation efficiency for binary neutron stars, neutron star–black hole, and binary black hole signals.

DOI: [10.1103/PhysRevD.90.024018](https://doi.org/10.1103/PhysRevD.90.024018)

PACS numbers: 04.30.-w, 97.60.Lf, 97.60.Jd

I. INTRODUCTION

The Advanced LIGO (aLIGO) gravitational-wave detectors are expected to come online in 2015 [1], with Advanced Virgo (AdV) following in 2016 [2]. They are expected to directly detect gravitational wave (GW) signals and, once they reach design sensitivity, could detect tens of events per year [3] from the coalescence of compact binaries composed of neutron stars (NSs) and/or black holes (BHs). Black holes in particular are expected to have large spin [4], so properly modeling spins will be especially important for neutron star–black hole (NS-BH) and binary black hole (BBH) systems. Spin-induced precession, caused when one or both spin vectors are misaligned with the orbital angular momentum, is a particularly challenging effect to model, as it induces amplitude and phase modulations, changes the relative strength of the two waveform polarizations, and complicates a spin-weighted spherical harmonic mode decomposition of the waveform [5,6].

In an important early work, Apostolatos *et al.* [7] laid out a “simple precession” model for the evolution and gravitational wave emission of compact binaries with generic spins. With the exception of rare cases of transitional precession, or the simpler special case of spins aligned or antialigned with the orbital angular momentum, the orbital angular momentum vector of the binary will precess on a

cone about the total angular momentum. Since this and other important early work [5,8], there have been a number of improvements to waveform models from spinning systems. This has included deriving and studying further post-Newtonian (PN) corrections to the waveform dynamics and phasing [9–12], amplitude [5,6] and precession equations [13], the development of spinning, precessing inspiral-merger-ringdown waveforms [14–17] and a frequency-domain precessing waveform model [18], and efforts to track precessional motion and disentangle it from other dynamical effects [19–23]. Despite all of these refinements, precessing waveforms still qualitatively match the Apostolatos *et al.* picture of the orbital angular momentum moving along a precession cone that slowly grows due to radiation reaction.

Although suboptimal, matched filter searches with non-spinning or aligned-spin templates could still detect GW signals from spinning, precessing binaries, albeit with significant parameter bias [24]. If a detection can be made, then the LALINFERENCE module of LAL [25] can be used for a focused parameter estimation follow-up effort [26]. LALINFERENCE is a suite of routines to perform Bayesian inference techniques such as Markov chain Monte Carlo (MCMC) and nested sampling on gravitational wave detector data. Bayesian inference has proven adept at sampling the 11–15-dimensional parameter space of circularized compact binary mergers with non-negligible spin [26–31].

Nonetheless, there is a significant computational cost to performing Bayesian inference. Typically, waveform

*bfarr@u.northwestern.edu

†evano@gravity.phys.uwm.edu

‡w.farr@bham.ac.uk

generation is the dominant cost. Performing parameter estimation with MCMC on real or simulated data typically requires generating several million compact binary coalescence (CBC) waveforms to produce ≈ 1000 independent samples from the posterior probability density on parameter space. Depending on the computational cost of generating the waveforms, this can take hours to weeks or longer to complete. The latency of such analyses must, at the very least, be low enough to keep up with potential advanced-detector trigger rates. If these analyses are completed within hours, then they can potentially play a critical role in the search for electromagnetic counterparts to binary neutron stars (BNS) or NS-BH GW signals.

The latency of MCMC analyses, assuming fixed waveform generation costs, are governed by the sampling efficiency of the chains, which are intimately tied to the parametrization and proposal distributions that are used. In an extreme limit, local one-dimensional jumps and distorted coordinates require the Markov chain to jump slowly through tightly correlated, twisting paths in parameter space.

In this work we demonstrate that MCMC-based parameter estimation is significantly more efficient when we adopt coordinates well adapted to the dynamics of the precessing binary. We parametrize the spin and orbital angular momentum degrees of freedom of a precessing binary with a set of angles describing the simple precession cone model of Apostolatos *et al.* [7]. We use the inclination of the *total* angular momentum to the line of sight, the azimuthal position of orbital angular momentum L_N on this cone at some reference point, three angles describing the orientation of the spins relative to L_N and each other at the reference point, and the magnitude of each spin. Critically, we also choose the GW frequency (i.e., twice the instantaneous orbital frequency) f_{ref} at which these angles are defined. Choosing such a reference frequency near the peak sensitivity of the detector also dramatically improves the convergence of Bayesian parameter estimation methods.

We find that the new parametrization and a suitable choice of f_{ref} can decrease the “mixing” or autocorrelation time in the stochastic parameter sampling by a factor of 3.7–11 for the cases considered here. We find the shortest autocorrelation times and best parameter constraints occur for $f_{\text{ref}} \approx 70$ –100 Hz. Previously, other authors have proposed specifying parameters at a fixed reference frequency and in a frame relative to the total angular momentum of the system, both for nonprecessing and single-spin binaries [32,33]. However, this work is the first to demonstrate that such a parametrization can improve the parameter estimation of precessing binaries.

This work is organized as follows. In Sec. II, we briefly review the dynamics of precessing binaries and the commonly used “radiation frame” coordinates, and we introduce our well-adapted coordinate system for generic precessing binaries. In Sec. III we very briefly describe

the parallel-tempered Markov chain Monte Carlo parameter estimation strategy adopted in this work, emphasizing why well-adapted coordinates improve its performance. In Sec. IV, we compare the results of parameter estimation calculations performed using radiation frame coordinates and our well-adapted coordinates. We show that MCMC calculations using the well-adapted coordinates converge more efficiently, providing reliable results in less time. Furthermore, our well-adapted coordinates correspond to physically pertinent, well-constrained observables, so the posteriors for astrophysically interesting angles such as spin tilts are obtained directly without further (expensive) postprocessing. We finish with some concluding remarks in Sec. V.

II. COORDINATES FOR PRECESSING BINARIES

A. Evolution equations for precessing PN binaries

Here for convenience we briefly review the formalism for generating PN precessing waveforms. For more details, we refer the reader to [34] (and references therein) for a summary of various nonspinning PN waveforms. For a generalization to spinning waveforms, the reader may refer to [5,6,35,36] among others.

The gravitational wave strain observed by a detector is given by¹

$$h(t) = F_+ h_+(t) + F_\times h_\times(t), \quad (2.1)$$

where F_+ and F_\times are antenna pattern functions describing the detector and $h_+(t)$ and $h_\times(t)$ are the gravitational wave polarizations. For a compact binary evolving along a series of quasicircular orbits, these functions have the following form at leading order:

$$h_+(t) = -\frac{2M\eta}{D_L} v(t)^2 (1 + (\hat{L} \cdot \hat{N})^2) \cos 2\phi(t), \quad (2.2)$$

$$h_\times(t) = -\frac{2M\eta}{D_L} v(t)^2 2(\hat{L} \cdot \hat{N}) \sin 2\phi(t). \quad (2.3)$$

Here D_L is the luminosity distance to the binary, M is the total mass, $\eta = m_1 m_2 / M^2$ is the symmetric mass ratio, \hat{L} is the direction of orbital angular momentum, \hat{N} is the direction of GW propagation, ϕ is the orbital phase of the binary, and $v = (2\pi M F_{\text{orb}})^{1/3}$ is the “characteristic velocity” PN expansion parameter (with F_{orb} the orbital frequency). Higher-order corrections to the polarizations valid for precessing binaries can be found in [5,6]. For our purposes, the important point is that these polarizations depend only on the masses, the inclination between the

¹We note that this expression assumes that the sky position of the source remains constant, i.e., that the Earth rotates by a negligible amount over the duration of the signal. As NS-BH GW signals evolve from 10 Hz to coalescence in $\lesssim 5$ min, this is true to a very good approximation.

orbital angular momentum and the line of sight, and the time-dependent phasing and frequency parameters $\phi(t)$ and $v(t)$.²

For compact binaries on quasicircular orbits, $\phi(t)$ and $v(t)$ can be computed via the energy balance equation. PN expansions are known for both the binding energy of the binary, E , and the gravitational wave luminosity (commonly called the “flux”), \mathcal{F} . We assume that emission of gravitational waves accounts for all of the loss of binding energy and a simple use of the chain rule provides a differential equation for the evolution of $v(t)$,

$$-\frac{dE}{dt} = \frac{dE}{dv} \frac{dv}{dt} = \mathcal{F} \Rightarrow \frac{dv}{dt} = -\frac{\mathcal{F}}{dE/dv}. \quad (2.4)$$

Since the derivative of the orbital phase is simply the (angular) orbital frequency, we trivially obtain a coupled differential equation for the phase $\phi(t)$,

$$\frac{d\phi}{dt} = 2\pi F_{\text{orb}} = \frac{v^3}{M}. \quad (2.5)$$

For nonspinning binaries, these are the only equations needed to evolve the orbital dynamics and compute a gravitational waveform. Equations (2.4) and (2.5) are integrated to compute $v(t)$ and $\phi(t)$, which are then plugged back into Eqs. (2.2) and (2.3) to obtain the waveform. There are a number of different ways to solve these differential equations which are equivalent up to the PN order of E and \mathcal{F} , but differ in truncation error at the next, unknown PN order. These different methods for solving the energy balance equation are known as PN *approximants*. For example, E and \mathcal{F} are known as Taylor series in v , so the right hand side of Eq. (2.4) is a rational function of v . One could keep it in this form (referred to as the TaylorT1 approximant), or reexpand it as a Taylor series (TaylorT4). In this work we will use TaylorT4, but our results are equally applicable to any PN approximant. See [34] for a summary of the various PN approximants, and [18,37] for two newly proposed approximants.

For precessing binaries, we note that there are spin corrections proportional to $\hat{S}_{1,2} \cdot \hat{L}$ and $\hat{S}_1 \cdot \hat{S}_2$ in E and \mathcal{F} . Additionally, we have already noted that $\hat{L} \cdot \hat{N}$ appears in the polarizations. The orientations of these vectors can change over time for precessing binaries, and so they must be computed as functions to time, along with $v(t)$ and $\phi(t)$. Their evolution is given by the precession equations [5,7].

$$\begin{aligned} \frac{d\hat{L}}{dt} = \frac{v^6}{2M^3} \left\{ \left[\left(4 + 3 \frac{m_2}{m_1} \right) \mathbf{S}_1 + \left(4 + 3 \frac{m_1}{m_2} \right) \mathbf{S}_2 \right] \right. \\ \left. - \frac{3v}{M^2 \eta} [(\mathbf{S}_2 \cdot \hat{L}) \mathbf{S}_1 + (\mathbf{S}_1 \cdot \hat{L}) \mathbf{S}_2] \right\} \times \hat{L}, \end{aligned} \quad (2.6)$$

²At higher order, the spins also enter the polarizations. Our main point is that Eqs. (2.4)–(2.8) are sufficient to compute the polarizations, and this is true at all PN orders.

$$\begin{aligned} \frac{d\mathbf{S}_1}{dt} = \frac{v^5}{2M} \left\{ \left(4 + 3 \frac{m_2}{m_1} \right) \hat{L} \right. \\ \left. + \frac{v}{M^2} [\mathbf{S}_2 - 3(\mathbf{S}_2 \cdot \hat{L}) \hat{L}] \right\} \times \mathbf{S}_1, \end{aligned} \quad (2.7)$$

$$\begin{aligned} \frac{d\mathbf{S}_2}{dt} = \frac{v^5}{2M} \left\{ \left(4 + 3 \frac{m_1}{m_2} \right) \hat{L} \right. \\ \left. + \frac{v}{M^2} [\mathbf{S}_1 - 3(\mathbf{S}_1 \cdot \hat{L}) \hat{L}] \right\} \times \mathbf{S}_2. \end{aligned} \quad (2.8)$$

To generate a gravitational wave signal from a precessing binary, we first solve for its orbit and spin [Eqs. (2.4)–(2.8)], and then substitute these kinematic quantities into Eqs. (2.2) and (2.3).

Higher-order polarization expressions valid for precessing binaries can be found in [6]. While in this work we report results using only leading-order polarizations, related work by [29,30] suggests our conclusions will also hold for higher-order polarizations.

B. Behavior of precessing waveforms

The behavior of precessing binaries was first laid out in significant qualitative and quantitative detail in [7]. To summarize, they find that almost all spinning binary configurations would undergo simple precession. This means that the direction of total angular momentum, \hat{J} , remains very nearly fixed and the orbital angular momentum vector moves on a cone about \hat{J} . The angle between \mathbf{L} and the total spin, $\mathbf{S} = \mathbf{S}_1 + \mathbf{S}_2$, will remain nearly constant throughout the binary evolution. The magnitude of \mathbf{L} will slowly decrease due to radiation reaction, and since its angle with \mathbf{S} does not change appreciably, this means the opening angle of the \mathbf{L} precession cone will slowly grow as the binary inspirals.

There are two types of special cases of spinning binary configurations which do not obey simple precession. First, if both spins are aligned and/or antialigned with the orbital angular momentum, then the binary will not precess and the orbital angular momentum will remain in a fixed direction. It is clear from Eqs. (2.6)–(2.8) that the direction of all of these vectors will be constant when they are all parallel. Second, transitional precession can occur if the binary has large spins which are nearly, but not perfectly, antialigned with the orbital angular momentum. In this case, the binary will initially be in a state where the orbital angular momentum is the dominant contribution to \mathbf{J} .³ Gravitational wave emission will decrease the magnitude of \mathbf{L} , but will not change the magnitude of \mathbf{S}_1 and \mathbf{S}_2 . If it decreases such that $|\mathbf{L}| \lesssim |\mathbf{S}|$, then the direction of \hat{J} will change from being near \hat{L} to near \hat{S} . This brief period when

³Since $\mathbf{L} \approx \mathbf{r} \times \mathbf{p}$, the orbital angular momentum will always dominate the spin angular momentum for sufficiently wide binary separations r .

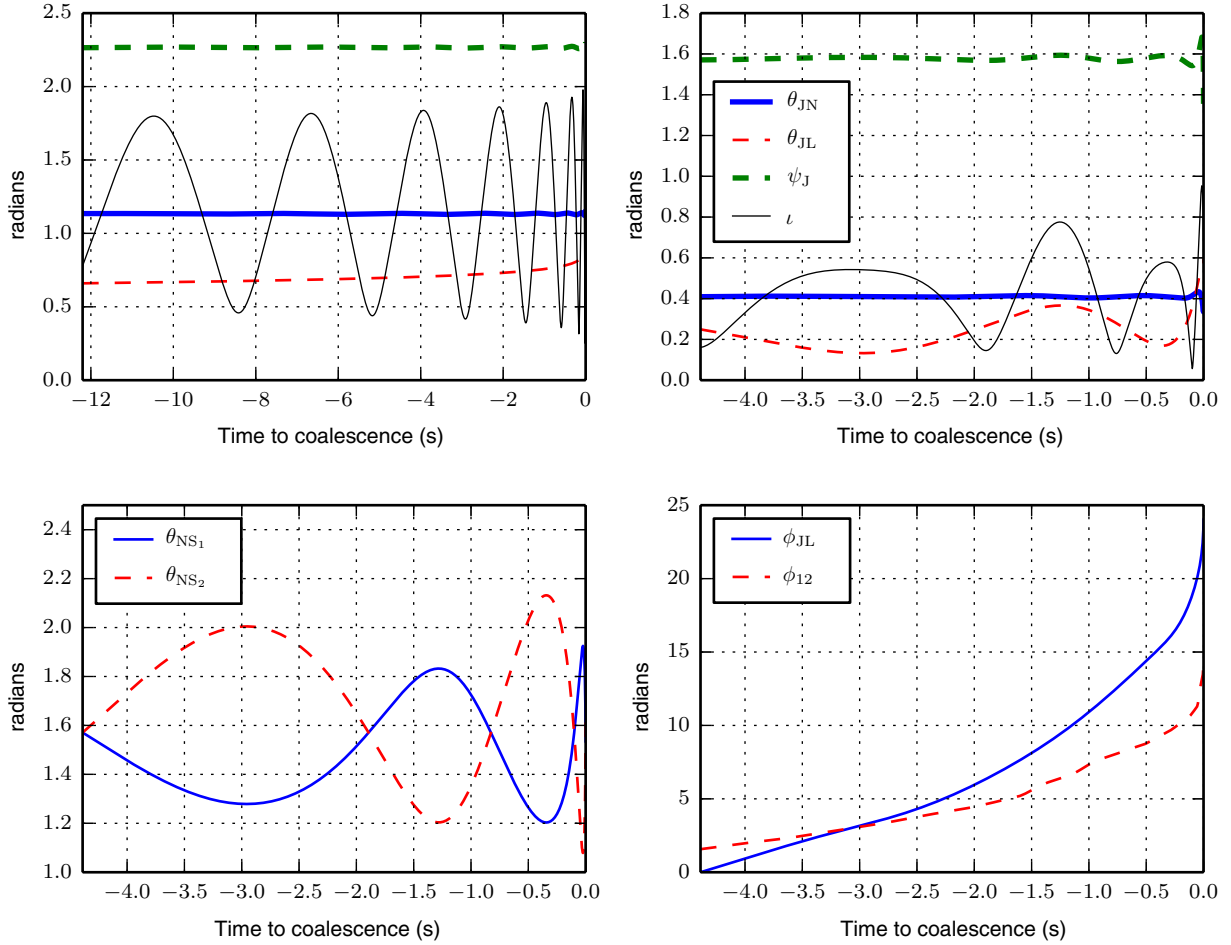


FIG. 1 (color online). *Fiducial binary orientation:* We plot the evolution of angles describing the orientation of our fiducial NS-BH and BBH binaries versus time to coalescence. For both the NS-BH (top left) and the BBH (top right) binaries, the radiation-frame angle i evolves significantly, while the system-frame angles ψ_J , θ_{JL} , θ_{JN} evolve slowly. For the single-spin NS-BH binary the spin maintains a constant tilt relative to \hat{L} (not shown). For the double-spin BBH binary, the tilts of the component spins relative to \hat{L} do vary by a few tenths of a radian (lower left panel). In the lower right panel we plot the azimuthal angles ϕ_{JL} describing the motion of \hat{L} moving in its cone about \hat{J} , and ϕ_{12} , the azimuthal separation of component spins measured relative to \hat{L} for the BBH binary.

$|\mathbf{J}| \approx 0$ and the direction of \hat{J} changes rapidly is known as transitional precession. Note that it is preceded and proceeded by periods of simple precession.

We note that the simple precession model presented in [7] primarily focused on the special cases of equal masses and/or single spin binaries. This was so the precession equations (2.6)–(2.8) would imply that $\mathbf{S}_1 \cdot \mathbf{S}_2 = \text{const}$ and simplify the behavior. However, even for a generic case with unequal masses and spins, the simple precession model is still a good *qualitative* description of the precessing binary's behavior. One difference is that the orbital angular momentum exhibits nutation. That is, \hat{L} bobs up and down as it moves on its cone about \hat{J} . The cone opening angle still tends to grow as orbital angular momentum is radiated away, but it is no longer a monotonic increase. Compare the single spin NS-BH binary to the double spin BBH binary in Fig. 1. The angle between \hat{J} and \hat{L} , θ_{JL} , nutates strongly in the latter but not in the former.

Note, however, that the angles θ_{JN} and ψ_J , the inclination and polarization angles of the total angular momentum, respectively, are nearly constant in each case, as \hat{J} remains essentially fixed. Additionally, Fig. 1 shows the tilt of the component spins and the total spin $\mathbf{S} = \mathbf{S}_1 + \mathbf{S}_2$ relative to \hat{L} . Note that these angles can change by a few tenths of a radian (with the total spin oscillating less than the individual spins), while for the single-spin case (not shown) $\hat{L} \cdot \hat{S}_1$ is constant throughout the evolution.

C. Previous coordinate conventions

As we have summarized, computing the orbital dynamics of a precessing post-Newtonian waveform involves numerically integrating the coupled set of ordinary differential equations (ODEs) in (2.4)–(2.8). Typically, one specifies a binary with parameters m_1 and m_2 , plus the initial Cartesian components of the spin and orbital angular

momentum vectors specified at the minimum frequency at which the signal enters the detector's sensitive band. This is indeed an obvious, natural choice to specify the initial conditions of the ODEs and is what has previously been used to generate CBC waveforms in the LAL software.

Actually, for precessing binaries LAL codes adopt a simplifying convention, without loss of generality, so that the orbital angular momentum vector can be specified by a single inclination angle, ι . In particular, the z axis of the Cartesian frame is taken to be the direction of propagation of the gravitational wave. The orbital angular momentum is assumed to lie in the x - z plane, such that $\hat{L} = (\sin \iota, 0, \cos \iota)$, and hence such that the projection of \hat{L} on the plane of the sky is in the direction $(\cos \psi_L, \sin \psi_L) = (1, 0)$. Therefore, the LAL waveform generation routines depend on the parameters

$$\{m_1, m_2, \iota, (\mathbf{S}_{1x}, \mathbf{S}_{1y}, \mathbf{S}_{1z}), (\mathbf{S}_{2x}, \mathbf{S}_{2y}, \mathbf{S}_{2z})\}. \quad (2.9)$$

This Cartesian parametrization, with the z axis along the direction of GW propagation, is commonly called the *radiation frame*. Previous parameter estimation efforts, such as [26–28,30], would estimate these radiation frame parameters.⁴

While the radiation frame is very convenient for generating waveforms, it has several drawbacks for parameter estimation. First of all, the Cartesian angular momentum coordinates of a binary, and hence their projection on the plane of the sky (i.e., ψ_L), can vary considerably over its evolution. Figure 1 shows how the inclination of the orbital plane (ι) varies across the observed signal, and the Cartesian components of \hat{S}_i (not shown) cover most of their allowed range $[-1, 1]$. Otherwise identical binaries that happen to be at two different points along the precession cone at the reference point will have very different Cartesian components. Thus, systems that produce similar GW signals are spread across the parameter space in complex ways, and parameter estimation analyses must map out these complicated correlations. Furthermore, the Cartesian components are specified at the low frequency limit, where detectors have poor sensitivity. Two binaries with similar vector components at f_{\min} might have rather different component values when the signal is in the sensitive band of the detector.

D. Nearly conserved coordinates

Motivated by the simple precession picture of Apostolatos *et al.*, we parametrize the binary configuration via a set of angles that describe the position and shape

⁴Additionally, the time and phase of some reference point t_{ref} , ϕ_{ref} , sky location (δ, α) , luminosity distance D_L , and polarization angle ψ_L are needed, but they merely describe the orientation of the binary relative to the observer and do not affect the orbital dynamics in any way.

of the precession cone, as well as where \hat{L} is along its cone. In particular, we parametrize a binary configuration with

$$\{m_1, m_2, \chi_1, \chi_2, \theta_{\text{JN}}, \theta_{\text{LS}_1}, \theta_{\text{LS}_2}, \phi_{12}, \phi_{\text{JL}}\}. \quad (2.10)$$

Here $0 \leq \chi_{1,2} \leq 1$ are the spin magnitudes, θ_{JN} is the inclination between the total angular momentum and the direction of propagation, $\theta_{\text{LS}_{1,2}}$ are the inclinations of each spin relative to \hat{L} (commonly referred to as tilts), ϕ_{12} is the azimuthal angle of $\hat{S}_2 - \hat{S}_1$ measured relative to \hat{L} ,⁵ and ϕ_{JL} is the azimuthal position of \hat{L} on its cone about \hat{J} (with \hat{N} setting the zero of azimuth). We refer to this parametrization as precessing system coordinates (or the “system frame” for short).

We note that this parametrization, like the standard Cartesian one, requires nine parameters to specify the orbital dynamics. The system frame parametrization captures all degrees of freedom in the binary but does not overdetermine it. Given values for all the parameters of either the system or the radiation frame, one can compute the values of the other parametrization through a series of Euler rotations. We have implemented such transformations in the LAL software. As a practical implementation, parameter estimation routines propose values for the parameters of Eq. (2.10); these are transformed into the parameters of Eq. (2.9) and passed to waveform generation routines which take input in terms of these parameters.

An advantage of the precessing system frame is that many of the parameters are nearly conserved for the duration that the signal is in band. It is well known the masses and spin magnitudes do not change significantly during inspiral, and Apostolatos *et al.* showed that θ_{JN} is very nearly constant. For the case of a single-spin binary, Apostolatos *et al.* showed θ_{LS_1} is constant at leading order. For a double-spin binary, the angles $\theta_{\text{LS}_1}, \theta_{\text{LS}_2}, \phi_{12}$ need not be conserved—indeed Fig. 1 shows the first two vary by a few tenths of a radian, while ϕ_{12} grows through approximately two full cycles. However, it should be noted that the BBH binary plotted in Fig. 1 was intentionally chosen as an extreme case. Many double spin binaries will have significantly less variation in these angles. For example, there are known to be *spin resonances* where both spins and \hat{L} get locked into a coplanar configuration [38]. Last, ϕ_{JL} is not constant, but does increase steadily and monotonically, essentially chirping on a precessional time scale. As we will see in Sec. IV, using these nearly conserved coordinates improves the convergence rates of Bayesian parameter estimation.

⁵Only the relative azimuthal difference between the spin vectors matters. The absolute azimuth, relative to, say \hat{N} , is implicitly set by ϕ_{JL} .

E. In-band coordinates

Precessional motion causes amplitude and phase modulations, and can change the polarization content of the observed waveform (i.e., the relative strength of h_+ and h_\times). For example, we see from Eqs. (2.2) and (2.3) that h_+ and h_\times polarizations will have equal strength when the binary is face on, but h_\times vanishes when the binary is edge on. The detectors will essentially be measuring the polarization content when the signal is in the “bucket,” or the most sensitive frequency band of the detector. They will necessarily be less sensitive when the binary is at lower frequencies near the seismic cutoff. Therefore, it stands to reason that it is more difficult to constrain the orientation of the binary at low frequencies, where there is little sensitivity, than at frequencies of peak sensitivity, which are typically $\sim 100\text{--}200$ Hz for ground-based detectors.

It is therefore unfortunate that waveform generation routines and parameter estimation jump proposals in LAL have until now specified initial conditions for binary orientation at the low frequency limit, where the sensitivity is worst. If the orientations of two binaries are similar at low frequencies, they need not be similar in the bucket. For example, if the masses and/or spin magnitudes are a bit different, they may precess at different rates and could be at different points along similar precession cones when in the bucket, thus having very different polarization content for certain observers. Or, they could move along very dissimilar precession cones that happen to be nearly tangent at a certain point at low frequency.

Fortunately, we note that the differential equations needed to evolve the orbital dynamics, Eqs. (2.4)–(2.8), can be integrated backward in time just as easily as forward. The same is true for any post-Newtonian waveform model. As a practical implementation, we choose some reference frequency, f_{ref} , and specify the “initial” conditions for the binary orientation at that gravitational-wave frequency. We then make two calls to evolve the orbital dynamics: one integrates Eqs. (2.4)–(2.8) backward in time until it reaches the minimum frequency of detector sensitivity f_{min} , and the other integrates forward in time to a frequency f_{end} , which can be the minimum energy circular orbit (MECO) or some other waveform stopping condition. We then stitch together the waveform time series from each integration to get a seamless waveform that covers the full frequency range $[f_{\text{min}}, f_{\text{end}}]$. As we have implemented it, this method of two-way integration agrees with the standard approach of forward integration from f_{min} to within numerical precision, and there is virtually no difference in the speed of waveform generation.

In principal, we can choose f_{ref} to be any value which the binary reaches before the MECO or other termination condition. However, as expected, we find in Sec. IV that parameter estimation codes are most efficient when choosing f_{ref} near where the detector has peak sensitivity.

III. PARAMETER ESTIMATION WITH MARKOV CHAIN MONTE CARLO

For CBC parameter estimation, MCMC methods are used to sample the full 15-dimensional posterior distribution as a function of the parameters describing the circularized compact binary merger. These methods employ serial Markov chains that stochastically wander the parameter space through the use of various proposal distributions. By accepting or rejecting proposed jumps according to the Metropolis-Hastings ratio [39,40], samples recorded by the Markov chains are distributed with a density proportional to the target posterior probability density. To ensure each sample is an independent draw from the posterior distribution, samples are first thinned based on the correlations present in the chain. Thus, assuming the likelihood function is equally expensive to compute at all times, the efficiency of the MCMC sampler will ultimately be decided by the maximum one-dimensional (1D) autocorrelation time (ACT), t_{max} , of the chains. We estimate the 1D autocorrelation time t for random variable X as the smallest s that satisfies

$$1 + \frac{2}{C_0} \sum_{\tau=1}^{Ms} C(\tau) < s, \quad (3.1)$$

where $C(\tau)$ is the autocorrelation function $C(\tau) = E[(X_t - \mu)(X_{t+\tau} - \mu)]/\sigma^2$, with μ and σ being the mean and standard deviation of X , respectively, and $C_0 = C(\tau = 0)$ is the zero-lag autocorrelation. M is a tunable parameter, ensuring the stability of the ACT estimate by requiring the length of the window used to estimate the ACT to be at least M times the estimated ACT. We have empirically found $M = 5$ to produce reliable ACT estimates.

To minimize ACTs, it is critical that the acceptance rates of jump proposals are balanced with the correlations they introduce. In general, the most efficient proposal distribution is the target distribution. The target distribution is typically unknown, and a collection of generally useful proposal distributions are used in its place. The most basic jump proposal typically used for MCMC sampling is a local Gaussian, centered on the current location of the chain. The width of this Gaussian in each dimension will affect the acceptance rates and ACTs. In the small-width limit jump acceptance rates approach 1; however, the very small steps of the chain greatly increase ACTs. In the large-width limit jump acceptance rates approach 0, also resulting in large ACTs due to many repeated samples in the chains. In the idealized case of an N -dimensional Gaussian target distribution composed of N independent one-dimensional Gaussians, it can be shown that the ideal acceptance rate of this proposal that minimizes ACTs is $\sim 23.4\%$ [41]. In the case of non-Gaussian target distributions, this acceptance rate is not necessarily optimal, and this proposal (even with optimized widths) can be very inefficient.

TABLE I. The precessing system parameters [see Eq. (2.10) and the surrounding text] for our BNS population and fiducial NS-BH and BBH binaries. For the NS-BH and BBH systems we provide the opening angle of the precession cone (θ_{JL}), the polarization angle of the total angular momentum (ψ_J), and the network SNR (ρ_{net}), for a three-detector network of early aLIGO-AdV detectors. A population of 15 BNS systems with masses and spins chosen uniformly from the specified ranges, orientations chosen isotropically, and the range of their SNRs in two-detector early aLIGO detectors are shown.

Binary type	m_1	m_2	χ_1	χ_2	θ_{JN}	θ_{LS_1}	θ_{LS_2}	ϕ_{12}	ϕ_{JL}	θ_{JL}	ψ_J	ρ_{net}
BNS	1.2–1.6	1.2–1.6	0–0.05	0–0.05	10.3–33.4
NS-BH	10	1.4	1	0	1.22	1.175	N/A	N/A	$\pi/4$	$\pi/4$	2.36	20.3
BBH	8	5	0.8	0.9	$\pi/6$	$\pi/2$	$\pi/2$	$\pi/2$	0	0.36	1.57	19.1

For CBC parameter estimation, the local Gaussian proposal distribution, though not the only one, is the proposal used most often. Since this proposal is optimal for Gaussian target distributions, the parametrization used for sampling should be chosen such that the posterior is as close to Gaussian as possible. For the purpose of determining the efficiency gains from the system frame, we will focus on the sampling of the intrinsic parameters: $\{\mathcal{M}, q, \iota, \chi_1, \theta_{NS_1}, \phi_1, \chi_2, \theta_{NS_2}, \phi_2\}$ for the radiation frame and $\{\mathcal{M}, q, \theta_{JN}, \phi_{JL}, \chi_1, \theta_{LS_1}, \chi_2, \theta_{LS_2}, \phi_{12}\}$ for the system frame, where $\mathcal{M} = (m_1 m_2)^{3/5} (m_1 + m_2)^{-1/5}$ is the chirp mass and q is the asymmetric mass ratio m_2/m_1 defined such that $0 < q \leq 1$.

IV. RESULTS

To assess the overall improvement in MCMC sampling efficiency, we have simulated GW inspiral signals from the three main types of compact binaries expected to be observed by advanced ground-based detectors: BNS, NS-BH, and BBH. The range of parameters studied are provided in Table I. For BNS systems we focus on the impact of the system frame on analyses during the first year of the advanced-detector era (2015), assuming Gaussian noise from the Hanford and Livingston detectors with “early” aLIGO sensitivity, as defined in [42].

For the fiducial NS-BH and BBH systems we move to the projected 2016 three-detector network, with “mid” aLIGO sensitivity for Hanford and Livingston and “early” AdV sensitivity for Virgo, as defined in [42]. For these systems we also consider the choice of reference frequency, and its impact on parameter constraints and sampling efficiency.

A. Binary neutron stars

With Virgo not coming online until 2016, the first year of the advanced-detector era will see only two operational detectors. Since binary neutron stars are observationally confirmed sources that aLIGO is expected to be sensitive to, we place particular emphasis on the impact of system frame analyses on BNS signals in the early advanced-detector era. To this end, we have randomly selected 15 BNS signals detected from an astrophysically distributed set of injections in a simulated two-detector network.

Motivated by the observed neutron star population to date, the injection population was chosen to have dimensionless spin magnitudes distributed uniformly below 0.05 [43], with isotropic spin orientations, and component masses drawn uniformly between 1.2 and 1.6 M_\odot . This injection set was distributed uniformly in the local universe, and 15 randomly selected signals detected by GSTLAL_INSPIRAL [44] were chosen for the purposes of this study [45].

We compare the maximum one-dimensional ACT from each MCMC chain sampling in the radiation frame ($f_{\text{ref}} = 40$ Hz) and system frame ($f_{\text{ref}} = 100$ Hz) parametrizations. We find that the precessing system frame gives ACTs which are shorter by 7.6, initially surprising due to the minimal effects of spin on the GW signal. Figure 2 shows the cumulative distribution of ACTs for the intrinsic parameters of the relevant frame across the 15 BNS systems. Even though these systems have very low spin, and therefore very little precession, a long-known degeneracy between mass ratio and spin magnitude [46–48] comes into play. In the simplified picture of a BNS system with spins aligned with the orbital angular momentum, changes to the waveform from increasing spin magnitude cannot be distinguished from increases in mass ratio, and vice versa. Because BNS systems are very close to equal mass, the posterior is highly skewed toward lower mass ratios, and thus higher spin magnitudes in the direction of the angular momentum. Support for antialigned spin would require support for higher mass ratios, which would be unphysical. Parametrically this is easily described in the system frame as $\cos(\theta_{LS_1}), \cos(\theta_{LS_2}) > 0$, resulting in an excluded region in $\theta_{LS_1} - \theta_{LS_2}$ space as seen in Fig. 3. In the radiation frame, however, this region of parameter space is nontrivially defined as a function of all four spin orientation angles, which no proposal used by the MCMC is able to navigate efficiently.

B. Neutron star–black hole and binary black hole systems

1. System frame efficiency

Based on current estimates from low-mass x-ray binaries, most black holes in binary systems are believed to have significant spin angular momentum [4]. Thus, even




FIG. 2 (color online). *Autocorrelation times*: Cumulative histogram of the largest intrinsic parameter ACT of each MCMC chain, for signals analyzed using the radiation frame (with parameters defined at 40 Hz) and precessing system frame (parameters defined at 100 Hz). *Left panel*: Results for the 15 BNS signals; using the system frame reduces ACTs by a median factor of 7.6. *Center panel*: Results for each analysis of the NS-BH binary described in Table I. The system frame is found to have a median improvement of a factor of 11 in efficiency over the radiation frame. *Right panel*: Results for each analysis of the BBH binary described in Table I. The system frame is found to have a median improvement of a factor of 3.7 in efficiency over the radiation frame.




FIG. 3 (color online). *Left panel*: MCMC samples from the posterior of a simulated BNS signal. The region where both spins have a negative \hat{L}_z component is excluded. This information is easily described, and sampled, in the system frame. *Right panel*: Posterior samples from the analyses of an NS-BH injection, centered on the mean. θ and ϕ correspond to θ_{NS_1} and ϕ_1 in the radiation frame, and θ_{LS_1} and ϕ_{JL} in the system frame. The strong, nonlinear correlation between spin parameters in the radiation frame is difficult to sample, resulting in large ACTs. The system frame parametrization eliminates this correlation, allowing for more efficient MCMC sampling.

slight misalignments for NS-BH and BBH systems will lead to precession of the orbital plane and modulation of the measured gravitational wave signal. For the case of NS-BH systems, the neutron star has negligible angular momentum compared to the black hole, and the system will undergo simple precession, as described in Sec. II B. For BBH systems, both components will have comparable angular momenta. In this regime the precession behavior becomes more complex, with $\hat{\mathbf{L}}$ nutating as it moves along its cone and the spins not maintaining a fixed orientation relative to $\hat{\mathbf{L}}$. We have chosen two fiducial NS-BH and BBH binaries with parameters given in Table I to assess the importance of well-chosen parameters for strongly precessing binaries.

For the NS-BH system, the spin of the BH is well constrained. In the radiation frame, shown in Fig. 3, the BH spin is constrained to a single mode with strong nonlinear

correlation between θ_{NS_1} and ϕ_1 . In the system frame, this constraint is mainly in θ_{LS_1} and ϕ_{JL} , but with little correlation between these parameters. Sampling correlations like that found in the radiation frame proves very inefficient, since the jump proposal used most often proposes jumps in one dimension at a time.

Both components of the BBH system in this study have significant angular momentum, making degeneracies stronger between them. For this particular system, the constraints placed on the spin parameters are much weaker than for the NS-BH system's primary spin. However, the system frame still isolates physical features in the waveform and reduces the correlation between spinning parameters. Figure 4 shows the 95% credible regions in the primary spin's orientation parameters. In the radiation frame, the posterior is highly structured, with regions of

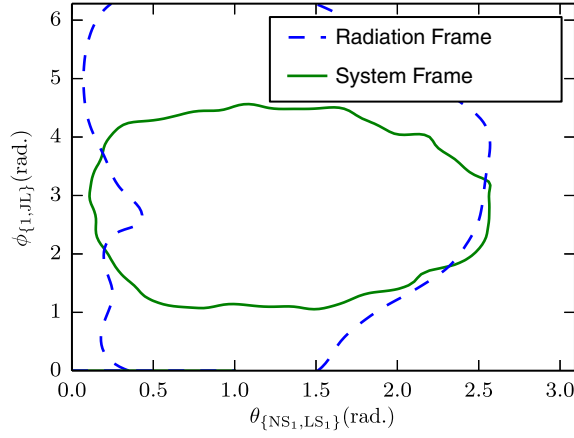


FIG. 4 (color online). The 95% posterior credible regions from analyses of a BBH injection. θ and ϕ correspond to θ_{NS_1} and ϕ_1 in the radiation frame, and θ_{LS_1} and ϕ_{JL} in the system frame. The posterior in the radiation frame parametrization has significant structure. In the system frame parameters, support is largely confined to a single uncorrelated mode that crosses the cyclic boundary of ϕ_{JL} , proving much more efficient to sample.

high correlation between parameters that make MCMC sampling inefficient. In the system frame, however, posterior support is confined to a single uncorrelated mode that stretches across the cyclic boundary of ϕ_{JL} , ideal for sampling efficiently with the proposal set employed.

To assess the improved sampling efficiency using the system frame, MCMC chains were used to analyze the fiducial NS-BH and BBH binaries described in Table I. In Fig. 2 we show the cumulative distributions of one-dimension ACTs for the intrinsic parameters in the parametrization relevant to the frame.

2. Reference frequency

It is clear the system frame parametrization, with evolving parameters defined at 100 Hz, is beneficial for MCMC sampling. To justify this choice in reference frequency, we have conducted a suite of analyses where the system frame parameters are specified at differing reference frequencies. Figure 5 shows the efficiency of MCMC analyses for 11 different choices of reference frequency. We find that a reference frequency near 100 Hz achieves the most efficient MCMC sampling.

By taking the posterior samples from an analysis with $f_{\text{ref}} = 40$ Hz, we can use the PN equations outlined in Sec. II to evolve the posterior samples to later times in the inspiral. We have done this for the NS-BH system to determine how the primary component's spin orientation is constrained over the course of the inspiral. Figure 5 shows the constraint on θ_{LS_1} is roughly uniform across the waveform, as we would expect under the simple precession evolution. The constraints on ϕ_{JL} , however, vary significantly over the inspiral with a minimum around ~ 70 Hz. The correlation between θ_{LS_1} and ϕ_{JL} has a less consistent evolution, but for this system has a minimum at ~ 140 Hz. Since the jump sizes for the Gaussian proposal are adapted to the width of the posterior in each dimension separately, the correlation is expected to impact the sampling efficiency more than the one-dimensional standard deviations. This is reflected by the ACTs in the left panel of Fig. 5, where the ACTs can be seen to oscillate roughly in accordance with the correlation in the right panel. This picture is likely to change for different systems, which is reflected by the BBH ACTs. Ultimately we find that choosing a reference frequency of ~ 100 Hz should reliably result in efficient sampling.

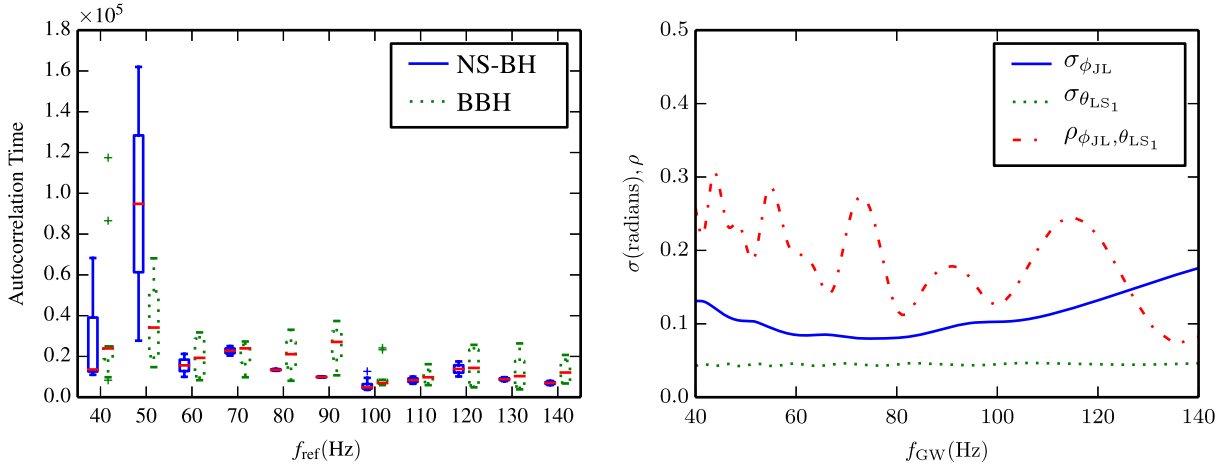


FIG. 5 (color online). *Left panel:* Box-and-whisker plots summarizing the distributions of maximum intrinsic-parameter ACTs for the fiducial NS-BH and BBH systems for the 11 tested reference frequencies. Boxes indicate the first and third quartiles, with interior lines indicating medians, whiskers showing the range of data within 1.5 inner-quartile widths, and +’s indicating points outside this range. *Right panel:* The standard deviations and correlation of the θ_{LS_1} and ϕ_{JL} distributions of the NS-BH fiducial system over the evolution of the waveform. θ_{LS_1} is found to have a relatively stationary distribution, while ϕ_{JL} is best constrained at ~ 70 Hz. The correlation between parameters, though highly variable, has a minimum near 140 Hz.

V. CONCLUSIONS

In this work we present an improved method for parametrizing precessing compact binary coalescence waveforms, which we have implemented in the LAL software library. We have compared the efficiency of MCMC sampling with both parametrizations for BNS, NS-BH, and BBH systems, and find improvements of factors of 7.6, 11, and 3.7, respectively. We have also determined the sampling efficiency and parameter constraints for system frame parameters for a range of reference frequencies. We find that using a reference frequency of ~ 100 Hz will ensure efficient sampling for both NS-BH and BBH signals.

While we assessed the improvement using a specific parameter estimation algorithm (LALINFERENCE_MCMC) and waveform model (a time domain SpinTaylorT4 implementation), this method can be expected to improve the performance of any underlying sampling method (e.g.,

nested sampling) and any preprocessing waveform implementation (e.g., frequency-domain templates [49] and preprocessing effective-one-body [50]). By moving to a more astrophysically intuitive parametrization, physical features in the waveform, such as a lack of precession, are described using fewer parameters. This reduces the correlation between model parameters and increases the efficiency of parameter estimation analyses. By additionally specifying the orientation of the binary and its components at a point near the detectors' peak sensitivity, the efficiency or parameter estimation can be further increased.

ACKNOWLEDGMENTS

B.F. was supported by NSF Fellowship No. DGE-0824162, with computing resources funded by Grant No. PHY-1126812. E.O. and R.O.S. would like to acknowledge support from NSF Grants No. PHY-0970074 and No. PHY-1307429.

-
- [1] G.M. Harry, *Classical Quantum Gravity* **27**, 084006 (2010).
 - [2] F. Acernese *et al.*, Virgo Technical Report No. 0027A-09, 2009.
 - [3] J. Abadie *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), *Classical Quantum Gravity* **27**, 173001 (2010).
 - [4] J.E. McClintock, R. Narayan, S.W. Davis, L. Gou, A. Kulkarni, J. A. Orosz, R. F. Penna, R. A. Remillard, and J. F. Steiner, *Classical Quantum Gravity* **28**, 114009 (2011).
 - [5] L. E. Kidder, *Phys. Rev. D* **52**, 821 (1995).
 - [6] K. G. Arun, A. Buonanno, G. Faye, and E. Ochsner, *Phys. Rev. D* **79**, 104023 (2009).
 - [7] T. A. Apostolatos, C. Cutler, G. J. Sussman, and K. S. Thorne, *Phys. Rev. D* **49**, 6274 (1994).
 - [8] L. E. Kidder, C. M. Will, and A. G. Wiseman, *Phys. Rev. D* **47**, R4183 (1993).
 - [9] G. Faye, L. Blanchet, and A. Buonanno, *Phys. Rev. D* **74**, 104033 (2006).
 - [10] L. Blanchet, A. Buonanno, and G. Faye, *Phys. Rev. D* **74**, 104034 (2006).
 - [11] S. Marsat, A. Bohe, G. Faye, and L. Blanchet, *Classical Quantum Gravity* **30**, 055007 (2013).
 - [12] A. Bohe, S. Marsat, and L. Blanchet, *Classical Quantum Gravity* **30**, 135009 (2013).
 - [13] A. Bohe, S. Marsat, G. Faye, and L. Blanchet, *Classical Quantum Gravity* **30**, 075017 (2013).
 - [14] R. Sturani, S. Fischetti, L. Cadonati, G. M. Guidi, J. Healy, D. Shoemaker, and A. Vicere', *J. Phys. Conf.* **243**, 012007 (2010).
 - [15] P. Schmidt, M. Hannam, and S. Husa, *Phys. Rev. D* **86**, 104063 (2012).
 - [16] M. Hannam, P. Schmidt, A. Bohé, L. Haegel, S. Husa, F. Ohme, G. Pratten, and M. Pürrer, [arXiv:1308.3271](https://arxiv.org/abs/1308.3271).
 - [17] A. Taracchini *et al.*, *Phys. Rev. D* **89**, 061502 (2014).
 - [18] A. Lundgren and R. O'Shaughnessy, *Phys. Rev. D* **89**, 044021 (2014).
 - [19] P. Schmidt, M. Hannam, S. Husa, and P. Ajith, *Phys. Rev. D* **84**, 024046 (2011).
 - [20] M. Boyle, R. Owen, and H. P. Pfeiffer, *Phys. Rev. D* **84**, 124011 (2011).
 - [21] R. O'Shaughnessy, B. Vaishnav, J. Healy, Z. Meeks, and D. Shoemaker, *Phys. Rev. D* **84**, 124002 (2011).
 - [22] E. Ochsner and R. O'Shaughnessy, *Phys. Rev. D* **86**, 104037 (2012).
 - [23] M. Boyle, *Phys. Rev. D* **87**, 104006 (2013).
 - [24] I. W. Harry, A. H. Nitz, D. A. Brown, A. P. Lundgren, E. Ochsner, and D. Keppel, *Phys. Rev. D* **89**, 024010 (2014).
 - [25] LIGO Scientific Collaboration, <https://www.lsc-group.phys.uwm.edu/daswg/projects/lal.html>.
 - [26] J. Aasi *et al.* (LIGO Collaboration, Virgo Collaboration), *Phys. Rev. D* **88**, 062001 (2013).
 - [27] M. V. van der Sluys, C. Röver, A. Stroeer, V. Raymond, I. Mandel, N. Christensen, V. Kalogera, R. Meyer, and A. Vecchio, *Astrophys. J. Lett.* **688**, L61 (2008).
 - [28] V. Raymond, M. V. van der Sluys, I. Mandel, V. Kalogera, C. Röver, and N. Christensen, *Classical Quantum Gravity* **26**, 114007 (2009).
 - [29] R. O'Shaughnessy, B. Farr, E. Ochsner, H.-S. Cho, C. Kim, and C.-H. Lee, *Phys. Rev. D* **89**, 064048 (2014).
 - [30] R. O'Shaughnessy, B. Farr, E. Ochsner, H.-S. Cho, V. Raymond, C. Kim, and C.-H. Lee, *Phys. Rev. D* **89**, 102005 (2014).

- [31] S. Vitale, R. Lynch, J. Veitch, V. Raymond, and R. Sturani, [arXiv:1403.0129 \[Phys. Rev. Lett. \(to be published\)\]](#).
- [32] D. A. Brown, A. Lundgren, and R. O'Shaughnessy, [Phys. Rev. D **86**, 064020 \(2012\)](#).
- [33] H.-S. Cho, E. Ochsner, R. O'Shaughnessy, C. Kim, and C.-H. Lee, [Phys. Rev. D **87**, 024004 \(2013\)](#).
- [34] A. Buonanno, B. R. Iyer, E. Ochsner, Y. Pan, and B. S. Sathyaprakash, [Phys. Rev. D **80**, 084043 \(2009\)](#).
- [35] Y. Pan, A. Buonanno, Y.-b. Chen, and M. Vallisneri, [Phys. Rev. D **69**, 104017 \(2004\)](#).
- [36] A. Buonanno, Y.-b. Chen, Y. Pan, and M. Vallisneri, [Phys. Rev. D **70**, 104003 \(2004\)](#).
- [37] A. H. Nitz, A. Lundgren, D. Brown, E. Ochsner, D. Keppel, and I. Harry, [Phys. Rev. D **88**, 124039 \(2013\)](#).
- [38] J. D. Schnittman, [Phys. Rev. D **70**, 124020 \(2004\)](#).
- [39] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, [J. Chem. Phys. **21**, 1087 \(1953\)](#).
- [40] W. K. Hastings, [Biometrika **57**, 97 \(1970\)](#).
- [41] A. Gelman, G. O. Roberts, and W. R. Gilks, in *Bayesian Statistics, 5 (Alicante, 1994)*, Oxford Science Publications (Oxford University Press, New York, 1996), pp. 599–607.
- [42] J. Aasi *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), [arXiv:1304.0670](#).
- [43] R. N. Manchester, G. B. Hobbs, A. Teoh, and M. Hobbs, [VizieR Online Data Catalog **7245**, 0 \(2005\)](#).
- [44] K. Cannon *et al.*, [Astrophys. J. **748**, 136 \(2012\)](#).
- [45] L. Singer *et al.*, [arXiv:1404.5623](#).
- [46] E. Baird, S. Fairhurst, M. Hannam, and P. Murphy, [Phys. Rev. D **87**, 024035 \(2013\)](#).
- [47] M. Hannam, D. A. Brown, S. Fairhurst, C. L. Fryer, and I. W. Harry, [Astrophys. J. Lett. **766**, L14 \(2013\)](#).
- [48] J. Aasi *et al.*, [Classical Quantum Gravity **31**, 115004 \(2014\)](#).
- [49] K. Chatziioannou, N. Cornish, A. Klein, and N. Yunes, [Phys. Rev. D **89**, 104023 \(2014\)](#).
- [50] Y. Pan, A. Buonanno, A. Taracchini, L. E. Kidder, A. H. Mroué, H. P. Pfeiffer, M. A. Scheel, and B. Szilágyi, [Phys. Rev. D **89**, 084006 \(2014\)](#).